Predictive Margins and Marginal Effects in Stata

Ben Jann

University of Bern, jann@soz.unibe.ch

11th German Stata Users Group meeting
Potsdam, June 7, 2013

Outline

- Motivation
- `margins` and `marginsplot`
- `regplot`

Motivation

- For a long time, regression tables have been the preferred way of communicating results from statistical models.
- However, interpretation of regression tables can be very challenging in the case of interaction effects, categorical variables, or nonlinear functional forms.
- Moreover, interpretational difficulties can be overwhelming in nonlinear models such as logistic regression. In these models the raw coefficients are often not of much interest; what we want to see for interpretation are effects on outcomes such as probabilities, not on „latent“ variables such as log odds.
- Fortunately, Stata has a number of handy commands such as `margins`, `contrasts`, and `marginsplot` for making sense of regression results.

Example: Factorial Survey on Just Incomes

- Mail survey among a random sample of the Swiss population ($N = 1945$). Written questionnaire in German, French and Italian.
  - Data collected in fall 2010 as part of a follow-up survey to the "Swiss Environmental Survey 2007"
    (see http://www.socio.ethz.ch/research/umweltsurvey/umweltsurvey2007)
- Respondents were asked to judge short text descriptions of (fictional) individuals (so called "vignettes"), in which certain elements are varied at random.
- For our research objective, we used vignettes describing men and women employing the following $2 \times 2 \times 2 \times 3$ design :
  - male vs. female
  - single without children vs. married without children
  - average work effort vs. above-average work effort
  - income levels: 5000 CHF, 5500 CHF, 6000 CHF

Example: The Vignette

In letzter Zeit wird viel über die Höhe von Löhnen in verschiedenen Berufen gesprochen. Wir interessieren uns für Ihre persönliche Einschätzung zu diesem Thema.

Stellen Sie sich die folgende Situation vor:

{Herr | Frau} Müller, 25-jährig, {allein stehend und ohne Kinder | verheiratet in kinderloser Ehe}, arbeitet als kaufmännische{r|} Angestellte{r|} im Rechnungswesen eines mittleren Dienstleistungsbetriebs und erbringt dort {überdurchschnittliche | durchschnittliche} Leistungen. {Sein | Ihr} monatliches Bruttoeinkommen beträgt {5'000 | 5'500 | 6'000} Franken.

Wie bewerten Sie das Einkommen dieser Person? Ist das Einkommen Ihrer Meinung nach gerecht oder ist es ungerechterweise zu hoch oder zu niedrig?

viel zu niedrig					gerecht					viel zu hoch
-5	**-4**	**-3**	**-2**	**-1**	**0**	**+1**	**+2**	**+3**	**+4**	**+5**

Example: The Data

```
. use vignettes
(2010 Vignette Study on Just Incomes)

. d

Contains data from vignettes.dta
  obs:         1,482                          2010 Vignette Study on Just Incomes
 vars:            13                          10 Jun 2013 11:16
 size:        51,870                          (_dta has notes)
-----------------------------------------------------------------------------------------------
              storage  display     value
variable name   type   format      label      variable label
-----------------------------------------------------------------------------------------------
vrating         double %13.0g      vrating    Vignette: rating (-5=much too low, 5=much too high)
vmale           byte   %8.0g       vmale      Vignette: male
vmarried        byte   %8.0g       vmarried   Vignette: married
veffort         byte   %25.0g      veffort    Vignette: above-average work effort
vinc            int    %10.0g                 Vignette: income (CHF per month)
rmale           byte   %8.0g       rmale      Respondent: male
rage            byte   %8.0g                  Respondent: age
reducyrs        double %10.0g                 Respondent: years of education
rright          byte   %8.0g                  Respondent: political orientation (0=left, 10=right)
rmarstat        byte   %8.0g       rmarstat   Respondent: marital status
rinc            byte   %13.0g      rinc       Respondent: income (CHF per month)
wt              double %10.0g                 sampling weights
strata          byte   %8.0g                  sampling strata
-----------------------------------------------------------------------------------------------
Sorted by:
```

Example: Analysis of the Vignette Data

- A simple linear regression model with the vignette responses as dependent variable – have fun interpreting!

```
. regress vrating vinc i.vmale i.vmarried i.veffort ///
>     vmale##rmale##c.reducyrs##c.reducyrs, vsquish noheader
```

vrating	Coef.	Std. Err.	t	P>|t|	[95% Conf.	Interval]
vinc	.0010069	.000094	10.71	0.000	.0008225	.0011913
1.vmale	-1.388362	3.140572	-0.44	0.659	-7.548853	4.772129
1.vmarried	-.1799115	.0771521	-2.33	0.020	-.3312517	-.0285713
1.veffort	-.6662772	.0771985	-8.63	0.000	-.8177084	-.5148459
1.rmale	-.3799532	3.355975	-0.11	0.910	-6.962975	6.203068
vmale#rmale						
1 1	3.984662	4.818208	0.83	0.408	-5.46665	13.43597
reducyrs	.1051598	.3327842	0.32	0.752	-.5476238	.7579435
vmale#c.reducyrs						
1	.0722016	.4839479	0.15	0.881	-.877102	1.021505
rmale#c.reducyrs						
1	.0129762	.5123413	0.03	0.980	-.9920235	1.017976
vmale#rmale#c.reducyrs						
1 1	-.4443146	.7331748	-0.61	0.545	-1.882497	.9938681
c.reducyrs#c.reducyrs	-.0063593	.0123638	-0.51	0.607	-.030612	.0178933
vmale#c.reducyrs#c.reducyrs						
1	.0010164	.0178947	0.06	0.955	-.0340855	.0361183
rmale#c.reducyrs#c.reducyrs						
1	.0013596	.0187994	0.07	0.942	-.0355169	.0382361
vmale#rmale#c.reducyrs#c.reducyrs						
1 1	.0105725	.0268027	0.39	0.693	-.0420033	.0631482
_cons	-4.825754	2.224811	-2.17	0.030	-9.189904	-.4616041

Example: Analysis of the Vignette Data

- Questions we might have about the regression output:
  - What are the overall effects of the vignette factors?
  - How does the effect of vignette factor "sex" depend on education and sex of the respondent?
  - What is the shape of the effect of education depending on sex?
  - Can we express effects in CHF?

Example: Binary Dependent Variable

- A logistic regression of whether income in vignette was judged as "too low" or not:

```
. generate byte toolow = vrating<0 if vrating<.

. logit toolow vinc i.vmale i.vmarried i.veffort

Iteration 0:   log likelihood = -726.94882
Iteration 1:   log likelihood = -660.31413
Iteration 2:   log likelihood = -656.56237
Iteration 3:   log likelihood = -656.55323
Iteration 4:   log likelihood = -656.55323

Logistic regression                               Number of obs   =       1482
                                                  LR chi2(4)      =     140.79
                                                  Prob > chi2     =     0.0000
Log likelihood = -656.55323                       Pseudo R2       =     0.0968

------------------------------------------------------------------------------
      toolow |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
        vinc |  -.0013663   .0001789    -7.64   0.000    -.0017169   -.0010157
     1.vmale |   .3994001   .1393606     2.87   0.004     .1262584    .6725417
  1.vmarried |   .2886296   .1389959     2.08   0.038     .0162027    .5610565
   1.veffort |   1.164184   .1466717     7.94   0.000     .8767125    1.451655
       _cons |   4.939563   .9606374     5.14   0.000     3.056748    6.822378
------------------------------------------------------------------------------
```

Example: Binary Dependent Variable

- Questions we might have about the logit output:
  - What the hell do these coefficients mean?
  - What is the conditional probability of "too low" depending on different levels of the factor variables?
  - What is the marginal effect of the vignette factors on the probability of "too low"?

