

margins — Adjusted predictions, predictive margins, and marginal effects

Description
Options

Quick start
Remarks and examples

Menu
Stored results

Syntax
Also see

Description

`margins` calculates statistics based on predictions of a previously fit model. These statistics can be calculated averaging over all covariates, or at fixed values of some covariates and averaged over the remaining covariates. After you fit a choice model, `margins` provides estimates such as marginal predicted choice probabilities, adjusted predictions, and marginal effects that allow you to easily interpret the results of a choice model.

Many possible margins can be calculated for choice models. Therefore, `margins` has special choice model options to select which outcomes are estimated or to select which alternatives are fixed or averaged within. These options are available after `cmlogit`, `cmmixlogit`, `cmxtmixlogit`, `cmmprobit`, and `cmroprobit`.

`margins` with the `contrast` option or with [contrast operators](#) performs contrasts (comparisons) of margins. After you fit a choice model, there are also special options to select contrasts for outcomes or for alternatives.

This entry focuses on the use of the special choice model options with `margins`. For the full capabilities of `margins`, see [\[R\] margins](#).

Quick start

Outcome probabilities, the average predicted probability of selecting each alternative

```
margins
```

Outcome probabilities for each level of factor-variable `a`

```
margins a
```

As above, but show results only for the probability of selecting the outcome labeled `Alt1`

```
margins a, outcome(Alt1)
```

Outcome probabilities setting alternative-specific variable `x1` to 50, 75, ..., 500 for each alternative, one at a time

```
margins, at(x1=(50(25)500))
```

As above, but set only the observations of `x1` corresponding to the alternative `Alt1` to the specified values

```
margins, at(x1=(50(25)500)) alternative(Alt1)
```

Outcome probabilities when `x1` is set to the specified values simultaneously across each alternative

```
margins, at(x1=(50(25)500)) alternative(simultaneous)
```

Average marginal effect of `x1` on the predicted probabilities of each outcome

```
margins, dydx(x1)
```

As above, but show results only for the outcome `Alt1` for a change in the alternative-specific variable `x1` at the observations corresponding to the alternative `Alt2`, and repeat for alternative `Alt3`

```
margins, dydx(x1) outcome(Alt1) alternative(Alt2 Alt3)
```

Contrasts

For each outcome, test equality of outcome probabilities across levels of a

```
margins a, contrast
```

Set alternative-specific variable `x1` to 4 at each of the alternatives identified by `altvar`, and test for differences from the previous alternative

```
margins, at(x1=4) contrast(alternativecontrast(ar.altvar))
```

Menu

Statistics > Postestimation

Syntax

```
margins [marginlist] [, options]
```

marginlist is a list of factor variables or interactions that appear in the current estimation results.

For the full syntax, see [\[R\] margins](#).

<i>options</i>	Description
<code>outcome(outcomes [, altsubpop])</code>	estimate margins for specified outcomes
<code>alternative(alts)</code>	estimate margins for specified alternatives for alternative-specific covariates
<code>alternative(simultaneous)</code>	estimate margins changing all alternatives simultaneously for alternative-specific covariates
<code>contrast</code>	joint tests of differences across levels of the elements of <i>marginlist</i>
<code>contrast(contrast_options)</code>	contrast the margins between the outcomes or alternatives as specified by <i>contrast_options</i>
<code>noesample</code>	do not restrict <code>margins</code> to the estimation sample
<i>other_margins_options</i>	see [R] margins for more options

<i>contrast_options</i>	Description
<code>outcomejoint</code>	joint test of differences across outcomes
<code>outcomecontrast(op[.outcome])</code>	apply the <i>op.</i> contrast operator to the outcomes
<code>alternativejoint</code>	joint test of differences across alternatives for alternative-specific covariates
<code>alternativecontrast(op[.altvar])</code>	apply the <i>op.</i> contrast operator to the levels of the alternatives for alternative-specific covariates

Options

Main

`outcome(outcomes [, altsubpop])` specifies that margins be estimated for the specified outcomes only. The default is to estimate margins for all outcomes.

outcomes is a list of one or more outcomes, which are the values of the alternatives variable; see [CM] [cmset](#). *outcomes* can be specified by

#1, #2, . . . , where #1 means the first level of the alternatives variable, #2 means the second level, etc.;

numeric values of the alternatives variable if it is a numeric variable;

value labels of the alternatives variable, enclosed in quotes if there are spaces in the value labels;

string values of the alternatives variable if it is a string variable, enclosed in quotes if there are spaces in the values; or

`_all` or `*` for all levels of the alternatives variable.

The suboption `altsubpop` applies only to samples with unbalanced choice sets. For balanced samples, the default is the same as specifying `altsubpop`. This option is used in conjunction with alternative-specific covariates and unbalanced choice sets to specify that calculations done for each alternative be restricted to the subpopulation of cases with that alternative in their choice set. The default treats the sample as if it were balanced with alternatives not in a choice set considered as alternatives with zero probability of being chosen. `altsubpop` is appropriate for unbalanced experimental designs in which decision makers were presented with different choice sets.

`alternative(alts)` applies only when one or more alternative-specific covariates are specified in an element of *marginlist*, in the `at()` option, or in one of the marginal effects options (`dydx()`, etc.). This option specifies that margins be estimated for the specified alternatives only. The default is to estimate margins for all alternatives. *alts* are specified in the same manner as in `outcome(outcomes)`.

`alternative(simultaneous)`, as with `alternative(alts)`, applies only when there are alternative-specific covariates in the specification of *margins*. By default, each alternative-specific covariate is changed (for example, set to a specified value) separately for each alternative, giving results for each alternative. This option specifies that each alternative-specific covariate is to be changed across all alternatives simultaneously to produce a single result.

For example, suppose *xvar* is an alternative-specific variable with alternatives *A*, *B*, and *C*, and *margins*, `at(xvar=1)` is specified. By default, *xvar* is first set to 1 for alternative *A* and kept at its sample values for *B* and *C*, then similarly for the alternative *B*, and then *C*, producing results for each of the three alternatives. The `alternative(simultaneous)` option sets *xvar* to 1 at each of the alternatives *A*, *B*, and *C* simultaneously, producing a single result for the alternatives as a group.

`contrast` applies only when *marginlist* is specified. If an element of *marginlist* contains only case-specific covariates, this option displays joint tests of differences among predicted probabilities across the levels of the element for each outcome. If the element contains alternative-specific covariates, this option displays joint tests of differences among predicted probabilities across the levels of the element for each outcome and alternative combination. It also displays a joint test of all the differences.

`contrast(outcomejoint)` displays a joint test of differences across all outcomes. It is a test of the null hypothesis: within each alternative, differences among predicted probabilities across levels of an element of *marginlist* are the same for each outcome.

`contrast(outcomecontrast(op[._outcome]))` applies the contrast operator *op.* to the outcomes. See the *op.* table in [R] **contrast** for a list of all contrast operators. The optional *._outcome* does nothing, but adding it will produce more readable code, showing what *op.* is operating on.

`contrast(alternativejoint)` applies only when there are alternative-specific covariates in the specification of *margins*. This option displays a joint test of differences across all alternatives. It is a test of the null hypothesis: within each outcome, differences among predicted probabilities across levels of an element of *marginlist* are the same for each alternative.

`contrast(alternativecontrast(op[.altvar]))` applies only when there are alternative-specific covariates in the specification of *margins*. This option applies the contrast operator *op.* to the alternatives. *altvar* is the name of the alternatives covariates used with *cmset*. The optional *.altvar* does nothing, but adding it will produce more readable code, showing what *op.* is operating on.

