

# Package ‘marginaleffects’

May 8, 2024

**Title** Predictions, Comparisons, Slopes, Marginal Means, and Hypothesis Tests

**Version** 0.20.1

**Description** Compute and plot predictions, slopes, marginal means, and comparisons (contrasts, risk ratios, odds, etc.) for over 100 classes of statistical and machine learning models in R. Conduct linear and non-linear hypothesis tests, or equivalence tests. Calculate uncertainty estimates using the delta method, bootstrapping, or simulation-based inference.

**License** GPL ( $\geq 3$ )

**Copyright** inst/COPYRIGHTS

**Encoding** UTF-8

**URL** <https://marginaleffects.com/>

**BugReports** <https://github.com/vincentarelbundock/marginaleffects/issues>

**RoxygenNote** 7.3.1

**Depends** R ( $\geq 3.6.0$ )

**Imports** checkmate, data.table, generics, insight ( $\geq 0.19.7$ ), methods, rlang, utils, Rcpp ( $\geq 1.0.0$ )

**Enhances** knitr

**LinkingTo** Rcpp, RcppEigen

**Suggests** AER, Amelia, afex, aod, bench, betareg, BH, bife, biglm, blme, boot, brglm2, brms, brmsmargins, broom, car, carData, causaldata, clarify, cjoint, collapse, conflicted, countrycode, covr, crch, DALEXtra, DCchoice, dbarts, distributional, dfix, dplyr, emmeans, equivalence, estimatr, fixest, flexsurv, fmeffects, fontquiver, future, future.apply, fw, gam, gamlss, gamlss.dist, geopack, ggdist, ggokabeito, ggplot2, ggrepel, glmmTMB, glmx, haven, here, itsadug, ivreg, kableExtra, lme4, lmerTest, logistf, magrittr, MatchIt, MASS, mclogit, MCMCglmm, missRanger, mgcv, mice, miceadds, mlogit, mlr3verse, modelbased, modelsummary, nlme, nnet, numDeriv, optmatch, ordinal, parameters, parsnip, partykit, patchwork, pkgdown, phylolm, plm, polyspline, poorman, posterior, pscl, purrr,

quantreg, Rchoice, REndo, rcmdcheck, remotes, rmarkdown, rms, robust, robustbase, robustlmm, rsample, rstanarm, rstantools, rsvg, sampleSelection, sandwich, scam, spelling, speedglm, survey, survival, svglite, systemfonts, tibble, tictoc, tidymodels, tidyr, tidyverse, tinysnapshot, tinytable (>= 0.2.0), tinytest, titanic, truncreg, tsModel, withr, workflows, yaml, xgboost, testthat (>= 3.0.0), altdoc

**Collate** 'RcppExports.R' 'backtransform.R' 'bootstrap\_boot.R' 'bootstrap\_fwb.R' 'bootstrap\_rsampl.R' 'bootstrap\_simulation.R' 'broom.R' 'by.R' 'ci.R' 'comparisons.R' 'complete\_levels.R' 'conformal.R' 'datagrid.R' 'deprecated.R' 'equivalence.R' 'get\_coef.R' 'get\_contrast\_data.R' 'get\_contrast\_data\_character.R' 'get\_contrast\_data\_factor.R' 'get\_contrast\_data\_logical.R' 'get\_contrast\_data\_numeric.R' 'get\_contrasts.R' 'get\_group\_names.R' 'get\_hypothesis.R' 'get\_jacobian.R' 'get\_model\_matrix.R' 'get\_model\_matrix\_attribute.R' 'get\_modeldata.R' 'get\_predict.R' 'get\_se\_delta.R' 'get\_term\_labels.R' 'get\_vcov.R' 'github\_issue.R' 'hush.R' 'hypotheses.R' 'hypotheses\_joint.R' 'hypothesis\_helper.R' 'imputation.R' 'inferences.R' 'mean\_or\_mode.R' 'methods.R' 'set\_coef.R' 'methods\_AER.R' 'methods\_MASS.R' 'methods\_MCMCglimm.R' 'methods\_Rchoice.R' 'methods\_afex.R' 'methods\_aod.R' 'methods\_betareg.R' 'methods\_bife.R' 'methods\_biglm.R' 'methods\_nnet.R' 'methods\_brglm2.R' 'sanity\_model.R' 'methods\_brms.R' 'methods\_crch.R' 'methods\_dataframe.R' 'methods\_dbarts.R' 'methods\_fixest.R' 'methods\_flexsurv.R' 'methods\_gamlss.R' 'methods\_glimmTMB.R' 'methods\_glmx.R' 'methods\_lme4.R' 'methods\_mclogit.R' 'methods\_mgcv.R' 'methods\_mhurdle.R' 'methods\_mlm.R' 'methods\_mlogit.R' 'methods\_mlr3.R' 'methods\_nlme.R' 'methods\_ordinal.R' 'methods\_plm.R' 'methods\_psc1.R' 'methods\_quantreg.R' 'methods\_rms.R' 'methods\_robustlmm.R' 'methods\_rstanarm.R' 'methods\_sampleSelection.R' 'methods\_scam.R' 'methods\_stats.R' 'methods\_survey.R' 'methods\_survival.R' 'methods\_tidymodels.R' 'methods\_tobit1.R' 'modelarchive.R' 'myTryCatch.R' 'package.R' 'plot.R' 'plot\_build.R' 'plot\_comparisons.R' 'plot\_predictions.R' 'plot\_slopes.R' 'posterior\_draws.R' 'predictions.R' 'print.R' 'recall.R' 'sanitize\_comparison.R' 'sanitize\_condition.R' 'sanitize\_conf\_level.R' 'sanitize\_hypothesis.R' 'sanitize\_interaction.R' 'sanitize\_newdata.R' 'sanitize\_numderiv.R' 'sanitize\_type.R' 'sanitize\_variables.R' 'sanitize\_vcov.R' 'sanity.R' 'sanity\_by.R' 'sanity\_dots.R' 'settings.R' 'slopes.R' 'sort.R' 'tinytest.R' 'type\_dictionary.R' 'unpack\_matrix\_cols.R' 'utils.R'

**Language** en-US

**Config/testthat/edition** 3

**NeedsCompilation** yes

**Author** Vincent Arel-Bundock [aut, cre, cph]

(<<https://orcid.org/0000-0003-2042-7063>>),

Noah Greifer [ctb] (<<https://orcid.org/0000-0003-3067-7154>>),

Etienne Bacher [ctb] (<<https://orcid.org/0000-0002-9271-5075>>)

**Maintainer** Vincent Arel-Bundock <vincent.arel-bundock@umontreal.ca>

**Repository** CRAN

**Date/Publication** 2024-05-08 12:10:03 UTC

## R topics documented:

comparisons . . . . .	3
datagrid . . . . .	16
hypotheses . . . . .	19
inferences . . . . .	25
plot_comparisons . . . . .	28
plot_predictions . . . . .	32
plot_slopes . . . . .	35
posterior_draws . . . . .	39
predictions . . . . .	40
print.marginaleffects . . . . .	51
slopes . . . . .	52
specify_hypothesis . . . . .	63

<b>Index</b>	<b>64</b>
--------------	-----------

---

comparisons	<i>Comparisons Between Predictions Made With Different Regressor Values</i>
-------------	---

---

### Description

Predict the outcome variable at different regressor values (e.g., college graduates vs. others), and compare those predictions by computing a difference, ratio, or some other function. `comparisons()` can return many quantities of interest, such as contrasts, differences, risk ratios, changes in log odds, lift, slopes, elasticities, etc.

- `comparisons()`: unit-level (conditional) estimates.
- `avg_comparisons()`: average (marginal) estimates.

`variables` identifies the focal regressors whose "effect" we are interested in. `comparison` determines how predictions with different regressor values are compared (difference, ratio, odds, etc.). The `newdata` argument and the `datagrid()` function control where statistics are evaluated in the predictor space: "at observed values", "at the mean", "at representative values", etc.

See the comparisons vignette and package website for worked examples and case studies:

- <https://marginaleffects.com/vignettes/comparisons.html>
- <https://marginaleffects.com/>

**Usage**

```
comparisons(  
  model,  
  newdata = NULL,  
  variables = NULL,  
  comparison = "difference",  
  type = NULL,  
  vcov = TRUE,  
  by = FALSE,  
  conf_level = 0.95,  
  transform = NULL,  
  cross = FALSE,  
  wts = FALSE,  
  hypothesis = NULL,  
  equivalence = NULL,  
  p_adjust = NULL,  
  df = Inf,  
  eps = NULL,  
  numberiv = "fdforward",  
  ...  
)
```

```
avg_comparisons(  
  model,  
  newdata = NULL,  
  variables = NULL,  
  type = NULL,  
  vcov = TRUE,  
  by = TRUE,  
  conf_level = 0.95,  
  comparison = "difference",  
  transform = NULL,  
  cross = FALSE,  
  wts = FALSE,  
  hypothesis = NULL,  
  equivalence = NULL,  
  p_adjust = NULL,  
  df = Inf,  
  eps = NULL,  
  numberiv = "fdforward",  
  ...  
)
```

**Arguments**

model	Model object
newdata	Grid of predictor values at which we evaluate the comparisons.

- Warning: Avoid modifying your dataset between fitting the model and calling a `marginalEffects` function. This can sometimes lead to unexpected results.
- NULL (default): Unit-level contrasts for each observed value in the dataset (empirical distribution). The dataset is retrieved using `insight::get_data()`, which tries to extract data from the environment. This may produce unexpected results if the original data frame has been altered since fitting the model.
- data frame: Unit-level contrasts for each row of the `newdata` data frame.
- string:
  - "mean": Contrasts at the Mean. Contrasts when each predictor is held at its mean or mode.
  - "median": Contrasts at the Median. Contrasts when each predictor is held at its median or mode.
  - "marginalmeans": Contrasts at Marginal Means.
  - "tukey": Contrasts at Tukey's 5 numbers.
  - "grid": Contrasts on a grid of representative numbers (Tukey's 5 numbers and unique values of categorical predictors).
- `datagrid()` call to specify a custom grid of regressors. For example:
  - `newdata = datagrid(cyl = c(4, 6))`: `cyl` variable equal to 4 and 6 and other regressors fixed at their means or modes.
  - `newdata = datagrid(mpg = fivenum)`: `mpg` variable held at Tukey's five numbers (using the `fivenum` function), and other regressors fixed at their means or modes.
  - See the Examples section and the `datagrid` documentation.
- `subset()` call with a single argument to select a subset of the dataset used to fit the model, ex: `newdata = subset(treatment == 1)`
- `dplyr::filter()` call with a single argument to select a subset of the dataset used to fit the model, ex: `newdata = filter(treatment == 1)`

## variables

## Focal variables

- NULL: compute comparisons for all the variables in the model object (can be slow).
- Character vector: subset of variables (usually faster).
- Named list: names identify the subset of variables of interest, and values define the type of contrast to compute. Acceptable values depend on the variable type:
  - Factor or character variables:
    - \* "reference": Each factor level is compared to the factor reference (base) level
    - \* "all": All combinations of observed levels
    - \* "sequential": Each factor level is compared to the previous factor level
    - \* "pairwise": Each factor level is compared to all other levels
    - \* "minmax": The highest and lowest levels of a factor.

- \* "revpairwise", "revreference", "revsequential": inverse of the corresponding hypotheses.
- \* Vector of length 2 with the two values to compare.
- \* Data frame with the same number of rows as newdata, with two columns of "lo" and "hi" values to compare.
- \* Function that accepts a vector and returns a data frame with two columns of "lo" and "hi" values to compare. See examples below.
- Logical variables:
  - \* NULL: contrast between TRUE and FALSE
  - \* Data frame with the same number of rows as newdata, with two columns of "lo" and "hi" values to compare.
  - \* Function that accepts a vector and returns a data frame with two columns of "lo" and "hi" values to compare. See examples below.
- Numeric variables:
  - \* Numeric of length 1: Forward contrast for a gap of  $x$ , computed between the observed value and the observed value plus  $x$ . Users can set a global option to get a "center" or "backward" contrast instead: `options(marginaleffects_contrast_direction="center")`
  - \* Numeric vector of length 2: Contrast between the largest and the smallest elements of the  $x$  vector.
  - \* Data frame with the same number of rows as newdata, with two columns of "lo" and "hi" values to compare.
  - \* Function that accepts a vector and returns a data frame with two columns of "lo" and "hi" values to compare. See examples below.
  - \* "iqr": Contrast across the interquartile range of the regressor.
  - \* "sd": Contrast across one standard deviation around the regressor mean.
  - \* "2sd": Contrast across two standard deviations around the regressor mean.
  - \* "minmax": Contrast between the maximum and the minimum values of the regressor.
- Examples:
  - \* `variables = list(gear = "pairwise", hp = 10)`
  - \* `variables = list(gear = "sequential", hp = c(100, 120))`
  - \* `variables = list(hp = \(x) data.frame(low = x - 5, high = x + 10))`
  - \* See the Examples section below for more.

comparison

How should pairs of predictions be compared? Difference, ratio, odds ratio, or user-defined functions.

- string: shortcuts to common contrast functions.
  - Supported shortcuts strings: `difference`, `differenceavg`, `differenceavgwts`, `dydx`, `eyex`, `eydx`, `dyex`, `dydxavg`, `eyexavg`, `eydxavg`, `dyexavg`, `dydxavgwts`, `eyexavgwts`, `eydxavgwts`, `dyexavgwts`, `ratio`, `ratioavg`, `ratioavgwts`, `lnratio`, `lnratioavg`, `lnratioavgwts`, `lnor`, `lnoravg`, `lnoravgwts`, `lift`, `liftavg`, `expdydx`, `expdydxavg`, `expdydxavgwts`

	<ul style="list-style-type: none"> <li>– See the Comparisons section below for definitions of each transformation.</li> <li>• function: accept two equal-length numeric vectors of adjusted predictions (hi and lo) and returns a vector of contrasts of the same length, or a unique numeric value. <ul style="list-style-type: none"> <li>– See the Transformations section below for examples of valid functions.</li> </ul> </li> </ul>
type	string indicates the type (scale) of the predictions used to compute contrasts or slopes. This can differ based on the model type, but will typically be a string such as: "response", "link", "probs", or "zero". When an unsupported string is entered, the model-specific list of acceptable values is returned in an error message. When type is NULL, the first entry in the error message is used by default.
vcov	Type of uncertainty estimates to report (e.g., for robust standard errors). Acceptable values: <ul style="list-style-type: none"> <li>• FALSE: Do not compute standard errors. This can speed up computation considerably.</li> <li>• TRUE: Unit-level standard errors using the default <code>vcov(model)</code> variance-covariance matrix.</li> <li>• String which indicates the kind of uncertainty estimates to return. <ul style="list-style-type: none"> <li>– Heteroskedasticity-consistent: "HC", "HC0", "HC1", "HC2", "HC3", "HC4", "HC4m", "HC5". See <code>?sandwich::vcovHC</code></li> <li>– Heteroskedasticity and autocorrelation consistent: "HAC"</li> <li>– Mixed-Models degrees of freedom: "satterthwaite", "kenward-roger"</li> <li>– Other: "NeweyWest", "KernHAC", "OPG". See the <code>sandwich</code> package documentation.</li> </ul> </li> <li>• One-sided formula which indicates the name of cluster variables (e.g., <code>~unit_id</code>). This formula is passed to the <code>cluster</code> argument of the <code>sandwich::vcovCL</code> function.</li> <li>• Square covariance matrix</li> <li>• Function which returns a covariance matrix (e.g., <code>stats::vcov(model)</code>)</li> </ul>
by	Aggregate unit-level estimates (aka, marginalize, average over). Valid inputs: <ul style="list-style-type: none"> <li>• FALSE: return the original unit-level estimates.</li> <li>• TRUE: aggregate estimates for each term.</li> <li>• Character vector of column names in <code>newdata</code> or in the data frame produced by calling the function without the <code>by</code> argument.</li> <li>• Data frame with a <code>by</code> column of group labels, and merging columns shared by <code>newdata</code> or the data frame produced by calling the same function without the <code>by</code> argument.</li> <li>• See examples below.</li> <li>• For more complex aggregations, you can use the <code>FUN</code> argument of the <code>hypotheses()</code> function. See that function's documentation and the Hypothesis Test vignettes on the <code>marginalEffects</code> website.</li> </ul>
conf_level	numeric value between 0 and 1. Confidence level to use to build a confidence interval.