Stata tools to answer these questions

- Stata commands `margins` and `marginsplot` can help us answer these questions.
- There's another useful command called `contrast`, but I am not going to talk about that.
- However, I will also show `marginscontplot` by Patrick Royston that will appear in one of the next issues of the Stata Journal.

What can margins do?

- `margins` computes so-called margins of responses.
  - A "margin" is a statistic computed from predictions from a model while manipulating the values of the covariates.
    - "conditional margin": a prediction from a model where all covariates are set to fixed values
    - "predictive margin": if some covariates are not fixed
  - Computed are *levels* of margins for different covariate values or *differences* in levels of margins if covariate values are changed (or even differences in differences). The later is often called *marginal effects*.
  - Continuous vs. discrete marginal effects
    - For a continuous covariate, `margins` computes the first derivative of the response with respect to the covariate.
    - For a discrete covariate, `margins` computes the effect of a discrete change of the covariate (discrete change effects).
  - MEM: marginal effects at the mean, AME: average marginal effects, MER: marginal effects at representative values

Technical note

- You must use Stata's factor variable notation in the estimation command for `margins` to be able to compute correct results (see `help fvvarlist`).
  - Use the `i.` operator for discrete variables.
  - Use the # and ## operators for interactions.
  - Use the `c.` for continuous variables involved in an interaction.

Answers: Overall effects of the vignette factors

- Predictive margins / adjusted predictions (levels)

```
. quietly regress vrating vinc i.vmale i.vmarried i.veffort ///
>     vmale##rmale##c.reducyrs##c.reducyrs, vsquish noheader

. margins vmale vmarried veffort

Predictive margins                              Number of obs   =       1482
Model VCE    : OLS

Expression   : Linear prediction, predict()

------------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       vmale |
          0  |   .5727241   .0541998    10.57   0.000     .4664945    .6789537
          1  |   .2916397   .0547894     5.32   0.000     .1842544     .399025
             |
    vmarried |
          0  |   .5250302    .054352     9.66   0.000     .4185023    .6315581
          1  |   .3451187   .0546468     6.32   0.000      .238013    .4522245
             |
     veffort |
          0  |   .7664507   .0543315    14.11   0.000     .6599629    .8729385
          1  |   .1001736   .0547003     1.83   0.067    -.0070371    .2073843
------------------------------------------------------------------------------
```

Answers: Overall effects of the vignette factors

- Interpretation of predictive margins for `vmale`:
  - If all respondents would have answered the *female* vignette (keeping the other vignette factors and the repondent's sex and education as they happen to be), then the average response would have been 0.57.
  - If all respondents would have answered the *male* vignette (keeping the other vignette factors and the repondent's sex and education as they happen to be), then the average response would have been 0.29.
  - This means that, keeping everything else constant, the same income is more likely to be judged as too low in the male vignette than in the female vignette.
  - To find out whether the difference is significant, we can use `margins` to compute contrasts or marginal effects.
    - ★ An alternative would be to specify the `post` option in the above command and then apply the `test` command (see below).

Answers: Overall effects of the vignette factors

- Contrasts (differences in levels): use the `r.` operator

```
. margins r.vmale r.vmarried r.veffort

Contrasts of predictive margins
Model VCE    : OLS

Expression   : Linear prediction, predict()
```

	df	chi2	P>chi2
vmale	1	13.30	0.0003
vmarried	1	5.44	0.0197
veffort	1	74.49	0.0000

	Contrast	Delta-method Std. Err.	[95% Conf.	Interval]
vmale (1 vs 0)	-.2810844	.0770733	-.4321454	-.1300234
vmarried (1 vs 0)	-.1799115	.0771521	-.3311268	-.0286961
veffort (1 vs 0)	-.6662772	.0771985	-.8175835	-.5149709

Answers: Overall effects of the vignette factors

- Marginal effects for discrete variables (discrete change effects): use the `dydx()` option

```
. margins, dydx(vmale vmarried veffort)

Average marginal effects                          Number of obs   =       1482
Model VCE    : OLS

Expression   : Linear prediction, predict()
dy/dx w.r.t. : 1.vmale 1.vmarried 1.veffort

------------------------------------------------------------------------------
             |            Delta-method
             |      dy/dx   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
     1.vmale |  -.2810844   .0770733    -3.65   0.000    -.4321454   -.1300234
  1.vmarried |  -.1799115   .0771521    -2.33   0.020    -.3311268   -.0286961
   1.veffort |  -.6662772   .0771985    -8.63   0.000    -.8175835   -.5149709
------------------------------------------------------------------------------
Note: dy/dx for factor levels is the discrete change from the base level.
```

Answers: Overall effects of the vignette factors

- Interpretation contrasts / marginal effects
  - We see that both commands yield the same results.
  - The effect of male vs. female sex in the vignette is an average decrease of 0.28 points on the response scale. This is simply the difference in the predictive margins computed above.
  - The difference is highly significant with a z-value of 3.65 or a $\chi^2(1)$-value of 13.3 (which is simply the square of the z-value because the test has 1 degree of freedom).
  - The 95% confidence interval of the effect is -0.43 to -0.13.

Answers: How does the effect of vignette sex depend on respondent's education and sex?

- Let's start with the interaction with respondent's sex.
- Predictive margins for the vignette sex by sex of the respondent can be computed as follows:

```
. margins vmale, over(rmale)

Predictive margins                              Number of obs   =       1482
Model VCE    : OLS

Expression   : Linear prediction, predict()
over         : rmale

------------------------------------------------------------------------------
             |            Delta-method
             |     Margin   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
 rmale#vmale |
        0 0  |    .580111   .0744694     7.79   0.000     .4341537    .7260683
        0 1  |   .2069453   .0758263     2.73   0.006     .0583285    .3555621
        1 0  |   .5645165   .0790644     7.14   0.000     .4095531    .7194799
        1 1  |   .3857447   .0792884     4.87   0.000     .2303422    .5411472
------------------------------------------------------------------------------
```

Answers: How does the effect of vignette sex depend on respondent's education and sex?

- We see that for female respondents, the difference in predictive margins for vignette sex is larger ($0.21 - 0.58 = -0.37$) than for male respondents ($0.39 - 0.56 = -0.17$).
  - Note that the difference could be due to differential educational levels of female and male respondents, because an interaction with education was included in the regression model and the predictive margins are averaged over female and male respondents as is. We could, for example, type

    ```
    . margins vmale rmale, at((omean) reducyrs)
    ```

    to find out whether controlling for education changes the picture (it does a bit, but not much) (`omean` sets the respondent's education to the overall mean across all observations).
- To find out whether effects for female and males are significant we can again resort to the `r.` contrast operator or the `dydx()` option.

Answers: How does the effect of vignette sex depend on respondent's education and sex?

- Using the `r.` contrast operator:

```
. margins r.vmale, over(rmale)

Contrasts of predictive margins
Model VCE    : OLS

Expression   : Linear prediction, predict()
over         : rmale

-------------------------------------------------
             |         df        chi2     P>chi2
-------------+-----------------------------------
vmale@rmale  |
 (1 vs 0) 0  |          1       12.32     0.0004
 (1 vs 0) 1  |          1        2.55     0.1105
      Joint  |          2       14.88     0.0006
-------------------------------------------------

--------------------------------------------------------------
             |            Delta-method
             |   Contrast   Std. Err.     [95% Conf. Interval]
-------------+------------------------------------------------
vmale@rmale  |
 (1 vs 0) 0  |  -.3731657    .1063105     -.5815304   -.164801
 (1 vs 0) 1  |  -.1787718    .1120037     -.3982952   .0407515
--------------------------------------------------------------
```

Answers: How does the effect of vignette sex depend on respondent's education and sex?

