Interpreting Models for Categorical and Count Outcomes

Rose Medeiros
StataCorp LLC

Stata Webinar
March 21, 2019

Contents

1 Introduction **1**
1.1 Goals . . . 1

2 Estimation **1**
2.1 Factor Variables . . . 1

3 Postestimation **7**
3.1 Tests of Coefficients . . . 7
3.2 Predictions . . . 10
3.3 Marginal Effects . . . 25
3.4 Other Models . . . 29

4 Conclusion **36**
4.1 Conclusion . . . 36

1 Introduction

1.1 Goals

Goals

- Learn how to fit models that include categorical variables and/or interactions using factor variable syntax
- Get an overview of tools available for investigating models
- Learn a bit about how Stata partitions model fitting and model testing tasks

---

2 Estimation

2.1 Factor Variables

A Logistic Regression Model

- We'll use data from the National Health and Nutrition Examination Survey (NHANES) for our examples

```
. webuse nhanes2
```

- We'll start with a model for high blood pressure (`highbp`) using age, body mass index (`bmi`) and sex (`female`)
- Before we fit the model, let's investigate the variables

```
. codebook highbp age bmi female

-----------------------------------------------------------------------------------------------
highbp                                         1 if bpsystol >= 140|bpdiast >= 90, 0 otherwise
-----------------------------------------------------------------------------------------------

                  type:  numeric (byte)

                 range:  [0,1]                        units:  1
         unique values:  2                        missing .:  0/10,351

            tabulation:  Freq.  Value
                         5,975  0
                         4,376  1

-----------------------------------------------------------------------------------------------
age                                                                               age in years
-----------------------------------------------------------------------------------------------

                  type:  numeric (byte)

                 range:  [20,74]                      units:  1
         unique values:  55                       missing .:  0/10,351

                  mean:   47.5797
              std. dev:   17.2148

           percentiles:        10%       25%       50%       75%       90%
                                24        31        49        63        69

-----------------------------------------------------------------------------------------------
bmi                                                                      Body Mass Index (BMI)
-----------------------------------------------------------------------------------------------

                  type:  numeric (float)

                 range:  [12.385596,61.129696]        units:  1.000e-07
         unique values:  9,941                    missing .:  0/10,351

                  mean:   25.5376
              std. dev:   4.91497

           percentiles:        10%       25%       50%       75%       90%
                           20.1037    22.142   24.8181   28.0267   31.7259

-----------------------------------------------------------------------------------------------
female                                                                        1=female, 0=male
-----------------------------------------------------------------------------------------------

                  type:  numeric (byte)

                 range:  [0,1]                        units:  1
         unique values:  2                        missing .:  0/10,351

            tabulation:  Freq.  Value
                         4,915  0
                         5,436  1
```

- Now we can fit the model

```
. logit highbp age bmi female
```

```
Iteration 0:   log likelihood = -7050.7655
Iteration 1:   log likelihood = -5859.5273
Iteration 2:   log likelihood = -5845.5355
Iteration 3:   log likelihood = -5845.4948
Iteration 4:   log likelihood = -5845.4948

Logistic regression                             Number of obs     =     10,351
                                                LR chi2(3)        =    2410.54
                                                Prob > chi2       =     0.0000
Log likelihood = -5845.4948                     Pseudo R2         =     0.1709

------------------------------------------------------------------------------
      highbp |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         age |   .0459846   .0013974    32.91   0.000     .0432457    .0487235
         bmi |   .1371553   .0050802    27.00   0.000     .1271983    .1471123
      female |  -.4824464   .0451382   -10.69   0.000    -.5709156   -.3939771
       _cons |   -5.84101   .1523939   -38.33   0.000    -6.139697   -5.542324
------------------------------------------------------------------------------
```

Working with Categorical Variables

- Now we would like to include `region` in the model, let's take a look at this variable

```
. codebook region

-----------------------------------------------------------------------------------------------
region                                                                     1=NE, 2=MW, 3=S, 4=W
-----------------------------------------------------------------------------------------------

                  type:  numeric (byte)
                 label:  region

                 range:  [1,4]                        units:  1
         unique values:  4                        missing .:  0/10,351

            tabulation:  Freq.   Numeric  Label
                         2,096         1  NE
                         2,774         2  MW
                         2,853         3  S
                         2,628         4  W
```

- `region` cannot simply be added to the list of covariates because it has 4 categories
- To include a categorical variable, put an `i.` in front of its name—this declares the variable to be a categorical variable, or in Stataese, a *factor variable*
- For example

```
. logit highbp age bmi i.female i.region

Iteration 0:   log likelihood = -7050.7655
Iteration 1:   log likelihood =  -5857.277
Iteration 2:   log likelihood = -5843.2102
Iteration 3:   log likelihood =  -5843.169
Iteration 4:   log likelihood =  -5843.169

Logistic regression                             Number of obs     =     10,351
                                                LR chi2(6)        =    2415.19
                                                Prob > chi2       =     0.0000
```

```
Log likelihood =  -5843.169                     Pseudo R2         =     0.1713

------------------------------------------------------------------------------
      highbp |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         age |   .0459318   .0013982    32.85   0.000     .0431914    .0486722
         bmi |   .1372797   .0050825    27.01   0.000     .1273182    .1472411
             |
      female |
          0  |          0  (base)
          1  |  -.4811765   .0451517   -10.66   0.000    -.5696723   -.3926807
             |
      region |
         NE  |          0  (base)
         MW  |  -.1324591   .0662441    -2.00   0.046    -.2622952    -.002623
          S  |  -.0887067   .0653787    -1.36   0.175    -.2168466    .0394331
          W  |  -.0403994   .0667441    -0.61   0.545    -.1712154    .0904166
             |
       _cons |  -5.772271   .1584937   -36.42   0.000    -6.082913   -5.461629
------------------------------------------------------------------------------
```

Niceities

- Starting in Stata 13, value labels associated with factor variables are displayed in the regression table
- We can tell Stata to show the base categories for our factor variables

  ```
  . set showbaselevels on
  ```

  ◇ This means the base category will always be clearly documented in the output

---

Factor Notation as Operators

- The `i.` operator can be applied to many variables at once:

```
. logit highbp age bmi i.(female region)

Iteration 0:   log likelihood = -7050.7655
Iteration 1:   log likelihood =  -5857.277
Iteration 2:   log likelihood = -5843.2102
Iteration 3:   log likelihood =  -5843.169
Iteration 4:   log likelihood =  -5843.169

Logistic regression                             Number of obs     =     10,351
                                                LR chi2(6)        =    2415.19
                                                Prob > chi2       =     0.0000
Log likelihood =  -5843.169                     Pseudo R2         =     0.1713

------------------------------------------------------------------------------
      highbp |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         age |   .0459318   .0013982    32.85   0.000     .0431914    .0486722
         bmi |   .1372797   .0050825    27.01   0.000     .1273182    .1472411
             |
      female |
          0  |          0  (base)
          1  |  -.4811765   .0451517   -10.66   0.000    -.5696723   -.3926807
             |
```

```
      region |
          NE  |          0  (base)
          MW  |  -.1324591   .0662441    -2.00   0.046    -.2622952    -.002623
           S  |  -.0887067   .0653787    -1.36   0.175    -.2168466    .0394331
           W  |  -.0403994   .0667441    -0.61   0.545    -.1712154    .0904166
              |
       _cons  |  -5.772271   .1584937   -36.42   0.000    -6.082913   -5.461629
------------------------------------------------------------------------------
```

- In other words, it understands the distributive property
  - This is useful when using variable ranges, for example
- For the curious, factor variable notation works with wildcards
  - If there were many variables starting with `u`, then `i.u*` would include them all as factor variables

---

Using Different Base Categories

- By default, the smallest-valued category is the base category
- This can be overridden within commands
  - `b#.` specifies the value # as the base
  - `b(##).` specifies the #'th largest value as the base
  - `b(first).` specifies the smallest value as the base
  - `b(last).` specifies the largest value as the base
  - `b(freq).` specifies the most prevalent value as the base
  - `bn.` specifies there should be no base
- The base can also be permanently changed using `fvset`; see `help fvset` for more information

---

Playing with the Base

- We can use `region=3` as the base class on the fly:

  ```
  . logit highbp age bmi i.female b3.region
  ```
- We can use the most prevalent category as the base

  ```
  . logit highbp age bmi i.female b(freq).region
  ```
- Factor variables can be distributed across many variables

  ```
  . logit highbp age bmi b(freq).(female region)
  ```
- The base category can be omitted (with some care here)

  ```
  . logit highbp age bmi i.female bn.region, noconstant
  ```
- We can also include a term for `region=4` only

  ```
  . logit highbp age bmi i.female 4.region
  ```