transform	string or function. Transformation applied to unit-level estimates and confidence intervals just before the function returns results. Functions must accept a vector and return a vector of the same length. Support string shortcuts: "exp", "ln"
cross	<ul style="list-style-type: none"> <li>• FALSE: Contrasts represent the change in adjusted predictions when one predictor changes and all other variables are held constant.</li> <li>• TRUE: Contrasts represent the changes in adjusted predictions when all the predictors specified in the <code>variables</code> argument are manipulated simultaneously (a "cross-contrast").</li> </ul>
wts	<p>logical, string or numeric: weights to use when computing average predictions, contrasts or slopes. These weights only affect the averaging in <code>avg_*</code>() or with the <code>by</code> argument, and not unit-level estimates. See <code>?weighted.mean</code></p> <ul style="list-style-type: none"> <li>• string: column name of the weights variable in <code>newdata</code>. When supplying a column name to <code>wts</code>, it is recommended to supply the original data (including the weights variable) explicitly to <code>newdata</code>.</li> <li>• numeric: vector of length equal to the number of rows in the original data or in <code>newdata</code> (if supplied).</li> <li>• FALSE: Equal weights.</li> <li>• TRUE: Extract weights from the fitted object with <code>insight::find_weights()</code> and use them when taking weighted averages of estimates. Warning: <code>newdata=datagrid()</code> returns a single average weight, which is equivalent to using <code>wts=FALSE</code></li> </ul>
hypothesis	<p>specify a hypothesis test or custom contrast using a numeric value, vector, or matrix, a string, a string formula, or a function.</p> <ul style="list-style-type: none"> <li>• Numeric: <ul style="list-style-type: none"> <li>– Single value: the null hypothesis used in the computation of Z and p (before applying transform).</li> <li>– Vector: Weights to compute a linear combination of (custom contrast between) estimates. Length equal to the number of rows generated by the same function call, but without the <code>hypothesis</code> argument.</li> <li>– Matrix: Each column is a vector of weights, as describe above, used to compute a distinct linear combination of (contrast between) estimates. The column names of the matrix are used as labels in the output.</li> </ul> </li> <li>• String formula to specify linear or non-linear hypothesis tests. If the term column uniquely identifies rows, terms can be used in the formula. Otherwise, use <code>b1</code>, <code>b2</code>, etc. to identify the position of each parameter. The <code>b*</code> wildcard can be used to test hypotheses on all estimates. If a named vector is used, the names are used as labels in the output. Examples: <ul style="list-style-type: none"> <li>– <code>hp = drat</code></li> <li>– <code>hp + drat = 12</code></li> <li>– <code>b1 + b2 + b3 = 0</code></li> <li>– <code>b* / b1 = 1</code></li> </ul> </li> <li>• String: <ul style="list-style-type: none"> <li>– "pairwise": pairwise differences between estimates in each row.</li> <li>– "reference": differences between the estimates in each row and the estimate in the first row.</li> </ul> </li> </ul>



	<ul style="list-style-type: none"> <li>- "sequential": difference between an estimate and the estimate in the next row.</li> <li>- "meandev": difference between an estimate and the mean of all estimates.</li> <li>- "meanotherdev": difference between an estimate and the mean of all other estimates, excluding the current one.</li> <li>- "revpairwise", "revreference", "revsequential": inverse of the corresponding hypotheses, as described above.</li> </ul>
	<ul style="list-style-type: none"> <li>• Function: <ul style="list-style-type: none"> <li>- Accepts an argument x: object produced by a <code>marginalEffects</code> function or a data frame with column <code>rowid</code> and <code>estimate</code></li> <li>- Returns a data frame with columns <code>term</code> and <code>estimate</code> (mandatory) and <code>rowid</code> (optional).</li> <li>- The function can also accept optional input arguments: <code>newdata</code>, <code>by</code>, <code>draws</code>.</li> <li>- This function approach will not work for Bayesian models or with bootstrapping. In those cases, it is easy to use <code>posterior_draws()</code> to extract and manipulate the draws directly.</li> </ul> </li> <li>• See the Examples section below and the vignette: <a href="https://marginaleffects.com/vignettes/hypothesis.ht">https://marginaleffects.com/vignettes/hypothesis.ht</a></li> </ul>
<code>equivalence</code>	Numeric vector of length 2: bounds used for the two-one-sided test (TOST) of equivalence, and for the non-inferiority and non-superiority tests. See Details section below.
<code>p_adjust</code>	Adjust p-values for multiple comparisons: "holm", "hochberg", "hommel", "bonferroni", "BH", "BY", or "fdr". See <a href="#">stats::p.adjust</a>
<code>df</code>	Degrees of freedom used to compute p values and confidence intervals. A single numeric value between 1 and Inf. When df is Inf, the normal distribution is used. When df is finite, the t distribution is used. See <a href="#">insight::get_df</a> for a convenient function to extract degrees of freedom. Ex: <code>slopes(model, df = insight::get_df(model))</code>
<code>eps</code>	NULL or numeric value which determines the step size to use when calculating numerical derivatives: $(f(x+eps)-f(x))/eps$ . When eps is NULL, the step size is 0.0001 multiplied by the difference between the maximum and minimum values of the variable with respect to which we are taking the derivative. Changing eps may be necessary to avoid numerical problems in certain models.
<code>numderiv</code>	string or list of strings indicating the method to use to for the numeric differentiation used in to compute delta method standard errors. <ul style="list-style-type: none"> <li>• "fdforward": finite difference method with forward differences</li> <li>• "fdcenter": finite difference method with central differences (default)</li> <li>• "richardson": Richardson extrapolation method</li> <li>• Extra arguments can be specified by passing a list to the <code>numDeriv</code> argument, with the name of the method first and named arguments following, ex: <code>numderiv=list("fdcenter", eps = 1e-5)</code>. When an unknown argument is used, <code>marginaleffects</code> prints the list of valid arguments for each method.</li> </ul>

... Additional arguments are passed to the `predict()` method supplied by the modeling package. These arguments are particularly useful for mixed-effects or bayesian models (see the online vignettes on the `marginalEffects` website). Available arguments can vary from model to model, depending on the range of supported arguments by each modeling package. See the "Model-Specific Arguments" section of the `?slopes` documentation for a non-exhaustive list of available arguments.

### Value

A data frame with one row per observation (per term/group) and several columns:

- `rowid`: row number of the `newdata` data frame
- `type`: prediction type, as defined by the `type` argument
- `group`: (optional) value of the grouped outcome (e.g., categorical outcome models)
- `term`: the variable whose marginal effect is computed
- `dydx`: slope of the outcome with respect to the term, for a given combination of predictor values
- `std.error`: standard errors computed by via the delta method.
- `p.value`: p value associated to the estimate column. The null is determined by the `hypothesis` argument (0 by default), and p values are computed before applying the `transform` argument.
- `s.value`: Shannon information transforms of p values. How many consecutive "heads" tosses would provide the same amount of evidence (or "surprise") against the null hypothesis that the coin is fair? The purpose of S is to calibrate the analyst's intuition about the strength of evidence encoded in p against a well-known physical phenomenon. See Greenland (2019) and Cole et al. (2020).
- `conf.low`: lower bound of the confidence interval (or equal-tailed interval for bayesian models)
- `conf.high`: upper bound of the confidence interval (or equal-tailed interval for bayesian models)

See `?print.marginalEffects` for printing options.

### Functions

- `avg_comparisons()`: Average comparisons

### Standard errors using the delta method

Standard errors for all quantities estimated by `marginalEffects` can be obtained via the delta method. This requires differentiating a function with respect to the coefficients in the model using a finite difference approach. In some models, the delta method standard errors can be sensitive to various aspects of the numeric differentiation strategy, including the step size. By default, the step size is set to  $1e-8$ , or to  $1e-4$  times the smallest absolute model coefficient, whichever is largest.

`marginalEffects` can delegate numeric differentiation to the `numDeriv` package, which allows more flexibility. To do this, users can pass arguments to the `numDeriv::jacobian` function through a global option. For example:

- `options(marginaleffects_numDeriv = list(method = "simple", method.args = list(eps = 1e-6)))`
- `options(marginaleffects_numDeriv = list(method = "Richardson", method.args = list(eps = 1e-5)))`
- `options(marginaleffects_numDeriv = NULL)`

See the "Standard Errors and Confidence Intervals" vignette on the `marginaleffects` website for more details on the computation of standard errors:

<https://marginaleffects.com/vignettes/uncertainty.html>

Note that the `inferences()` function can be used to compute uncertainty estimates using a bootstrap or simulation-based inference. See the vignette:

<https://marginaleffects.com/vignettes/bootstrap.html>

### Model-Specific Arguments

Some model types allow model-specific arguments to modify the nature of marginal effects, predictions, marginal means, and contrasts. Please report other package-specific `predict()` arguments on Github so we can add them to the table below.

<https://github.com/vincentarelbundock/marginaleffects/issues>

Package	Class	Argument	Documentation
brms	brmsfit	<code>ndraws</code>	<a href="#">brms::posterior_predict</a>
		<code>re_formula</code>	<a href="#">brms::posterior_predict</a>
lme4	merMod	<code>re.form</code>	<a href="#">lme4::predict.merMod</a>
		<code>allow.new.levels</code>	<a href="#">lme4::predict.merMod</a>
glmmTMB	glmmTMB	<code>re.form</code>	<a href="#">glmmTMB::predict.glmmTMB</a>
		<code>allow.new.levels</code>	<a href="#">glmmTMB::predict.glmmTMB</a>
		<code>zitype</code>	<a href="#">glmmTMB::predict.glmmTMB</a>
mgcv	bam	<code>exclude</code>	<a href="#">mgcv::predict.bam</a>
robustlmm	rlmerMod	<code>re.form</code>	<a href="#">robustlmm::predict.rlmerMod</a>
		<code>allow.new.levels</code>	<a href="#">robustlmm::predict.rlmerMod</a>
MCMCglmm	MCMCglmm	<code>ndraws</code>	

### comparison argument functions

The following transformations can be applied by supplying one of the shortcut strings to the `comparison` argument. `hi` is a vector of adjusted predictions for the "high" side of the contrast. `lo` is a vector of adjusted predictions for the "low" side of the contrast. `y` is a vector of adjusted predictions for the original data. `x` is the predictor in the original data. `eps` is the step size to use to compute derivatives and elasticities.

Shortcut	Function
<code>difference</code>	$\backslash(\text{hi}, \text{lo}) \text{ hi} - \text{lo}$
<code>differenceavg</code>	$\backslash(\text{hi}, \text{lo}) \text{ mean}(\text{hi} - \text{lo})$
<code>dydx</code>	$\backslash(\text{hi}, \text{lo}, \text{eps}) (\text{hi} - \text{lo})/\text{eps}$
<code>eyex</code>	$\backslash(\text{hi}, \text{lo}, \text{eps}, \text{y}, \text{x}) (\text{hi} - \text{lo})/\text{eps} * (\text{x}/\text{y})$
<code>eydx</code>	$\backslash(\text{hi}, \text{lo}, \text{eps}, \text{y}, \text{x}) ((\text{hi} - \text{lo})/\text{eps})/\text{y}$

dyex	$\backslash(\text{hi}, \text{lo}, \text{eps}, \text{x}) ((\text{hi} - \text{lo})/\text{eps}) * \text{x}$
dydxavg	$\backslash(\text{hi}, \text{lo}, \text{eps}) \text{mean}((\text{hi} - \text{lo})/\text{eps})$
eyexavg	$\backslash(\text{hi}, \text{lo}, \text{eps}, \text{y}, \text{x}) \text{mean}((\text{hi} - \text{lo})/\text{eps} * (\text{x}/\text{y}))$
eydxavg	$\backslash(\text{hi}, \text{lo}, \text{eps}, \text{y}, \text{x}) \text{mean}(((\text{hi} - \text{lo})/\text{eps})/\text{y})$
dyexavg	$\backslash(\text{hi}, \text{lo}, \text{eps}, \text{x}) \text{mean}(((\text{hi} - \text{lo})/\text{eps}) * \text{x})$
ratio	$\backslash(\text{hi}, \text{lo}) \text{hi}/\text{lo}$
ratioavg	$\backslash(\text{hi}, \text{lo}) \text{mean}(\text{hi})/\text{mean}(\text{lo})$
lnratio	$\backslash(\text{hi}, \text{lo}) \log(\text{hi}/\text{lo})$
lnratioavg	$\backslash(\text{hi}, \text{lo}) \log(\text{mean}(\text{hi})/\text{mean}(\text{lo}))$
lnor	$\backslash(\text{hi}, \text{lo}) \log((\text{hi}/(1 - \text{hi})) / (\text{lo}/(1 - \text{lo})))$
lnoravg	$\backslash(\text{hi}, \text{lo}) \log((\text{mean}(\text{hi}) / (1 - \text{mean}(\text{hi}))) / (\text{mean}(\text{lo}) / (1 - \text{mean}(\text{lo}))))$
lift	$\backslash(\text{hi}, \text{lo}) (\text{hi} - \text{lo})/\text{lo}$
liftavg	$\backslash(\text{hi}, \text{lo}) (\text{mean}(\text{hi} - \text{lo})) / \text{mean}(\text{lo})$
expdydx	$\backslash(\text{hi}, \text{lo}, \text{eps}) ((\exp(\text{hi}) - \exp(\text{lo})) / \exp(\text{eps})) / \text{eps}$
expdydxavg	$\backslash(\text{hi}, \text{lo}, \text{eps}) \text{mean}(((\exp(\text{hi}) - \exp(\text{lo})) / \exp(\text{eps})) / \text{eps})$

### Bayesian posterior summaries

By default, credible intervals in bayesian models are built as equal-tailed intervals. This can be changed to a highest density interval by setting a global option:

```
options("marginaleffects_posterior_interval" = "eti")
```

```
options("marginaleffects_posterior_interval" = "hdi")
```

By default, the center of the posterior distribution in bayesian models is identified by the median. Users can use a different summary function by setting a global option:

```
options("marginaleffects_posterior_center" = "mean")
```

```
options("marginaleffects_posterior_center" = "median")
```

When estimates are averaged using the `by` argument, the `tidy()` function, or the `summary()` function, the posterior distribution is marginalized twice over. First, we take the average *across* units but *within* each iteration of the MCMC chain, according to what the user requested in `by` argument or `tidy()/summary()` functions. Then, we identify the center of the resulting posterior using the function supplied to the `"marginaleffects_posterior_center"` option (the median by default).

### Equivalence, Inferiority, Superiority

$\theta$  is an estimate,  $\sigma_\theta$  its estimated standard error, and  $[a, b]$  are the bounds of the interval supplied to the equivalence argument.

Non-inferiority:

- $H_0: \theta \leq a$
- $H_1: \theta > a$
- $t = (\theta - a) / \sigma_\theta$
- p: Upper-tail probability

Non-superiority:

- $H_0: \theta \geq b$

- $H_1: \theta < b$
- $t = (\theta - b)/\sigma_\theta$
- p: Lower-tail probability

Equivalence: Two One-Sided Tests (TOST)

- p: Maximum of the non-inferiority and non-superiority p values.

Thanks to Russell V. Lenth for the excellent emmeans package and documentation which inspired this feature.

### Prediction types

The `type` argument determines the scale of the predictions used to compute quantities of interest with functions from the `marginalEffects` package. Admissible values for `type` depend on the model object. When users specify an incorrect value for `type`, `marginalEffects` will raise an informative error with a list of valid `type` values for the specific model object. The first entry in the list in that error message is the default `type`.

The `invlink(link)` is a special `type` defined by `marginalEffects`. It is available for some (but not all) models and functions. With this `link` type, we first compute predictions on the link scale, then we use the inverse link function to backtransform the predictions to the response scale. This is useful for models with non-linear link functions as it can ensure that confidence intervals stay within desirable bounds, ex: 0 to 1 for a logit model. Note that an average of estimates with `type="invlink(link)"` will not always be equivalent to the average of estimates with `type="response"`.

Some of the most common `type` values are:

response, link, E, Ep, average, class, conditional, count, cum.prob, cumhaz, cumprob, density, detection, disp, ev, expected, expvalue, fitted, hazard, invlink(link), latent, latent\_N, linear, linear.predictor, linpred, location, lp, mean, numeric, p, ppd, pr, precision, prediction, prob, probability, probs, quantile, risk, rmst, scale, survival, unconditional, utility, variance, xb, zero, zlink, zprob

### Order of operations

Behind the scenes, the arguments of `marginalEffects` functions are evaluated in this order:

1. newdata
2. variables
3. comparison and slopes
4. by
5. vcov
6. hypothesis
7. transform

## Parallel computation

The `slopes()` and `comparisons()` functions can use parallelism to speed up computation. Operations are parallelized for the computation of standard errors, at the model coefficient level. There is always considerable overhead when using parallel computation, mainly involved in passing the whole dataset to the different processes. Thus, parallel computation is most likely to be useful when the model includes many parameters and the dataset is relatively small.

Warning: In many cases, parallel processing will not be useful at all.