- Using the `dydx()` option:

```
. margins, dydx(vmale) over(rmale)

Average marginal effects                          Number of obs   =       1482
Model VCE    : OLS

Expression   : Linear prediction, predict()
dy/dx w.r.t. : 1.vmale
over         : rmale

------------------------------------------------------------------------------
             |            Delta-method
             |      dy/dx   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
1.vmale      |
       rmale |
          0  |  -.3731657   .1063105    -3.51   0.000    -.5815304    -.164801
          1  |  -.1787718   .1120037    -1.60   0.110    -.3982952    .0407515
------------------------------------------------------------------------------
Note: dy/dx for factor levels is the discrete change from the base level.
```

Answers: How does the effect of vignette sex depend on respondent's education and sex?

- Mechanics of the `over()` option:
  - Specifying over is equivalent to running `margins` on subpopulations. That is, the results above could also be computed by typing

```
. margins if rmale==0, dydx(vmale)
. margins if rmale==1, dydx(vmale)
```

  - Beware, however, that using `if` can lead to biased standard errors in complex samples. A saver approach is to use the `subpop()` option:

```
. margins, dydx(vmale) subpop(if rmale==0)
. margins, dydx(vmale) subpop(if rmale==1)
```

Answers: How does the effect of vignette sex depend on respondent's education and sex?

- What we really want to know is whether the effect of the vignette sex is different for female respondents and for male respondents.
- We could test this, for example, as follows:

```
. estimates store lin

. margins, dydx(vmale) over(rmale) coeflegend post

Average marginal effects                          Number of obs   =       1482
Model VCE    : OLS

Expression   : Linear prediction, predict()
dy/dx w.r.t. : 1.vmale
over         : rmale

------------------------------------------------------------------------------
             |      dy/dx  Legend
-------------+----------------------------------------------------------------
1.vmale      |
       rmale |
          0  |  -.3731657  _b[1.vmale:0bn.rmale]
          1  |  -.1787718  _b[1.vmale:1.rmale]
------------------------------------------------------------------------------
Note: dy/dx for factor levels is the discrete change from the base level.

. test _b[1.vmale:0bn.rmale] = _b[1.vmale:1.rmale]

 ( 1)  [1.vmale]0bn.rmale - [1.vmale]1.rmale = 0

           chi2(  1) =    1.58
         Prob > chi2 =    0.2083

. est restore lin
(results lin are active now)
```

Answers: How does the effect of vignette sex depend on respondent's education and sex?

- A more direct approach is to have `margins` compute an estimate for the difference in differences by adding the `r.` operator within the `over()` option.
- Either type ...

```
. margins r.vmale, over(r.rmale)

Contrasts of predictive margins
Model VCE    : OLS

Expression   : Linear prediction, predict()
over         : rmale

------------------------------------------------
             |         df        chi2     P>chi2
-------------+----------------------------------
 rmale#vmale |          1        1.58     0.2083
------------------------------------------------

--------------------------------------------------------------
             |            Delta-method
             |   Contrast   Std. Err.     [95% Conf. Interval]
-------------+------------------------------------------------
 rmale#vmale |
(1 vs 0) (1 vs 0) |   .1943939   .1544914    -.1084037    .4971914
--------------------------------------------------------------
```

Answers: How does the effect of vignette sex depend on respondent's education and sex?

- . . . or type

```
. margins, dydx(vmale) over(r.rmale)

Contrasts of average marginal effects
Model VCE    : OLS

Expression   : Linear prediction, predict()
dy/dx w.r.t. : 1.vmale
over         : rmale

------------------------------------------------
             |         df        chi2     P>chi2
-------------+----------------------------------
0b.vmale     |
       rmale |  (omitted)
-------------+----------------------------------
1.vmale      |
       rmale |          1        1.58     0.2083
------------------------------------------------

--------------------------------------------------------------
             |   Contrast Delta-method
             |      dy/dx   Std. Err.     [95% Conf. Interval]
-------------+------------------------------------------------
1.vmale      |
       rmale |
    (1 vs 0) |   .1943939   .1544914     -.1084037    .4971914
--------------------------------------------------------------
Note: dy/dx for factor levels is the discrete change from the
      base level.
```

Answers: How does the effect of vignette sex depend on respondent's education and sex?

- Our conclusion from the above results would be that the effects of the vignette sex appears to be a bit stronger for female, but the difference is not significant.
- Again, note that there is a third variable involved in the interaction (education), so that part of the differences between the effects for female and male respondents might be due to different educational level.
- To see how the effect of the vignette sex changes by education, we could type

```
. margins r.vmale, at(reducyrs=(8(1)18))
. marginsplot, yline(0)
```

or

```
. margins, dydx(vmale) at(reducyrs=(8(1)18))
. marginsplot, yline(0)
```

Answers: How does the effect of vignette sex depend on respondent's education and sex?

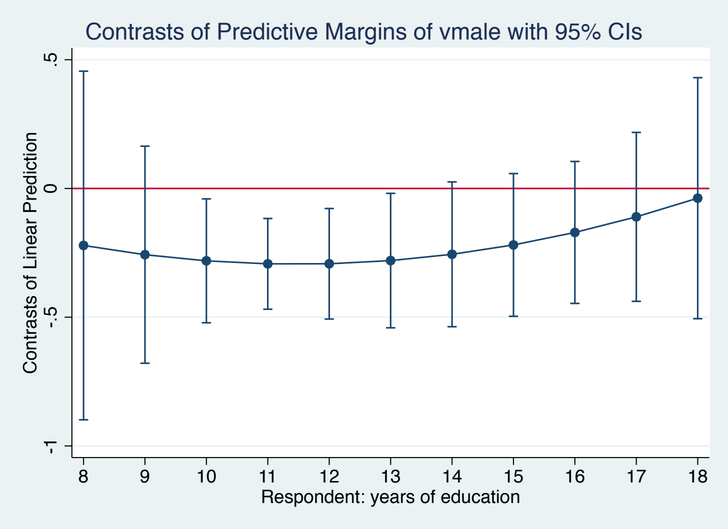

Answers: How does the effect of vignette sex depend on respondent's education and sex?

- We see that, after an initial increase, the effect of the vignette sex diminishes (i.e. gets closer to zero) as education increases.
- Still we are looking only at a two-way interaction (vignette sex by education). To explore the full three-way interaction specified in the model we have to go one step further.
- For example, to see how the effect of vignette sex depends on education by sex of respondent, we could type

```
. margins r.vmale, over(rmale) at(reducyrs=(8(1)18))
. marginsplot, yline(0)
```

or

```
. margins, dydx(vmale) over(rmale) at(reducyrs=(8(1)18))
. marginsplot, yline(0)
```

Answers: How does the effect of vignette sex depend on respondent's education and sex?

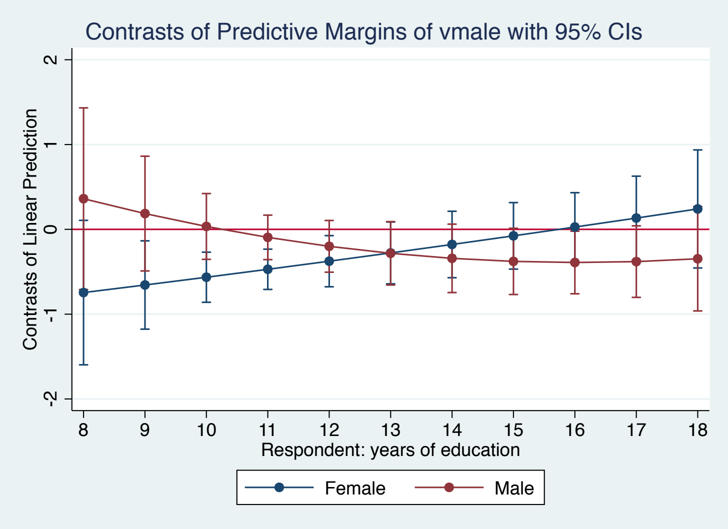

Answers: How does the effect of vignette sex depend on respondent's education and sex?