Specifying Interactions

- Factor variables are also used for specifying interactions
  - This is where they really shine
- To include both main effects and interaction terms in a model, put ## between the variables
- To include only the interaction terms, put # between the terms
- Variables involved in interactions are treated as categorical by default
  - Prefix a variable with `c.` to specify that a variable is continuous
- Here is our model with an interaction between age and `female`

```
. logit highbp bmi c.age##female i.region

Iteration 0:   log likelihood = -7050.7655
Iteration 1:   log likelihood = -5824.3249
Iteration 2:   log likelihood = -5795.4621
Iteration 3:   log likelihood = -5795.4025
Iteration 4:   log likelihood = -5795.4025

Logistic regression                             Number of obs     =     10,351
                                                LR chi2(7)        =    2510.73
                                                Prob > chi2       =     0.0000
Log likelihood = -5795.4025                     Pseudo R2         =     0.1780

------------------------------------------------------------------------------
      highbp |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         bmi |   .1378163    .005139    26.82   0.000     .1277441    .1478886
         age |   .0334439   .0018514    18.06   0.000     .0298151    .0370727
             |
      female |
          0  |          0  (base)
          1  |  -1.883645   .1530275   -12.31   0.000    -2.183574   -1.583717
             |
female#c.age |
          1  |   .0276653   .0028606     9.67   0.000     .0220585     .033272
             |
      region |
         NE  |          0  (base)
         MW  |  -.1359488   .0664206    -2.05   0.041    -.2661308   -.0057668
          S  |  -.0902012   .0655982    -1.38   0.169    -.2187713    .0383689
          W  |  -.0379412    .066944    -0.57   0.571     -.169149    .0932666
             |
       _cons |  -5.176679   .1687139   -30.68   0.000    -5.507352   -4.846006
------------------------------------------------------------------------------
```

Some Factor Variable Notes

- If you plan to look at marginal effects of any kind, it is best to
  - Explicitly mark all categorical variables with `i.`
  - Specify all interactions using # or ##
  - Specify powers of a variable as interactions of the variable with itself

- There can be up to 8 categorical and 8 continuous interactions in one expression
  - Have fun with the interpretation

---

3 Postestimation

3.1 Tests of Coefficients

Introduction to Postestimation

- In Stata jargon, postestimation commands are commands that can be run after a model is fit, for example
  - Predictions
  - Additional hypothesis tests
  - Checks of assumptions
- We'll explore postestimation tools that can be used to help interpret model results
  - The main example here is after `logit` models, but these tools can be used with most estimation commands
- The usefulness of specific tools will depend on the types of hypotheses you wish to examine

---

Finding the Coefficient Names

- Some postestimation commands require that you know the names used to store the coefficients
- To see these names we can replay the model showing the *coefficient legend*

```
. logit, coeflegend

Logistic regression                             Number of obs     =     10,351
                                                LR chi2(7)        =    2510.73
                                                Prob > chi2       =     0.0000
Log likelihood = -5795.4025                     Pseudo R2         =     0.1780

------------------------------------------------------------------------------
      highbp |      Coef.  Legend
-------------+----------------------------------------------------------------
         bmi |   .1378163  _b[bmi]
         age |   .0334439  _b[age]
             |
      female |
          0  |          0  _b[0b.female]
          1  |  -1.883645  _b[1.female]
             |
female#c.age |
          1  |   .0276653  _b[1.female#c.age]
             |
      region |
         NE  |          0  _b[1b.region]
         MW  |  -.1359488  _b[2.region]
          S  |  -.0902012  _b[3.region]
          W  |  -.0379412  _b[4.region]
             |
       _cons |  -5.176679  _b[_cons]
------------------------------------------------------------------------------
```

- From here, we can see the full specification of the factor levels:
  - _b[2.region] corresponds to region=2 which is “MW” or midwest
  - _b[3.region] corresponds to region=3 which is “S” or south
- The coefficient for the female by age interaction is stored as _b[1.female#c.age]

Joint Tests

- The test command performs a Wald test of the specified null hypothesis
  - The default test is that the listed terms are equal to 0
- test takes a list of terms, which may be variable names, but can also be terms associated with factor variables
- To specify a joint test of the null hypothesis that the coefficients for the levels of region are all equal to 0

```
. test 2.region 3.region 4.region

 ( 1)  [highbp]2.region = 0
 ( 2)  [highbp]3.region = 0
 ( 3)  [highbp]4.region = 0

           chi2(  3) =    4.96
         Prob > chi2 =    0.1744
```

Testing Sets of Coefficients

- If you are testing a large number of terms, typing them all out can be laborious
- testparm also performs Wald tests, but it accepts lists of variables, rather than coefficients in the model
- For example, to test all coefficients associated with i.region

```
. testparm i.region

 ( 1)  [highbp]2.region = 0
 ( 2)  [highbp]3.region = 0
 ( 3)  [highbp]4.region = 0

           chi2(  3) =    4.96
         Prob > chi2 =    0.1744
```

Likelihood Ratio Tests

- Likelihood ratio tests provide an alternative method of testing sets of coefficients
- To test the coefficients associated with region we need to store our model results. The name is arbitrary, we’ll call them m1

```
. estimates store m1
```

- Now we can rerun our model without region

```
. logit highbp bmi c.age##female if e(sample)
```

```
Iteration 0:   log likelihood = -7050.7655
Iteration 1:   log likelihood =  -5826.855
Iteration 2:   log likelihood = -5797.9206
Iteration 3:   log likelihood = -5797.8856
Iteration 4:   log likelihood = -5797.8856

Logistic regression                             Number of obs     =     10,351
                                                LR chi2(4)        =    2505.76
                                                Prob > chi2       =     0.0000
Log likelihood = -5797.8856                     Pseudo R2         =     0.1777

------------------------------------------------------------------------------
      highbp |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         bmi |   .1376855   .0051366    26.80   0.000      .127618     .147753
         age |   .0335286   .0018501    18.12   0.000     .0299025    .0371548
             |
      female |
          0  |          0  (base)
          1  |  -1.882479   .1530115   -12.30   0.000    -2.182376   -1.582582
             |
female#c.age |
          1  |    .027615   .0028601     9.66   0.000     .0220092    .0332207
             |
       _cons |  -5.247536   .1628488   -32.22   0.000    -5.566713   -4.928358
------------------------------------------------------------------------------
```

- Adding `if e(sample)` makes sure the same sample, what Stata calls the *estimation sample*, is used for both models

---

Likelihood Ratio Tests (Continued)

- Now we store the second set of estimates

```
. estimates store m2
```

- And use the `lrtest` command to perform the likelihood ratio test

```
. lrtest m1 m2

Likelihood-ratio test                                 LR chi2(3)  =      4.97
(Assumption: m2 nested in m1)                         Prob > chi2 =    0.1743
```

- We'll restore the results from `m1` which includes `region` even though the terms are not collectively significant

```
. estimates restore m1

(results m1 are active now)
```

- Now it's as though we just ran the model stored as `m1`

Tests of Differences

- `test` can also be used to the equality of coefficients

```
. test 3.region = 4.region

 ( 1)  [highbp]3.region - [highbp]4.region = 0

           chi2(  1) =    0.71
         Prob > chi2 =    0.3978
```

- A likelihood ratio test can also be used; see `help constraint` for information on setting the necessary constraints
- The `lincom` command calculates linear combinations of coefficients, along with standard errors, hypothesis tests, and confidence intervals
- For example, to obtain the difference in coefficients

```
. lincom 3.region - 4.region

 ( 1)  [highbp]3.region - [highbp]4.region = 0

------------------------------------------------------------------------------
      highbp |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         (1) |    -.05226   .0618078    -0.85   0.398    -.1734012    .0688811
------------------------------------------------------------------------------
```

3.2 Predictions

What are margins?

- Stata defines margins as "statistics calculated from predictions of a previously fit model at fixed values of some covariates and averaging or otherwise integrating over the remaining covariates."
  - Also known as counterfactuals, or when we fix a categorical variable, potential outcomes
- What sorts of predictions does `margins` work with?
  - Predicted means, probabilities, and counts
  - Derivatives
  - Elasticities
- We'll also see contrasts and pairwise comparisons of the above

---

Average Predictions

- Let's start with `margins` in its most basic form

```
. margins

Predictive margins                              Number of obs     =     10,351
Model VCE    : OIM

Expression   : Pr(highbp), predict()

------------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       _cons |   .4227611   .0042898    98.55   0.000     .4143533    .4311689
------------------------------------------------------------------------------
```

- What happened here?
  1. The predicted probability of `highbp=1` was calculated for each case, using each case's observed values of `bmi`, `age`, `female`, and `region`
  2. The average of those predictions was calculated and displayed
- Unless we tell it to do otherwise, `margins` works with the estimation sample

---

Predictions at the Average

- An alternative is to calculate the predicted probability fixing all the covariates at some value, often the mean

```
. margins, atmeans

Adjusted predictions                              Number of obs     =     10,351
Model VCE    : OIM

Expression   : Pr(highbp), predict()
at           : bmi             =     25.5376 (mean)
               age             =    47.57965 (mean)
               0.female        =    .4748333 (mean)
               1.female        =    .5251667 (mean)
               1.region        =    .2024925 (mean)
               2.region        =    .2679934 (mean)
               3.region        =    .2756255 (mean)
               4.region        =    .2538885 (mean)

------------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       _cons |   .3929783   .0056167    69.97   0.000     .3819697    .4039869
------------------------------------------------------------------------------
```

- What happened here?