`noesample` specifies that *margins* not restrict its computations to the estimation sample used by the previous estimation command. If the estimation command used casewise deletion (the default), *margins* with `noesample` also omits missing values casewise. If the estimation command used alternativewise deletion (option `altwise`), alternativewise deletion is also used by *margins*.

other_margins_options; see [R] **margins** for additional options.

Remarks and examples

[stata.com](https://www.stata.com)

Remarks are presented under the following headings:

Introduction

Estimating margins for case-specific variables

Estimating margins for alternative-specific variables

The `altsubpop` suboption for unbalanced choice sets

More on unbalanced choice sets

The `outcomecontrast()` and `alternativecontrast()` suboptions

Graphing margins results

Introduction

Before reading this entry, read [CM] **Intro 1**. There you will learn about many of the common questions you can answer using *margins* after choice models and about some of the special options that are specific to *cm* commands. This entry explores even more of the choice model options for *margins* and delves deeper into the types of hypotheses you can test. Here we also provide advice on using the more advanced options, such as those for handling unbalanced choice sets.

The special choice model options for *margins* can be used after

```

cmlogit
cmmixlogit
cmxtmixlogit
cmmprobit
cmroprobit

```

margins has the same capabilities when used after any of these commands. All examples of *margins* shown in this entry will work after any of these commands.

The special choice model options described in this entry cannot, however, be used with `margins` after other choice models. `cmrologit` does not have explicitly identified alternatives, so you use the standard syntax of `margins` after this command; see [R] `margins`. `nlogit` does not allow `margins` because of the structure of the hierarchical model it fits.

`margins` produces estimates based on predictions. After `cm` estimation commands, two types of predictions can be used with `margins`, predicted probabilities or linear-form predictions. Predicted probabilities are the default and likely the only one you will use.

For choice models, a predicted probability is the probability of a decision maker choosing one of the possible alternatives, and these probabilities sum to one across the alternatives. (Note that for rank-ordered alternatives modeled by `cmprobit`, `margins` bases its results on predictions calculated using the `predict` option `pr1`, which gives the probability of ranking an alternative as first. So `margins` after `cmprobit` behaves just as it does after the `cm` estimators for models in which a single alternative is chosen.)

It is important to understand the difference between the use of the word “outcome” and the use of the word “alternative” in a `margins` specification or in output from `margins`. Whenever the word “outcome” is used in `margins`, it refers to what alternative is chosen.

Whenever the word “alternative” is used in a `margins` specification, it means do something special by alternative when `margins` operates on alternative-specific variables. So when you see “alternative” used with `margins`, do not think of alternatives as outcomes, think of manipulating alternative-specific variables at the observations corresponding to the alternatives. If you specify one of the `alternative*` options when there are only case-specific variables in the `margins` specification, it does nothing. It is simply ignored, no error message is given. You will not see the word `alternative` used in `margins` output; instead, you will see the name of the alternatives variable that you specified with `cmset`.

Let’s explain what `margins` does after `cm` estimators using a simple example. Suppose `cost_cat` is a case-specific categorical variable, and we included `i.cost_cat` as a `casevar()` in our `cm` estimation. If we now type

```
. margins cost_cat
```

the output will show the average predicted probability of selecting each alternative (as a possible outcome) at each of the levels of `cost_cat`. If there are k levels of `cost_cat` and n possible alternatives, there will be $k \times n$ predicted probabilities in the output from `margins`.

If, however, `cost_cat` were an alternative-specific variable, then `margins` would display $k \times n \times n$ predicted probabilities. Alternative-specific variables are variables that vary across both alternatives and cases, and each alternative-specific variable can be thought of as n different variables, one for each alternative. (Indeed, it is only because `cm` commands require data in long format that each alternative-specific variable is stored as a single Stata variable. If the storage design had been wide format, there would be n Stata variables for each alternative-specific variable.)

Suppose `cost_cat` is alternative specific and its levels are 1, 2, 3, and 4. Suppose the possible alternatives are `car`, `bus`, and `train`. `margins` first sets the level of `cost_cat` to 1 for the observations corresponding to alternative `car`. At the observations corresponding to the other alternatives, `cost_cat` is, by default, kept at its observed value. Then average predicted probabilities are calculated for each of the outcomes `car`, `bus`, and `train`. Then `cost_cat` is set to 2, still at alternative `car` only, and three more probabilities are calculated. This process continues until all $k \times n \times n = 4 \times 3 \times 3 = 36$ predicted probabilities are estimated. So the default output of `margins` can contain an overwhelming number of predicted probabilities for alternative-specific variables when n is not small. Some of the special options for `margins` after the `cm` estimation commands are there for the purpose of reducing the number of probabilities estimated and displayed.

The `outcome(levles)` option restricts the probabilities estimated by `margins` to the probabilities of the decision maker choosing only the alternatives in `levles`. This restriction works for both case-specific and alternative-specific variables.

The `alternative(levles)` option applies only to alternative-specific variables. With this option, `margins` changes only the variable at the observations corresponding to the alternatives in `levles`. For example, if we typed

```
. margins cost_cat, outcome(Car) alternative(Bus)
```

`margins` would estimate the average predicted probability of choosing `car` for `cost_cat` set to each of its levels, 1, 2, 3, and 4, at observations corresponding to the alternative `bus` only.

`margins` uses numerical derivatives to calculate standard errors. These computations can take a long time when your data have lots of cases, when there are lots of alternatives, or when there are lots of levels in the covariates in the `margins` specification. This is another reason to use `outcome()` and `alternative()`. Restricting the estimates to a smaller set of possibilities reduces computation time.

When `margins` is taking a long time to calculate estimates, you may want to first run `margins` with its `nose` option, which skips the standard error calculations.

```
. margins ..., ... nose
```

You can check the output and confirm that your specification of the `margins` command is what you want. Then, you can run it again without `nose` to get the standard errors.

Estimating margins for case-specific variables

When all the variables in the `margins` specification are case-specific variables, there is less output and it is easier to interpret.

► Example 1: Only case-specific variables

In [example 1](#) of [\[CM\] cmclogit](#), we fit a model of consumer choice of automobile. The alternatives are nationality of the automobile manufacturer: American, Japanese, European, and Korean. There is one alternative-specific variable in the model, `dealers`, that contains the number of dealerships of each nationality in the consumer's city. The case-specific variables are `gender` and `income`, the consumer's income in thousands of dollars.