To activate parallel computation, users must load the `future.apply` package, call `plan()` function, and set a global option. For example:

```
library(future.apply)
plan("multicore", workers = 4)
options(marginaleffects_parallel = TRUE)
```

```
slopes(model)
```

To disable parallelism in `marginaleffects` altogether, you can set a global option:

```
options(marginaleffects_parallel = FALSE)
```

## References

- Greenland S. 2019. "Valid P-Values Behave Exactly as They Should: Some Misleading Criticisms of P-Values and Their Resolution With S-Values." *The American Statistician*. 73(S1): 106–114.
- Cole, Stephen R, Jessie K Edwards, and Sander Greenland. 2020. "Surprise!" *American Journal of Epidemiology* 190 (2): 191–93. <https://doi.org/10.1093/aje/kwaa136>

## Examples

```
library(marginaleffects)

# Linear model
tmp <- mtcars
tmp$am <- as.logical(tmp$am)
mod <- lm(mpg ~ am + factor(cyl), tmp)
avg_comparisons(mod, variables = list(cyl = "reference"))
avg_comparisons(mod, variables = list(cyl = "sequential"))
avg_comparisons(mod, variables = list(cyl = "pairwise"))

# GLM with different scale types
mod <- glm(am ~ factor(gear), data = mtcars)
avg_comparisons(mod, type = "response")
avg_comparisons(mod, type = "link")

# Contrasts at the mean
comparisons(mod, newdata = "mean")
```

```

# Contrasts between marginal means
comparisons(mod, newdata = "marginalmeans")

# Contrasts at user-specified values
comparisons(mod, newdata = datagrid(am = 0, gear = tmp$gear))
comparisons(mod, newdata = datagrid(am = unique, gear = max))

m <- lm(mpg ~ hp + drat + factor(cyl) + factor(am), data = mtcars)
comparisons(m, variables = "hp", newdata = datagrid(FUN_factor = unique, FUN_numeric = median))

# Numeric contrasts
mod <- lm(mpg ~ hp, data = mtcars)
avg_comparisons(mod, variables = list(hp = 1))
avg_comparisons(mod, variables = list(hp = 5))
avg_comparisons(mod, variables = list(hp = c(90, 100)))
avg_comparisons(mod, variables = list(hp = "iqr"))
avg_comparisons(mod, variables = list(hp = "sd"))
avg_comparisons(mod, variables = list(hp = "minmax"))

# using a function to specify a custom difference in one regressor
dat <- mtcars
dat$new_hp <- 49 * (dat$hp - min(dat$hp)) / (max(dat$hp) - min(dat$hp)) + 1
modlog <- lm(mpg ~ log(new_hp) + factor(cyl), data = dat)
fdiff <- \(x) data.frame(x, x + 10)
avg_comparisons(modlog, variables = list(new_hp = fdiff))

# Adjusted Risk Ratio: see the contrasts vignette
mod <- glm(vs ~ mpg, data = mtcars, family = binomial)
avg_comparisons(mod, comparison = "lnratioavg", transform = exp)

# Adjusted Risk Ratio: Manual specification of the `comparison`
avg_comparisons(
  mod,
  comparison = function(hi, lo) log(mean(hi) / mean(lo)),
  transform = exp)
# cross contrasts
mod <- lm(mpg ~ factor(cyl) * factor(gear) + hp, data = mtcars)
avg_comparisons(mod, variables = c("cyl", "gear"), cross = TRUE)

# variable-specific contrasts
avg_comparisons(mod, variables = list(gear = "sequential", hp = 10))

# hypothesis test: is the `hp` marginal effect at the mean equal to the `drat` marginal effect
mod <- lm(mpg ~ wt + drat, data = mtcars)

comparisons(
  mod,
  newdata = "mean",
  hypothesis = "wt = drat")

# same hypothesis test using row indices
comparisons(

```

```

    mod,
    newdata = "mean",
    hypothesis = "b1 - b2 = 0")

# same hypothesis test using numeric vector of weights
comparisons(
  mod,
  newdata = "mean",
  hypothesis = c(1, -1))

# two custom contrasts using a matrix of weights
lc <- matrix(c(
  1, -1,
  2, 3),
  ncol = 2)
comparisons(
  mod,
  newdata = "mean",
  hypothesis = lc)

# Effect of a 1 group-wise standard deviation change
# First we calculate the SD in each group of `cyl`
# Second, we use that SD as the treatment size in the `variables` argument
library(dplyr)
mod <- lm(mpg ~ hp + factor(cyl), mtcars)
tmp <- mtcars %>%
  group_by(cyl) %>%
  mutate(hp_sd = sd(hp))
avg_comparisons(mod,
  variables = list(hp = function(x) data.frame(x, x + tmp$hp_sd)),
  by = "cyl")

# `by` argument
mod <- lm(mpg ~ hp * am * vs, data = mtcars)
comparisons(mod, by = TRUE)

mod <- lm(mpg ~ hp * am * vs, data = mtcars)
avg_comparisons(mod, variables = "hp", by = c("vs", "am"))

library(nnet)
mod <- multinom(factor(gear) ~ mpg + am * vs, data = mtcars, trace = FALSE)
by <- data.frame(
  group = c("3", "4", "5"),
  by = c("3,4", "3,4", "5"))
comparisons(mod, type = "probs", by = by)

```



## Description

Generate a data grid of user-specified values for use in the `newdata` argument of the `predictions()`, `comparisons()`, and `slopes()` functions. This is useful to define where in the predictor space we want to evaluate the quantities of interest. Ex: the predicted outcome or slope for a 37 year old college graduate.

## Usage

```
datagrid(
  ...,
  model = NULL,
  newdata = NULL,
  by = NULL,
  grid_type = "mean_or_mode",
  response = FALSE,
  FUN_character = NULL,
  FUN_factor = NULL,
  FUN_logical = NULL,
  FUN_numeric = NULL,
  FUN_integer = NULL,
  FUN_binary = NULL,
  FUN_other = NULL
)
```

## Arguments

<code>...</code>	<p>named arguments with vectors of values or functions for user-specified variables.</p> <ul style="list-style-type: none"> <li>• Functions are applied to the variable in the <code>model</code> dataset or <code>newdata</code>, and must return a vector of the appropriate type.</li> <li>• Character vectors are automatically transformed to factors if necessary. +The output will include all combinations of these variables (see Examples below.)</li> </ul>
<code>model</code>	Model object
<code>newdata</code>	<code>data.frame</code> (one and only one of the <code>model</code> and <code>newdata</code> arguments can be used.)
<code>by</code>	character vector with grouping variables within which <code>FUN_*</code> functions are applied to create "sub-grids" with unspecified variables.
<code>grid_type</code>	<p>character. Determines the functions to apply to each variable. The defaults can be overridden by defining individual variables explicitly in <code>...</code>, or by supplying a function to one of the <code>FUN_*</code> arguments.</p> <ul style="list-style-type: none"> <li>• "mean_or_mode": Character, factor, logical, and binary variables are set to their modes. Numeric, integer, and other variables are set to their means.</li> <li>• "balanced": Each unique level of character, factor, logical, and binary variables are preserved. Numeric, integer, and other variables are set to their means. Warning: When there are many variables and many levels per variable, a balanced grid can be very large. In those cases, it is better to use</li> </ul>

	grid_type="mean_or_mode" and to specify the unique levels of a subset of named variables explicitly.
	<ul style="list-style-type: none"> <li>• "counterfactual": the entire dataset is duplicated for each combination of the variable values specified in ... Variables not explicitly supplied to datagrid() are set to their observed values in the original dataset.</li> </ul>
response	Logical should the response variable be included in the grid, even if it is not specified explicitly.
FUN_character	the function to be applied to character variables.
FUN_factor	the function to be applied to factor variables. This only applies if the variable in the original data is a factor. For variables converted to factor in a model-fitting formula, for example, FUN_character is used.
FUN_logical	the function to be applied to logical variables.
FUN_numeric	the function to be applied to numeric variables.
FUN_integer	the function to be applied to integer variables.
FUN_binary	the function to be applied to binary variables.
FUN_other	the function to be applied to other variable types.

## Details

If datagrid is used in a predictions(), comparisons(), or slopes() call as the newdata argument, the model is automatically inserted in the model argument of datagrid() call, and users do not need to specify either the model or newdata arguments. The same behavior will occur when the value supplied to newdata= is a function call which starts with "datagrid". This is intended to allow users to create convenience shortcuts like:

```
library(marginaleffects)
mod <- lm(mpg ~ am + vs + factor(cyl) + hp, mtcars)
datagrid_bal <- function(...) datagrid(..., grid_type = "balanced")
predictions(model, newdata = datagrid_bal(cyl = 4))
```

If users supply a model, the data used to fit that model is retrieved using the insight::get\_data function.

## Value

A data.frame in which each row corresponds to one combination of the named predictors supplied by the user via the ... dots. Variables which are not explicitly defined are held at their mean or mode.

## Examples

```
# The output only has 2 rows, and all the variables except `hp` are at their
# mean or mode.
datagrid(newdata = mtcars, hp = c(100, 110))

# We get the same result by feeding a model instead of a data.frame
mod <- lm(mpg ~ hp, mtcars)
```

```

datagrid(model = mod, hp = c(100, 110))

# Use in `marginaleffects` to compute "Typical Marginal Effects". When used
# in `slopes()` or `predictions()` we do not need to specify the
# `model` or `newdata` arguments.
slopes(mod, newdata = datagrid(hp = c(100, 110)))

# datagrid accepts functions
datagrid(hp = range, cyl = unique, newdata = mtcars)
comparisons(mod, newdata = datagrid(hp = fivenum))

# The full dataset is duplicated with each observation given counterfactual
# values of 100 and 110 for the `hp` variable. The original `mtcars` includes
# 32 rows, so the resulting dataset includes 64 rows.
dg <- datagrid(newdata = mtcars, hp = c(100, 110), grid_type = "counterfactual")
nrow(dg)

# We get the same result by feeding a model instead of a data.frame
mod <- lm(mpg ~ hp, mtcars)
dg <- datagrid(model = mod, hp = c(100, 110), grid_type = "counterfactual")
nrow(dg)

```

---

hypotheses

*(Non-)Linear Tests for Null Hypotheses, Joint Hypotheses, Equivalence, Non Superiority, and Non Inferiority*

---

## Description

Uncertainty estimates are calculated as first-order approximate standard errors for linear or non-linear functions of a vector of random variables with known or estimated covariance matrix. In that sense, `hypotheses` emulates the behavior of the excellent and well-established `car::deltaMethod` and `car::linearHypothesis` functions, but it supports more models; requires fewer dependencies; expands the range of tests to equivalence and superiority/inferiority; and offers convenience features like robust standard errors.

To learn more, read the hypothesis tests vignette, visit the package website, or scroll down this page for a full list of vignettes:

- <https://marginaleffects.com/vignettes/hypothesis.html>
- <https://marginaleffects.com/>

Warning #1: Tests are conducted directly on the scale defined by the `type` argument. For some models, it can make sense to conduct hypothesis or equivalence tests on the "link" scale instead of the "response" scale which is often the default.

Warning #2: For hypothesis tests on objects produced by the `marginaleffects` package, it is safer to use the `hypothesis` argument of the original function. Using `hypotheses()` may not work in certain environments, in lists, or when working programmatically with `*apply` style functions.

Warning #3: The tests assume that the hypothesis expression is (approximately) normally distributed, which for non-linear functions of the parameters may not be realistic. More reliable confidence intervals can be obtained using the `inferences()` function with `method = "boot"`.

**Usage**

```

hypotheses(
  model,
  hypothesis = NULL,
  vcov = NULL,
  conf_level = 0.95,
  df = NULL,
  equivalence = NULL,
  joint = FALSE,
  joint_test = "f",
  numderiv = "fdforward",
  ...
)

```

**Arguments**

- |            |  |
|------------|--|
| model      | Model object or object generated by the <code>comparisons()</code> , <code>slopes()</code> , or <code>predictions()</code> functions.  |
| hypothesis | <p>specify a hypothesis test or custom contrast using a numeric value, vector, or matrix, a string, a string formula, or a function.</p> <ul style="list-style-type: none"> <li>• Numeric: <ul style="list-style-type: none"> <li>– Single value: the null hypothesis used in the computation of Z and p (before applying transform).</li> <li>– Vector: Weights to compute a linear combination of (custom contrast between) estimates. Length equal to the number of rows generated by the same function call, but without the <code>hypothesis</code> argument.</li> <li>– Matrix: Each column is a vector of weights, as describe above, used to compute a distinct linear combination of (contrast between) estimates. The column names of the matrix are used as labels in the output.</li> </ul> </li> <li>• String formula to specify linear or non-linear hypothesis tests. If the term column uniquely identifies rows, terms can be used in the formula. Otherwise, use <code>b1</code>, <code>b2</code>, etc. to identify the position of each parameter. The <code>b*</code> wildcard can be used to test hypotheses on all estimates. If a named vector is used, the names are used as labels in the output. Examples: <ul style="list-style-type: none"> <li>– <code>hp = drat</code></li> <li>– <code>hp + drat = 12</code></li> <li>– <code>b1 + b2 + b3 = 0</code></li> <li>– <code>b* / b1 = 1</code></li> </ul> </li> <li>• String: <ul style="list-style-type: none"> <li>– "pairwise": pairwise differences between estimates in each row.</li> <li>– "reference": differences between the estimates in each row and the estimate in the first row.</li> <li>– "sequential": difference between an estimate and the estimate in the next row.</li> <li>– "meandev": difference between an estimate and the mean of all estimates.</li> </ul> </li> </ul> |

	<ul style="list-style-type: none"> <li>– "meanotherdev": difference between an estimate and the mean of all other estimates, excluding the current one.</li> <li>– "revpairwise", "revreference", "revsequential": inverse of the corresponding hypotheses, as described above.</li> <li>• Function: <ul style="list-style-type: none"> <li>– Accepts an argument <code>x</code>: object produced by a <code>marginalEffects</code> function or a data frame with column <code>rowid</code> and <code>estimate</code></li> <li>– Returns a data frame with columns <code>term</code> and <code>estimate</code> (mandatory) and <code>rowid</code> (optional).</li> <li>– The function can also accept optional input arguments: <code>newdata</code>, <code>by</code>, <code>draws</code>.</li> <li>– This function approach will not work for Bayesian models or with bootstrapping. In those cases, it is easy to use <code>posterior_draws()</code> to extract and manipulate the draws directly.</li> </ul> </li> <li>• See the Examples section below and the vignette: <a href="https://marginaleffects.com/vignettes/hypothesis.ht">https://marginaleffects.com/vignettes/hypothesis.ht</a></li> </ul>
<code>vcov</code>	<p>Type of uncertainty estimates to report (e.g., for robust standard errors). Acceptable values:</p> <ul style="list-style-type: none"> <li>• FALSE: Do not compute standard errors. This can speed up computation considerably.</li> <li>• TRUE: Unit-level standard errors using the default <code>vcov(model)</code> variance-covariance matrix.</li> <li>• String which indicates the kind of uncertainty estimates to return. <ul style="list-style-type: none"> <li>– Heteroskedasticity-consistent: "HC", "HC0", "HC1", "HC2", "HC3", "HC4", "HC4m", "HC5". See <code>?sandwich::vcovHC</code></li> <li>– Heteroskedasticity and autocorrelation consistent: "HAC"</li> <li>– Mixed-Models degrees of freedom: "satterthwaite", "kenward-roger"</li> <li>– Other: "NeweyWest", "KernHAC", "OPG". See the <code>sandwich</code> package documentation.</li> </ul> </li> <li>• One-sided formula which indicates the name of cluster variables (e.g., <code>~unit_id</code>). This formula is passed to the <code>cluster</code> argument of the <code>sandwich::vcovCL</code> function.</li> <li>• Square covariance matrix</li> <li>• Function which returns a covariance matrix (e.g., <code>stats::vcov(model)</code>)</li> </ul>
<code>conf_level</code>	numeric value between 0 and 1. Confidence level to use to build a confidence interval.
<code>df</code>	Degrees of freedom used to compute p values and confidence intervals. A single numeric value between 1 and Inf. When using <code>joint_test="f"</code> , the <code>df</code> argument should be a numeric vector of length 2.
<code>equivalence</code>	Numeric vector of length 2: bounds used for the two-one-sided test (TOST) of equivalence, and for the non-inferiority and non-superiority tests. See Details section below.
<code>joint</code>	<p>Joint test of statistical significance. The null hypothesis value can be set using the <code>hypothesis</code> argument.</p> <ul style="list-style-type: none"> <li>• FALSE: Hypotheses are not tested jointly.</li> </ul>

	<ul style="list-style-type: none"> <li>• TRUE: All parameters are tested jointly.</li> <li>• String: A regular expression to match parameters to be tested jointly. <code>grep(joint, perl = TRUE)</code></li> <li>• Character vector of parameter names to be tested. Characters refer to the names of the vector returned by <code>coef(object)</code>.</li> <li>• Integer vector of indices. Which parameters positions to test jointly.</li> </ul>
<code>joint_test</code>	A character string specifying the type of test, either "f" or "chisq". The null hypothesis is set by the <code>hypothesis</code> argument, with default null equal to 0 for all parameters.
<code>numberiv</code>	string or list of strings indicating the method to use to for the numeric differentiation used in to compute delta method standard errors. <ul style="list-style-type: none"> <li>• "fdforward": finite difference method with forward differences</li> <li>• "fdcenter": finite difference method with central differences (default)</li> <li>• "richardson": Richardson extrapolation method</li> <li>• Extra arguments can be specified by passing a list to the <code>numDeriv</code> argument, with the name of the method first and named arguments following, ex: <code>numberiv=list("fdcenter", eps = 1e-5)</code>. When an unknown argument is used, <code>marginaleffects</code> prints the list of valid arguments for each method.</li> </ul>
...	Additional arguments are passed to the <code>predict()</code> method supplied by the modeling package. These arguments are particularly useful for mixed-effects or bayesian models (see the online vignettes on the <code>marginaleffects</code> website). Available arguments can vary from model to model, depending on the range of supported arguments by each modeling package. See the "Model-Specific Arguments" section of the <code>?slopes</code> documentation for a non-exhaustive list of available arguments.