  1. The mean of each independent variable was calculated
  2. The predicted probability of `highbp=1` was calculated using the means from step 1

Predictions at Each Level of a Factor Variable

- Adding a factor variable specifies that the predictions be repeated at each level of the variable, for example

```
. margins region

Predictive margins                                Number of obs     =     10,351
Model VCE    : OIM

Expression   : Pr(highbp), predict()

------------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
      region |
         NE  |   .4362592   .0095422    45.72   0.000     .4175568    .4549616
         MW  |   .4103455   .0083278    49.27   0.000     .3940234    .4266677
          S  |   .4190352   .0081188    51.61   0.000     .4031226    .4349477
          W  |   .4290013   .0085434    50.21   0.000     .4122565    .4457461
------------------------------------------------------------------------------
```

- What happened here?

  1. The predicted probability is calculated treating all cases as if `region=1` and using each case's observed values of `bmi`, `age`, and `female`
  2. The mean of the predictions from step 1 is calculated
  3. Repeat steps 1 and 2 for each value of `region`

Multiple Factor Variables

- We can obtain margins for multiple variables

```
. margins region female

Predictive margins                              Number of obs     =     10,351
Model VCE    : OIM

Expression   : Pr(highbp), predict()

------------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
      region |
         NE  |   .4362592   .0095422    45.72   0.000     .4175568    .4549616
         MW  |   .4103455   .0083278    49.27   0.000     .3940234    .4266677
          S  |   .4190352   .0081188    51.61   0.000     .4031226    .4349477
          W  |   .4290013   .0085434    50.21   0.000     .4122565    .4457461
             |
      female |
          0  |   .4692315    .006393    73.40   0.000     .4567014    .4817616
          1  |   .3766361   .0057397    65.62   0.000     .3653866    .3878857
------------------------------------------------------------------------------
```

- Or combinations of values, for example each combination of region and female

```
. margins region#female

Predictive margins                              Number of obs     =     10,351
Model VCE    : OIM

Expression   : Pr(highbp), predict()

-------------------------------------------------------------------------------
              |            Delta-method
              |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
--------------+----------------------------------------------------------------
region#female |
        NE#0  |   .4839466   .0112276    43.10   0.000      .461941    .5059522
        NE#1  |   .3889131   .0097392    39.93   0.000     .3698246    .4080015
        MW#0  |   .4556986   .0100844    45.19   0.000     .4359337    .4754636
        MW#1  |   .3652888   .0086372    42.29   0.000     .3483602    .3822173
         S#0  |   .4651826   .0099214    46.89   0.000     .4457369    .4846282
         S#1  |   .3731942   .0084524    44.15   0.000     .3566278    .3897605
         W#0  |   .4760455   .0102535    46.43   0.000      .455949     .496142
         W#1  |   .3822812   .0088891    43.01   0.000     .3648589    .3997034
-------------------------------------------------------------------------------
```

- We can graph the resulting predictions using the marginsplot command

---

Graphing Predicted Probabilities

- For example to graph the last set of margins

```
. marginsplot
```

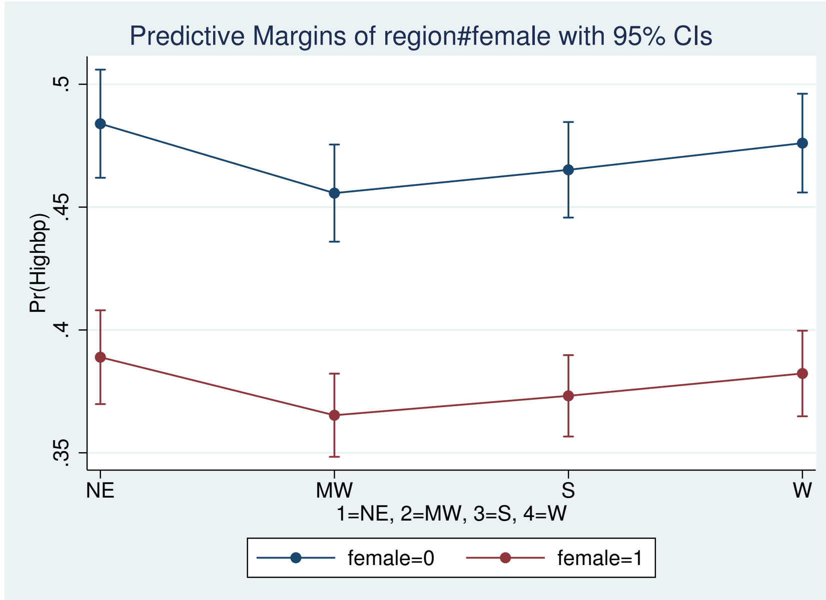

Predictions at Specified Values of Covariates

- The `at()` option is used to specify values at which margins should be calculated
- To obtain the average predicted probability setting age=40 specify

```
. margins, at(age=40)

Predictive margins                              Number of obs     =     10,351
Model VCE    : OIM

Expression   : Pr(highbp), predict()
at           : age             =          40

------------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       _cons |   .3287856   .0053346    61.63   0.000       .31833    .3392413
------------------------------------------------------------------------------
```

- `at()` accepts number lists, so we can obtain predictions setting age to 20, 30, ..., 70

```
. margins, at(age=(20(10)70)) vsquish

Predictive margins                              Number of obs     =     10,351
Model VCE    : OIM

Expression   : Pr(highbp), predict()
1._at        : age             =          20
2._at        : age             =          30
3._at        : age             =          40
4._at        : age             =          50
5._at        : age             =          60
6._at        : age             =          70

------------------------------------------------------------------------------
```

```
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         _at |
          1  |   .1826464    .006436    28.38   0.000      .170032    .1952608
          2  |    .247219   .0060245    41.04   0.000     .2354113    .2590268
          3  |   .3287856   .0053346    61.63   0.000       .31833    .3392413
          4  |    .425936   .0050064    85.08   0.000     .4161236    .4357485
          5  |   .5326646   .0059775    89.11   0.000     .5209488    .5443804
          6  |   .6387994   .0076524    83.48   0.000     .6238009    .6537979
------------------------------------------------------------------------------
```

- The `vsquish` option reduces the amount of vertical space the header for `margins` takes up

---

Graphing Across Values of Continuous Variables

```
. marginsplot
```

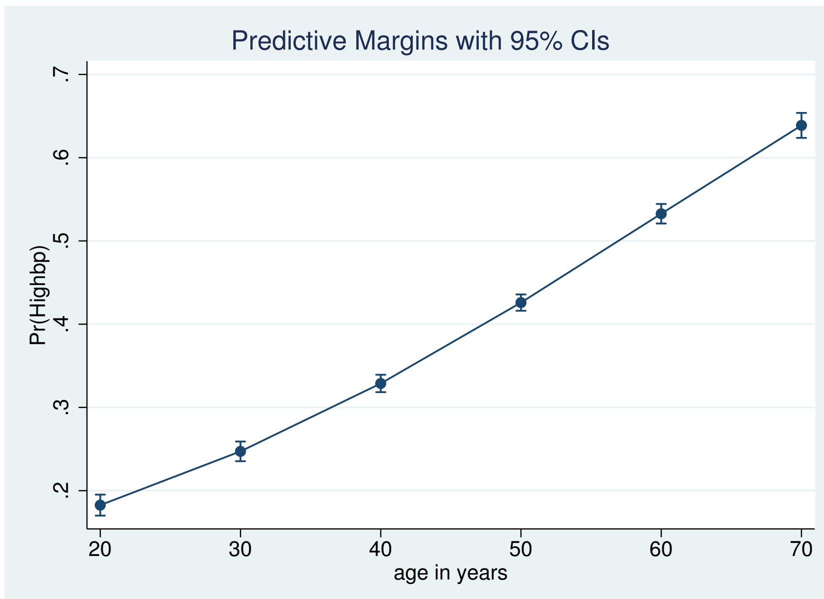

Specifying Values of Multiple Variables

- We can specify values of multiple variables using `at()`
- If we set values of all the independent variables in our model, we can ask very specific questions
- For example, what is the predicted probability of high blood pressure for an male who is age 40, with a bmi of 25 and living in the midwest (`region=2`)? What is the predicted probability if the person is female?

```
. margins female, at(age=40 bmi=25 region=2)

Adjusted predictions                            Number of obs     =     10,351
Model VCE    : OIM

Expression   : Pr(highbp), predict()
at           : bmi             =          25
               age             =          40
```

```
               region          =           2

------------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
      female |
          0  |   .3706418   .0118974    31.15   0.000     .3473232    .3939603
          1  |   .2130731   .0096757    22.02   0.000      .194109    .2320372
------------------------------------------------------------------------------
```

- We can use the contrast operator `r.` to compare the predicted probabilities for males and females

```
. margins r.female, at(age=40 bmi=25 region=2)

Contrasts of adjusted predictions
Model VCE    : OIM

Expression   : Pr(highbp), predict()
at           : bmi             =          25
               age             =          40
               region          =           2

------------------------------------------------
             |         df        chi2     P>chi2
-------------+----------------------------------
      female |          1      200.44     0.0000
------------------------------------------------

--------------------------------------------------------------
             |            Delta-method
             |   Contrast   Std. Err.     [95% Conf. Interval]
-------------+------------------------------------------------
      female |
   (1 vs 0)  |  -.1575687   .0111296     -.1793822   -.1357551
--------------------------------------------------------------
```