We load the data and `cmset` them. For this example, we create a categorical variable, `income_cat`, that contains quartiles of income. We specify it as a case-specific variable along with `gender` and fit a `cmclogit` model:

```
. use https://www.stata-press.com/data/r17/carchoice
(Car choice data)
. cmset consumerid car
note: alternatives are unbalanced across choice sets; choice sets of
      different sizes found.
      Case ID variable: consumerid
      Alternatives variable: car
. xtile income_cat = income, nquantiles(4)
```

```
. cmclogit purchase dealers, casevars(i.gender i.income_cat)
```

```
Iteration 0: log likelihood = -960.62626
Iteration 1: log likelihood = -950.15551
Iteration 2: log likelihood = -949.74982
Iteration 3: log likelihood = -949.74885
Iteration 4: log likelihood = -949.74885
```

```
Conditional logit choice model      Number of obs      =      3,075
Case ID variable: consumerid        Number of cases    =         862
Alternatives variable: car           Alts per case: min =         3
                                       avg =         3.6
                                       max =         4
                                       Wald chi2(13)     =      48.28
Log likelihood = -949.74885         Prob > chi2        =      0.0000
```

purchase	Coefficient	Std. err.	z	P> z	[95% conf. interval]	
car						
dealers	.04461	.0263695	1.69	0.091	-.0070733	.0962932
American	(base alternative)					
Japanese						
gender						
Male	-.4224753	.1906288	-2.22	0.027	-.7961009	-.0488497
income_cat						
2	.0946656	.242739	0.39	0.697	-.3810941	.5704254
3	.5016051	.2215662	2.26	0.024	.0673433	.935867
4	.4315078	.2280086	1.89	0.058	-.0153808	.8783964
_cons	-.0483645	.2470234	-0.20	0.845	-.5325214	.4357924
European						
gender						
Male	.6949682	.2814506	2.47	0.014	.1433352	1.246601
income_cat						
2	.496032	.3314079	1.50	0.134	-.1535156	1.14558
3	.8082547	.2960338	2.73	0.006	.2280392	1.38847
4	.9865233	.2939615	3.36	0.001	.4103694	1.562677
_cons	-1.969437	.3723598	-5.29	0.000	-2.699249	-1.239625
Korean						
gender						
Male	.0257312	.4913433	0.05	0.958	-.937284	.9887464
income_cat						
2	-.638151	.5053717	-1.26	0.207	-1.628661	.3523593
3	-.8951082	.5090436	-1.76	0.079	-1.892815	.102599
4	-1.082247	.5441119	-1.99	0.047	-2.148687	-.0158072
_cons	-.8522904	.5679136	-1.50	0.133	-1.965381	.2607998

We now use margins to get average predicted probabilities for the different levels of `income_cat`.

```
. margins income_cat
Predictive margins                                Number of obs = 3,075
Model VCE: OIM
Expression: Pr(car|1 selected), predict()
```

	Delta-method				
	Margin	std. err.	z	P> z	[95% conf. interval]
<code>_outcome#</code>					
<code>income_cat</code>					
American#1	.4941932	.0345185	14.32	0.000	.4265382 .5618482
American#2	.4677756	.0369774	12.65	0.000	.3953013 .5402499
American#3	.385185	.0338328	11.38	0.000	.3188739 .4514961
American#4	.3850117	.0335619	11.47	0.000	.3192316 .4507918
Japanese#1	.3237821	.0328847	9.85	0.000	.2593293 .3882349
Japanese#2	.3375662	.0339412	9.95	0.000	.2710427 .4040897
Japanese#3	.4165363	.0340384	12.24	0.000	.3498223 .4832504
Japanese#4	.3891954	.0332832	11.69	0.000	.3239615 .4544294
European#1	.0971881	.0181104	5.37	0.000	.0616924 .1326838
European#2	.1506216	.0282786	5.33	0.000	.0951966 .2060467
European#3	.1698688	.0255855	6.64	0.000	.1197222 .2200154
European#4	.2021599	.0277523	7.28	0.000	.1477663 .2565534
Korean#1	.0848366	.0194847	4.35	0.000	.0466472 .123026
Korean#2	.0440366	.0150923	2.92	0.004	.0144562 .073617
Korean#3	.0284098	.0113046	2.51	0.012	.0062532 .0505665
Korean#4	.023633	.0103463	2.28	0.022	.0033546 .0439114

Rows are labeled first by the outcome category, the alternative that is chosen; and second by the value of `income_cat`. The first column in the body of the table gives predicted probabilities of the outcome. If we have a random or otherwise representative sample, these are the expected probabilities based on our model. The 0.494 next to `American#1` is the expected probability of a consumer with `income_cat = 1` buying an American car. Said differently, we expect 49.4% of individuals to buy American cars if they are in the first income quartile and have the same distribution of dealers and gender that we observe in the data.

The probability is computed by setting `income_cat = 1` for all persons, leaving the other variables unchanged, calculating the predicted probability for each outcome with these altered data, and then averaging the probabilities. And so on for the other levels of `income_cat`. So, to be precise, what is calculated by `margins` is an average predicted probability.

Examination of the table shows that as income increases, consumers are less likely to buy American cars and less likely to buy Korean cars but more likely to buy European cars. The association between Japanese cars and income categories is less clear. Admittedly, it would be easier to spot these changes if we plotted the probabilities instead of looking at the values reported in this table. To see a plot, you can simply type `marginsplot`, as we demonstrate in [CM] [Intro 1](#) and in [example 4](#) below.

We can test for the effect of `income_cat` on the probability of selecting each nationality of car. To get the joint test of any difference in expected probabilities across income levels for each outcome, we specify the option `contrast`.


```
. margins income_cat, contrast
Contrasts of predictive margins          Number of obs = 3,075
Model VCE: OIM
Expression: Pr(car|1 selected), predict()
```

	df	chi2	P>chi2
income_cat@outcome			
American	3	8.50	0.0368
Japanese	3	5.33	0.1493
European	3	11.93	0.0076
Korean	3	9.00	0.0293
Joint	9	27.52	0.0011

We see that the joint test of any difference across income levels has the smallest p -value for European cars (0.0076). The joint test of the effect of income is nonsignificant for Japanese cars.

We can use an *op.* contrast operator to see differences (contrasts). The contrast operators most typically used in this context are *r.*, differences from the reference level; *a.*, differences from the next level (adjacent contrasts); and *ar.*, differences from the previous level (reverse adjacent contrasts). See the *op.* table in [R] [contrast](#) for a list of all contrast operators.

Here we use the *ar.* contrast operator to estimate differences in expected probabilities between each level of `income_cat` and its previous level. We also specify the `outcome()` option to restrict the results to the probability of buying a European car.

```
. margins ar.income_cat, outcome(European)
Contrasts of predictive margins          Number of obs = 3,075
Model VCE: OIM
Expression: Pr(car|1 selected), predict()
Outcome:    European
```

	df	chi2	P>chi2
income_cat			
(2 vs 1)	1	2.42	0.1199
(3 vs 2)	1	0.25	0.6152
(4 vs 3)	1	0.73	0.3934
Joint	3	11.93	0.0076

	Delta-method		
	Contrast	std. err.	[95% conf. interval]
income_cat			
(2 vs 1)	.0534335	.0343615	-.0139138 .1207808
(3 vs 2)	.0192472	.0382871	-.0557941 .0942885
(4 vs 3)	.032291	.0378334	-.0418611 .1064432

From the second table, the line labeled (2 vs 1) estimates that the probability of selecting a European car increases by 0.053 when we go from the first income category to the second. The p -value for this difference is given in the first table on the line (2 vs 1), $p = 0.12$, and we see that this difference is not significant at the 5% level. In fact, none of these reverse adjacent contrasts are significant. The joint significance reported in the first table is 0.0076 and is, of course, the same as what was calculated by the previous `margins` command.

Instead of testing whether income has an effect on the expected probability of selecting one nationality of car, we might want to test whether the effects of income are different for different

nationalities of cars. The option `contrast(outcomecontrast(op._outcome))` can be used to get tests of the differences between outcomes in the differences of expected probabilities across income levels.

For instance, when we typed `margins income_cat`, we saw that the expected probability of selecting a Japanese car was 0.338 for the second income category and was 0.324 for the first income category, a difference of 0.014. If we look at the probabilities of selecting an American car, we get 0.468 for the second income category and 0.494 for the first income category, a difference of -0.026 . The differences of 0.014 and -0.026 have opposite signs, but are they statistically different from each other? We could ask the same question for differences in the third versus first income categories and for differences in the fourth versus first income categories. The `contrast(outcomecontrast())` option gives a joint test of all of these differences—a test of whether `income_cat` has the same effect on the probability of selecting a Japanese car as it has on the probability of selecting an American car.