- The pattern of the effect of the vignette sex appears to be somewhat different for female respondents and male respondents.
- To better see how the difference in the effect of the vignette sex between female and male respondents changes with education, we could type

```
. margins r.vmale, over(r.rmale) at(reducyrs=(8(1)18))
. marginsplot, yline(0)
```

or

```
. margins, dydx(vmale) over(r.rmale) at(reducyrs=(8(1)18))
. marginsplot, yline(0)
```

Answers: How does the effect of vignette sex depend on respondent's education and sex?

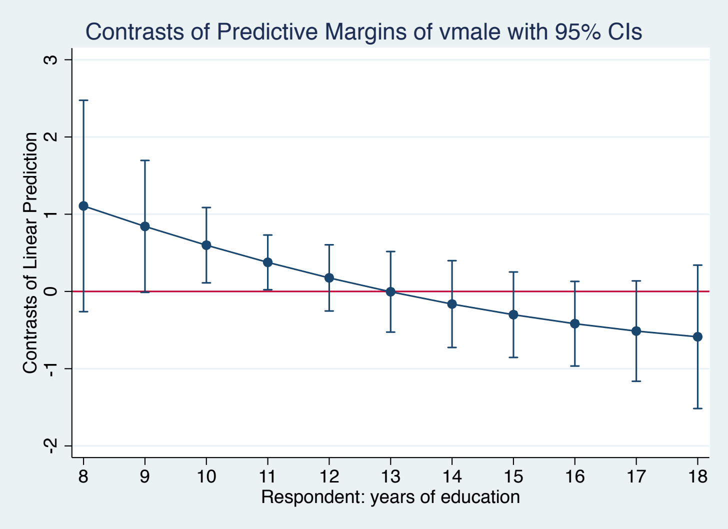

Answers: How does the effect of vignette sex depend on respondent's education and sex?

- In this graph, a difference-in-differences estimate is displayed for each educational level. We see that the difference in effects between female and male respondents is first positive, but then declines and becomes negative after 13 years of education

Answers: Shape of the effect of education

- Up to now, we looked at how the effect of the vignette sex changes depending on education. However, we might also be interested in the main effect of education on the responses.
- To see how the response level changes with education we can simply type

```
. margins, at(reducyrs=(8(1)18))
. marginsplot
```

- Furthermore, for a picture of how the effect of education changes with educational level, type

```
. margins, at(reducyrs=(8(1)18)) contrast(atcontrast(ar._at))
. marginsplot, yline(0)
```

  The `contrast(atcontrast(ar._at))` option causes `margins` to compute contrasts between predictive margins across adjacent educational levels.

Answers: Shape of the effect of education

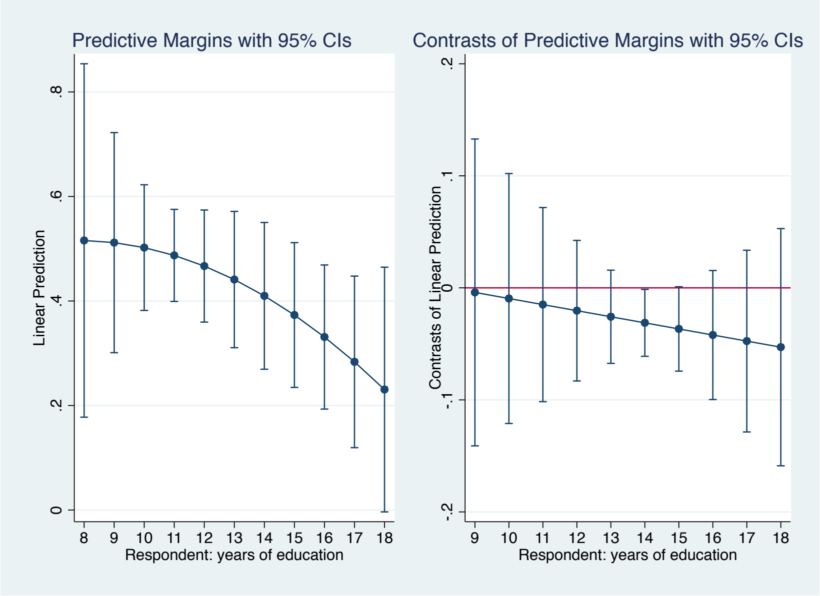

Answers: Shape of the effect of education

- Respondents with lower educational level are more likely to judge the income in the vignette as too high than respondents with higher educational level. Furthermore, the (negative) effect of education is getting steeper with additional education.
- A variant of the plot on the left can also be quickly produced by Patrick Royston's `marginscontplot`:

```
. marginscontplot reducyrs, ci
```

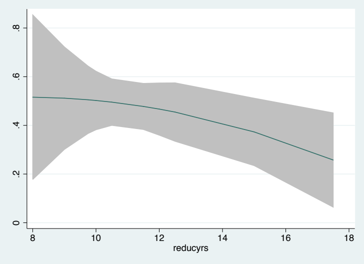

Answers: Shape of the effect of education

- To compare the shape of the effect of education between female and male respondents, we could type something like

```
. margins, at(reducyrs=(8(1)18)) over(rmale)
. marginsplot, name(a, replace)
. margins, at(reducyrs=(8(1)18)) over(r.rmale)
. marginsplot, yline(0) nodraw name(b, replace)
. graph combine a b
```

Answers: Shape of the effect of education

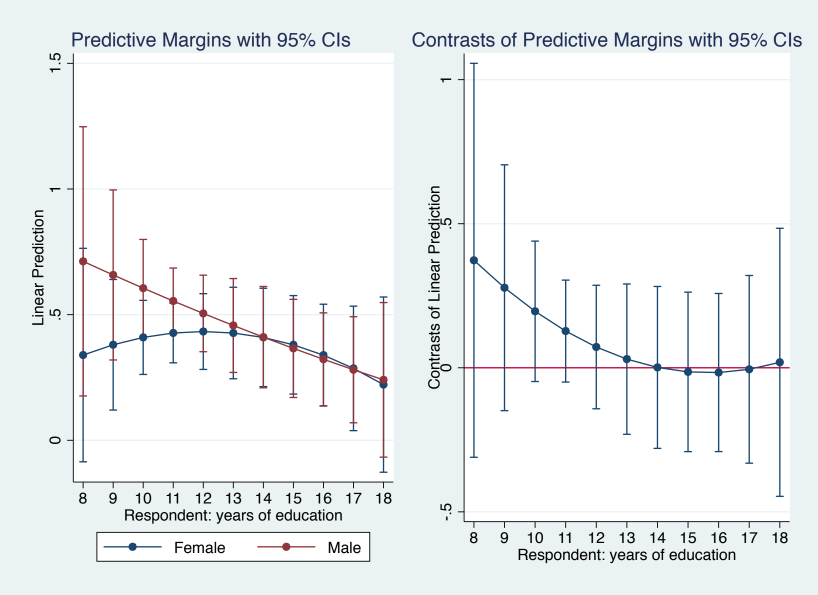

Answers: Shape of the effect of education

- In the left plot, the predictive margins by educational level are shown for female respondents and male respondents.
- In the right plot, the difference between the two curves is plotted. This shows how the effect of the sex of the respondent changes with educational level.