- We'll see more on contrasts below

Specifying Ranges of Multiple Variables

- We can also specify ranges of values for multiple variables, for example multiple values of age and `bmi`

```
. margins, at(age=(20(10)70) bmi=(20(10)40))
```

- We can also combine the use of factor and continuous variables, for example

```
. margins female, at(age=(20(10)70)) vsquish

Predictive margins                              Number of obs     =     10,351
Model VCE    : OIM

Expression   : Pr(highbp), predict()
1._at        : age             =          20
2._at        : age             =          30
3._at        : age             =          40
4._at        : age             =          50
5._at        : age             =          60
6._at        : age             =          70
```

```
-----------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+---------------------------------------------------------------
  _at#female |
        1 0  |   .2728133   .0110381    24.72   0.000     .2511789    .2944477
        1 1  |   .0997683   .0065837    15.15   0.000     .0868644    .1126722
        2 0  |   .3369363   .0093499    36.04   0.000     .3186108    .3552617
        2 1  |   .1629921   .0074737    21.81   0.000      .148344    .1776402
        3 0  |   .4076871   .0075866    53.74   0.000     .3928176    .4225566
        3 1  |   .2537634   .0074437    34.09   0.000     .2391741    .2683527
        4 0  |   .4826887   .0070403    68.56   0.000       .46889    .4964874
        4 1  |   .3718821   .0071293    52.16   0.000      .357909    .3858552
        5 0  |   .5588757   .0084852    65.87   0.000     .5422451    .5755063
        5 1  |   .5079403   .0083938    60.51   0.000     .4914886    .5243919
        6 0  |   .6329264   .0108508    58.33   0.000     .6116592    .6541935
        6 1  |   .6442392   .0106744    60.35   0.000     .6233177    .6651607
-----------------------------------------------------------------------------
```

More Plots

```
. marginsplot, legend(order(3 "Males" 4 "Females"))
```

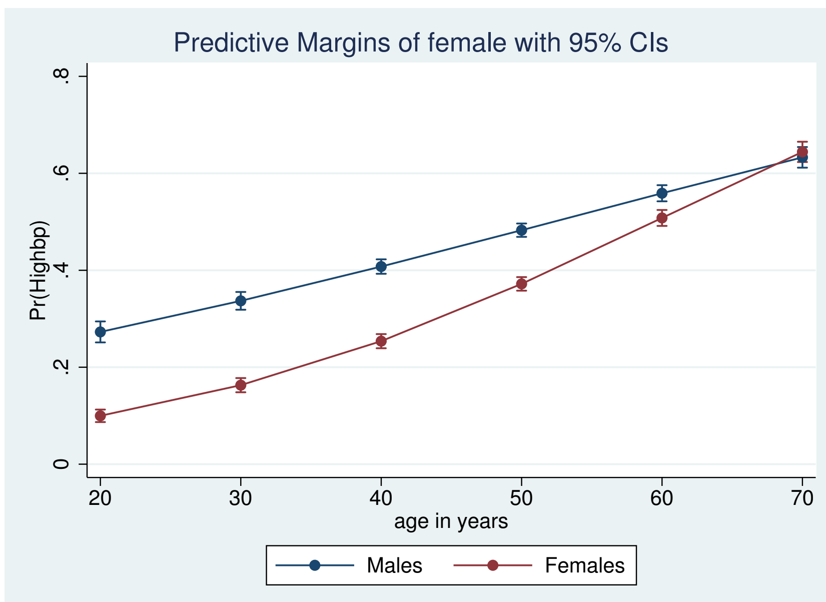

◇ The standard errors are drawn before the lines for the predictions, so we want the legend to show the third and fourth plots

---

More Predictions

- We can use `at()` with the `generate()` suboption to answer different sorts of questions
- For example, what would the averaged predicted probability be if everyone aged 5 years, while their values `female` and `region` remained the same?
- The `generate(age+5)` requests margins calculated at each observations value of age plus 5

```
. margins, at(age=generate(age+5))

Predictive margins                              Number of obs     =     10,351
Model VCE    : OIM

Expression   : Pr(highbp), predict()
at           : age             = age+5

------------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       _cons |   .4672688    .004476   104.39   0.000      .458496    .4760416
------------------------------------------------------------------------------
```

- We can specify `at()` multiple times, to obtain predictions under different scenarios

```
. margins, at(age=generate(age)) ///
     at(age=generate(age+5)) at(age=generate(age+10))

Predictive margins                              Number of obs     =     10,351
Model VCE    : OIM

Expression   : Pr(highbp), predict()

1._at        : age             = age

2._at        : age             = age+5

3._at        : age             = age+10

------------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         _at |
          1  |   .4227611   .0042898    98.55   0.000     .4143533    .4311689
          2  |   .4672688    .004476   104.39   0.000      .458496    .4760416
          3  |    .512185   .0048335   105.97   0.000     .5027115    .5216585
------------------------------------------------------------------------------
```

Predictions Over Groups

- The `over()` option produces predictions averaging within groups defined by the factor variable, for example, `female`

```
. margins, over(female)

Predictive margins                              Number of obs     =     10,351
Model VCE    : OIM

Expression   : Pr(highbp), predict()
over         : female

------------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
      female |
          0  |   .4687691   .0066113    70.90   0.000     .4558112    .4817269
          1  |   .3811626    .005567    68.47   0.000     .3702516    .3920737
------------------------------------------------------------------------------
```

- What happened here?
    1. The predicted probability for each case is calculated, using the case's observed values on all variables
    2. The average predicted probability is calculated using only cases where `female=0`
    3. Repeat step 2 using only cases where `female=1`

Pairwise Comparisons of Predictions

- Earlier we obtained average predicted probabilities at each level of `region` using

```
. margins region
```

- For pairwise comparisons of these margins we can add the `pwcompare` option

```
. margins region, pwcompare

Pairwise comparisons of predictive margins
Model VCE    : OIM

Expression   : Pr(highbp), predict()

--------------------------------------------------------------
             |            Delta-method         Unadjusted
             |   Contrast   Std. Err.     [95% Conf. Interval]
-------------+------------------------------------------------
      region |
    MW vs NE |  -.0259137   .0126665     -.0507396   -.0010878
     S vs NE |   -.017224   .0125288     -.0417801     .007332
     W vs NE |  -.0072579   .0128075     -.0323601    .0178443
     S vs MW |   .0086896   .0116321     -.0141089    .0314882
     W vs MW |   .0186558   .0119339     -.0047343    .0420459
      W vs S |   .0099661   .0117862     -.0131345    .0330667
--------------------------------------------------------------
```

- Adding the groups option will allow us to see which levels are statistically distinguishable

```
. margins region, pwcompare(groups)

Pairwise comparisons of predictive margins
Model VCE    : OIM

Expression   : Pr(highbp), predict()

-------------------------------------------------
             |            Delta-method Unadjusted
             |     Margin   Std. Err.      Groups
-------------+-----------------------------------
      region |
         NE  |   .4362592   .0095422           B
         MW  |   .4103455   .0083278          A
          S  |   .4190352   .0081188          AB
          W  |   .4290013   .0085434          AB
-------------------------------------------------
Note: Margins sharing a letter in the group label
      are not significantly different at the 5%
      level.
```

- The `pwcompare()` option can be used to specify other suboptions; see `help margins pwcompare` for more information

---

Contrasts of Predictions

- The `margins` command allows *contrast operators* which are used to request comparisons of the margins
  - In this case the margins are predicted probabilities
- For example, to compare average predicted probabilities setting `female=0` versus `female=1` add the `r.` prefix

```
. margins r.female

Contrasts of predictive margins
Model VCE    : OIM

Expression   : Pr(highbp), predict()

------------------------------------------------
             |         df        chi2     P>chi2
-------------+----------------------------------
      female |          1      116.16     0.0000
------------------------------------------------

--------------------------------------------------------------
             |            Delta-method
             |   Contrast   Std. Err.     [95% Conf. Interval]
-------------+------------------------------------------------
      female |
   (1 vs 0)  |  -.0925953   .0085912     -.1094338   -.0757569
--------------------------------------------------------------
```

- We can use the `@` operator to contrast `female` at each level of `region`

```
. margins r.female@region

Contrasts of predictive margins
Model VCE    : OIM

Expression   : Pr(highbp), predict()

-------------------------------------------------
              |         df        chi2     P>chi2
--------------+----------------------------------
female@region |
 (1 vs 0) NE  |          1      117.89     0.0000
 (1 vs 0) MW  |          1      109.28     0.0000
  (1 vs 0) S  |          1      112.04     0.0000
  (1 vs 0) W  |          1      115.96     0.0000
       Joint  |          4      119.65     0.0000
-------------------------------------------------

---------------------------------------------------------------
              |            Delta-method
              |   Contrast   Std. Err.     [95% Conf. Interval]
--------------+------------------------------------------------
female@region |
 (1 vs 0) NE  |  -.0950335   .0087525     -.1121881   -.0778789
 (1 vs 0) MW  |  -.0904099   .0086485     -.1073606   -.0734592
  (1 vs 0) S  |  -.0919884   .0086906     -.1090216   -.0749552
  (1 vs 0) W  |  -.0937643   .0087074     -.1108305   -.0766982
---------------------------------------------------------------
```