### Joint hypothesis tests

The test statistic for the joint Wald test is calculated as  $(R * \theta_{\hat{}} - r)' * \text{inv}(R * V_{\hat{}} * R') * (R * \theta_{\hat{}} - r) / Q$ , where  $\theta_{\hat{}}$  is the vector of estimated parameters,  $V_{\hat{}}$  is the estimated covariance matrix,  $R$  is a  $Q \times P$  matrix for testing  $Q$  hypotheses on  $P$  parameters,  $r$  is a  $Q \times 1$  vector for the null hypothesis, and  $Q$  is the number of rows in  $R$ . If the test is a Chi-squared test, the test statistic is not normalized.

The p-value is then calculated based on either the F-distribution (for F-test) or the Chi-squared distribution (for Chi-squared test). For the F-test, the degrees of freedom are  $Q$  and  $(n - P)$ , where  $n$  is the sample size and  $P$  is the number of parameters. For the Chi-squared test, the degrees of freedom are  $Q$ .

### Equivalence, Inferiority, Superiority

$\theta$  is an estimate,  $\sigma_{\theta}$  its estimated standard error, and  $[a, b]$  are the bounds of the interval supplied to the equivalence argument.

Non-inferiority:

- $H_0: \theta \leq a$

- $H_1: \theta > a$
- $t = (\theta - a)/\sigma_\theta$
- p: Upper-tail probability

Non-superiority:

- $H_0: \theta \geq b$
- $H_1: \theta < b$
- $t = (\theta - b)/\sigma_\theta$
- p: Lower-tail probability

Equivalence: Two One-Sided Tests (TOST)

- p: Maximum of the non-inferiority and non-superiority p values.

Thanks to Russell V. Lenth for the excellent emmeans package and documentation which inspired this feature.

## Examples

```
library(marginaleffects)
mod <- lm(mpg ~ hp + wt + factor(cyl), data = mtcars)

hypotheses(mod)

# Test of equality between coefficients
hypotheses(mod, hypothesis = "hp = wt")

# Non-linear function
hypotheses(mod, hypothesis = "exp(hp + wt) = 0.1")

# Robust standard errors
hypotheses(mod, hypothesis = "hp = wt", vcov = "HC3")

# b1, b2, ... shortcuts can be used to identify the position of the
# parameters of interest in the output of
hypotheses(mod, hypothesis = "b2 = b3")

# wildcard
hypotheses(mod, hypothesis = "b* / b2 = 1")

# term names with special characters have to be enclosed in backticks
hypotheses(mod, hypothesis = "`factor(cyl)6` = `factor(cyl)8`")

mod2 <- lm(mpg ~ hp * drat, data = mtcars)
hypotheses(mod2, hypothesis = "`hp:drat` = drat")

# predictions(), comparisons(), and slopes()
mod <- glm(am ~ hp + mpg, data = mtcars, family = binomial)
cmp <- comparisons(mod, newdata = "mean")
hypotheses(cmp, hypothesis = "b1 = b2")
```

```

mfx <- slopes(mod, newdata = "mean")
hypotheses(cmp, hypothesis = "b2 = 0.2")

pre <- predictions(mod, newdata = datagrid(hp = 110, mpg = c(30, 35)))
hypotheses(pre, hypothesis = "b1 = b2")

# The `hypothesis` argument can be used to compute standard errors for fitted values
mod <- glm(am ~ hp + mpg, data = mtcars, family = binomial)

f <- function(x) predict(x, type = "link", newdata = mtcars)
p <- hypotheses(mod, hypothesis = f)
head(p)

f <- function(x) predict(x, type = "response", newdata = mtcars)
p <- hypotheses(mod, hypothesis = f)
head(p)

# Complex aggregation
# Step 1: Collapse predicted probabilities by outcome level, for each individual
# Step 2: Take the mean of the collapsed probabilities by group and `cyl`
library(dplyr)
library(MASS)
library(dplyr)

dat <- transform(mtcars, gear = factor(gear))
mod <- polr(gear ~ factor(cyl) + hp, dat)

aggregation_fun <- function(x) {
  predictions(x, vcov = FALSE) |>
  mutate(group = ifelse(group %in% c("3", "4"), "3 & 4", "5")) |>
  summarize(estimate = sum(estimate), .by = c("rowid", "cyl", "group")) |>
  summarize(estimate = mean(estimate), .by = c("cyl", "group")) |>
  rename(term = cyl)
}

hypotheses(mod, hypothesis = aggregation_fun)

# Equivalence, non-inferiority, and non-superiority tests
mod <- lm(mpg ~ hp + factor(gear), data = mtcars)
p <- predictions(mod, newdata = "median")
hypotheses(p, equivalence = c(17, 18))

mfx <- avg_slopes(mod, variables = "hp")
hypotheses(mfx, equivalence = c(-.1, .1))

cmp <- avg_comparisons(mod, variables = "gear", hypothesis = "pairwise")
hypotheses(cmp, equivalence = c(0, 10))

# joint hypotheses: character vector
model <- lm(mpg ~ as.factor(cyl) * hp, data = mtcars)
hypotheses(model, joint = c("as.factor(cyl)6:hp", "as.factor(cyl)8:hp"))

# joint hypotheses: regular expression

```



```

hypotheses(model, joint = "cyl")

# joint hypotheses: integer indices
hypotheses(model, joint = 2:3)

# joint hypotheses: different null hypotheses
hypotheses(model, joint = 2:3, hypothesis = 1)
hypotheses(model, joint = 2:3, hypothesis = 1:2)

# joint hypotheses: marginaleffects object
cmp <- avg_comparisons(model)
hypotheses(cmp, joint = "cyl")

```

---

inferences	<i>(EXPERIMENTAL) Bootstrap, Conformal, and Simulation-Based Inference</i>
------------	--

---

### Description

Warning: This function is experimental. It may be renamed, the user interface may change, or the functionality may migrate to arguments in other `marginaleffects` functions.

Apply this function to a `marginaleffects` object to change the inferential method used to compute uncertainty estimates.

### Usage

```

inferences(
  x,
  method,
  R = 1000,
  conf_type = "perc",
  conformal_test = NULL,
  conformal_calibration = NULL,
  conformal_score = "residual_abs",
  ...
)

```

### Arguments

<code>x</code>	Object produced by one of the core <code>marginaleffects</code> functions.
<code>method</code>	String <ul style="list-style-type: none"> <li>• "delta": delta method standard errors</li> <li>• "boot" package</li> <li>• "fwb": fractional weighted bootstrap</li> <li>• "rsample" package</li> <li>• "simulation" from a multivariate normal distribution (Krinsky &amp; Robb, 1986)</li> </ul>

	<ul style="list-style-type: none"> <li>• "mi" multiple imputation for missing data</li> <li>• "conformal_split": prediction intervals using split conformal prediction (see Angelopoulos &amp; Bates, 2022)</li> <li>• "conformal_cv+": prediction intervals using cross-validation+ conformal prediction (see Barber et al., 2020)</li> </ul>
R	Number of resamples, simulations, or cross-validation folds.
conf_type	String: type of bootstrap interval to construct. <ul style="list-style-type: none"> <li>• boot: "perc", "norm", "basic", or "bca"</li> <li>• fwb: "perc", "norm", "basic", "bc", or "bca"</li> <li>• rsample: "perc" or "bca"</li> <li>• simulation: argument ignored.</li> </ul>
conformal_test	Data frame of test data for conformal prediction.
conformal_calibration	Data frame of calibration data for split conformal prediction (method="conformal_split").
conformal_score	String. Warning: The type argument in predictions() must generate predictions which are on the same scale as the outcome variable. Typically, this means that type must be "response" or "probs". <ul style="list-style-type: none"> <li>• "residual_abs" or "residual_sq" for regression tasks (numeric outcome)</li> <li>• "softmax" for classification tasks (when predictions() returns a group columns, such as multinomial or ordinal logit models).</li> </ul>
...	<ul style="list-style-type: none"> <li>• If method="boot", additional arguments are passed to boot::boot().</li> <li>• If method="fwb", additional arguments are passed to fwb::fwb().</li> <li>• If method="rsample", additional arguments are passed to rsample::bootstraps().</li> <li>• Additional arguments are ignored for all other methods.</li> </ul>

## Details

When method="simulation", we conduct simulation-based inference following the method discussed in Krinsky & Robb (1986):

1. Draw  $R$  sets of simulated coefficients from a multivariate normal distribution with mean equal to the original model's estimated coefficients and variance equal to the model's variance-covariance matrix (classical, "HC3", or other).
2. Use the  $R$  sets of coefficients to compute  $R$  sets of estimands: predictions, comparisons, slopes, or hypotheses.
3. Take quantiles of the resulting distribution of estimands to obtain a confidence interval and the standard deviation of simulated estimates to estimate the standard error.

When method="fwb", drawn weights are supplied to the model fitting function's weights argument; if the model doesn't accept non-integer weights, this method should not be used. If weights were included in the original model fit, they are extracted by `weights()` and multiplied by the drawn weights. These weights are supplied to the `wts` argument of the estimation function (e.g., `comparisons()`).

**Value**

A `marginaleffects` object with simulation or bootstrap resamples and objects attached.

**References**

- Krinsky, I., and A. L. Robb. 1986. "On Approximating the Statistical Properties of Elasticities." *Review of Economics and Statistics* 68 (4): 715–9.
- King, Gary, Michael Tomz, and Jason Wittenberg. "Making the most of statistical analyses: Improving interpretation and presentation." *American journal of political science* (2000): 347-361
- Dowd, Bryan E., William H. Greene, and Edward C. Norton. "Computation of standard errors." *Health services research* 49.2 (2014): 731-750.
- Angelopoulos, Anastasios N., and Stephen Bates. 2022. "A Gentle Introduction to Conformal Prediction and Distribution-Free Uncertainty Quantification." *arXiv*. <https://doi.org/10.48550/arXiv.2107.07511>.
- Barber, Rina Foygel, Emmanuel J. Candès, Aaditya Ramdas, and Ryan J. Tibshirani. 2020. "Predictive Inference with the Jackknife+." *arXiv*. <http://arxiv.org/abs/1905.02928>.

**Examples**

```
## Not run:
library(marginaleffects)
library(magrittr)
set.seed(1024)
mod <- lm(Sepal.Length ~ Sepal.Width * Species, data = iris)

# bootstrap
avg_predictions(mod, by = "Species") %>%
  inferences(method = "boot")

avg_predictions(mod, by = "Species") %>%
  inferences(method = "rsample")

# Fractional (bayesian) bootstrap
avg_slopes(mod, by = "Species") %>%
  inferences(method = "fwb") %>%
  posterior_draws("rvar") %>%
  data.frame()

# Simulation-based inference
slopes(mod) %>%
  inferences(method = "simulation") %>%
  head()

## End(Not run)
```

---

plot\_comparisons      *Plot Conditional or Marginal Comparisons*

---

### Description

Plot comparisons on the y-axis against values of one or more predictors (x-axis, colors/shapes, and facets).

The `by` argument is used to plot marginal comparisons, that is, comparisons made on the original data, but averaged by subgroups. This is analogous to using the `by` argument in the `comparisons()` function.

The `condition` argument is used to plot conditional comparisons, that is, comparisons made on a user-specified grid. This is analogous to using the `newdata` argument and `datagrid()` function in a `comparisons()` call. All variables whose values are not specified explicitly are treated as usual by `datagrid()`, that is, they are held at their mean or mode (or rounded mean for integers). This includes grouping variables in mixed-effects models, so analysts who fit such models may want to specify the groups of interest using the `condition` argument, or supply model-specific arguments to compute population-level estimates. See details below.

See the "Plots" vignette and website for tutorials and information on how to customize plots:

- <https://marginaleffects.com/vignettes/plot.html>
- <https://marginaleffects.com>

### Usage

```
plot_comparisons(
  model,
  variables = NULL,
  condition = NULL,
  by = NULL,
  newdata = NULL,
  type = "response",
  vcov = NULL,
  conf_level = 0.95,
  wts = FALSE,
  comparison = "difference",
  transform = NULL,
  rug = FALSE,
  gray = FALSE,
  draw = TRUE,
  ...
)
```

### Arguments

<code>model</code>	Model object
<code>variables</code>	Name of the variable whose contrast we want to plot on the y-axis.

condition	<p>Conditional slopes</p> <ul style="list-style-type: none"> <li>• Character vector (max length 4): Names of the predictors to display.</li> <li>• Named list (max length 4): List names correspond to predictors. List elements can be: <ul style="list-style-type: none"> <li>– Numeric vector</li> <li>– Function which returns a numeric vector or a set of unique categorical values</li> <li>– Shortcut strings for common reference values: "minmax", "quartile", "threenum"</li> </ul> </li> <li>• 1: x-axis. 2: color/shape. 3: facet (wrap if no fourth variable, otherwise cols of grid). 4: facet (rows of grid).</li> <li>• Numeric variables in positions 2 and 3 are summarized by Tukey's five numbers <code>?stats::fivenum</code>.</li> </ul>
by	<p>Aggregate unit-level estimates (aka, marginalize, average over). Valid inputs:</p> <ul style="list-style-type: none"> <li>• FALSE: return the original unit-level estimates.</li> <li>• TRUE: aggregate estimates for each term.</li> <li>• Character vector of column names in <code>newdata</code> or in the data frame produced by calling the function without the <code>by</code> argument.</li> <li>• Data frame with a <code>by</code> column of group labels, and merging columns shared by <code>newdata</code> or the data frame produced by calling the same function without the <code>by</code> argument.</li> <li>• See examples below.</li> <li>• For more complex aggregations, you can use the <code>FUN</code> argument of the <code>hypotheses()</code> function. See that function's documentation and the Hypothesis Test vignettes on the <code>marginalEffects</code> website.</li> </ul>
newdata	<p>When <code>newdata</code> is NULL, the grid is determined by the <code>condition</code> argument. When <code>newdata</code> is not NULL, the argument behaves in the same way as in the <code>comparisons()</code> function.</p>
type	<p>string indicates the type (scale) of the predictions used to compute contrasts or slopes. This can differ based on the model type, but will typically be a string such as: "response", "link", "probs", or "zero". When an unsupported string is entered, the model-specific list of acceptable values is returned in an error message. When <code>type</code> is NULL, the first entry in the error message is used by default.</p>
vcov	<p>Type of uncertainty estimates to report (e.g., for robust standard errors). Acceptable values:</p> <ul style="list-style-type: none"> <li>• FALSE: Do not compute standard errors. This can speed up computation considerably.</li> <li>• TRUE: Unit-level standard errors using the default <code>vcov(model)</code> variance-covariance matrix.</li> <li>• String which indicates the kind of uncertainty estimates to return. <ul style="list-style-type: none"> <li>– Heteroskedasticity-consistent: "HC", "HC0", "HC1", "HC2", "HC3", "HC4", "HC4m", "HC5". See <code>?sandwich::vcovHC</code></li> <li>– Heteroskedasticity and autocorrelation consistent: "HAC"</li> </ul> </li> </ul>

	<ul style="list-style-type: none"> <li>– Mixed-Models degrees of freedom: "satterthwaite", "kenward-roger"</li> <li>– Other: "NeweyWest", "KernHAC", "OPG". See the sandwich package documentation.</li> <li>• One-sided formula which indicates the name of cluster variables (e.g., <code>~unit_id</code>). This formula is passed to the <code>cluster</code> argument of the <code>sandwich::vcovCL</code> function.</li> <li>• Square covariance matrix</li> <li>• Function which returns a covariance matrix (e.g., <code>stats::vcov(model)</code>)</li> </ul>
<code>conf_level</code>	numeric value between 0 and 1. Confidence level to use to build a confidence interval.
<code>wts</code>	<p>logical, string or numeric: weights to use when computing average predictions, contrasts or slopes. These weights only affect the averaging in <code>avg_*</code>() or with the <code>by</code> argument, and not unit-level estimates. See <code>?weighted.mean</code></p> <ul style="list-style-type: none"> <li>• string: column name of the weights variable in <code>newdata</code>. When supplying a column name to <code>wts</code>, it is recommended to supply the original data (including the weights variable) explicitly to <code>newdata</code>.</li> <li>• numeric: vector of length equal to the number of rows in the original data or in <code>newdata</code> (if supplied).</li> <li>• FALSE: Equal weights.</li> <li>• TRUE: Extract weights from the fitted object with <code>insight::find_weights()</code> and use them when taking weighted averages of estimates. Warning: <code>newdata=datagrid()</code> returns a single average weight, which is equivalent to using <code>wts=FALSE</code></li> </ul>
<code>comparison</code>	<p>How should pairs of predictions be compared? Difference, ratio, odds ratio, or user-defined functions.</p> <ul style="list-style-type: none"> <li>• string: shortcuts to common contrast functions. <ul style="list-style-type: none"> <li>– Supported shortcuts strings: <code>difference</code>, <code>differenceavg</code>, <code>differenceavgwts</code>, <code>dydx</code>, <code>eyex</code>, <code>eydx</code>, <code>dyex</code>, <code>dydxavg</code>, <code>eyexavg</code>, <code>eydxavg</code>, <code>dyexavg</code>, <code>dydxavgwts</code>, <code>eyexavgwts</code>, <code>eydxavgwts</code>, <code>dyexavgwts</code>, <code>ratio</code>, <code>ratioavg</code>, <code>ratioavgwts</code>, <code>lnratio</code>, <code>lnratioavg</code>, <code>lnratioavgwts</code>, <code>lnor</code>, <code>lnoravg</code>, <code>lnoravgwts</code>, <code>lift</code>, <code>liftavg</code>, <code>expdydx</code>, <code>expdydxavg</code>, <code>expdydxavgwts</code></li> <li>– See the Comparisons section below for definitions of each transformation.</li> </ul> </li> <li>• function: accept two equal-length numeric vectors of adjusted predictions (<code>hi</code> and <code>lo</code>) and returns a vector of contrasts of the same length, or a unique numeric value. <ul style="list-style-type: none"> <li>– See the Transformations section below for examples of valid functions.</li> </ul> </li> </ul>
<code>transform</code>	string or function. Transformation applied to unit-level estimates and confidence intervals just before the function returns results. Functions must accept a vector and return a vector of the same length. Support string shortcuts: "exp", "ln"
<code>rug</code>	TRUE displays tick marks on the axes to mark the distribution of raw data.
<code>gray</code>	FALSE grayscale or color plot
<code>draw</code>	TRUE returns a <code>ggplot2</code> plot. FALSE returns a <code>data.frame</code> of the underlying data.