Answers: Can we express the effects in CHF?

- As we can see in the regression output on slide 7, the vignette income has an effect of 0.0010069 on the judgement. An increase in income by 1 CHF increases the expected judgement by 0.0010069 points.
- This means, that 1 point on the judgement scale is worth about 993 CHF:

```
. display 1/_b[vinc]
993.15639
```

- You can use the `expression()` option in margins to compute predictive margins and marginal effects with respect to a rescaled outcome so that, in our case, all effects are expressed in CHF.
- `margins` will take care of the details and also provide consistent standard errors.

Answers: Can we express the effects in CHF?

- Example: Effects of vignette factors

```
. margins, dydx(vmale vmarried veffort) expression(xb()/-_b[vinc])

Average marginal effects                          Number of obs   =       1482
Model VCE     : OLS

Expression    : xb()/-_b[vinc]
dy/dx w.r.t. : 1.vmale 1.vmarried 1.veffort

------------------------------------------------------------------------------
             |            Delta-method
             |      dy/dx   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
     1.vmale |   279.1608   80.52689     3.47   0.001      121.331    436.9906
  1.vmarried |   178.6802   78.74228     2.27   0.023      24.3482    333.0123
   1.veffort |   661.7174   98.82199     6.70   0.000     468.0299     855.405
------------------------------------------------------------------------------
Note: dy/dx for factor levels is the discrete change from the base level.

. marginsplot, horizontal recast(scatter) recastci(rspike) ///
>        xline(0) xlabel(, grid) ylabel(, nogrid)

  Variables that uniquely identify margins: _deriv
```

- Interpretation: Males "should" get 279 CHF more per month than females, the married should get 179 CHF more, the hard-working should get 661 CHF more.

Answers: Can we express the effects in CHF?

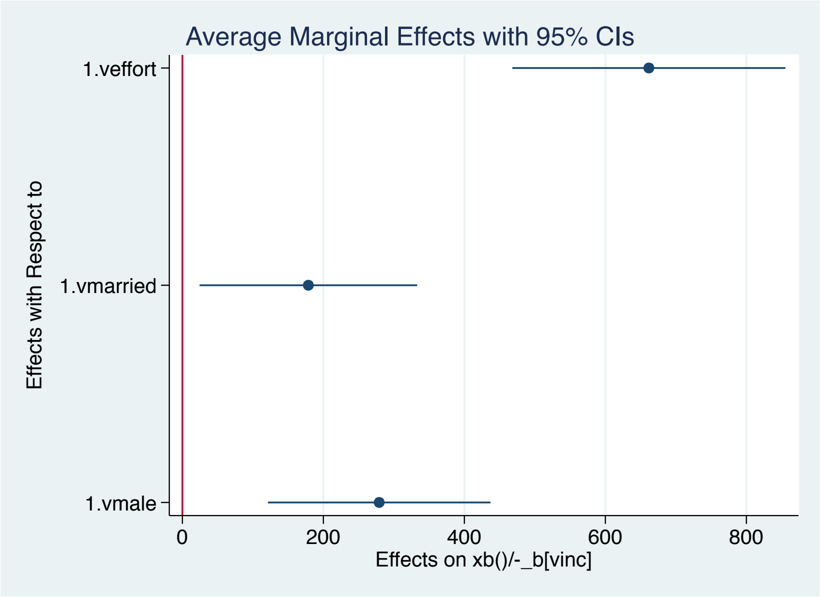

Answers: Can we express the effects in CHF?

- Example: "Just" income levels for females and males by education of respondent

```
. margins vmale, at(reducyrs = (8(1)18)) expression(xb()/-_b[vinc] + vinc)
. marginsplot
```

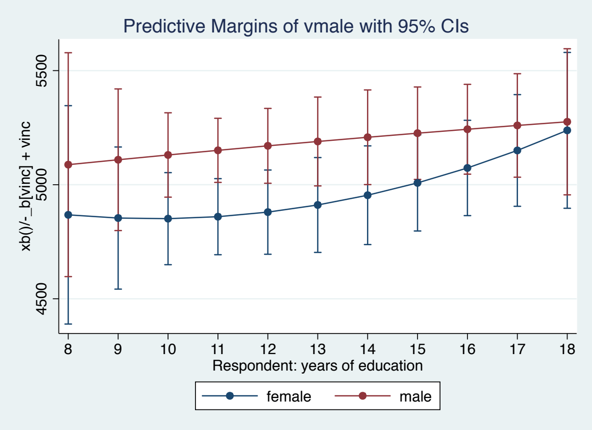

Answers: Interpretation of a logit model

- What is the conditional probability of "too low" depending on different levels of the factor variables?

```
. quietly logit toolow vinc i.vmale i.vmarried i.veffort

. margins vmale vmarried veffort

Predictive margins                                Number of obs   =       1482
Model VCE    : OIM

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

	Margin	Delta-method Std. Err.	z	P>|z|	[95% Conf.	Interval]
vmale						
0	.1652877	.012963	12.75	0.000	.1398807	.1906946
1	.2215201	.0145671	15.21	0.000	.192969	.2500711
vmarried						
0	.1725705	.0133054	12.97	0.000	.1464925	.1986485
1	.2131927	.0142477	14.96	0.000	.1852677	.2411177
veffort						
0	.1115456	.0113453	9.83	0.000	.0893093	.1337819
1	.2749636	.015868	17.33	0.000	.2438629	.3060642

Answers: Interpretation of a logit model

- What is the (average) marginal effect of the vignette factors on the probability of "too low"?

```
. margins, dydx(vmale vmarried veffort)

Average marginal effects                          Number of obs   =       1482
Model VCE    : OIM

Expression   : Pr(toolow), predict()
dy/dx w.r.t. : 1.vmale 1.vmarried 1.veffort

------------------------------------------------------------------------------
             |            Delta-method
             |      dy/dx   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
     1.vmale |   .0562324   .0195033     2.88   0.004     .0180066    .0944582
  1.vmarried |   .0406222   .0194988     2.08   0.037     .0024053    .0788392
   1.veffort |   .1634179   .0195074     8.38   0.000     .1251842    .2016517
------------------------------------------------------------------------------
Note: dy/dx for factor levels is the discrete change from the base level.
```

- We see that, for example, the average marginal effect of the vignette sex is 5 percentage points. That is, everything else equal, we would expect a 5 percentage point increase in the proportion of respondents who judge the vignette income as too low if we change the vignette sex from female to male.

What margins and marginsplot can't do

- The default for continuous variables is to compute marginal effects as first derivatives. Discrete change effects for continuous variables can only be computed for special cases (e.g. min to max, in steps across the scale).
  - It would be nice to be able to compute discrete change effects around observed values (e.g +/- half a standard deviation).
- Marginal effects for transformed covariates can only be computed in special cases (e.g. quadratic).
- `margins` can only deal with one equation at the time in models with multiple outcomes (e.g. `mlogit`).
- `margins` can be excessively slow on big datasets.
- `marginsplot` can only handle results from `margins` and can only display one set of results at the time.

New command: regplot

- A new command called `regplot` provides a solution for the last problem.