We use the `r.` `contrast` operator to get differences between outcomes relative to the first level of the alternatives variable, which is `American`.

```
. margins income_cat, contrast(outcomecontrast(r._outcome))
Contrasts of predictive margins                Number of obs = 3,075
Model VCE: OIM
Expression: Pr(car|1 selected), predict()
```

	df	chi2	P>chi2
_outcome#income_cat			
(Japanese vs American) (joint)	3	8.07	0.0446
(European vs American) (joint)	3	13.16	0.0043
(Korean vs American) (joint)	3	2.28	0.5167
Joint	9	27.52	0.0011

To be clear about what is being tested: The test labeled “(Japanese vs American) (joint)” is a test of the null hypothesis,

$$\begin{aligned} \Pr(\text{Japanese}\#2) - \Pr(\text{Japanese}\#1) &= \Pr(\text{American}\#2) - \Pr(\text{American}\#1) \\ \Pr(\text{Japanese}\#3) - \Pr(\text{Japanese}\#1) &= \Pr(\text{American}\#3) - \Pr(\text{American}\#1) \\ \Pr(\text{Japanese}\#4) - \Pr(\text{Japanese}\#1) &= \Pr(\text{American}\#4) - \Pr(\text{American}\#1) \end{aligned}$$

where `Japanese#2` denotes the outcome of choosing a Japanese car for `income_cat = 2`, etc.

At the 5% level, there are significant differences in the effect of income on the probability of selecting Japanese versus American cars and in the probability of selecting European versus American cars.

We do not need to perform all of these tests at once. If we are interested only in testing the difference in effect of income on Japanese versus American outcomes, we can use the `r(1/2).` operator to restrict the outcome levels.

```
. margins income_cat, contrast(outcomecontrast(r(1/2)._out))
Contrasts of predictive margins                Number of obs = 3,075
Model VCE: OIM
Expression: Pr(car|1 selected), predict()
```

	df	chi2	P>chi2
_outcome#income_cat	3	8.07	0.0446

The result is the same as it was in the previous `margins` output for the test of Japanese versus American outcomes.



Estimating margins for alternative-specific variables

For alternative-specific variables, we can explore even more possibilities using `margins`. For instance, we can estimate the effect of changing the value of an alternative-specific variable at only one of the alternatives, or we could change its value across all alternatives. As we discussed earlier, even when an alternative-specific variable is changed only at one value of the alternatives, it creates changes in the predicted probabilities of selecting an outcome for all the possible outcomes. To handle this additional complexity, the option `alternative()` is extremely useful when we want to test hypotheses about alternative-specific variables that involve only one (or a subset) of the alternatives. We demonstrate this below.

► Example 2: Alternative-specific variables

We continue with the same `cmlogit` choice model on the nationality of car purchased. The model was fit with an alternative-specific variable `dealers`, which contains the number of dealerships of each nationality of car in the individual's community.

Would increasing the number of dealerships for a certain nationality of car affect the likelihood of more people buying that car? And at what nationality's expense? For example, if increasing the number of Korean dealerships increases the probability of buying a Korean car, then the probability of buying an American, a Japanese, or a European car must go down—and we would like to know which one has the biggest decrease in probability. (It is possible one of these could also go up, but because the changes in probabilities must sum to zero, one of the changes must be negative if the change for Korean cars is positive.)

`margins` can answer these questions—based on the fitted `cmlogit` model.

`margins` by default produces a lot of output as we discussed earlier. Here there are four outcomes in this model, one for each nationality of car. There are also four ways to change the alternative-specific variable `dealers`. We can change it just for Korean dealerships, or just for American dealerships, or Japanese, or European. Restricting the change to a particular alternative—in this case, Korean—is what we want.

We can run `margins` with two `at()` options, the first with the `dealers` set to the original value and the second with `dealers` increased by one. Because we specify the option `alternative(Korean)`, only `dealers` corresponding to the alternative Korean are increased.

```
. margins, at(dealers=generate(dealers)) at(dealers=generate(dealers+1))
> alternative(Korean)

Predictive margins                                Number of obs = 3,075
Model VCE: OIM

Expression: Pr(car|1 selected), predict()
Alternative: Korean

1._at: dealers = dealers
2._at: dealers = dealers+1
```

	Delta-method			P> z	[95% conf. interval]	
	Margin	std. err.	z			
_outcome#_at						
American#1	.4361949	.0167567	26.03	0.000	.4033524	.4690374
American#2	.4352739	.0167586	25.97	0.000	.4024276	.4681202
Japanese#1	.3665893	.0162132	22.61	0.000	.3348121	.3983666
Japanese#2	.3659044	.016203	22.58	0.000	.3341472	.3976617
European#1	.1508121	.0120258	12.54	0.000	.1272419	.1743822
European#2	.1505359	.0120085	12.54	0.000	.1269998	.1740721
Korean#1	.0464037	.0069344	6.69	0.000	.0328124	.059995
Korean#2	.0482858	.007271	6.64	0.000	.0340348	.0625367

These results indicate that if we add one Korean dealership in each community, we would expect the percentage of individuals purchasing a Korean car to go from 4.64% to 4.82%. Said another way, we expect the probability of buying a Korean car to increase from 0.0464 to 0.0482.

We run margins again with the `contrast(atcontrast(op))` option to estimate each of the differences in expected probabilities. Here `op` represents a [contrast operator](#). There are only two `at()`'s to contrast, so any pairwise operator will do. We use the `r` operator.

```
. margins, at(dealers=generate(dealers)) at(dealers=generate(dealers+1))
> alternative(Korean) contrast(atcontrast(r))

Contrasts of predictive margins                    Number of obs = 3,075
Model VCE: OIM

Expression: Pr(car|1 selected), predict()
Alternative: Korean

1._at: dealers = dealers
2._at: dealers = dealers+1
```

	df	chi2	P>chi2
_at@_outcome			
(2 vs 1) American	1	2.65	0.1036
(2 vs 1) Japanese	1	2.63	0.1051
(2 vs 1) European	1	2.49	0.1148
(2 vs 1) Korean	1	2.64	0.1041
Joint	3	2.67	0.4454

	Delta-method		
	Contrast	std. err.	[95% conf. interval]
_at@_outcome			
(2 vs 1) American	-.000921	.0005658	-.00203 .0001879
(2 vs 1) Japanese	-.0006849	.0004226	-.0015131 .0001433
(2 vs 1) European	-.0002761	.0001751	-.0006194 .0000671
(2 vs 1) Korean	.001882	.001158	-.0003875 .0041516

Based on these results, we expect the probability of buying a Korean car to increase by 0.0019. Most of this increase came at the expense of American and Japanese cars. We expect that the probability of buying an American car will decrease by 0.0009, and the probability of buying a Japanese car will decrease by 0.0007. The probability of buying a European car decreases only by 0.0003.



The altsubpop suboption for unbalanced choice sets

Not everyone in this dataset has Korean as one of his or her possible choices for a car. The choice sets are unbalanced. This can be seen by running `cmchoiceset`.

```
. cmchoiceset, generate(choiceset)
Tabulation of choice-set possibilities
```

Choice set	Freq.	Percent	Cum.
1 2 3	380	42.94	42.94
1 2 3 4	505	57.06	100.00
Total	885	100.00	

Note: Total is number of cases.