... Additional arguments are passed to the `predict()` method supplied by the modeling package. These arguments are particularly useful for mixed-effects or bayesian models (see the online vignettes on the `marginalEffects` website). Available arguments can vary from model to model, depending on the range of supported arguments by each modeling package. See the "Model-Specific Arguments" section of the `?slopes` documentation for a non-exhaustive list of available arguments.

## Value

A `ggplot2` object

## Model-Specific Arguments

Some model types allow model-specific arguments to modify the nature of marginal effects, predictions, marginal means, and contrasts. Please report other package-specific `predict()` arguments on Github so we can add them to the table below.

<https://github.com/vincentarelbundock/marginalEffects/issues>

Package	Class	Argument	Documentation
brms	brmsfit	ndraws	<a href="#">brms::posterior_predict</a>
		re_formula	<a href="#">brms::posterior_predict</a>
lme4	merMod	re.form	<a href="#">lme4::predict.merMod</a>
		allow.new.levels	<a href="#">lme4::predict.merMod</a>
glmmTMB	glmmTMB	re.form	<a href="#">glmmTMB::predict.glmmTMB</a>
		allow.new.levels	<a href="#">glmmTMB::predict.glmmTMB</a>
		zitype	<a href="#">glmmTMB::predict.glmmTMB</a>
mgcv	bam	exclude	<a href="#">mgcv::predict.bam</a>
robustlmm	rlmerMod	re.form	<a href="#">robustlmm::predict.rlmerMod</a>
		allow.new.levels	<a href="#">robustlmm::predict.rlmerMod</a>
MCMCglmm	MCMCglmm	ndraws	

## Examples

```
mod <- lm(mpg ~ hp * drat * factor(am), data = mtcars)

plot_comparisons(mod, variables = "hp", condition = "drat")

plot_comparisons(mod, variables = "hp", condition = c("drat", "am"))

plot_comparisons(mod, variables = "hp", condition = list("am", "drat" = 3:5))

plot_comparisons(mod, variables = "am", condition = list("hp", "drat" = range))

plot_comparisons(mod, variables = "am", condition = list("hp", "drat" = "threenum"))
```

---

plot\_predictions      *Plot Conditional or Marginal Predictions*

---

## Description

Plot predictions on the y-axis against values of one or more predictors (x-axis, colors/shapes, and facets).

The `by` argument is used to plot marginal predictions, that is, predictions made on the original data, but averaged by subgroups. This is analogous to using the `by` argument in the `predictions()` function.

The `condition` argument is used to plot conditional predictions, that is, predictions made on a user-specified grid. This is analogous to using the `newdata` argument and `datagrid()` function in a `predictions()` call. All variables whose values are not specified explicitly are treated as usual by `datagrid()`, that is, they are held at their mean or mode (or rounded mean for integers). This includes grouping variables in mixed-effects models, so analysts who fit such models may want to specify the groups of interest using the `condition` argument, or supply model-specific arguments to compute population-level estimates. See details below.

See the "Plots" vignette and website for tutorials and information on how to customize plots:

- <https://marginaleffects.com/vignettes/plot.html>
- <https://marginaleffects.com>

## Usage

```
plot_predictions(
  model,
  condition = NULL,
  by = NULL,
  newdata = NULL,
  type = NULL,
  vcov = NULL,
  conf_level = 0.95,
  wts = FALSE,
  transform = NULL,
  points = 0,
  rug = FALSE,
  gray = FALSE,
  draw = TRUE,
  ...
)
```

## Arguments

<code>model</code>	Model object
<code>condition</code>	Conditional predictions



	<ul style="list-style-type: none"> <li>• Character vector (max length 4): Names of the predictors to display.</li> <li>• Named list (max length 4): List names correspond to predictors. List elements can be: <ul style="list-style-type: none"> <li>– Numeric vector</li> <li>– Function which returns a numeric vector or a set of unique categorical values</li> <li>– Shortcut strings for common reference values: "minmax", "quartile", "threenum"</li> </ul> </li> <li>• 1: x-axis. 2: color/shape. 3: facet (wrap if no fourth variable, otherwise cols of grid). 4: facet (rows of grid).</li> <li>• Numeric variables in positions 2 and 3 are summarized by Tukey's five numbers <code>?stats::fivenum</code></li> </ul>
by	<p>Marginal predictions</p> <ul style="list-style-type: none"> <li>• Character vector (max length 3): Names of the categorical predictors to marginalize across.</li> <li>• 1: x-axis. 2: color. 3: facets.</li> </ul>
newdata	<p>When newdata is NULL, the grid is determined by the condition argument. When newdata is not NULL, the argument behaves in the same way as in the <code>predictions()</code> function.</p>
type	<p>string indicates the type (scale) of the predictions used to compute contrasts or slopes. This can differ based on the model type, but will typically be a string such as: "response", "link", "probs", or "zero". When an unsupported string is entered, the model-specific list of acceptable values is returned in an error message. When type is NULL, the first entry in the error message is used by default.</p>
vcov	<p>Type of uncertainty estimates to report (e.g., for robust standard errors). Acceptable values:</p> <ul style="list-style-type: none"> <li>• FALSE: Do not compute standard errors. This can speed up computation considerably.</li> <li>• TRUE: Unit-level standard errors using the default <code>vcov(model)</code> variance-covariance matrix.</li> <li>• String which indicates the kind of uncertainty estimates to return. <ul style="list-style-type: none"> <li>– Heteroskedasticity-consistent: "HC", "HC0", "HC1", "HC2", "HC3", "HC4", "HC4m", "HC5". See <code>?sandwich::vcovHC</code></li> <li>– Heteroskedasticity and autocorrelation consistent: "HAC"</li> <li>– Mixed-Models degrees of freedom: "satterthwaite", "kenward-roger"</li> <li>– Other: "NeweyWest", "KernHAC", "OPG". See the <code>sandwich</code> package documentation.</li> </ul> </li> <li>• One-sided formula which indicates the name of cluster variables (e.g., <code>~unit_id</code>). This formula is passed to the <code>cluster</code> argument of the <code>sandwich::vcovCL</code> function.</li> <li>• Square covariance matrix</li> <li>• Function which returns a covariance matrix (e.g., <code>stats::vcov(model)</code>)</li> </ul>
conf_level	<p>numeric value between 0 and 1. Confidence level to use to build a confidence interval.</p>

wt	logical, string or numeric: weights to use when computing average predictions, contrasts or slopes. These weights only affect the averaging in <code>avg_*</code> () or with the <code>by</code> argument, and not unit-level estimates. See <code>?weighted.mean</code> <ul style="list-style-type: none"> <li>• string: column name of the weights variable in <code>newdata</code>. When supplying a column name to <code>wt</code>, it is recommended to supply the original data (including the weights variable) explicitly to <code>newdata</code>.</li> <li>• numeric: vector of length equal to the number of rows in the original data or in <code>newdata</code> (if supplied).</li> <li>• FALSE: Equal weights.</li> <li>• TRUE: Extract weights from the fitted object with <code>insight::find_weights()</code> and use them when taking weighted averages of estimates. Warning: <code>newdata=datagrid()</code> returns a single average weight, which is equivalent to using <code>wt=FALSE</code></li> </ul>
transform	A function applied to unit-level adjusted predictions and confidence intervals just before the function returns results. For bayesian models, this function is applied to individual draws from the posterior distribution, before computing summaries.
points	Number between 0 and 1 which controls the transparency of raw data points. 0 (default) does not display any points.
rug	TRUE displays tick marks on the axes to mark the distribution of raw data.
gray	FALSE grayscale or color plot
draw	TRUE returns a <code>ggplot2</code> plot. FALSE returns a <code>data.frame</code> of the underlying data.
...	Additional arguments are passed to the <code>predict()</code> method supplied by the modeling package. These arguments are particularly useful for mixed-effects or bayesian models (see the online vignettes on the <code>marginaleffects</code> website). Available arguments can vary from model to model, depending on the range of supported arguments by each modeling package. See the "Model-Specific Arguments" section of the <code>?slopes</code> documentation for a non-exhaustive list of available arguments.

## Value

A `ggplot2` object or data frame (if `draw=FALSE`)

## Model-Specific Arguments

Some model types allow model-specific arguments to modify the nature of marginal effects, predictions, marginal means, and contrasts. Please report other package-specific `predict()` arguments on Github so we can add them to the table below.

<https://github.com/vincentarelbundock/marginaleffects/issues>

Package	Class	Argument	Documentation
brms	brmsfit	ndraws	<a href="#">brms::posterior_predict</a>
		re_formula	<a href="#">brms::posterior_predict</a>
lme4	merMod	re.form	<a href="#">lme4::predict.merMod</a>
		allow.new.levels	<a href="#">lme4::predict.merMod</a>

glmmTMB	glmmTMB	re.form	<a href="#">glmmTMB::predict.glmmTMB</a>
		allow.new.levels	<a href="#">glmmTMB::predict.glmmTMB</a>
		zitype	<a href="#">glmmTMB::predict.glmmTMB</a>
mgcv	bam	exclude	<a href="#">mgcv::predict.bam</a>
robustlmm	rlmerMod	re.form	<a href="#">robustlmm::predict.rlmerMod</a>
		allow.new.levels	<a href="#">robustlmm::predict.rlmerMod</a>
MCMCglmm	MCMCglmm	ndraws	

## Prediction types

The `type` argument determines the scale of the predictions used to compute quantities of interest with functions from the `marginalEffects` package. Admissible values for `type` depend on the model object. When users specify an incorrect value for `type`, `marginalEffects` will raise an informative error with a list of valid `type` values for the specific model object. The first entry in the list in that error message is the default `type`.

The `invlink(link)` is a special `type` defined by `marginalEffects`. It is available for some (but not all) models and functions. With this link `type`, we first compute predictions on the link scale, then we use the inverse link function to backtransform the predictions to the response scale. This is useful for models with non-linear link functions as it can ensure that confidence intervals stay within desirable bounds, ex: 0 to 1 for a logit model. Note that an average of estimates with `type="invlink(link)"` will not always be equivalent to the average of estimates with `type="response"`.

Some of the most common `type` values are:

`response`, `link`, `E`, `Ep`, `average`, `class`, `conditional`, `count`, `cum.prob`, `cumhaz`, `cumprob`, `density`, `detection`, `disp`, `ev`, `expected`, `expvalue`, `fitted`, `hazard`, `invlink(link)`, `latent`, `latent_N`, `linear`, `linear.predictor`, `linpred`, `location`, `lp`, `mean`, `numeric`, `p`, `ppd`, `pr`, `precision`, `prediction`, `prob`, `probability`, `probs`, `quantile`, `risk`, `rmst`, `scale`, `survival`, `unconditional`, `utility`, `variance`, `xb`, `zero`, `zlink`, `zprob`

## Examples

```
mod <- lm(mpg ~ hp + wt, data = mtcars)
plot_predictions(mod, condition = "wt")

mod <- lm(mpg ~ hp * wt * am, data = mtcars)
plot_predictions(mod, condition = c("hp", "wt"))

plot_predictions(mod, condition = list("hp", wt = "threenum"))

plot_predictions(mod, condition = list("hp", wt = range))
```

## Description

Plot slopes on the y-axis against values of one or more predictors (x-axis, colors/shapes, and facets).

The `by` argument is used to plot marginal slopes, that is, slopes made on the original data, but averaged by subgroups. This is analogous to using the `by` argument in the `slopes()` function.

The `condition` argument is used to plot conditional slopes, that is, slopes computed on a user-specified grid. This is analogous to using the `newdata` argument and `datagrid()` function in a `slopes()` call. All variables whose values are not specified explicitly are treated as usual by `datagrid()`, that is, they are held at their mean or mode (or rounded mean for integers). This includes grouping variables in mixed-effects models, so analysts who fit such models may want to specify the groups of interest using the `condition` argument, or supply model-specific arguments to compute population-level estimates. See details below. See the "Plots" vignette and website for tutorials and information on how to customize plots:

- <https://marginaleffects.com/vignettes/plot.html>
- <https://marginaleffects.com>

## Usage

```
plot_slopes(
  model,
  variables = NULL,
  condition = NULL,
  by = NULL,
  newdata = NULL,
  type = "response",
  vcov = NULL,
  conf_level = 0.95,
  wts = FALSE,
  slope = "dydx",
  rug = FALSE,
  gray = FALSE,
  draw = TRUE,
  ...
)
```

## Arguments

<code>model</code>	Model object
<code>variables</code>	Name of the variable whose marginal effect (slope) we want to plot on the y-axis.
<code>condition</code>	Conditional slopes <ul style="list-style-type: none"> <li>• Character vector (max length 4): Names of the predictors to display.</li> <li>• Named list (max length 4): List names correspond to predictors. List elements can be:           <ul style="list-style-type: none"> <li>– Numeric vector</li> <li>– Function which returns a numeric vector or a set of unique categorical values</li> </ul> </li> </ul>

- Shortcut strings for common reference values: "minmax", "quartile", "threenum"
  - 1: x-axis. 2: color/shape. 3: facet (wrap if no fourth variable, otherwise cols of grid). 4: facet (rows of grid).
  - Numeric variables in positions 2 and 3 are summarized by Tukey's five numbers `?stats::fivenum`.
- by** Aggregate unit-level estimates (aka, marginalize, average over). Valid inputs:
- FALSE: return the original unit-level estimates.
  - TRUE: aggregate estimates for each term.
  - Character vector of column names in `newdata` or in the data frame produced by calling the function without the `by` argument.
  - Data frame with a `by` column of group labels, and merging columns shared by `newdata` or the data frame produced by calling the same function without the `by` argument.
  - See examples below.
  - For more complex aggregations, you can use the `FUN` argument of the `hypotheses()` function. See that function's documentation and the Hypothesis Test vignettes on the `marginalEffects` website.
- newdata** When `newdata` is NULL, the grid is determined by the `condition` argument. When `newdata` is not NULL, the argument behaves in the same way as in the `slopes()` function.
- type** string indicates the type (scale) of the predictions used to compute contrasts or slopes. This can differ based on the model type, but will typically be a string such as: "response", "link", "probs", or "zero". When an unsupported string is entered, the model-specific list of acceptable values is returned in an error message. When `type` is NULL, the first entry in the error message is used by default.
- vcov** Type of uncertainty estimates to report (e.g., for robust standard errors). Acceptable values:
- FALSE: Do not compute standard errors. This can speed up computation considerably.
  - TRUE: Unit-level standard errors using the default `vcov(model)` variance-covariance matrix.
  - String which indicates the kind of uncertainty estimates to return.
    - Heteroskedasticity-consistent: "HC", "HC0", "HC1", "HC2", "HC3", "HC4", "HC4m", "HC5". See `?sandwich::vcovHC`
    - Heteroskedasticity and autocorrelation consistent: "HAC"
    - Mixed-Models degrees of freedom: "satterthwaite", "kenward-roger"
    - Other: "NeweyWest", "KernHAC", "OPG". See the `sandwich` package documentation.
  - One-sided formula which indicates the name of cluster variables (e.g., `~unit_id`). This formula is passed to the `cluster` argument of the `sandwich::vcovCL` function.
  - Square covariance matrix

	<ul style="list-style-type: none"> <li>Function which returns a covariance matrix (e.g., <code>stats::vcov(model)</code>)</li> </ul>
<code>conf_level</code>	numeric value between 0 and 1. Confidence level to use to build a confidence interval.
<code>wts</code>	<p>logical, string or numeric: weights to use when computing average predictions, contrasts or slopes. These weights only affect the averaging in <code>avg_*</code>() or with the <code>by</code> argument, and not unit-level estimates. See <code>?weighted.mean</code></p> <ul style="list-style-type: none"> <li>string: column name of the weights variable in <code>newdata</code>. When supplying a column name to <code>wts</code>, it is recommended to supply the original data (including the weights variable) explicitly to <code>newdata</code>.</li> <li>numeric: vector of length equal to the number of rows in the original data or in <code>newdata</code> (if supplied).</li> <li>FALSE: Equal weights.</li> <li>TRUE: Extract weights from the fitted object with <code>insight::find_weights()</code> and use them when taking weighted averages of estimates. Warning: <code>newdata=datagrid()</code> returns a single average weight, which is equivalent to using <code>wts=FALSE</code></li> </ul>
<code>slope</code>	<p>string indicates the type of slope or (semi-)elasticity to compute:</p> <ul style="list-style-type: none"> <li>"dydx": <math>dY/dX</math></li> <li>"eyex": <math>dY/dX * Y / X</math></li> <li>"eydx": <math>dY/dX * Y</math></li> <li>"dyex": <math>dY/dX / X</math></li> <li>Y is the predicted value of the outcome; X is the observed value of the predictor.</li> </ul>
<code>rug</code>	TRUE displays tick marks on the axes to mark the distribution of raw data.
<code>gray</code>	FALSE grayscale or color plot
<code>draw</code>	TRUE returns a <code>ggplot2</code> plot. FALSE returns a <code>data.frame</code> of the underlying data.
<code>...</code>	Additional arguments are passed to the <code>predict()</code> method supplied by the modeling package. These arguments are particularly useful for mixed-effects or bayesian models (see the online vignettes on the <code>marginaleffects</code> website). Available arguments can vary from model to model, depending on the range of supported arguments by each modeling package. See the "Model-Specific Arguments" section of the <code>?slopes</code> documentation for a non-exhaustive list of available arguments.