- This reports the differences in predicted probabilities when `female=1` versus `female=0` at each level of `region`

Contrasts of Predictions (Continued)

- To perform contrasts at different values of a continuous variable use the `at()` option

```
. margins r.female, at(age=(20(10)70)) vsquish

Contrasts of predictive margins
Model VCE    : OIM

Expression   : Pr(highbp), predict()
1._at        : age             =          20
2._at        : age             =          30
3._at        : age             =          40
4._at        : age             =          50
5._at        : age             =          60
6._at        : age             =          70

------------------------------------------------
             |         df        chi2     P>chi2
-------------+----------------------------------
  female@_at |
 (1 vs 0) 1  |          1      182.15     0.0000
 (1 vs 0) 2  |          1      211.82     0.0000
 (1 vs 0) 3  |          1      209.80     0.0000
 (1 vs 0) 4  |          1      122.51     0.0000
 (1 vs 0) 5  |          1       18.36     0.0000
 (1 vs 0) 6  |          1        0.56     0.4552
      Joint  |          6   123716.83     0.0000
------------------------------------------------

--------------------------------------------------------------
             |            Delta-method
             |   Contrast   Std. Err.     [95% Conf. Interval]
-------------+------------------------------------------------
  female@_at |
 (1 vs 0) 1  |   -.173045   .0128218     -.1981752   -.1479147
 (1 vs 0) 2  |  -.1739442   .0119516     -.1973689   -.1505195
 (1 vs 0) 3  |  -.1539237   .0106268     -.1747518   -.1330956
 (1 vs 0) 4  |  -.1108066   .0100111      -.130428   -.0911851
 (1 vs 0) 5  |  -.0509354   .0118889     -.0742372   -.0276335
 (1 vs 0) 6  |   .0113128   .0151483     -.0183773     .041003
--------------------------------------------------------------
```

- The output gives tests of the differences in predicted probabilities for `female=1` versus `female=0` at each of the specified values of age
  - The joint test is statistically significant
  - The differences get smaller in absolute value as age increases

---

Plotting Contrasts

```
. marginsplot, yline(0)
```

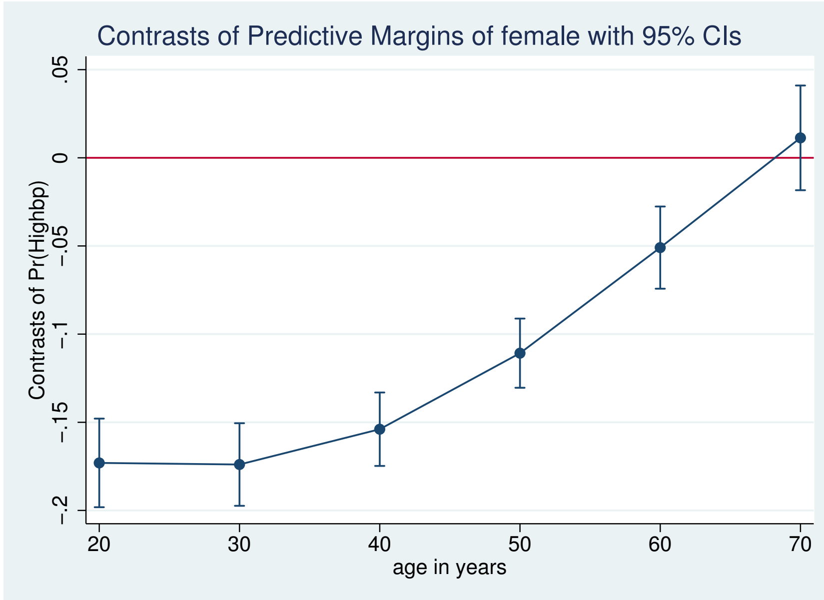

Contrast Operators

- A few common contrast operators are
  - ◇ `r.` differences from the base (a.k.a. reference) level
  - ◇ `a.` differences from the next (adjacent) level
  - ◇ `ar.` differences from the previous level (reverse adjacent)
  - ◇ `g.` differences from the balanced grand mean
  - ◇ `gw.` differences from the observeration-weighted grand mean
  - ◇ There are also operators for Helmert contrats and contrasts using orthogonal polynomials for balanced and unbalanced cases

`contrast` suboptions

- So far we've obtained contrasts using *contrast operators*, but `margins` also allows a `contrast()` option
- The `contrast()` option is particularly useful for specifying options to `contrast`
- For example, to obtain contrasts for continuous variables the `atcontrast()` suboption is used
  - ◇ The `effects` suboption requests a table showing the contrasts along with confidence intervals and p-values
  - ◇ In `atcontrast(a)` the a contrast operator requests comparisons of adjacent categories

```
. margins, at(age=(20(10)70)) contrast(atcontrast(a) effects) vsquish

Contrasts of predictive margins
Model VCE    : OIM

Expression   : Pr(highbp), predict()
1._at        : age             =          20
2._at        : age             =          30
3._at        : age             =          40
4._at        : age             =          50
```

```
5._at        : age            =          60
6._at        : age            =          70

------------------------------------------------
             |         df        chi2     P>chi2
-------------+----------------------------------
         _at |
   (1 vs 2)  |          1     1947.70     0.0000
   (2 vs 3)  |          1     1565.44     0.0000
   (3 vs 4)  |          1     1191.96     0.0000
   (4 vs 5)  |          1     1064.35     0.0000
   (5 vs 6)  |          1     1301.80     0.0000
      Joint  |          5   278027.30     0.0000
------------------------------------------------

------------------------------------------------------------------------------
             |            Delta-method
             |   Contrast   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         _at |
   (1 vs 2)  |  -.0645726   .0014631   -44.13   0.000    -.0674403   -.0617049
   (2 vs 3)  |  -.0815666   .0020616   -39.57   0.000    -.0856072    -.077526
   (3 vs 4)  |  -.0971504   .0028139   -34.52   0.000    -.1026656   -.0916352
   (4 vs 5)  |  -.1067286   .0032714   -32.62   0.000    -.1131405   -.1003167
   (5 vs 6)  |  -.1061348   .0029416   -36.08   0.000    -.1119002   -.1003693
------------------------------------------------------------------------------
```

Contrasts with `generate()`

- Earlier we used the `generate()` suboption to obtain predicted probabilities modifying the observed values
- Specifically, we obtained predicted probabilities using each case's observed value of age and each case's observed value $+5$ years

```
. margins, at(age=generate(age)) at(age=generate(age+5))

Predictive margins                              Number of obs     =     10,351
Model VCE    : OIM

Expression   : Pr(highbp), predict()

1._at        : age            = age

2._at        : age            = age+5

------------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         _at |
          1  |   .4227611   .0042898    98.55   0.000     .4143533    .4311689
          2  |   .4672688    .004476   104.39   0.000      .458496    .4760416
------------------------------------------------------------------------------
```

- Using the `contrast` option, we can compare the two

```
. margins, at(age=generate(age)) ///
     at(age=generate(age+5)) contrast(atcontrast(r))
```

```
Contrasts of predictive margins
Model VCE    : OIM

Expression   : Pr(highbp), predict()

1._at        : age             = age

2._at        : age             = age+5

------------------------------------------------
             |         df        chi2     P>chi2
-------------+----------------------------------
         _at |          1     1728.47     0.0000
------------------------------------------------

--------------------------------------------------------------
             |            Delta-method
             |   Contrast   Std. Err.     [95% Conf. Interval]
-------------+------------------------------------------------
         _at |
   (2 vs 1)  |   .0445077   .0010705      .0424095    .0466059
--------------------------------------------------------------
```

Contrasts of Differences

- We can also request contrasts of contrasts by combining contrast operators
- For example, to compare the differences between males and females across levels of `region` use

```
. margins r.female#r.region

Contrasts of predictive margins
Model VCE    : OIM

Expression   : Pr(highbp), predict()

--------------------------------------------------------
                     |         df        chi2     P>chi2
---------------------+----------------------------------
       female#region |
(1 vs 0) (MW vs NE)  |          1        4.11     0.0426
 (1 vs 0) (S vs NE)  |          1        1.88     0.1703
 (1 vs 0) (W vs NE)  |          1        0.32     0.5709
              Joint  |          3        4.83     0.1851
--------------------------------------------------------

----------------------------------------------------------------------
                     |            Delta-method
                     |   Contrast   Std. Err.     [95% Conf. Interval]
---------------------+------------------------------------------------
       female#region |
(1 vs 0) (MW vs NE)  |   .0046236   .0022806      .0001537    .0090935
 (1 vs 0) (S vs NE)  |   .0030451   .0022208     -.0013077    .0073979
 (1 vs 0) (W vs NE)  |   .0012692   .0022396     -.0031203    .0056586
----------------------------------------------------------------------
```

Adjusting for Multiple Comparisons

- Use of `contrast` and `pwcompare` can result in a large number of hypothesis tests
- The `mcompare()` option can be used to adjust p-values and confidence intervals for multiple comparisons within factor variable terms
- The available methods are
  - ◇ `noadjust`
  - ◇ `bonferroni`
  - ◇ `sidak`
  - ◇ `scheffe`