```
regplot models, options
```

- `regplot` uses `marginsplot` internally, but it can be applied to any estimation results, be they from margins or not.

This command has been published as `coefplot`.

- Multiple estimation results can be combined in a single graph.
- By default, `regplot` creates a horizontal "dot plot" of the coefficients found in `e(b)` and includes spikes for confidence intervals. I call this a "regression plot". Others sometimes call it an "airplane plot".

New command: regplot

- How to specify *models*:
  - Models in separate plots:

    `(`*modelname*`,` *options*`) || (`*modelname*`,` *options*`) ...`

  - Several models in one plot:

    `(`*modelname*`,` *options*`) (`*modelname*`,` *options*`) ...`

    This command has been published as `coefplot`**.**

  - Append models:

    `(`*modelname*`,` *options* `\` *modelname*`,` *options*`) ...`

  - A combination of the above.

Plot of raw coefficients

- `regplot` can be applied directly after an estimation command to produce a plot of the estimated coefficients.

```
. logit toolow vinc i.vmale i.vmarried i.veffort, nolog

Logistic regression                               Number of obs   =       1482
                                                  LR chi2(4)      =     140.79
                                                  Prob > chi2     =     0.0000
Log likelihood = -656.55323                       Pseudo R2       =     0.0968

------------------------------------------------------------------------------
      toolow |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
        vinc |  -.0013663   .0001789    -7.64   0.000    -.0017169   -.0010157
     1.vmale |   .3994001   .1393606     2.87   0.004     .1262584    .6725417
  1.vmarried |   .2886296   .1389959     2.08   0.038     .0162027    .5610565
   1.veffort |   1.164184   .1466717     7.94   0.000     .8767125    1.451655
       _cons |   4.939563   .9606374     5.14   0.000     3.056748    6.822378
------------------------------------------------------------------------------

. regplot ., keep(1.*) xline(0) xtitle(Log odds)
```

This command has been published as `coefplot`.

Plot of raw coefficients

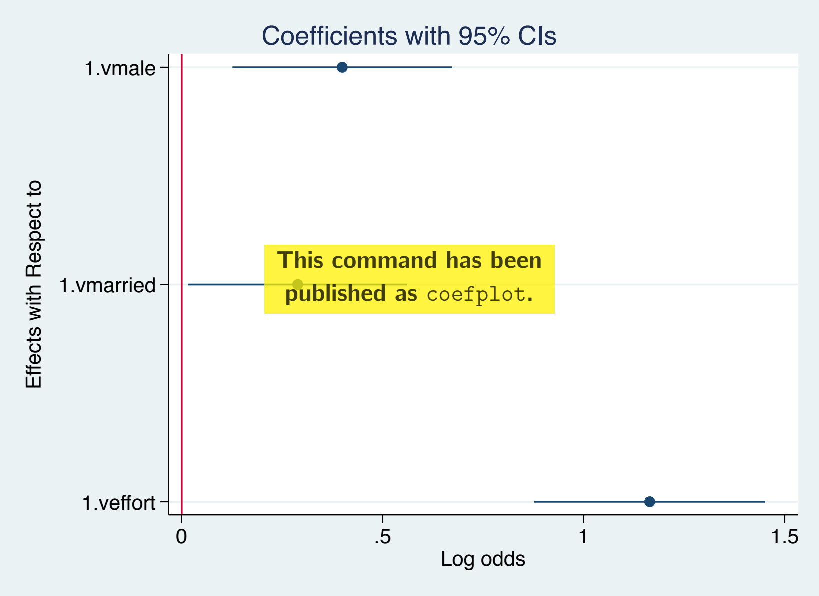

This command has been published as `coefplot`.

Label the coefficients

- Coefficients are positioned on the y-axis at 1, 2, 3, ... (from top). Hence, you can use `ylabel()` to define custom labels.

```
. regplot ., keep(1.*) xline(0) xtitle(Log odds) ///
>     ylabel(1 "male" 2 "married" 3 `""above-average" "effort""´) ///
>     yscale(range(0.75 3.25))
```

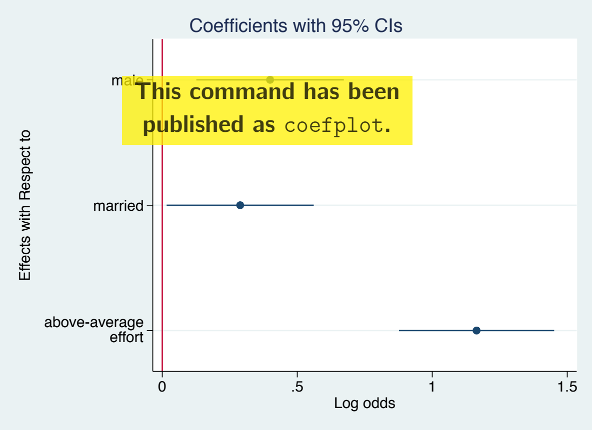

Plot results from margins

- `regplot` can also be applied after `margins`, if the `post` option is specified with `margins`:

```
. margins, dydx(vmale vmarried veffort) post

Average marginal effects                          Number of obs   =       1482
Model VCE    : OIM

Expression   : Pr(toolow), predict()
dy/dx w.r.t. : 1.vmale 1.vmarried 1.veffort

------------------------------------------------------------------------------
             |            Delta-method
             |      dy/dx   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
     1.vmale |   .0562324   .0195033     2.88   0.004     .0180066    .0944582
  1.vmarried |   .0406222   .0194988     2.08   0.037     .0024053    .0788392
   1.veffort |   .1634179   .0195074     8.38   0.000     .1251842    .2016517
------------------------------------------------------------------------------
Note: dy/dx for factor levels is the discrete change from the base level.

. regplot ., xline(0) keep(1.*) title(Average marginal effects with 95% CIs)
```

This command has been published as `coefplot`.

Plot results from margins

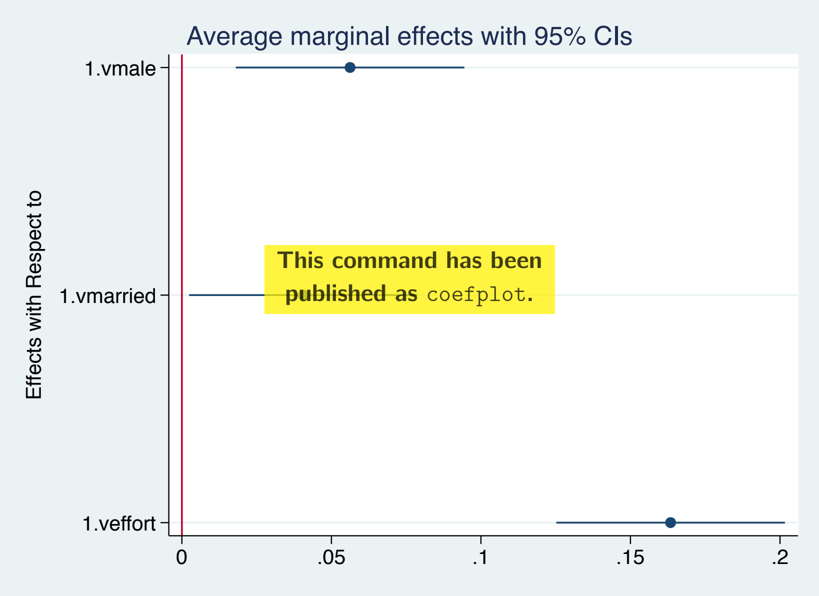

This command has been published as `coefplot`**.**

Combine raw coefficients and AMEs

- `regplot` can be applied to stored models. For example, we could combine raw coefficients from `logit` and corresponding average marginal effects from `margins`