### Value

A `ggplot2` object

### Model-Specific Arguments

Some model types allow model-specific arguments to modify the nature of marginal effects, predictions, marginal means, and contrasts. Please report other package-specific `predict()` arguments on Github so we can add them to the table below.

<https://github.com/vincentarelbundock/marginaleffects/issues>

Package	Class	Argument	Documentation
brms	brmsfit	ndraws	<a href="#">brms::posterior_predict</a>
		re_formula	<a href="#">brms::posterior_predict</a>
lme4	merMod	re.form	<a href="#">lme4::predict.merMod</a>
		allow.new.levels	<a href="#">lme4::predict.merMod</a>
glmmTMB	glmmTMB	re.form	<a href="#">glmmTMB::predict.glmmTMB</a>
		allow.new.levels	<a href="#">glmmTMB::predict.glmmTMB</a>
		zitype	<a href="#">glmmTMB::predict.glmmTMB</a>
mgcv	bam	exclude	<a href="#">mgcv::predict.bam</a>
robustlmm	rlmerMod	re.form	<a href="#">robustlmm::predict.rlmerMod</a>
		allow.new.levels	<a href="#">robustlmm::predict.rlmerMod</a>
MCMCglmm	MCMCglmm	ndraws	

## Examples

```
library(marginaleffects)
mod <- lm(mpg ~ hp * drat * factor(am), data = mtcars)

plot_slopes(mod, variables = "hp", condition = "drat")

plot_slopes(mod, variables = "hp", condition = c("drat", "am"))

plot_slopes(mod, variables = "hp", condition = list("am", "drat" = 3:5))

plot_slopes(mod, variables = "am", condition = list("hp", "drat" = range))

plot_slopes(mod, variables = "am", condition = list("hp", "drat" = "threenum"))
```

---

posterior\_draws      *Extract Posterior Draws or Bootstrap Resamples from marginaleffects Objects*

---

## Description

Extract Posterior Draws or Bootstrap Resamples from marginaleffects Objects

## Usage

```
posterior_draws(x, shape = "long")
```

## Arguments

- |       |   |
|-------|---|
| x     | An object produced by a marginaleffects package function, such as <code>predictions()</code> , <code>avg_slopes()</code> , <code>hypotheses()</code> , etc. |
| shape | string indicating the shape of the output format: <ul style="list-style-type: none"> <li>• "long": long format data frame</li> </ul>                        |

- "DxP": Matrix with draws as rows and parameters as columns
- "PxD": Matrix with draws as rows and parameters as columns
- "rvar": Random variable datatype (see posterior package documentation).

### Value

A data.frame with drawid and draw columns.

---

predictions

*Predictions*

---

### Description

Outcome predicted by a fitted model on a specified scale for a given combination of values of the predictor variables, such as their observed values, their means, or factor levels (a.k.a. "reference grid").

- `predictions()`: unit-level (conditional) estimates.
- `avg_predictions()`: average (marginal) estimates.

The `newdata` argument and the `datagrid()` function can be used to control where statistics are evaluated in the predictor space: "at observed values", "at the mean", "at representative values", etc.

See the predictions vignette and package website for worked examples and case studies:

- <https://marginaleffects.com/vignettes/predictions.html>
- <https://marginaleffects.com/>

### Usage

```
predictions(
  model,
  newdata = NULL,
  variables = NULL,
  vcov = TRUE,
  conf_level = 0.95,
  type = NULL,
  by = FALSE,
  byfun = NULL,
  wts = FALSE,
  transform = NULL,
  hypothesis = NULL,
  equivalence = NULL,
  p_adjust = NULL,
  df = Inf,
  numberiv = "fdforward",
  ...
)
```



```

)

avg_predictions(
  model,
  newdata = NULL,
  variables = NULL,
  vcov = TRUE,
  conf_level = 0.95,
  type = NULL,
  by = TRUE,
  byfun = NULL,
  wts = FALSE,
  transform = NULL,
  hypothesis = NULL,
  equivalence = NULL,
  p_adjust = NULL,
  df = Inf,
  numberiv = "fdforward",
  ...
)

```

## Arguments

model	Model object
newdata	<p>Grid of predictor values at which we evaluate predictions.</p> <ul style="list-style-type: none"> <li>• Warning: Please avoid modifying your dataset between fitting the model and calling a <code>margineffects</code> function. This can sometimes lead to unexpected results.</li> <li>• NULL (default): Unit-level predictions for each observed value in the dataset (empirical distribution). The dataset is retrieved using <code>insight::get_data()</code>, which tries to extract data from the environment. This may produce unexpected results if the original data frame has been altered since fitting the model.</li> <li>• string: <ul style="list-style-type: none"> <li>– "mean": Predictions at the Mean. Predictions when each predictor is held at its mean or mode.</li> <li>– "median": Predictions at the Median. Predictions when each predictor is held at its median or mode.</li> <li>– "marginalmeans": Predictions at Marginal Means. See Details section below.</li> <li>– "tukey": Predictions at Tukey's 5 numbers.</li> <li>– "grid": Predictions on a grid of representative numbers (Tukey's 5 numbers and unique values of categorical predictors).</li> </ul> </li> <li>• <code>datagrid()</code> call to specify a custom grid of regressors. For example: <ul style="list-style-type: none"> <li>– <code>newdata = datagrid(cyl = c(4, 6))</code>: <code>cyl</code> variable equal to 4 and 6 and other regressors fixed at their means or modes.</li> <li>– See the Examples section and the <code>datagrid()</code> documentation.</li> </ul> </li> </ul>

- `subset()` call with a single argument to select a subset of the dataset used to fit the model, ex: `newdata = subset(treatment == 1)`
- `dplyr::filter()` call with a single argument to select a subset of the dataset used to fit the model, ex: `newdata = filter(treatment == 1)`

variables

Counterfactual variables.

- Output:
  - `predictions()`: The entire dataset is replicated once for each unique combination of `variables`, and predictions are made.
  - `avg_predictions()`: The entire dataset is replicated, predictions are made, and they are marginalized by `variables` categories.
  - Warning: This can be expensive in large datasets.
  - Warning: Users who need "conditional" predictions should use the `newdata` argument instead of `variables`.
- Input:
  - NULL: computes one prediction per row of `newdata`
  - Character vector: the dataset is replicated once of every combination of unique values of the variables identified in `variables`.
  - Named list: names identify the subset of variables of interest and their values. For numeric variables, the `variables` argument supports functions and string shortcuts:
    - \* A function which returns a numeric value
    - \* Numeric vector: Contrast between the 2nd element and the 1st element of the `x` vector.
    - \* "iqr": Contrast across the interquartile range of the regressor.
    - \* "sd": Contrast across one standard deviation around the regressor mean.
    - \* "2sd": Contrast across two standard deviations around the regressor mean.
    - \* "minmax": Contrast between the maximum and the minimum values of the regressor.
    - \* "threenum": mean and 1 standard deviation on both sides
    - \* "fivenum": Tukey's five numbers

vcov

Type of uncertainty estimates to report (e.g., for robust standard errors). Acceptable values:

- FALSE: Do not compute standard errors. This can speed up computation considerably.
- TRUE: Unit-level standard errors using the default `vcov(model)` variance-covariance matrix.
- String which indicates the kind of uncertainty estimates to return.
  - Heteroskedasticity-consistent: "HC", "HC0", "HC1", "HC2", "HC3", "HC4", "HC4m", "HC5". See `?sandwich::vcovHC`
  - Heteroskedasticity and autocorrelation consistent: "HAC"
  - Mixed-Models degrees of freedom: "satterthwaite", "kenward-roger"

	<ul style="list-style-type: none"> <li>- Other: "NeweyWest", "KernHAC", "OPG". See the sandwich package documentation.</li> <li>• One-sided formula which indicates the name of cluster variables (e.g., <code>~unit_id</code>). This formula is passed to the <code>cluster</code> argument of the <code>sandwich::vcovCL</code> function.</li> <li>• Square covariance matrix</li> <li>• Function which returns a covariance matrix (e.g., <code>stats::vcov(model)</code>)</li> </ul>
<code>conf_level</code>	numeric value between 0 and 1. Confidence level to use to build a confidence interval.
<code>type</code>	string indicates the type (scale) of the predictions used to compute contrasts or slopes. This can differ based on the model type, but will typically be a string such as: "response", "link", "probs", or "zero". When an unsupported string is entered, the model-specific list of acceptable values is returned in an error message. When <code>type</code> is NULL, the first entry in the error message is used by default.
<code>by</code>	Aggregate unit-level estimates (aka, marginalize, average over). Valid inputs: <ul style="list-style-type: none"> <li>• FALSE: return the original unit-level estimates.</li> <li>• TRUE: aggregate estimates for each term.</li> <li>• Character vector of column names in <code>newdata</code> or in the data frame produced by calling the function without the <code>by</code> argument.</li> <li>• Data frame with a <code>by</code> column of group labels, and merging columns shared by <code>newdata</code> or the data frame produced by calling the same function without the <code>by</code> argument.</li> <li>• See examples below.</li> <li>• For more complex aggregations, you can use the <code>FUN</code> argument of the <code>hypotheses()</code> function. See that function's documentation and the Hypothesis Test vignettes on the <code>marginalEffects</code> website.</li> </ul>
<code>byfun</code>	A function such as <code>mean()</code> or <code>sum()</code> used to aggregate estimates within the subgroups defined by the <code>by</code> argument. NULL uses the <code>mean()</code> function. Must accept a numeric vector and return a single numeric value. This is sometimes used to take the sum or mean of predicted probabilities across outcome or predictor levels. See examples section.
<code>wts</code>	logical, string or numeric: weights to use when computing average predictions, contrasts or slopes. These weights only affect the averaging in <code>avg_*</code> or with the <code>by</code> argument, and not unit-level estimates. See <code>?weighted.mean</code> <ul style="list-style-type: none"> <li>• string: column name of the weights variable in <code>newdata</code>. When supplying a column name to <code>wts</code>, it is recommended to supply the original data (including the weights variable) explicitly to <code>newdata</code>.</li> <li>• numeric: vector of length equal to the number of rows in the original data or in <code>newdata</code> (if supplied).</li> <li>• FALSE: Equal weights.</li> <li>• TRUE: Extract weights from the fitted object with <code>insight::find_weights()</code> and use them when taking weighted averages of estimates. Warning: <code>newdata=datagrid()</code> returns a single average weight, which is equivalent to using <code>wts=FALSE</code></li> </ul>

- transform** A function applied to unit-level adjusted predictions and confidence intervals just before the function returns results. For bayesian models, this function is applied to individual draws from the posterior distribution, before computing summaries.
- hypothesis** specify a hypothesis test or custom contrast using a numeric value, vector, or matrix, a string, a string formula, or a function.
- Numeric:
    - Single value: the null hypothesis used in the computation of Z and p (before applying transform).
    - Vector: Weights to compute a linear combination of (custom contrast between) estimates. Length equal to the number of rows generated by the same function call, but without the hypothesis argument.
    - Matrix: Each column is a vector of weights, as describe above, used to compute a distinct linear combination of (contrast between) estimates. The column names of the matrix are used as labels in the output.
  - String formula to specify linear or non-linear hypothesis tests. If the term column uniquely identifies rows, terms can be used in the formula. Otherwise, use b1, b2, etc. to identify the position of each parameter. The b\* wildcard can be used to test hypotheses on all estimates. If a named vector is used, the names are used as labels in the output. Examples:
    - `hp = drat`
    - `hp + drat = 12`
    - `b1 + b2 + b3 = 0`
    - `b* / b1 = 1`
  - String:
    - "pairwise": pairwise differences between estimates in each row.
    - "reference": differences between the estimates in each row and the estimate in the first row.
    - "sequential": difference between an estimate and the estimate in the next row.
    - "meandev": difference between an estimate and the mean of all estimates.
    - "meanotherdev": difference between an estimate and the mean of all other estimates, excluding the current one.
    - "revpairwise", "revreference", "revsequential": inverse of the corresponding hypotheses, as described above.
  - Function:
    - Accepts an argument x: object produced by a `marginalEffects` function or a data frame with column `rowid` and `estimate`
    - Returns a data frame with columns `term` and `estimate` (mandatory) and `rowid` (optional).
    - The function can also accept optional input arguments: `newdata`, `by`, `draws`.
    - This function approach will not work for Bayesian models or with bootstrapping. In those cases, it is easy to use `posterior_draws()` to extract and manipulate the draws directly.

	<ul style="list-style-type: none"> <li>• See the Examples section below and the vignette: <a href="https://marginaleffects.com/vignettes/hypothesis.ht">https://marginaleffects.com/vignettes/hypothesis.ht</a></li> </ul>
equivalence	Numeric vector of length 2: bounds used for the two-one-sided test (TOST) of equivalence, and for the non-inferiority and non-superiority tests. See Details section below.
p_adjust	Adjust p-values for multiple comparisons: "holm", "hochberg", "hommel", "bonferroni", "BH", "BY", or "fdr". See <a href="#">stats::p.adjust</a>
df	Degrees of freedom used to compute p values and confidence intervals. A single numeric value between 1 and Inf. When df is Inf, the normal distribution is used. When df is finite, the t distribution is used. See <a href="#">insight::get_df</a> for a convenient function to extract degrees of freedom. Ex: <code>slopes(model, df = insight::get_df(model))</code>
numderiv	string or list of strings indicating the method to use to for the numeric differentiation used in to compute delta method standard errors. <ul style="list-style-type: none"> <li>• "fdforward": finite difference method with forward differences</li> <li>• "fdcenter": finite difference method with central differences (default)</li> <li>• "richardson": Richardson extrapolation method</li> <li>• Extra arguments can be specified by passing a list to the numDeriv argument, with the name of the method first and named arguments following, ex: <code>numderiv=list("fdcenter", eps = 1e-5)</code>. When an unknown argument is used, <code>marginaleffects</code> prints the list of valid arguments for each method.</li> </ul>
...	Additional arguments are passed to the <code>predict()</code> method supplied by the modeling package. These arguments are particularly useful for mixed-effects or bayesian models (see the online vignettes on the <code>marginaleffects</code> website). Available arguments can vary from model to model, depending on the range of supported arguments by each modeling package. See the "Model-Specific Arguments" section of the <code>?slopes</code> documentation for a non-exhaustive list of available arguments.

## Value

A data.frame with one row per observation and several columns:

- `rowid`: row number of the `newdata` data frame
- `type`: prediction type, as defined by the `type` argument
- `group`: (optional) value of the grouped outcome (e.g., categorical outcome models)
- `estimate`: predicted outcome
- `std.error`: standard errors computed using the delta method.
- `p.value`: p value associated to the `estimate` column. The null is determined by the `hypothesis` argument (0 by default), and p values are computed before applying the `transform` argument. For models of class `feglm`, `Gam`, `glm` and `negbin`, p values are computed on the link scale by default unless the `type` argument is specified explicitly.
- `s.value`: Shannon information transforms of p values. How many consecutive "heads" tosses would provide the same amount of evidence (or "surprise") against the null hypothesis that the coin is fair? The purpose of S is to calibrate the analyst's intuition about the strength of

evidence encoded in `p` against a well-known physical phenomenon. See Greenland (2019) and Cole et al. (2020).

- `conf.low`: lower bound of the confidence interval (or equal-tailed interval for bayesian models)
- `conf.high`: upper bound of the confidence interval (or equal-tailed interval for bayesian models)

See `?print.marginaleffects` for printing options.