Using `mcompare()`

- To apply Bonferroni's adjustment to an earlier contrast

```
. margins r.female@region, mcompare(bonferroni)

Contrasts of predictive margins
Model VCE    : OIM

Expression   : Pr(highbp), predict()

------------------------------------------------------------
              |                                   Bonferroni
              |         df        chi2     P>chi2     P>chi2
--------------+---------------------------------------------
female@region |
 (1 vs 0) NE  |          1      117.89     0.0000     0.0000
 (1 vs 0) MW  |          1      109.28     0.0000     0.0000
  (1 vs 0) S  |          1      112.04     0.0000     0.0000
  (1 vs 0) W  |          1      115.96     0.0000     0.0000
       Joint  |          4      119.65     0.0000
------------------------------------------------------------
Note: Bonferroni-adjusted p-values are reported for tests on
      individual contrasts only.

----------------------------
              |    Number of
              |  Comparisons
--------------+-------------
female@region |            4
----------------------------

---------------------------------------------------------------
              |            Delta-method         Bonferroni
              |   Contrast   Std. Err.     [95% Conf. Interval]
--------------+------------------------------------------------
female@region |
 (1 vs 0) NE  |  -.0950335   .0087525     -.1168946   -.0731723
 (1 vs 0) MW  |  -.0904099   .0086485     -.1120112   -.0688085
  (1 vs 0) S  |  -.0919884   .0086906     -.1136949   -.0702819
  (1 vs 0) W  |  -.0937643   .0087074     -.1155128   -.0720159
---------------------------------------------------------------
```

- Specifying adjusted p-values with the `pwcompare` option

```
. margins region, mcompare(sidak) pwcompare

Pairwise comparisons of predictive margins
Model VCE    : OIM

Expression   : Pr(highbp), predict()

---------------------------
             |    Number of
             |  Comparisons
-------------+-------------
      region |            6
---------------------------

--------------------------------------------------------------
             |            Delta-method           Sidak
             |   Contrast   Std. Err.     [95% Conf. Interval]
-------------+------------------------------------------------
      region |
   MW vs NE  |  -.0259137   .0126665     -.0592398    .0074124
    S vs NE  |   -.017224   .0125288     -.0501878    .0157398
    W vs NE  |  -.0072579   .0128075     -.0409548     .026439
    S vs MW  |   .0086896   .0116321      -.021915    .0392943
    W vs MW  |   .0186558   .0119339     -.0127429    .0500544
     W vs S  |   .0099661   .0117862     -.0210439    .0409762
--------------------------------------------------------------
```

3.3 Marginal Effects

Marginal Effects

- In a straightforward linear model, the marginal effect of a variable is the coefficient b

$$y = b_0 + b_1 x_1 + b_2 x_2 + e$$

- In more complex models, this is no longer true
  - ◇ models with interactions
  - ◇ models with polynomial terms
  - ◇ generalized linear models when the margin is not on the linear scale
- For example, in a logistic regression model, the marginal effect of covariates is not constant on the probability scale
- `margins` can be used to estimate the margins of the derivative of a response

---

A Closer Look at Slopes

- Here is a graph of predicted probabilities across values of `bmi`

```
. margins, at(bmi=(12(5)62))

. marginsplot
```

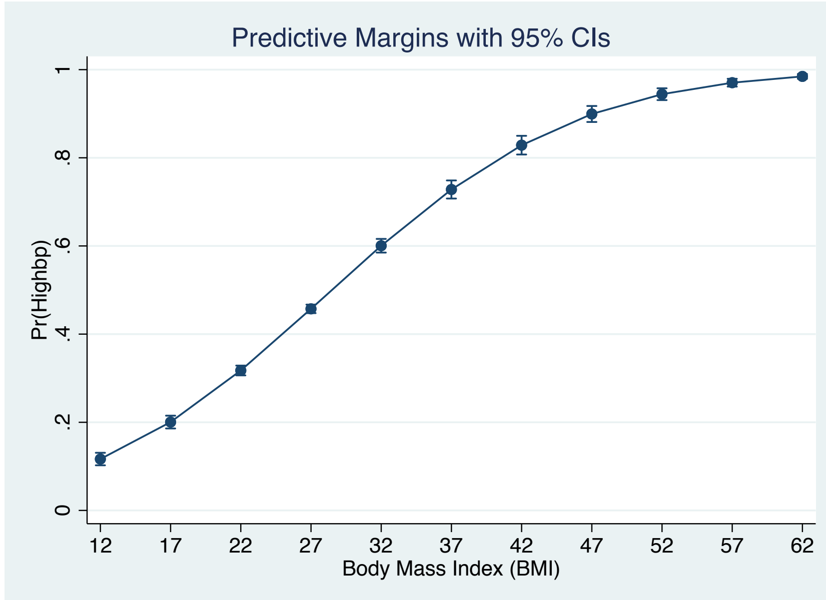

Average Marginal Effects

- The slope of bmi is not constant, but we might want to know what it is on average
- We can obtain the average marginal effect of bmi

```
. margins, dydx(bmi)

Average marginal effects                        Number of obs     =     10,351
Model VCE    : OIM

Expression   : Pr(highbp), predict()
dy/dx w.r.t. : bmi

------------------------------------------------------------------------------
             |            Delta-method
             |      dy/dx   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         bmi |   .0262514    .000852    30.81   0.000     .0245816    .0279212
------------------------------------------------------------------------------
```

- What happened here?
  1. Calculate the derivative of the predicted probability with respect to bmi for each observaton
  2. Calculate the average of derivatives from step 1
- We can do the same for all variables in our model

```
. margins, dydx(*)

Average marginal effects                        Number of obs     =     10,351
Model VCE    : OIM

Expression   : Pr(highbp), predict()
dy/dx w.r.t. : bmi age 1.female 2.region 3.region 4.region

------------------------------------------------------------------------------
```

```
             |            Delta-method
             |      dy/dx   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         bmi |   .0262514    .000852    30.81   0.000     .0245816    .0279212
         age |   .0088181    .0002145   41.11   0.000     .0083976    .0092385
             |
      female |
          0  |          0  (base)
          1  |  -.0925953    .0085912  -10.78   0.000    -.1094338   -.0757569
             |
      region |
         NE  |          0  (base)
         MW  |  -.0259137    .0126665   -2.05   0.041    -.0507396   -.0010878
          S  |   -.017224    .0125288   -1.37   0.169    -.0417801     .007332
          W  |  -.0072579    .0128075   -0.57   0.571    -.0323601    .0178443
------------------------------------------------------------------------------
Note: dy/dx for factor levels is the discrete change from the base level.
```

Marginal Effects Over the Response Surface

- It can also be informative to estimate the marginal effect of x at different values of x
- For example, we can obtain the derviative with respect to age at age=20, 30, ..., 70

```
. margins, dydx(age) at(age=(20(10)70)) vsquish

Average marginal effects                        Number of obs     =     10,351
Model VCE    : OIM

Expression   : Pr(highbp), predict()
dy/dx w.r.t. : age
1._at        : age             =          20
2._at        : age             =          30
3._at        : age             =          40
4._at        : age             =          50
5._at        : age             =          60
6._at        : age             =          70

------------------------------------------------------------------------------
             |            Delta-method
             |      dy/dx   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
age          |
         _at |
          1  |   .0056454    .0001263    44.70   0.000     .0053978    .0058929
          2  |   .0072988    .0001734    42.09   0.000     .0069589    .0076387
          3  |   .0089942     .000245    36.71   0.000      .008514    .0094744
          4  |   .0103355    .0003148    32.83   0.000     .0097184    .0109526
          5  |   .0108342    .0003262    33.21   0.000     .0101949    .0114736
          6  |   .0102041    .0002508    40.69   0.000     .0097125    .0106957
------------------------------------------------------------------------------
```

- Here we do something similar, setting `female=0` and then `female=1`

```
. margins female, dydx(age) at(age=(20(10)70)) vsquish

Average marginal effects                        Number of obs     =     10,351
Model VCE    : OIM

Expression   : Pr(highbp), predict()
```

```
dy/dx w.r.t. : age
1._at        : age             =          20
2._at        : age             =          30
3._at        : age             =          40
4._at        : age             =          50
5._at        : age             =          60
6._at        : age             =          70

------------------------------------------------------------------------------
             |            Delta-method
             |      dy/dx   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
age          |
  _at#female |
        1 0  |   .0060242   .0002192    27.48   0.000     .0055945    .0064538
        1 1  |   .0050964   .0001457    34.98   0.000     .0048108     .005382
        2 0  |   .0067761   .0003143    21.56   0.000     .0061601     .007392
        2 1  |   .0076423   .0001587    48.17   0.000     .0073313    .0079532
        3 0  |   .0073341   .0003896    18.82   0.000     .0065704    .0080978
        3 1  |   .0105163   .0002922    35.99   0.000     .0099436     .011089
        4 0  |   .0076144   .0004244    17.94   0.000     .0067825    .0084463
        4 1  |   .0129499   .0004576    28.30   0.000     .0120531    .0138467
        5 0  |   .0075668    .000407    18.59   0.000      .006769    .0083645
        5 1  |   .0139526   .0005002    27.89   0.000     .0129722    .0149331
        6 0  |   .0071918     .00034    21.15   0.000     .0065255    .0078581
        6 1  |   .0129829   .0003617    35.90   0.000      .012274    .0136917
------------------------------------------------------------------------------
```