The value of 4 for the alternatives variable `car` represents Korean. So we see that 380 persons out of a total of 885 do not have Korean as a choice.

How does `margins` handle the cases in which a particular alternative is not part of the choice set for the case? By default, `margins` considers an alternative that is missing from a choice set to have zero probability of being chosen. This makes sense if we are looking at the change in buying a Korean car for a change in a variable like income. There are no Korean dealerships in the community, so even if income changes, there is still no way for a person in that community to purchase a Korean car (ignoring the possibility of buying a Korean car in another community).

The `outcome()` option of `margins` has a suboption `altsubpop`, which changes the way `margins` handles alternatives that are not present in the case's choice set. When `altsubpop` is specified, the results for each outcome are restricted to those cases that have that outcome in their choice set. Here is what we get when we use `altsubpop`.

```
. margins, at(dealers=generate(dealers)) at(dealers=generate(dealers+1))
> alternative(Korean) outcome(_all, altsubpop)
Predictive margins                                Number of obs = 3,075
Model VCE: OIM
Expression: Pr(car|1 selected), predict()
Alternative: Korean
1._at: dealers = dealers
2._at: dealers = dealers+1
```

_outcome#_at	Delta-method				[95% conf. interval]	
	Margin	std. err.	z	P> z		
American#1	.4361949	.0167567	26.03	0.000	.4033524	.4690374
American#2	.4352739	.0167586	25.97	0.000	.4024276	.4681202
Japanese#1	.3665893	.0162132	22.61	0.000	.3348121	.3983666
Japanese#2	.3659044	.016203	22.58	0.000	.3341472	.3976617
European#1	.1508121	.0120258	12.54	0.000	.1272419	.1743822
European#2	.1505359	.0120085	12.54	0.000	.1269998	.1740721
Korean#1	.0817996	.0122239	6.69	0.000	.0578412	.105758
Korean#2	.0851172	.0128173	6.64	0.000	.0599959	.1102386

We find that for the subpopulation of individuals who had Korean in their choice set, the expected probability of selecting a Korean car is 0.0818. For the same subpopulation, if we increase the number of Korean dealerships by 1, the probability of selecting a Korean car goes to 0.0851. We can use the `contrast()` option to estimate the effect.

```
. margins, at(dealers=generate(dealers)) at(dealers=generate(dealers+1))
> alternative(Korean) contrast(atcontrast(r)) outcome(_all, altsubpop)
Contrasts of predictive margins                                Number of obs = 3,075
Model VCE: OIM
Expression: Pr(car|1 selected), predict()
Alternative: Korean
1._at: dealers = dealers
2._at: dealers = dealers+1
```

	df	chi2	P>chi2
_at@_outcome			
(2 vs 1) American	1	2.65	0.1036
(2 vs 1) Japanese	1	2.63	0.1051
(2 vs 1) European	1	2.49	0.1148
(2 vs 1) Korean	1	2.64	0.1041
Joint	3	2.67	0.4454

	Delta-method		
	Contrast	std. err.	[95% conf. interval]
_at@_outcome			
(2 vs 1) American	-.000921	.0005658	-.00203 .0001879
(2 vs 1) Japanese	-.0006849	.0004226	-.0015131 .0001433
(2 vs 1) European	-.0002761	.0001751	-.0006194 .0000671
(2 vs 1) Korean	.0033176	.0020412	-.0006831 .0073184

The change in the expected probability of buying a Korean car is estimated at 0.0033 for this subpopulation, a considerable difference from the previous estimate of 0.0019. This is not surprising. The earlier estimate considered those cases without Korean in their choice set as having a zero probability of buying a Korean car even when we added a Korean dealer. The estimate of 0.0033 ignores those cases with zero probability.

Note also that changes in probabilities no longer sum to zero. The number of cases for each estimate varies, so we would not expect them to sum to zero. The estimates for the Korean outcome are only for the subpopulation of individuals who had Korean in their choice set, while all other estimates are for the full population.

Another way to handle unbalanced choice sets is to use the `subpop()` option (or the `over()` option) with `margins` and use an indicator variable for the different choice sets. We created such a variable and called it `choiceset` when we ran `cmchoiceset` earlier. See [CM] [cmchoiceset](#). Here is `margins` using `subpop()` to restrict the sample to the cases that have all four alternatives.

```

. margins, at(dealers=generate(dealers)) at(dealers=generate(dealers+1))
> alternative(Korean) contrast(atcontrast(r) nowald)
> subpop(if choicest=="1 2 3 4":choicest)

Contrasts of predictive margins                Number of obs   = 3,075
Model VCE: OIM                               Subpop. no. obs = 1,956

Expression: Pr(car|1 selected), predict()
Alternative: Korean

1._at: dealers = dealers
2._at: dealers = dealers+1
    
```

	Delta-method			
	Contrast	std. err.	[95% conf. interval]	
_at@_outcome (2 vs 1) American	-.0016235	.0009974	-.0035784	.0003313
(2 vs 1) Japanese	-.0012073	.0007449	-.0026672	.0002526
(2 vs 1) European	-.0004868	.0003087	-.0010918	.0001183
(2 vs 1) Korean	.0033176	.0020412	-.0006831	.0073184

We now have results only for those cases having all four choices, and the changes in probabilities sum to zero.

For observational data, the default behavior of `margins` is likely what you want. If an alternative was not in a decision maker's choice set, how can changing a covariate make it possible to choose that alternative? The assumption is that the reason the alternative is not in the choice set is that the alternative does not exist for that decision maker under any condition. If the choices for commuters are car, train, or bus, but there is no train service in a commuter's community, then the probability of that commuter taking a train to work is zero.

Imagine, however, a different type of study in which, by design, individuals were not offered a particular alternative. Suppose, for example, a marketer is testing consumer preferences among six different types of breakfast cereal. He or she thinks that giving each consumer six different cereals to taste would be overwhelming. So the marketer gives each consumer only four cereals. Is it reasonable to keep the probability of picking a cereal not offered fixed at zero when looking at estimates for the entire sample? If it had been offered to someone to whom it was not, he or she might have chosen it. Using `altsubpop` in this case seems not only reasonable but also, perhaps, essential. We want to make comparisons only among those alternatives that persons were able to choose between.

More on unbalanced choice sets

When we are looking at changing the number of Korean dealerships to see its effect on buying cars of different nationalities, neither the default, the `altsubpop` suboption, nor `subpop()` really does what we want. We want to increase the number of Korean dealerships in all communities, including those who currently do not have any Korean dealerships in their community. We now show how this could be done.

First, we use the command `expand` to add observations to the cases that do not have Korean (`car = 4`) in their choice set. We generate a variable, `new`, that flags the newly created observations. See [\[D\] expand](#) for details.

```

. expand 2 if choicest=="1 2 3":choicest & car==3, gen(new)
(380 observations created)
    
```

Second, for new observations, we set the alternatives variable `car` equal to 4 (representing Korean) and `dealers` equal to 0. Then, we run `cmchoicest` to confirm we did what we wanted.

```
. replace car = 4 if new==1
(380 real changes made)
. replace dealers = 0 if new==1
(380 real changes made)
. cmchoiceset
Tabulation of choice-set possibilities
```

Choice set	Freq.	Percent	Cum.
1 2 3 4	885	100.00	100.00
Total	885	100.00	

Note: Total is number of cases.