```
. quietly logit toolow vinc i.vmale i.vmarried i.veffort
. estimates store raw
. quietly margins, dydx(vmale vmarried veffort) post
. estimates store ame
. regplot (raw, bylabel(Log Odds)) || (ame, bylabel(AMEs)) ///
>     , xline(0) keep(1.*) byopts(xrescale)
```

This command has been published as `coefplot`.

Combine raw coefficients and AMEs

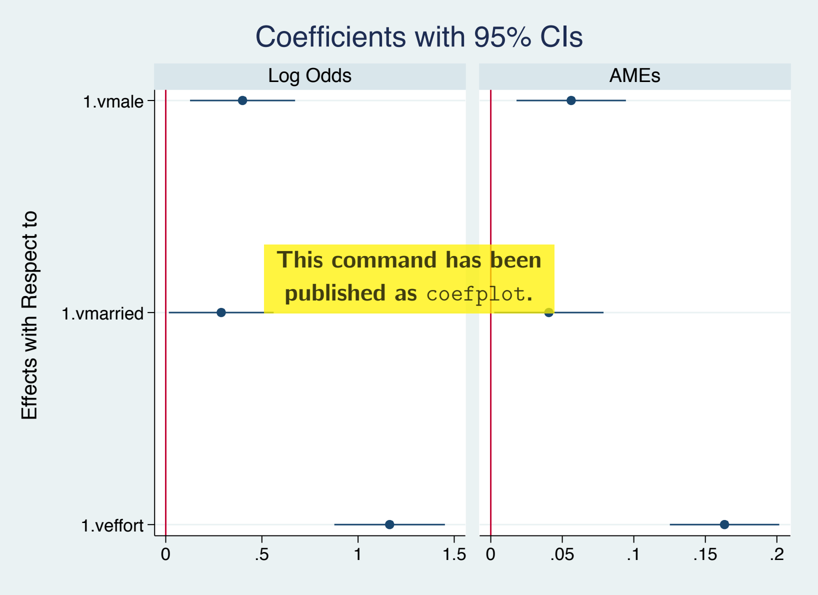

This command has been published as `coefplot`.

Combine predictive margins and contrasts

- We could also predictive margins and contrasts/marginal effects in one graph.

```
. quietly logit toolow vinc i.vmale i.vmarried i.veffort

. estimates store raw

. quietly margins vmale vmarried veffort, post

. estimates store margins

. estimates restore raw
(results raw are active now)

. quietly margins, dydx(vmale vmarried veffort) post

. estimates store contrasts

. regplot (margins, rename(0bn. "") drop(1.*)   label(p(x=0))) ///
>         (margins, rename(1. "")   drop(0bn.*) label(p(x=1)) ///
>             bylabel(predictive margins)) || ///
>         (contrasts, rename(1. "") drop(0.*) label(p(x=1)-p(x=0)) ///
>             bylabel(contrasts)) , norecycle legend(rows(1))
```

This command has been published as `coefplot`.

Combine predictive margins and contrasts

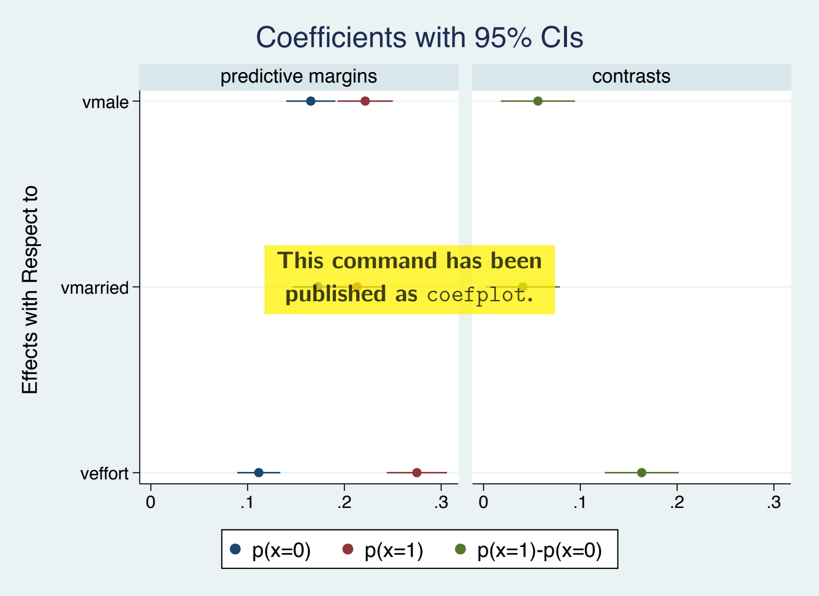

This command has been published as `coefplot`.

Append models

- Use \ to append models in one plot.

```
. quietly logit toolow vinc i.vmale i.vmarried i.veffort

. estimates store logit

. quietly margins vmale, post

. estimates store vmale

. estimates restore logit
(results logit are active now)

. quietly margins vmarried, post

. estimates store vmarried

. estimates restore logit
(results logit are active now)

. quietly margins veffort, post

. estimates store veffort

. regplot (vmale \ veffort \ vmarried), title(Appended in same plot) ///
>     nodraw name(a, replace)

. regplot (vmale) (veffort) (vmarried), title(As separate (overlayed) plots) ///
>     legend(off) nodraw name(b, replace)

. graph combine a b
```

This command has been published as `coefplot`.

Append models

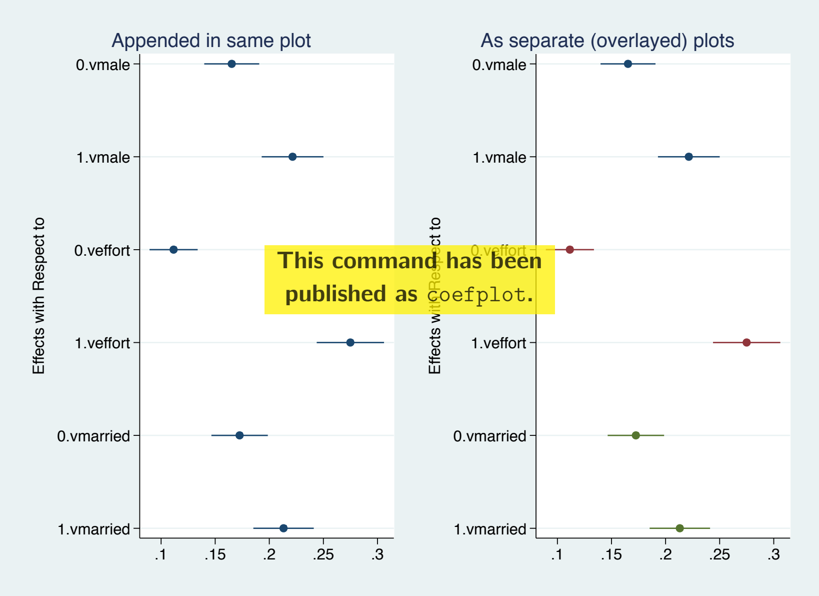

This command has been published as `coefplot`.

Plot coefficients from mlogit

- In multiequation models, use the `eq()` option to select the equation to be plotted.