Plots of Marginal Effects

- We can, of course, plot these marginal effects, to see how they change with different values of `female` and age

```
. marginsplot
```

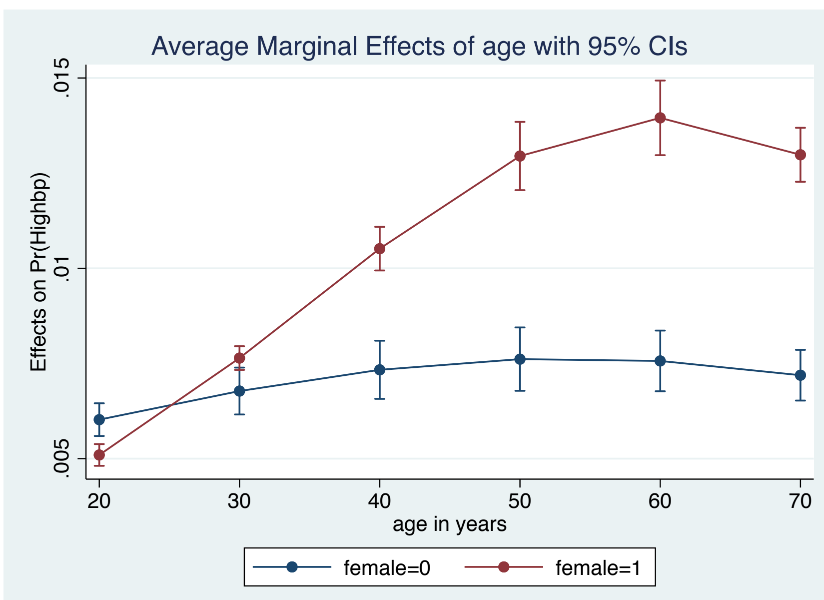

3.4 Other Models

margins with Other Estimation Commands

- `margins` works after most estimation commands
- The default prediction for `margins` is the same as the default prediction for `predict` after a given command
- See `help` *command* `postestimation` for information on postestimation commands and their defaults after a given command
- You can specify different predictions from `margins` using the `predict()` option

---

Modeling Household Size

- For the next set of examples we will model the number of individuals in a household (`houssiz`) using a Poisson model
- Our model will include covariates age, age^2, `region`, `rural`, and a `region` by `rural` interaction
- We've been working with age and `region` but we'll take a look at the new variables

```
. codebook houssiz rural

-----------------------------------------------------------------------------------------------
houssiz                                                             # persons in household, 1-14
-----------------------------------------------------------------------------------------------

                  type:  numeric (byte)

                 range:  [1,14]                       units:  1
         unique values:  14                       missing .:  0/10,351

                  mean:   2.94377
              std. dev:   1.69516

           percentiles:        10%       25%       50%       75%       90%
                                 1         2         2         4         5

-----------------------------------------------------------------------------------------------
rural                                                                          1=rural, 0=urban
-----------------------------------------------------------------------------------------------

                  type:  numeric (byte)

                 range:  [0,1]                        units:  1
         unique values:  2                        missing .:  0/10,351

            tabulation:  Freq.  Value
                         6,548  0
                         3,803  1
```

- Now we can fit our model

```
. poisson houssiz i.region##i.rural age c.age#c.age

Iteration 0:   log likelihood = -18385.275
Iteration 1:   log likelihood = -18385.272
Iteration 2:   log likelihood = -18385.272

Poisson regression                              Number of obs     =     10,351
```

```
                                                  LR chi2(9)        =    1780.26
                                                  Prob > chi2       =     0.0000
Log likelihood = -18385.272                       Pseudo R2         =     0.0462

------------------------------------------------------------------------------
     houssiz |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
      region |
         NE  |          0  (base)
         MW  |  -.0586473   .0204129    -2.87   0.004    -.0986558   -.0186387
          S  |   .0021845    .021345     0.10   0.918    -.0396509     .04402
          W  |  -.0305816   .0208232    -1.47   0.142    -.0713943    .0102311
             |
       rural |
          0  |          0  (base)
          1  |   .0441422   .0278741     1.58   0.113    -.0104901    .0987745
             |
region#rural |
       MW#1  |   .0474625    .036487     1.30   0.193    -.0240508    .1189758
        S#1  |  -.0013947   .0352449    -0.04   0.968    -.0704734    .0676839
        W#1  |   .0300379   .0366293     0.82   0.412    -.0417541     .10183
             |
         age |   .0561718   .0025069    22.41   0.000     .0512584    .0610852
             |
 c.age#c.age |  -.0007312   .0000272   -26.87   0.000    -.0007845   -.0006779
             |
       _cons |   .2472973   .0539633     4.58   0.000     .1415311    .3530634
------------------------------------------------------------------------------
```

`margins` after `poisson`

- `predict`'s default after `poisson` is the predicted count
- To obtain the average predicted count, using the observed values of all covarites use

```
. margins

Predictive margins                                Number of obs     =     10,351
Model VCE    : OIM

Expression   : Predicted number of events, predict()

------------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       _cons |   2.943774    .016864   174.56   0.000     2.910721    2.976826
------------------------------------------------------------------------------
```

- As before, we can request predicted counts at specified values of factor variables

```
. margins region#rural

Predictive margins                                Number of obs     =     10,351
Model VCE    : OIM

Expression   : Predicted number of events, predict()

------------------------------------------------------------------------------
             |            Delta-method
```

```
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
region#rural |
        NE#0  |   2.942144   .0441807    66.59   0.000     2.855552    3.028737
        NE#1  |   3.074926   .0722057    42.59   0.000     2.933405    3.216447
        MW#0  |   2.774558   .0383527    72.34   0.000     2.699388    2.849728
        MW#1  |   3.040725   .0579537    52.47   0.000     2.927138    3.154312
         S#0  |   2.948578   .0447353    65.91   0.000     2.860899    3.036258
         S#1  |   3.077355   .0472768    65.09   0.000     2.984695    3.170016
         W#0  |   2.853531   .0411629    69.32   0.000     2.772853    2.934209
         W#1  |   3.073255   .0580446    52.95   0.000     2.959489     3.18702
------------------------------------------------------------------------------
```

- And continuous variables

```
. margins, at(age=(20(10)70)) vsquish

Predictive margins                              Number of obs     =     10,351
Model VCE    : OIM

Expression   : Predicted number of events, predict()
1._at        : age             =          20
2._at        : age             =          30
3._at        : age             =          40
4._at        : age             =          50
5._at        : age             =          60
6._at        : age             =          70

------------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         _at |
          1  |    2.94187   .0438937    67.02   0.000      2.85584      3.0279
          2  |   3.579277   .0276575   129.41   0.000     3.525069    3.633484
          3  |   3.762318   .0326109   115.37   0.000     3.698402    3.826234
          4  |   3.416678   .0306675   111.41   0.000     3.356571    3.476785
          5  |   2.680655   .0216814   123.64   0.000      2.63816     2.72315
          6  |   1.817047   .0254912    71.28   0.000     1.767085    1.867009
------------------------------------------------------------------------------
```

Plotting Predicted Counts

```
. marginsplot
```

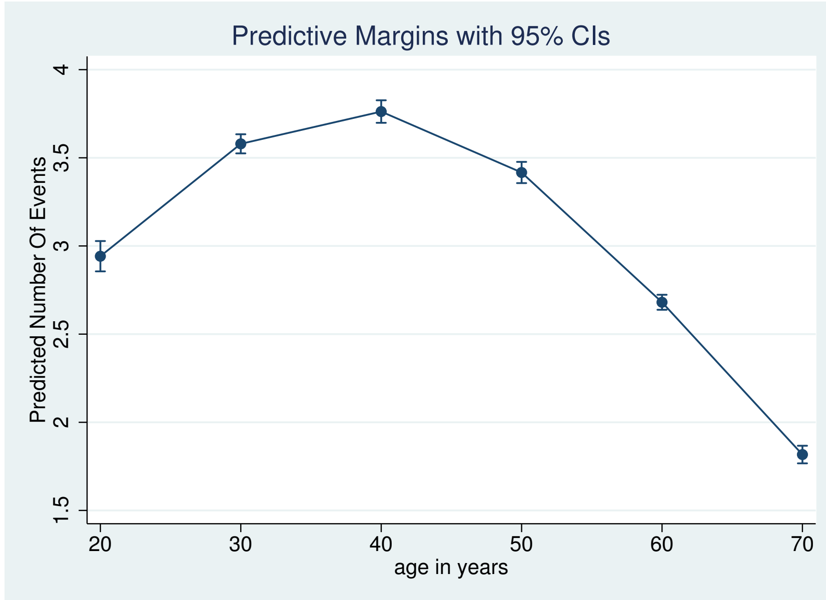

Other Margins

- After `poisson`, `margins` can be used to predict the following
  - ◇ `n` number of events; the default
  - ◇ `ir` incidence rate, `exp(xb)`, `n` when the exposure variable $= 1$
  - ◇ `pr(`*n*`)` probability that `y=n`
  - ◇ `pr(`*a*,*b*`)` probability that `a` $\leq$ `y` $\leq$ `b`
  - ◇ `xb` the linear predcition
- Predicted probability that `houssiz=1`

```
. margins rural, predict(pr(1))

Predictive margins                              Number of obs     =     10,351
Model VCE    : OIM

Expression   : Pr(houssiz=1), predict(pr(1))

------------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       rural |
          0  |   .1714666    .0020282    84.54   0.000     .1674915    .1754417
          1  |   .1541716    .0025566    60.30   0.000     .1491608    .1591823
------------------------------------------------------------------------------
```