## Functions

- `avg_predictions()`: Average predictions

## Standard errors using the delta method

Standard errors for all quantities estimated by `marginaleffects` can be obtained via the delta method. This requires differentiating a function with respect to the coefficients in the model using a finite difference approach. In some models, the delta method standard errors can be sensitive to various aspects of the numeric differentiation strategy, including the step size. By default, the step size is set to  $1e-8$ , or to  $1e-4$  times the smallest absolute model coefficient, whichever is largest.

`marginaleffects` can delegate numeric differentiation to the `numDeriv` package, which allows more flexibility. To do this, users can pass arguments to the `numDeriv::jacobian` function through a global option. For example:

- `options(marginaleffects_numDeriv = list(method = "simple", method.args = list(eps = 1e-6)))`
- `options(marginaleffects_numDeriv = list(method = "Richardson", method.args = list(eps = 1e-5)))`
- `options(marginaleffects_numDeriv = NULL)`

See the "Standard Errors and Confidence Intervals" vignette on the `marginaleffects` website for more details on the computation of standard errors:

<https://marginaleffects.com/vignettes/uncertainty.html>

Note that the `inferences()` function can be used to compute uncertainty estimates using a bootstrap or simulation-based inference. See the vignette:

<https://marginaleffects.com/vignettes/bootstrap.html>

## Model-Specific Arguments

Some model types allow model-specific arguments to modify the nature of marginal effects, predictions, marginal means, and contrasts. Please report other package-specific `predict()` arguments on Github so we can add them to the table below.

<https://github.com/vincentarelbundock/marginaleffects/issues>

Package	Class	Argument	Documentation
<code>brms</code>	<code>brmsfit</code>	<code>ndraws</code> <code>re_formula</code>	<a href="#">brms::posterior_predict</a> <a href="#">brms::posterior_predict</a>

lme4	merMod	re.form allow.new.levels	<a href="#">lme4::predict.merMod</a> <a href="#">lme4::predict.merMod</a>
glmmTMB	glmmTMB	re.form allow.new.levels zitype	<a href="#">glmmTMB::predict.glmmTMB</a> <a href="#">glmmTMB::predict.glmmTMB</a> <a href="#">glmmTMB::predict.glmmTMB</a>
mgcv	bam	exclude	<a href="#">mgcv::predict.bam</a>
robustlmm	rlmerMod	re.form allow.new.levels	<a href="#">robustlmm::predict.rlmerMod</a> <a href="#">robustlmm::predict.rlmerMod</a>
MCMCglmm	MCMCglmm	ndraws	

### Bayesian posterior summaries

By default, credible intervals in bayesian models are built as equal-tailed intervals. This can be changed to a highest density interval by setting a global option:

```
options("marginaleffects_posterior_interval" = "eti")
options("marginaleffects_posterior_interval" = "hdi")
```

By default, the center of the posterior distribution in bayesian models is identified by the median. Users can use a different summary function by setting a global option:

```
options("marginaleffects_posterior_center" = "mean")
options("marginaleffects_posterior_center" = "median")
```

When estimates are averaged using the `by` argument, the `tidy()` function, or the `summary()` function, the posterior distribution is marginalized twice over. First, we take the average *across* units but *within* each iteration of the MCMC chain, according to what the user requested in `by` argument or `tidy()/summary()` functions. Then, we identify the center of the resulting posterior using the function supplied to the `"marginaleffects_posterior_center"` option (the median by default).

### Equivalence, Inferiority, Superiority

$\theta$  is an estimate,  $\sigma_\theta$  its estimated standard error, and  $[a, b]$  are the bounds of the interval supplied to the equivalence argument.

Non-inferiority:

- $H_0: \theta \leq a$
- $H_1: \theta > a$
- $t = (\theta - a) / \sigma_\theta$
- p: Upper-tail probability

Non-superiority:

- $H_0: \theta \geq b$
- $H_1: \theta < b$
- $t = (\theta - b) / \sigma_\theta$
- p: Lower-tail probability

Equivalence: Two One-Sided Tests (TOST)

- p: Maximum of the non-inferiority and non-superiority p values.

Thanks to Russell V. Lenth for the excellent `emmeans` package and documentation which inspired this feature.

### Prediction types

The `type` argument determines the scale of the predictions used to compute quantities of interest with functions from the `marginalEffects` package. Admissible values for `type` depend on the model object. When users specify an incorrect value for `type`, `marginalEffects` will raise an informative error with a list of valid `type` values for the specific model object. The first entry in the list in that error message is the default `type`.

The `invlink(link)` is a special `type` defined by `marginalEffects`. It is available for some (but not all) models and functions. With this `link` `type`, we first compute predictions on the link scale, then we use the inverse link function to backtransform the predictions to the response scale. This is useful for models with non-linear link functions as it can ensure that confidence intervals stay within desirable bounds, ex: 0 to 1 for a logit model. Note that an average of estimates with `type="invlink(link)"` will not always be equivalent to the average of estimates with `type="response"`.

Some of the most common `type` values are:

`response`, `link`, `E`, `Ep`, `average`, `class`, `conditional`, `count`, `cum.prob`, `cumhaz`, `cumprob`, `density`, `detection`, `disp`, `ev`, `expected`, `expvalue`, `fitted`, `hazard`, `invlink(link)`, `latent`, `latent_N`, `linear`, `linear.predictor`, `linpred`, `location`, `lp`, `mean`, `numeric`, `p`, `ppd`, `pr`, `precision`, `prediction`, `prob`, `probability`, `probs`, `quantile`, `risk`, `rmst`, `scale`, `survival`, `unconditional`, `utility`, `variance`, `xb`, `zero`, `zlink`, `zprob`

### Order of operations

Behind the scenes, the arguments of `marginalEffects` functions are evaluated in this order:

1. `newdata`
2. `variables`
3. `comparison` and `slopes`
4. `by`
5. `vcov`
6. `hypothesis`
7. `transform`

### Parallel computation

The `slopes()` and `comparisons()` functions can use parallelism to speed up computation. Operations are parallelized for the computation of standard errors, at the model coefficient level. There is always considerable overhead when using parallel computation, mainly involved in passing the whole dataset to the different processes. Thus, parallel computation is most likely to be useful when the model includes many parameters and the dataset is relatively small.

Warning: In many cases, parallel processing will not be useful at all.

To activate parallel computation, users must load the `future.apply` package, call `plan()` function, and set a global option. For example:



```
library(future.apply)
plan("multicore", workers = 4)
options(marginaleffects_parallel = TRUE)

slopes(model)
```

To disable parallelism in `marginaleffects` altogether, you can set a global option:

```
options(marginaleffects_parallel = FALSE)
```

## References

- Greenland S. 2019. "Valid P-Values Behave Exactly as They Should: Some Misleading Criticisms of P-Values and Their Resolution With S-Values." *The American Statistician*. 73(S1): 106–114.
- Cole, Stephen R, Jessie K Edwards, and Sander Greenland. 2020. "Surprise!" *American Journal of Epidemiology* 190 (2): 191–93. <https://doi.org/10.1093/aje/kwaa136>

## Examples

```
# Adjusted Prediction for every row of the original dataset
mod <- lm(mpg ~ hp + factor(cyl), data = mtcars)
pred <- predictions(mod)
head(pred)

# Adjusted Predictions at User-Specified Values of the Regressors
predictions(mod, newdata = datagrid(hp = c(100, 120), cyl = 4))

m <- lm(mpg ~ hp + drat + factor(cyl) + factor(am), data = mtcars)
predictions(m, newdata = datagrid(FUN_factor = unique, FUN_numeric = median))

# Average Adjusted Predictions (AAP)
library(dplyr)
mod <- lm(mpg ~ hp * am * vs, mtcars)

avg_predictions(mod)

predictions(mod, by = "am")

# Conditional Adjusted Predictions
plot_predictions(mod, condition = "hp")

# Counterfactual predictions with the `variables` argument
# the `mtcars` dataset has 32 rows

mod <- lm(mpg ~ hp + am, data = mtcars)
p <- predictions(mod)
head(p)
nrow(p)
```

```
# average counterfactual predictions
avg_predictions(mod, variables = "am")

# counterfactual predictions obtained by replicating the entire for different
# values of the predictors
p <- predictions(mod, variables = list(hp = c(90, 110)))
nrow(p)

# hypothesis test: is the prediction in the 1st row equal to the prediction in the 2nd row
mod <- lm(mpg ~ wt + drat, data = mtcars)

predictions(
  mod,
  newdata = datagrid(wt = 2:3),
  hypothesis = "b1 = b2")

# same hypothesis test using row indices
predictions(
  mod,
  newdata = datagrid(wt = 2:3),
  hypothesis = "b1 - b2 = 0")

# same hypothesis test using numeric vector of weights
predictions(
  mod,
  newdata = datagrid(wt = 2:3),
  hypothesis = c(1, -1))

# two custom contrasts using a matrix of weights
lc <- matrix(c(
  1, -1,
  2, 3),
  ncol = 2)
predictions(
  mod,
  newdata = datagrid(wt = 2:3),
  hypothesis = lc)

# `by` argument
mod <- lm(mpg ~ hp * am * vs, data = mtcars)
predictions(mod, by = c("am", "vs"))

library(nnet)
nom <- multinom(factor(gear) ~ mpg + am * vs, data = mtcars, trace = FALSE)

# first 5 raw predictions
predictions(nom, type = "probs") |> head()

# average predictions
avg_predictions(nom, type = "probs", by = "group")
```

```

by <- data.frame(
  group = c("3", "4", "5"),
  by = c("3,4", "3,4", "5"))

predictions(nom, type = "probs", by = by)

# sum of predicted probabilities for combined response levels
mod <- multinom(factor(cyl) ~ mpg + am, data = mtcars, trace = FALSE)
by <- data.frame(
  by = c("4,6", "4,6", "8"),
  group = as.character(c(4, 6, 8)))
predictions(mod, newdata = "mean", byfun = sum, by = by)

```

---

print.marginaleffects *Print marginaleffects objects*

---

## Description

This function controls the text which is printed to the console when one of the core `marginaleffects` functions is called and the object is returned: `predictions()`, `comparisons()`, `slopes()`, `hypotheses()`, `avg_predictions()`, `avg_comparisons()`, `avg_slopes()`.

All of those functions return standard data frames. Columns can be extracted by name, `predictions(model)$estimate`, and all the usual data manipulation functions work out-of-the-box: `colnames()`, `head()`, `subset()`, `dplyr::filter()`, `dplyr::arrange()`, etc.

Some of the data columns are not printed by default. You can disable pretty printing and print the full results as a standard data frame using the `style` argument or by applying `as.data.frame()` on the object. See examples below.

## Usage

```

## S3 method for class 'marginaleffects'
print(
  x,
  style = getOption("marginaleffects_print_style", default = "summary"),
  digits = getOption("marginaleffects_print_digits", default = 3),
  p_eps = getOption("marginaleffects_print_p_eps", default = 0.001),
  topn = getOption("marginaleffects_print_topn", default = 5),
  nrows = getOption("marginaleffects_print_nrows", default = 30),
  ncols = getOption("marginaleffects_print_ncols", default = 30),
  type = getOption("marginaleffects_print_type", default = TRUE),
  column_names = getOption("marginaleffects_print_column_names", default = TRUE),
  ...
)

```

**Arguments**

x	An object produced by one of the <code>marginaleffects</code> package functions.
style	"summary", "data.frame", or "tinytable"
digits	The number of digits to display.
p_eps	p values smaller than this number are printed in "<0.001" style.
topn	The number of rows to be printed from the beginning and end of tables with more than <code>nrows</code> rows.
nrows	The number of rows which will be printed before truncation.
ncols	The maximum number of column names to display at the bottom of the printed output.
type	boolean: should the type be printed?
column_names	boolean: should the column names be printed?
...	Other arguments are currently ignored.

**Examples**

```
library(marginaleffects)
mod <- lm(mpg ~ hp + am + factor(gear), data = mtcars)
p <- predictions(mod, by = c("am", "gear"))
p

subset(p, am == 1)

print(p, style = "data.frame")

data.frame(p)
```

---

slopes

*Slopes (aka Partial derivatives, Marginal Effects, or Trends)*


---

**Description**

Partial derivative of the regression equation with respect to a regressor of interest.

- `slopes()`: unit-level (conditional) estimates.
- `avg_slopes()`: average (marginal) estimates.

The `newdata` argument and the `datagrid()` function can be used to control where statistics are evaluated in the predictor space: "at observed values", "at the mean", "at representative values", etc.

See the `slopes` vignette and package website for worked examples and case studies:

- <https://marginaleffects.com/vignettes/slopes.html>
- <https://marginaleffects.com/>

**Usage**

```
slopes(  
  model,  
  newdata = NULL,  
  variables = NULL,  
  type = NULL,  
  by = FALSE,  
  vcov = TRUE,  
  conf_level = 0.95,  
  slope = "dydx",  
  wts = FALSE,  
  hypothesis = NULL,  
  equivalence = NULL,  
  p_adjust = NULL,  
  df = Inf,  
  eps = NULL,  
  numderiv = "fdforward",  
  ...  
)
```

```
avg_slopes(  
  model,  
  newdata = NULL,  
  variables = NULL,  
  type = NULL,  
  by = TRUE,  
  vcov = TRUE,  
  conf_level = 0.95,  
  slope = "dydx",  
  wts = FALSE,  
  hypothesis = NULL,  
  equivalence = NULL,  
  p_adjust = NULL,  
  df = Inf,  
  eps = NULL,  
  numderiv = "fdforward",  
  ...  
)
```

**Arguments**

model	Model object
newdata	Grid of predictor values at which we evaluate the slopes. <ul style="list-style-type: none"><li>• Warning: Please avoid modifying your dataset between fitting the model and calling a <code>marginaleffects</code> function. This can sometimes lead to unexpected results.</li><li>• NULL (default): Unit-level slopes for each observed value in the dataset (empirical distribution). The dataset is retrieved using <code>insight::get_data()</code>,</li></ul>

which tries to extract data from the environment. This may produce unexpected results if the original data frame has been altered since fitting the model.

- `datagrid()` call to specify a custom grid of regressors. For example:
  - `newdata = datagrid(cyl = c(4, 6))`: `cyl` variable equal to 4 and 6 and other regressors fixed at their means or modes.
  - See the Examples section and the `datagrid()` documentation.
- `subset()` call with a single argument to select a subset of the dataset used to fit the model, ex: `newdata = subset(treatment == 1)`
- `dplyr::filter()` call with a single argument to select a subset of the dataset used to fit the model, ex: `newdata = filter(treatment == 1)`
- string:
  - "mean": Marginal Effects at the Mean. Slopes when each predictor is held at its mean or mode.
  - "median": Marginal Effects at the Median. Slopes when each predictor is held at its median or mode.
  - "marginalmeans": Marginal Effects at Marginal Means. See Details section below.
  - "tukey": Marginal Effects at Tukey's 5 numbers.
  - "grid": Marginal Effects on a grid of representative numbers (Tukey's 5 numbers and unique values of categorical predictors).

variables

Focal variables

- NULL: compute slopes or comparisons for all the variables in the model object (can be slow).
- Character vector: subset of variables (usually faster).

type

string indicates the type (scale) of the predictions used to compute contrasts or slopes. This can differ based on the model type, but will typically be a string such as: "response", "link", "probs", or "zero". When an unsupported string is entered, the model-specific list of acceptable values is returned in an error message. When type is NULL, the first entry in the error message is used by default.

by

Aggregate unit-level estimates (aka, marginalize, average over). Valid inputs:

- FALSE: return the original unit-level estimates.
- TRUE: aggregate estimates for each term.
- Character vector of column names in `newdata` or in the data frame produced by calling the function without the `by` argument.
- Data frame with a `by` column of group labels, and merging columns shared by `newdata` or the data frame produced by calling the same function without the `by` argument.
- See examples below.
- For more complex aggregations, you can use the `FUN` argument of the `hypotheses()` function. See that function's documentation and the Hypothesis Test vignettes on the `marginaleffects` website.

vcov

Type of uncertainty estimates to report (e.g., for robust standard errors). Acceptable values:

- FALSE: Do not compute standard errors. This can speed up computation considerably.
  - TRUE: Unit-level standard errors using the default `vcov(model)` variance-covariance matrix.
  - String which indicates the kind of uncertainty estimates to return.
    - Heteroskedasticity-consistent: "HC", "HC0", "HC1", "HC2", "HC3", "HC4", "HC4m", "HC5". See `?sandwich::vcovHC`
    - Heteroskedasticity and autocorrelation consistent: "HAC"
    - Mixed-Models degrees of freedom: "satterthwaite", "kenward-roger"
    - Other: "NeweyWest", "KernHAC", "OPG". See the `sandwich` package documentation.
  - One-sided formula which indicates the name of cluster variables (e.g., `~unit_id`). This formula is passed to the `cluster` argument of the `sandwich::vcovCL` function.
  - Square covariance matrix
  - Function which returns a covariance matrix (e.g., `stats::vcov(model)`)
- `conf_level` numeric value between 0 and 1. Confidence level to use to build a confidence interval.
- `slope` string indicates the type of slope or (semi-)elasticity to compute:
- "dydx":  $dY/dX$
  - "eyex":  $dY/dX * Y / X$
  - "eydx":  $dY/dX * Y$
  - "dyex":  $dY/dX / X$
  - Y is the predicted value of the outcome; X is the observed value of the predictor.
- `wts` logical, string or numeric: weights to use when computing average predictions, contrasts or slopes. These weights only affect the averaging in `avg_*()` or with the `by` argument, and not unit-level estimates. See `?weighted.mean`
- string: column name of the weights variable in `newdata`. When supplying a column name to `wts`, it is recommended to supply the original data (including the weights variable) explicitly to `newdata`.
  - numeric: vector of length equal to the number of rows in the original data or in `newdata` (if supplied).
  - FALSE: Equal weights.
  - TRUE: Extract weights from the fitted object with `insight::find_weights()` and use them when taking weighted averages of estimates. Warning: `newdata=datagrid()` returns a single average weight, which is equivalent to using `wts=FALSE`
- `hypothesis` specify a hypothesis test or custom contrast using a numeric value, vector, or matrix, a string, a string formula, or a function.
- Numeric:
    - Single value: the null hypothesis used in the computation of Z and p (before applying transform).
    - Vector: Weights to compute a linear combination of (custom contrast between) estimates. Length equal to the number of rows generated by the same function call, but without the `hypothesis` argument.