We can now estimate the probability of selecting each nationality of car after adding a Korean dealership to all communities. We run `margins` with the `noesample` option because we are now doing predictions outside the estimation sample.

```
. margins, at(dealers=generate(dealers+1)) alternative(Korean) noesample
Predictive margins                                Number of obs = 3,448
Model VCE: OIM
Expression: Pr(car|1 selected), predict()
Alternative: Korean
At: dealers = dealers+1
```

	Delta-method				
	Margin	std. err.	z	P> z	[95% conf. interval]
_outcome					
American	.4177746	.0169372	24.67	0.000	.3845782 .450971
Japanese	.3532169	.0160733	21.98	0.000	.3217138 .38472
European	.1454895	.0117044	12.43	0.000	.1225493 .1684298
Korean	.0835189	.0124662	6.70	0.000	.0590857 .1079522

We find that if we add one Korean dealership in each community, including those that had no dealerships originally, the expected probability of selecting a Korean car is 0.0835. When we used `altsubpop` and considered only the subpopulation of communities that had a Korean car in their choice set, this expected probability was just slightly larger, 0.0851.

► Example 3: margins, contrast

Testing contrasts with `margins` after the `cm` estimation commands can be overwhelming, especially with alternative-specific variables, because there are so many possibilities. Contrasts can be made between different outcomes (for both case-specific and alternative-specific variables). Contrasts can also be made between the alternatives at which alternative-specific variables are changed.

This example explains how you can read the output from `margins` to understand exactly what is being tested. For those familiar with manually defining and testing contrasts, we also show you how to find the contrast matrix to see exactly what `margins` is testing for any contrast it performs. In this example, we use `margins` after `cmmprobit` but recall what we said earlier. All the special options of `margins` for use after `cm` estimation commands work in the same way after `cmlogit`, `cmmprobit`, `cmroprobit`, `cmmixlogit`, and `cmxtmixlogit`. So this example applies to all of these commands.

We use data from [example 1](#) in [\[CM\] cmmprobit](#). The data represent individuals' choices of travel mode: `air`, `train`, `bus`, or `car`. There are two alternative-specific variables: `travelcost`, a measure of generalized cost of travel; and `termtime`, time spent in the terminal. The variable `income` gives household income and is case specific.

We load the data and `cmset` them. The command `xtile` is used to make a categorical variable, `cost_cat`, that contains tertiles of `travelcost`. This is the alternative-specific categorical variable we will use with `margins`. We then fit our `cmmprobit` model.

```
. use https://www.stata-press.com/data/r17/travel, clear
(Modes of travel)
. cmset id mode
      Case ID variable: id
      Alternatives variable: mode
. xtile cost_cat = travelcost, nquantiles(3)
. cmmprobit choice i.cost_cat termtime, casevars(income)
note: variable 2.cost_cat has 70 cases that are not alternative-specific;
      there is no within-case variability.
note: variable 3.cost_cat has 113 cases that are not alternative-specific;
      there is no within-case variability.
      (iteration log omitted)
Multinomial probit choice model                Number of obs      =      840
Case ID variable: id                          Number of cases    =      210
Alternatives variable: mode                   Alts per case: min =       4
                                                avg =      4.0
                                                max =       4
Integration sequence:      Hammersley
Integration points:        601                Wald chi2(6)       =      36.54
Log simulated-likelihood = -190.38007         Prob > chi2        =      0.0000
```

choice	Coefficient	Std. err.	z	P> z	[95% conf. interval]	
mode						
cost_cat						
2	.2155852	.2214832	0.97	0.330	-.2185139	.6496843
3	-.3822223	.2805587	-1.36	0.173	-.9321073	.1676628
termtime	-.043736	.008679	-5.04	0.000	-.0607466	-.0267255
Air						
(base alternative)						
Train						
income	-.0340284	.0092383	-3.68	0.000	-.0521352	-.0159216
_cons	.4632755	.3902352	1.19	0.235	-.3015715	1.228122
Bus						
income	-.0136801	.008338	-1.64	0.101	-.0300223	.0026622
_cons	-.1941561	.4585625	-0.42	0.672	-1.092922	.7046098
Car						
income	-.0039416	.0082461	-0.48	0.633	-.0201036	.0122204
_cons	-2.100911	.7742118	-2.71	0.007	-3.618338	-.5834833
/ln12_2	-.3689019	.3215014	-1.15	0.251	-.999033	.2612293
/ln13_3	-.4548538	.3229243	-1.41	0.159	-1.087774	.1780661
/12_1	1.075525	.2157021	4.99	0.000	.6527562	1.498293
/13_1	.9613768	.2451866	3.92	0.000	.4808199	1.441934
/13_2	.6096931	.2902669	2.10	0.036	.0407804	1.178606

```
(mode=Air is the alternative normalizing location)
(mode=Train is the alternative normalizing scale)
. estimates store ourmodel
```

Sometimes, we want to test a single, simple hypothesis. For instance, we could test whether changing the cost of the train alternative affects the probability of selecting the air outcome. We can request this test and also estimate the differences in expected probabilities across cost categories by typing

```
. margins r.cost_cat, alternative(Train) outcome(Air)
Contrasts of predictive margins                                Number of obs = 840
Model VCE: OIM
Expression: Pr(mode), predict()
Alternative: Train
Outcome:    Air
```

	df	chi2	P>chi2
cost_cat			
(2 vs 1)	1	1.00	0.3184
(3 vs 1)	1	1.52	0.2176
Joint	2	6.49	0.0390

	Delta-method		
	Contrast	std. err.	[95% conf. interval]
cost_cat			
(2 vs 1)	-.0153519	.015386	-.045508 .0148042
(3 vs 1)	.0226503	.0183692	-.0133527 .0586533

The joint test of the effect of `cost_cat` for the train alternative on the probability of selecting air travel is significant, with a *p*-value of 0.0390.

What if we wanted to perform all such tests—all tests of the effect of `cost_cat` for each alternative on the expected probabilities of each outcome? We specify the contrast option with margins:

```
. margins cost_cat, contrast
Contrasts of predictive margins                                Number of obs = 840
Model VCE: OIM
Expression: Pr(mode), predict()
```

	df	chi2	P>chi2
cost_cat@_outcome#mode			
Air Air	2	15.45	0.0004
Air Train	2	6.49	0.0390
Air Bus	2	5.59	0.0612
Air Car	2	10.96	0.0042
Train Air	2	6.92	0.0315
Train Train	2	9.75	0.0076
Train Bus	2	8.87	0.0119
Train Car	2	7.11	0.0286
Bus Air	2	5.68	0.0584
Bus Train	2	8.36	0.0153
Bus Bus	2	10.45	0.0054
Bus Car	2	7.77	0.0206
Car Air	2	11.09	0.0039
Car Train	2	6.64	0.0361
Car Bus	2	7.47	0.0238
Car Car	2	11.85	0.0027
Joint	14	307.62	0.0000

Now each row of the output gives a test of the effect of `cost_cat` on the probability of selecting each outcome by each alternative, but we have to look carefully to understand which hypothesis is being tested in each line. Let's look at the second row of the output. This tests the null hypothesis that the cost of the `train` alternative has no effect on the probability of selecting the `air` outcome. This is the same hypothesis we tested with our previous `margins` command.