- Predicted probability that $3 \leq$ `houssiz` ≤ 5

```
. margins region#rural, predict(pr(3,5))

Predictive margins                              Number of obs     =     10,351
Model VCE    : OIM

Expression   : Pr(3<=houssiz<=5), predict(pr(3,5))
```

```
---------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+-------------------------------------------------------------
region#rural |
        NE#0 |   .4557062   .0047091    96.77   0.000     .4464765     .464936
        NE#1 |   .4682528   .0063677    73.54   0.000     .4557723    .4807332
        MW#0 |   .4365671   .0049383    88.41   0.000     .4268883    .4462459
        MW#1 |   .4652407    .005386    86.38   0.000     .4546843    .4757971
         S#0 |   .4563673   .0047189    96.71   0.000     .4471185    .4656162
         S#1 |    .468461    .004296   109.05   0.000      .460041    .4768809
         W#0 |   .4460472    .004858    91.82   0.000     .4365256    .4555688
         W#1 |   .4681091   .0051371    91.12   0.000     .4580405    .4781777
---------------------------------------------------------------------------
```

Multiple Responses

- Starting in Stata 14, `margins` can compute margins for multiple responses at the same time
  - ◇ After, for example, `ologit`, `mlogit`, `mvreg`
- To demonstrate this, we'll model self-rated health in a different version of the NHANES dataset

```
. webuse nhanes2f
. codebook health

------------------------------------------------------------------------------------------
health                                                                1=poor,..., 5=excellent
------------------------------------------------------------------------------------------

                  type:  numeric (byte)
                 label:  hlthgrp

                 range:  [1,5]                        units:  1
         unique values:  5                        missing .:  2/10,337

            tabulation:  Freq.   Numeric  Label
                           729         1  poor
                         1,670         2  fair
                         2,938         3  average
                         2,591         4  good
                         2,407         5  excellent
                             2         .
```

- Our model is

```
. ologit health i.female age c.age#c.age

Iteration 0:   log likelihood = -15764.397
Iteration 1:   log likelihood =  -15042.53
Iteration 2:   log likelihood = -15036.362
Iteration 3:   log likelihood = -15036.355
Iteration 4:   log likelihood = -15036.355

Ordered logistic regression                     Number of obs     =     10,335
                                                LR chi2(3)        =    1456.09
                                                Prob > chi2       =     0.0000
Log likelihood = -15036.355                     Pseudo R2         =     0.0462
```

```
------------------------------------------------------------------------------
      health |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
      female |
          0  |          0  (base)
          1  |  -.1223788   .0355107    -3.45   0.001    -.1919786    -.052779
             |
         age |  -.0251916   .0076063    -3.31   0.001    -.0400997   -.0102834
             |
 c.age#c.age |    -.00016   .0000812    -1.97   0.049    -.0003191   -9.73e-07
-------------+----------------------------------------------------------------
       /cut1 |  -4.442363   .1659171                     -4.767554   -4.117171
       /cut2 |  -2.975821   .1632372                      -3.29576   -2.655882
       /cut3 |  -1.573015   .1618158                     -1.890168   -1.255862
       /cut4 |  -.3384551   .1606298                     -.6532838   -.0236264
------------------------------------------------------------------------------
```

Specifying the Response

- By default `margins` will produce the average predicted probability of each value of `health`

```
. margins

Predictive margins                              Number of obs     =     10,335
Model VCE    : OIM

1._predict   : Pr(health==1), predict(pr outcome(1))
2._predict   : Pr(health==2), predict(pr outcome(2))
3._predict   : Pr(health==3), predict(pr outcome(3))
4._predict   : Pr(health==4), predict(pr outcome(4))
5._predict   : Pr(health==5), predict(pr outcome(5))

------------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
    _predict |
          1  |   .0709472   .0024959    28.43   0.000     .0660554     .075839
          2  |   .1643302   .0035781    45.93   0.000     .1573172    .1713432
          3  |   .2868785   .0044083    65.08   0.000     .2782384    .2955187
          4  |   .2474815    .004184    59.15   0.000      .239281     .255682
          5  |   .2303626   .0039468    58.37   0.000      .222627    .2380981
------------------------------------------------------------------------------
```

- To request a single outcome we can use `predict(outcome(#))`

```
. margins, predict(outcome(2))

Predictive margins                              Number of obs     =     10,335
Model VCE    : OIM

Expression   : Pr(health==2), predict(outcome(2))

------------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       _cons |   .1643302   .0035781    45.93   0.000     .1573172    .1713432
------------------------------------------------------------------------------
```

- For multiple responses from a single command, repeat the predict() option

```
. margins, predict(outcome(1)) predict(outcome(2))

Predictive margins                              Number of obs     =     10,335
Model VCE    : OIM

1._predict   : Pr(health==1), predict(outcome(1))
2._predict   : Pr(health==2), predict(outcome(2))

------------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
    _predict |
          1  |   .0709472   .0024959    28.43   0.000     .0660554     .075839
          2  |   .1643302   .0035781    45.93   0.000     .1573172    .1713432
------------------------------------------------------------------------------
```

- To obtain predictions across values of age

```
. margins, at(age=(20(10)70)) pr(out(1)) pr(out(2)) vsquish

Predictive margins                              Number of obs     =     10,335
Model VCE    : OIM

1._predict   : Pr(health==1), predict(out(1))
2._predict   : Pr(health==2), predict(out(2))
1._at        : age             =          20
2._at        : age             =          30
3._at        : age             =          40
4._at        : age             =          50
5._at        : age             =          60
6._at        : age             =          70

------------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
_predict#_at |
        1 1  |   .0217005   .0013107    16.56   0.000     .0191315    .0242695
        1 2  |   .0299861    .001366    21.95   0.000     .0273087    .0326635
        1 3  |   .0425874   .0019332    22.03   0.000     .0387984    .0463763
        1 4  |   .0619896   .0026898    23.05   0.000     .0567177    .0672616
        1 5  |   .0920429   .0034083    27.01   0.000     .0853627    .0987231
        1 6  |   .1383404   .0055654    24.86   0.000     .1274324    .1492485
        2 1  |   .0659885   .0032038    20.60   0.000     .0597092    .0722678
        2 2  |   .0881333   .0027672    31.85   0.000     .0827097    .0935568
        2 3  |   .1189848   .0036317    32.76   0.000     .1118668    .1261029
        2 4  |   .1605636   .0045152    35.56   0.000      .151714    .1694132
        2 5  |   .2130434   .0049117    43.37   0.000     .2034167    .2226701
        2 6  |   .2717448   .0066991    40.56   0.000     .2586149    .2848748
------------------------------------------------------------------------------
```

Plots with Multiple Responses

```
. marginsplot, legend(order(3 "Poor" 4 "Fair"))
```

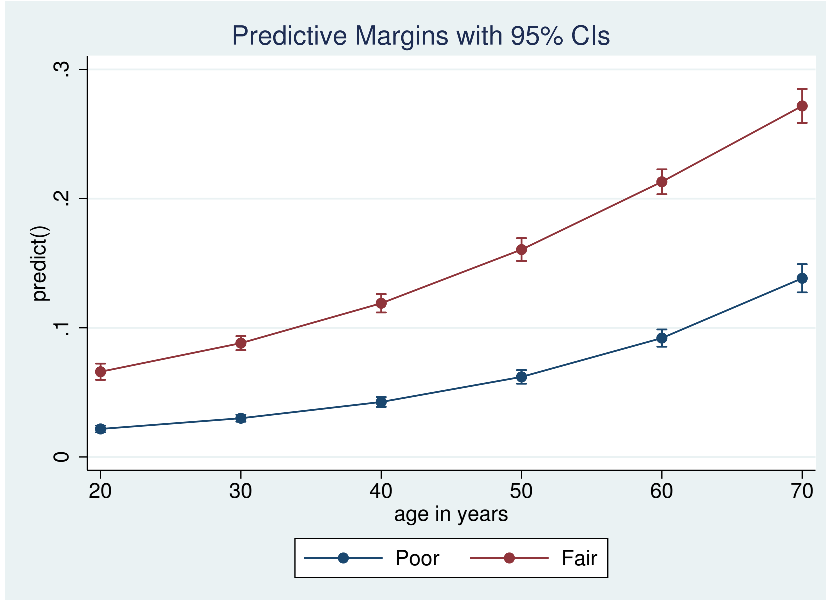

4 Conclusion

4.1 Conclusion

Conclusion

- We've seen how to obtain a variety of predictions and marginal effects after regression models
- We now know how to perform contrasts of predictions and marginal effects
- We've also seen how to graph these results

Index