```
. mlogit rmarstat rage reducyrs rright rinc, nolog

Multinomial logistic regression                   Number of obs   =       1482
                                                  LR chi2(12)     =     477.71
                                                  Prob > chi2     =     0.0000
Log likelihood = -1380.6739                       Pseudo R2       =     0.1475

------------------------------------------------------------------------------
    rmarstat |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
single       |
        rage |  -.0705039   .0055962   -12.60   0.000    -.0814724   -.0595355
    reducyrs |   .0531021   .0278531     1.91   0.057    -.0014889    .1076931
      rright |  -.0924911   .0340842    -2.71   0.007    -.1592948   -.0256873
        rinc |  -.0604214   .0522317    -1.16   0.247    -.1627937    .0419509
       _cons |   2.533033   .4611503     5.49   0.000     1.629195    3.436871
-------------+----------------------------------------------------------------
married      |  (base outcome)
-------------+----------------------------------------------------------------
divorced     |
        rage |   .0190981   .0064624     2.96   0.003      .006432    .0317642
    reducyrs |  -.0109678   .0386512    -0.28   0.777    -.0867228    .0647873
      rright |  -.1147781   .0449214    -2.56   0.011    -.2028225   -.0267338
        rinc |  -.0691778   .0695684    -0.99   0.320    -.2055293    .0671736
       _cons |  -1.927327    .633678    -3.04   0.002    -3.169313   -.6853408
-------------+----------------------------------------------------------------
widowed      |
        rage |   .1021225   .0101098    10.10   0.000     .0823077    .1219373
    reducyrs |  -.1693922   .0577359    -2.93   0.003    -.2825526   -.0562319
      rright |  -.1726326   .0580334    -2.97   0.003     -.286376   -.0588892
        rinc |   .0955222   .0920018     1.04   0.299     -.084798    .2758423
       _cons |  -6.098965   .9666954    -6.31   0.000    -7.993653   -4.204277
------------------------------------------------------------------------------

. estimates store marstat

. regplot (marstat, eq(single))   ///
>         (marstat, eq(married))  ///
>         (marstat, eq(divorced)) ///
>         (marstat, eq(widowed)), nocons xline(0)
```

This command has been published as `coefplot`.

Plot coefficients from mlogit

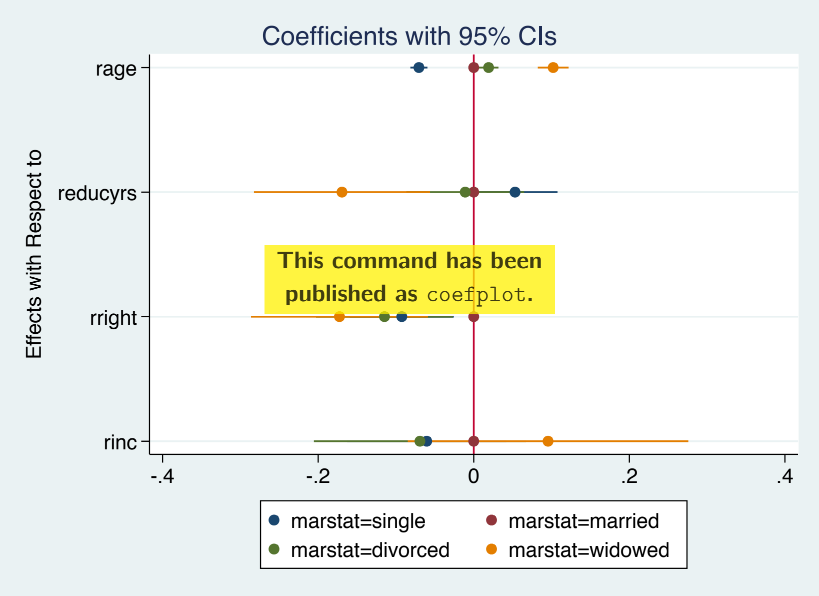

This command has been published as `coefplot`.

Plot marginal effects from mlogit

- To plot marginal effects for an `mlogit` model you have to run `margins` for each outcome.

```
. quietly margins, dydx(*) predict(outcome(single)) post

. estimates store single

. estimates restore marstat
(results marstat are active now)

. quietly margins, dydx(*) predict(outcome(married)) post

. estimates store married

. estimates restore marstat
(results marstat are active now)

. quietly margins, dydx(*) predict(outcome(divorced)) post

. estimates store divorced

. estimates restore marstat
(results marstat are active now)

. quietly margins, dydx(*) predict(outcome(widowed)) post

. estimates store widowed

. regplot single married divorced widowed, legend(row(1)) xline(0) ///
>     title(AMEs with 95% CIs)
```

This command has been published as `coefplot`.

Plot marginal effects from mlogit

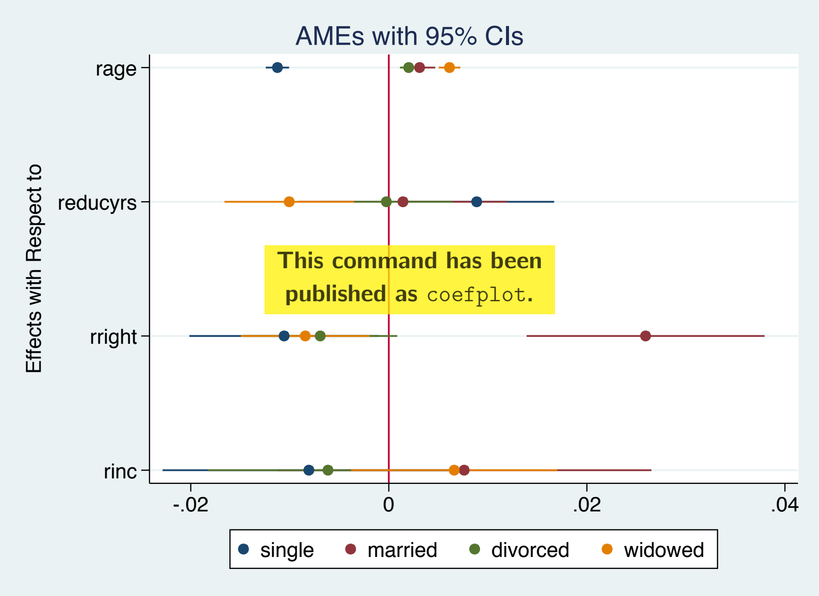

Plot marginal effects from mlogit

```
. regplot single || married || divorced || widowed, xline(0) ///
>     byopt(rows(1) title(AMEs with 95% CIs))
```

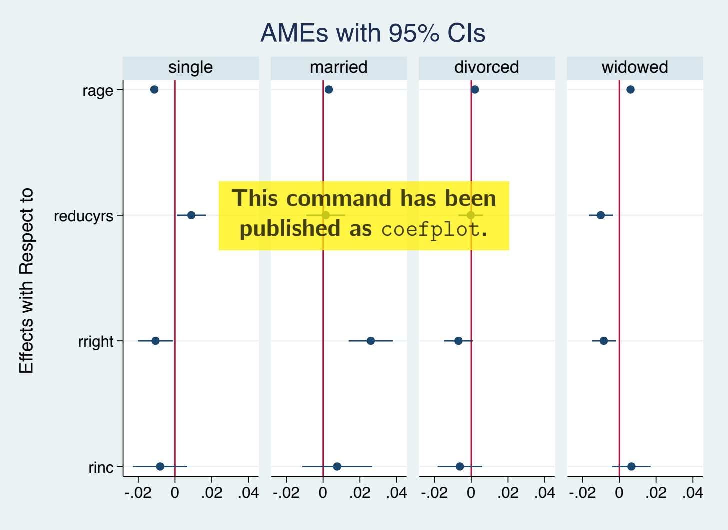

This command has been published as `coefplot`**.**

References

- Rising, B. (2012). Working in the margins to plot a clear course. Presentation at the 10th German Stata Users Group Meeting in Berlin.
- Rising, B. (2012). How to get an edge with margins and marginsplot. Presentation at the 2012 UK Stata Users Group Meeting in London.
- Royston, P. (forthcoming). marginscontplot: plotting the marginal effects of continuous predictors. *The Stata Journal*.
- Williams, R. (2012). Using the margins command to estimate and interpret adjusted predictions and marginal effects. *The Stata Journal* 12(2):308–331.