Having run the previous `margins` command, we could easily spot the row in this output that tested the same hypothesis. But if we ran only the `margins cost_cat, contrast` command, how would we determine what the hypothesis on a given row is? Recall what was said earlier about the use of “outcome” and “alternative” in `margins` specifications. The key for the labels on the table is `cost_cat@_outcome#mode`. The first part of the key, `cost_cat@`, means we are testing differences across `cost_cat`. The second part of the key, `_outcome#mode`, is where the differences are being tested. `_outcome` is the alternative hypothetically chosen. `mode`, which is the alternatives variable, gives the alternative at which the value of `cost_cat` is being changed.

The joint test shown in the last row is a test of the null hypothesis: within each outcome by alternative group, expected probabilities across levels of `cost_cat` are the same; that is, `cost_cat` has no effect anywhere.

We can duplicate the results manually using the `test` command and the contrast coefficients that `margins` uses. `margins` stores in `r(L)` the matrix of contrasts that are tested.

```
. matrix list r(L)
r(L) [32,48]
           1._outcome#  1._outcome#           1._outcome#
                2.mode#    2.mode#                2.mode#
...  1.cost_cat  2.cost_cat  ...  3.cost_cat
...
2.cost_cat@
1._outcome#
  2.mode           -1           1           0
...
3.cost_cat@
1._outcome#
  2.mode           -1           0           1
```

It is a huge matrix. We will again focus on the two-degrees-of-freedom test reported in the second line of the previous `margins` output that is labeled `air train`, so we show only the relevant portion of `r(L)` here. Our test is based on rows of this matrix that include the `@1.outcome#2.mode` in the row label. We know this because our alternatives variable `mode` is coded as 1 for `air` (our `_outcome`) and 2 for `train` (our `mode`). The nonzero elements within each row define a single contrast. In the first row of `r(L)` that we displayed here, it shows the contrast of the expected probabilities for the first and second levels of `cost_cat`. The contrast in the second row compares the expected probabilities for the first and third levels of `cost_cat`. The column labels show the syntax that we can later use to perform our own test.

We now run `margins` with the `post` option to save the results of `margins` as if it were an estimation command.

```
. margins cost_cat, post
Predictive margins                                Number of obs = 840
Model VCE: OIM
Expression: Pr(mode), predict()

-----+-----
|                                     Delta-method
|      Margin      std. err.      z    P>|z|    [95% conf. interval]
-----+-----
|  _outcome#
|  mode#
|  cost_cat
|  Air#Air#1      .2713238   .0427612    6.35  0.000    .1875135   .3551341
|  Air#Air#2      .3165716   .0324305    9.76  0.000    .2530089   .3801342
|  Air#Air#3      .2004557   .0340447    5.89  0.000    .1337293   .2671821
|  Air#Train#1    .2764299   .0300901    9.19  0.000    .2174544   .3354053
|  Air#Train#2    .261078   .0283491    9.21  0.000    .2055147   .3166413
|  Air#Train#3    .2990802   .0292282   10.23  0.000    .2417939   .3563665
|
|  (output omitted)
-----+-----

. estimates store ourmargins
```

Our test for an effect of the cost of the train travel on the probability of selecting the air travel is just a test for a difference in the expected probabilities labeled `air#train#1`, `air#train#2`, and `air#train#3` above.

We use `test` with the syntax we saw in `r(L)`.

```
. test (1._outcome#2.mode#1.cost_cat = 1._outcome#2.mode#2.cost_cat)
>      (1._outcome#2.mode#1.cost_cat = 1._outcome#2.mode#3.cost_cat)

( 1) 1bn._outcome#2.mode#1bn.cost_cat - 1bn._outcome#2.mode#2.cost_cat = 0
( 2) 1bn._outcome#2.mode#1bn.cost_cat - 1bn._outcome#2.mode#3.cost_cat = 0

      chi2( 2) =      6.49
      Prob > chi2 =      0.0390
```

We duplicated the second row of output from `margins cost_cat, contrast`.

4

The `outcomecontrast()` and `alternativecontrast()` suboptions

`margins` has two other options, `contrast(outcomecontrast(op))` and `contrast(alternativecontrast(op))`, that perform joint tests of hypothesis after fitting choice models. We use `contrast(outcomecontrast(op))` to test for differences across outcomes; we demonstrate this below. We use `contrast(alternativecontrast(op))` to test for differences across alternatives.

Continuing our example, we use the `contrast(outcomecontrast(op))` option to test for differences in the effects of `cost_cat` across outcomes. Before we can run `margins`, we must get our `cmmprobit` estimation results back (because the estimation results currently active are those from `margins`).

```
. estimates restore ourmodel
(results ourmodel are active now)
```

Now we can test whether changing the cost of the `air` travel alternative has the same effect on the probability of selecting the `train` outcome as it has on the probability of selecting the `air` outcome.

```
. margins r.cost_cat, alternative(Air) outcome(Train Air)
> contrast(outcomecontrast(r))

Contrasts of predictive margins                                Number of obs = 840
Model VCE: OIM
Expression: Pr(mode), predict()
Alternative: Air
```

	df	chi2	P>chi2	
_outcome#cost_cat				
(Train vs Air) (2 vs 1)	1	1.01	0.3157	
(Train vs Air) (3 vs 1)	1	1.87	0.1711	
Joint	2	14.12	0.0009	

	Delta-method			
	Contrast	std. err.	[95% conf. interval]	
_outcome#cost_cat				
(Train vs Air) (2 vs 1)	-.058856	.0586622	- .1738319	.0561199
(Train vs Air) (3 vs 1)	.0912034	.0666299	- .0393888	.2217957

There is a lot going on in this `margins` command. By specifying `r.cost_cat`, we requested differences in expected probabilities when comparing levels of the `cost_cat` variable. The `alternative(Air)` option tells `margins` that we want to estimate these differences only when changing the cost of the `air` alternative. The `outcome(Train Air)` option specifies that we want to estimate only these differences in expected probabilities of selecting the `train` outcome and the `air` outcome. Finally, `contrast(outcomecontrast(r))` says that we want to test whether the differences are the same for `train` travel and for `air` travel. Thus, we are testing whether the effect of the cost of `air` travel is the same on the probability of selecting `train` travel as it is on the probability of selecting `air` travel.

In the output from this command, the contrast labeled `(train vs air) (2 vs 1)` is the difference in the effect of changing the `cost_cat` of `air` travel from 1 to 2 on the probability of selecting `train` versus `air` travel. The p -value reported in the first table for this test is 0.3157. We do not have evidence that changing the cost of `air` travel from the first tertile to the second has different effects on the probabilities of selecting `train` and `air` travel. Similarly, looking at the lines labeled `(train vs air) (3 vs 1)`, we find no evidence that the effect of changing the cost of `air` travel from the first tertile to the third tertile has different effects on the probabilities of selecting `train` and `air` travel.

The joint test is provided in the last line of the top table in the output. With a p -value of 0.0009, we reject the null hypothesis that the effects of the `cost_cat` of `air` travel on the probability of selecting `train` travel are the same as the effects of the `cost_cat` of `air` travel on the probability of selecting `air` travel.

If we are interested in all tests comparing the effects of the `cost_cat` of one alternative on the probabilities of selecting two different outcomes, we can run `margins` again but without the `alternative()` and `outcome()` options.

```
. margins cost_cat, contrast(outcomecontrast(r))
Contrasts of predictive margins                                Number of obs = 840
Model VCE: OIM
Expression: Pr(mode), predict()
```

	df	chi2	P>chi2
_outcome#cost_cat@mode			
(Train vs Air) (joint) Air	2	14.12	0.0009
(Train vs Air) (joint) Train	2	9.73	0.0077
(Train vs Air) (joint) Bus	2	3.96	0.1379
(Train vs Air) (joint) Car	2	1.13	0.5683
(Bus vs Air) (joint) Air	2	14.88	0.0006
(Bus vs Air) (joint) Train	2	0.88	0.6440
(Bus vs Air) (joint) Bus	2	10.42	0.0055
(Bus vs Air) (joint) Car	2	1.23	0.5401
(Car vs Air) (joint) Air	2	15.06	0.0005
(Car vs Air) (joint) Train	2	1.47	0.4793
(Car vs Air) (joint) Bus	2	3.30	0.1918
(Car vs Air) (joint) Car	2	12.59	0.0018
Joint	14	1112.36	0.0000

The first line in the output matches the joint test reported in our previous `margins` command. The interpretations of the remaining rows are similar. For instance, in the second row, `(train vs air) (joint) train`, we test whether the effects of the cost of `train` travel on the probability of selecting `train` travel are the same as the effects of the cost of `train` travel on the probability of selecting `air` travel.

Again, we can list the `r(L)` matrix to see how the contrasts for each of these joint tests were formulated.

```
. matrix list r(L)
(output omitted)
```

We restore the `margins` estimation results and run `test` using the formulation of the contrasts we saw in `r(L)`.

```
. estimates restore ourmargins
(results ourmargins are active now)
> test ( 1._outcome#1.mode#1.cost_cat - 1._outcome#1.mode#2.cost_cat
>       = 2._outcome#1.mode#1.cost_cat - 2._outcome#1.mode#2.cost_cat )
>       ( 1._outcome#1.mode#1.cost_cat - 1._outcome#1.mode#3.cost_cat
>       = 2._outcome#1.mode#1.cost_cat - 2._outcome#1.mode#3.cost_cat )
( 1) 1bn._outcome#1bn.mode#1bn.cost_cat - 1bn._outcome#1bn.mode#2.cost_cat -
2._outcome#1bn.mode#1bn.cost_cat + 2._outcome#1bn.mode#2.cost_cat = 0
( 2) 1bn._outcome#1bn.mode#1bn.cost_cat - 1bn._outcome#1bn.mode#3.cost_cat -
2._outcome#1bn.mode#1bn.cost_cat + 2._outcome#1bn.mode#3.cost_cat = 0
      chi2( 2) = 14.12
      Prob > chi2 = 0.0009
```

We have duplicated the first row of the output from `margins`. Its interpretation is now clear. It is a test of the null hypothesis that for the alternative `air` (which is `mode = 1`), differences in predicted probabilities across levels of `cost_cat` are the same for the outcome `train` as they are for the outcome `air`. We say “for the alternative `air`”, meaning that the observations corresponding to the alternative `air` are the observations where `cost_cat` is set to 1, 2, or 3, and predicted probabilities are calculated for these values. At the observations corresponding to other alternatives, `cost_cat` is kept at its observed values. We could set `cost_cat` to its mean (or median, etc.) at these other alternatives using the option `at((mean) cost_cat)`. See [\[R\] margins](#).

We can follow the same steps after using the `contrast(alternativecontrast(op))` option with `margins` to be sure we understand what contrasts are being tested there.

Graphing margins results

After any `margins` command, you can use `marginsplot` to create a graph of the estimated probabilities or contrasts. See [R] [marginsplot](#) for information on using this command.

▷ Example 4: marginsplot

Here we give an example of using `marginsplot` after `margins` to graph the expected probabilities of selecting four outcomes across a range of values of a continuous case-specific variable. The model was fit using `cmmprobit`, but this example is perfectly general and works after any command that supports `margins`.

We continue with the data from [example 1](#) in [CM] `cmmprobit` and use the model fit in the [previous example](#).

We want to see how income affects the choices, controlling for travel cost and terminal time. The range of `income` is 2 to 72, and we call `margins` with an `at()` option that calculates probabilities at many points over the full range of income.

```
. margins, at(income=(2 5(5)70 72))
Predictive margins                                Number of obs = 840
Model VCE: OIM
Expression: Pr(mode), predict()
1._at: income = 2
      (output omitted)
16._at: income = 72
```

	Delta-method		z	P> z	[95% conf. interval]	
	Margin	std. err.				
<code>_outcome#_at</code>						
Air# 1	.1721153	.0456076	3.77	0.000	.0827261	.2615045
Air# 2	.1827666	.0443583	4.12	0.000	.0958259	.2697072
(output omitted)						
Car#15	.4052022	.0722242	5.61	0.000	.2636454	.5467589
Car#16	.4100206	.0758821	5.40	0.000	.2612945	.5587467

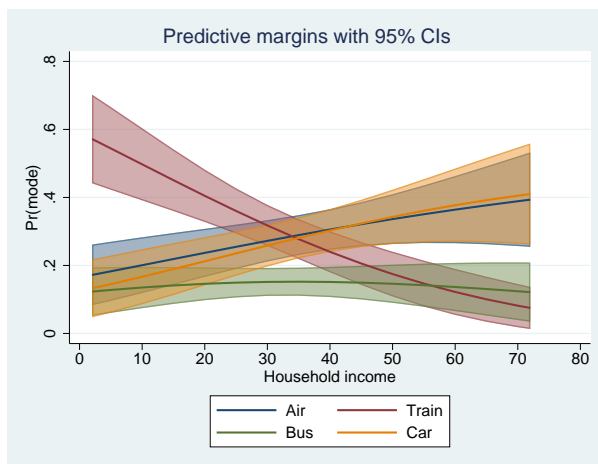
We could simply type

```
. marginsplot
```

to visualize the results.

Here we include the option `recast(line)` to smooth the plotting of the lines and the option `recastci(rarea)` to make the confidence intervals curves with shaded fill. The option `ciopts(color(%50))` makes the fill of the confidence intervals semitransparent.

```
. marginsplot, xlabel(0(10)80) recast(line) recastci(rarea)
> ciopts(color(%50)) plotopts(lwidth(medthick))
Variables that uniquely identify margins: income _outcome
```



From the graph, we see that the expected probability of choosing air travel increases with increasing income. The probability of choosing car travel also increases with increasing income. In fact, its probability is almost the same as the probability for air travel at all values of income. The probability of choosing bus travel changes little by income. The probability of choosing train travel has the biggest change over the range of income. At `income = 2`, the expected probability of choosing train travel is 54%. At `income = 72`, the expected probability of choosing train travel is only 9%.

◀

For more examples of `marginsplot` after CM commands, see [\[CM\] Intro 1](#).

Stored results

In addition to the results shown in [\[R\] margins](#), `margins` after `cm` estimators stores the following in `r()`:

Scalars

`r(k_alt)` number of levels of alternatives variable

Macros

`r(altvar)` name of alternatives variable
`r(alt#)` #th level of alternatives variable

Matrices

`r(altvals)` vector containing levels of alternatives variable

`margins` with the `post` option also stores the following in `e()`:

Scalars

`e(k_alt)` number of levels of alternatives variable

Macros

`e(altvar)` name of alternatives variable
`e(alt#)` #th level of alternatives variable

Matrices

`e(altvals)` vector containing levels of alternatives variable

Also see

- [R] [contrast](#) — Contrasts and linear hypothesis tests after estimation
- [R] [margins, contrast](#) — Contrasts of margins
- [R] [margins, pwcompare](#) — Pairwise comparisons of margins
- [R] [margins postestimation](#) — Postestimation tools for margins
- [U] [20 Estimation and postestimation commands](#)