- Matrix: Each column is a vector of weights, as describe above, used to compute a distinct linear combination of (contrast between) estimates. The column names of the matrix are used as labels in the output.
  - String formula to specify linear or non-linear hypothesis tests. If the term column uniquely identifies rows, terms can be used in the formula. Otherwise, use `b1`, `b2`, etc. to identify the position of each parameter. The `b*` wildcard can be used to test hypotheses on all estimates. If a named vector is used, the names are used as labels in the output. Examples:
    - `hp = drat`
    - `hp + drat = 12`
    - `b1 + b2 + b3 = 0`
    - `b* / b1 = 1`
  - String:
    - "pairwise": pairwise differences between estimates in each row.
    - "reference": differences between the estimates in each row and the estimate in the first row.
    - "sequential": difference between an estimate and the estimate in the next row.
    - "meandev": difference between an estimate and the mean of all estimates.
    - "meanotherdev": difference between an estimate and the mean of all other estimates, excluding the current one.
    - "revpairwise", "revreference", "revsequential": inverse of the corresponding hypotheses, as described above.
  - Function:
    - Accepts an argument `x`: object produced by a `marginalEffects` function or a data frame with column `rowid` and `estimate`
    - Returns a data frame with columns `term` and `estimate` (mandatory) and `rowid` (optional).
    - The function can also accept optional input arguments: `newdata`, `by`, `draws`.
    - This function approach will not work for Bayesian models or with bootstrapping. In those cases, it is easy to use `posterior_draws()` to extract and manipulate the draws directly.
  - See the Examples section below and the vignette: <https://marginaleffects.com/vignettes/hypothesis.ht>
- `equivalence` Numeric vector of length 2: bounds used for the two-one-sided test (TOST) of equivalence, and for the non-inferiority and non-superiority tests. See Details section below.
- `p_adjust` Adjust p-values for multiple comparisons: "holm", "hochberg", "hommel", "bonferroni", "BH", "BY", or "fdr". See [stats::p.adjust](#)
- `df` Degrees of freedom used to compute p values and confidence intervals. A single numeric value between 1 and `Inf`. When `df` is `Inf`, the normal distribution is used. When `df` is finite, the t distribution is used. See [insight::get\\_df](#) for a convenient function to extract degrees of freedom. Ex: `slopes(model, df = insight::get_df(model))`



eps	NULL or numeric value which determines the step size to use when calculating numerical derivatives: $(f(x+eps)-f(x))/eps$ . When eps is NULL, the step size is 0.0001 multiplied by the difference between the maximum and minimum values of the variable with respect to which we are taking the derivative. Changing eps may be necessary to avoid numerical problems in certain models.
numberiv	string or list of strings indicating the method to use to for the numeric differentiation used in to compute delta method standard errors. <ul style="list-style-type: none"> <li>• "fdforward": finite difference method with forward differences</li> <li>• "fdcenter": finite difference method with central differences (default)</li> <li>• "richardson": Richardson extrapolation method</li> <li>• Extra arguments can be specified by passing a list to the numDeriv argument, with the name of the method first and named arguments following, ex: <code>numberiv=list("fdcenter", eps = 1e-5)</code>. When an unknown argument is used, <code>marginaleffects</code> prints the list of valid arguments for each method.</li> </ul>
...	Additional arguments are passed to the <code>predict()</code> method supplied by the modeling package. These arguments are particularly useful for mixed-effects or bayesian models (see the online vignettes on the <code>marginaleffects</code> website). Available arguments can vary from model to model, depending on the range of supported arguments by each modeling package. See the "Model-Specific Arguments" section of the <code>?slopes</code> documentation for a non-exhaustive list of available arguments.

## Details

A "slope" or "marginal effect" is the partial derivative of the regression equation with respect to a variable in the model. This function uses automatic differentiation to compute slopes for a vast array of models, including non-linear models with transformations (e.g., polynomials). Uncertainty estimates are computed using the delta method.

Numerical derivatives for the `slopes` function are calculated using a simple epsilon difference approach:  $\partial Y / \partial X = (f(X + \varepsilon/2) - f(X - \varepsilon/2)) / \varepsilon$ , where `f` is the `predict()` method associated with the model class, and  $\varepsilon$  is determined by the `eps` argument.

## Value

A data frame with one row per observation (per term/group) and several columns:

- `rowid`: row number of the `newdata` data frame
- `type`: prediction type, as defined by the `type` argument
- `group`: (optional) value of the grouped outcome (e.g., categorical outcome models)
- `term`: the variable whose marginal effect is computed
- `dydx`: slope of the outcome with respect to the term, for a given combination of predictor values
- `std.error`: standard errors computed by via the delta method.

- `p.value`: p value associated to the estimate column. The null is determined by the hypothesis argument (0 by default), and p values are computed before applying the `transform` argument. For models of class `feglm`, `Gam`, `glm` and `negbin`, p values are computed on the link scale by default unless the `type` argument is specified explicitly.
- `s.value`: Shannon information transforms of p values. How many consecutive "heads" tosses would provide the same amount of evidence (or "surprise") against the null hypothesis that the coin is fair? The purpose of `S` is to calibrate the analyst's intuition about the strength of evidence encoded in `p` against a well-known physical phenomenon. See Greenland (2019) and Cole et al. (2020).
- `conf.low`: lower bound of the confidence interval (or equal-tailed interval for bayesian models)
- `conf.high`: upper bound of the confidence interval (or equal-tailed interval for bayesian models)

See `?print.marginaleffects` for printing options.

## Functions

- `avg_slopes()`: Average slopes

## Standard errors using the delta method

Standard errors for all quantities estimated by `marginaleffects` can be obtained via the delta method. This requires differentiating a function with respect to the coefficients in the model using a finite difference approach. In some models, the delta method standard errors can be sensitive to various aspects of the numeric differentiation strategy, including the step size. By default, the step size is set to  $1e-8$ , or to  $1e-4$  times the smallest absolute model coefficient, whichever is largest.

`marginaleffects` can delegate numeric differentiation to the `numDeriv` package, which allows more flexibility. To do this, users can pass arguments to the `numDeriv::jacobian` function through a global option. For example:

- `options(marginaleffects_numDeriv = list(method = "simple", method.args = list(eps = 1e-6)))`
- `options(marginaleffects_numDeriv = list(method = "Richardson", method.args = list(eps = 1e-5)))`
- `options(marginaleffects_numDeriv = NULL)`

See the "Standard Errors and Confidence Intervals" vignette on the `marginaleffects` website for more details on the computation of standard errors:

<https://marginaleffects.com/vignettes/uncertainty.html>

Note that the `inferences()` function can be used to compute uncertainty estimates using a bootstrap or simulation-based inference. See the vignette:

<https://marginaleffects.com/vignettes/bootstrap.html>

### Model-Specific Arguments

Some model types allow model-specific arguments to modify the nature of marginal effects, predictions, marginal means, and contrasts. Please report other package-specific `predict()` arguments on Github so we can add them to the table below.

<https://github.com/vincentarelbundock/marginaleffects/issues>

Package	Class	Argument	Documentation
brms	brmsfit	ndraws	<a href="#">brms::posterior_predict</a>
		re_formula	<a href="#">brms::posterior_predict</a>
lme4	merMod	re.form	<a href="#">lme4::predict.merMod</a>
		allow.new.levels	<a href="#">lme4::predict.merMod</a>
glmmTMB	glmmTMB	re.form	<a href="#">glmmTMB::predict.glmmTMB</a>
		allow.new.levels	<a href="#">glmmTMB::predict.glmmTMB</a>
		zitype	<a href="#">glmmTMB::predict.glmmTMB</a>
mgcv	bam	exclude	<a href="#">mgcv::predict.bam</a>
robustlmm	rlmerMod	re.form	<a href="#">robustlmm::predict.rlmerMod</a>
		allow.new.levels	<a href="#">robustlmm::predict.rlmerMod</a>
MCMCglmm	MCMCglmm	ndraws	

### Bayesian posterior summaries

By default, credible intervals in bayesian models are built as equal-tailed intervals. This can be changed to a highest density interval by setting a global option:

```
options("marginaleffects_posterior_interval" = "eti")
options("marginaleffects_posterior_interval" = "hdi")
```

By default, the center of the posterior distribution in bayesian models is identified by the median. Users can use a different summary function by setting a global option:

```
options("marginaleffects_posterior_center" = "mean")
options("marginaleffects_posterior_center" = "median")
```

When estimates are averaged using the `by` argument, the `tidy()` function, or the `summary()` function, the posterior distribution is marginalized twice over. First, we take the average *across* units but *within* each iteration of the MCMC chain, according to what the user requested in `by` argument or `tidy()/summary()` functions. Then, we identify the center of the resulting posterior using the function supplied to the `"marginaleffects_posterior_center"` option (the median by default).

### Equivalence, Inferiority, Superiority

$\theta$  is an estimate,  $\sigma_\theta$  its estimated standard error, and  $[a, b]$  are the bounds of the interval supplied to the equivalence argument.

Non-inferiority:

- $H_0: \theta \leq a$
- $H_1: \theta > a$
- $t = (\theta - a)/\sigma_\theta$

- p: Upper-tail probability

Non-superiority:

- $H_0: \theta \geq b$
- $H_1: \theta < b$
- $t = (\theta - b)/\sigma_\theta$
- p: Lower-tail probability

Equivalence: Two One-Sided Tests (TOST)

- p: Maximum of the non-inferiority and non-superiority p values.

Thanks to Russell V. Lenth for the excellent `emmeans` package and documentation which inspired this feature.

### Prediction types

The `type` argument determines the scale of the predictions used to compute quantities of interest with functions from the `marginaleffects` package. Admissible values for `type` depend on the model object. When users specify an incorrect value for `type`, `marginaleffects` will raise an informative error with a list of valid `type` values for the specific model object. The first entry in the list in that error message is the default `type`.

The `invlink(link)` is a special `type` defined by `marginaleffects`. It is available for some (but not all) models and functions. With this `link` type, we first compute predictions on the link scale, then we use the inverse link function to backtransform the predictions to the response scale. This is useful for models with non-linear link functions as it can ensure that confidence intervals stay within desirable bounds, ex: 0 to 1 for a logit model. Note that an average of estimates with `type="invlink(link)"` will not always be equivalent to the average of estimates with `type="response"`.

Some of the most common `type` values are:

response, link, E, Ep, average, class, conditional, count, cum.prob, cumhaz, cumprob, density, detection, disp, ev, expected, expvalue, fitted, hazard, invlink(link), latent, latent\_N, linear, linear.predictor, linpred, location, lp, mean, numeric, p, ppd, pr, precision, prediction, prob, probability, probs, quantile, risk, rmst, scale, survival, unconditional, utility, variance, xb, zero, zlink, zprob

### Parallel computation

The `slopes()` and `comparisons()` functions can use parallelism to speed up computation. Operations are parallelized for the computation of standard errors, at the model coefficient level. There is always considerable overhead when using parallel computation, mainly involved in passing the whole dataset to the different processes. Thus, parallel computation is most likely to be useful when the model includes many parameters and the dataset is relatively small.

Warning: In many cases, parallel processing will not be useful at all.

To activate parallel computation, users must load the `future.apply` package, call `plan()` function, and set a global option. For example:

```
library(future.apply)
plan("multicore", workers = 4)
options(marginaleffects_parallel = TRUE)

slopes(model)
```

To disable parallelism in `marginaleffects` altogether, you can set a global option:

```
options(marginaleffects_parallel = FALSE)
```

### Order of operations

Behind the scenes, the arguments of `marginaleffects` functions are evaluated in this order:

1. `newdata`
2. `variables`
3. `comparison` and `slopes`
4. `by`
5. `vcov`
6. `hypothesis`
7. `transform`

### References

- Greenland S. 2019. "Valid P-Values Behave Exactly as They Should: Some Misleading Criticisms of P-Values and Their Resolution With S-Values." *The American Statistician*. 73(S1): 106–114.
- Cole, Stephen R, Jessie K Edwards, and Sander Greenland. 2020. "Surprise!" *American Journal of Epidemiology* 190 (2): 191–93. <https://doi.org/10.1093/aje/kwaa136>

### Examples

```
# Unit-level (conditional) Marginal Effects
mod <- glm(am ~ hp * wt, data = mtcars, family = binomial)
mfx <- slopes(mod)
head(mfx)

# Average Marginal Effect (AME)
avg_slopes(mod, by = TRUE)

# Marginal Effect at the Mean (MEM)
slopes(mod, newdata = datagrid())

# Marginal Effect at User-Specified Values
# Variables not explicitly included in `datagrid()` are held at their means
slopes(mod, newdata = datagrid(hp = c(100, 110)))
```

```

# Group-Average Marginal Effects (G-AME)
# Calculate marginal effects for each observation, and then take the average
# marginal effect within each subset of observations with different observed
# values for the `cyl` variable:
mod2 <- lm(mpg ~ hp * cyl, data = mtcars)
avg_slopes(mod2, variables = "hp", by = "cyl")

# Marginal Effects at User-Specified Values (counterfactual)
# Variables not explicitly included in `datagrid()` are held at their
# original values, and the whole dataset is duplicated once for each
# combination of the values in `datagrid()`
mfx <- slopes(mod,
              newdata = datagrid(hp = c(100, 110),
                                grid_type = "counterfactual"))
head(mfx)

# Heteroskedasticity robust standard errors
mfx <- slopes(mod, vcov = sandwich::vcovHC(mod))
head(mfx)

# hypothesis test: is the `hp` marginal effect at the mean equal to the `drat` marginal effect
mod <- lm(mpg ~ wt + drat, data = mtcars)

slopes(
  mod,
  newdata = "mean",
  hypothesis = "wt = drat")

# same hypothesis test using row indices
slopes(
  mod,
  newdata = "mean",
  hypothesis = "b1 - b2 = 0")

# same hypothesis test using numeric vector of weights
slopes(
  mod,
  newdata = "mean",
  hypothesis = c(1, -1))

# two custom contrasts using a matrix of weights
lc <- matrix(c(
  1, -1,
  2, 3),
  ncol = 2)
colnames(lc) <- c("Contrast A", "Contrast B")
slopes(
  mod,
  newdata = "mean",
  hypothesis = lc)

```

---

specify\_hypothesis      *(EXPERIMENTAL) Complex aggregation and test functions for the hypothesis argument*

---

### Description

Warning: This function is experimental. It may be renamed, the user interface may change, or the functionality may migrate to arguments in other `marginaleffects` functions.

This function creates aggregation and test functions for use with the `hypothesis` argument in `marginaleffects` functions like `predictions()`, `slopes()`, and `comparisons()`. The benefit of this function is that it handles a lot of the "boilerplate" code such as label creation and transformations by subgroups.

### Usage

```
specify_hypothesis(
  hypothesis = "reference",
  label = NULL,
  label_columns = NULL,
  by = c("term", "group", "contrast")
)
```

### Arguments

<code>hypothesis</code>	String or Function. Compute a test statistic. <ul style="list-style-type: none"> <li>String: "reference" or "sequential"</li> <li>Function: Accepts a single argument named <code>estimate</code> and returns a numeric vector.</li> </ul>
<code>label</code>	Function. Accepts a vector of row labels and combines them to create hypothesis labels.
<code>label_columns</code>	Character vector. Column names to use for hypothesis labels. Default is <code>c("group", "term", "rowid", attr(x, "variables_datagrid"), attr(x, "by"))</code> .
<code>by</code>	Character vector. Variable names which indicate subgroups in which the hypothesis function should be applied.

### Value

`specify_hypothesis()` is a "function factory", which means that executing it will return a function suitable for use in the `hypothesis` argument of a `marginaleffects` function.

# Index

avg\_comparisons (comparisons), 3  
avg\_predictions (predictions), 40  
avg\_slopes (slopes), 52

brms::posterior\_predict, 11, 31, 34, 39,  
46, 59

car::deltaMethod, 19  
car::linearHypothesis, 19  
comparisons, 3

datagrid, 5, 16  
datagrid(), 5, 41, 54  
dplyr::filter(), 5, 42, 54

glmmTMB::predict.glmmTMB, 11, 31, 35, 39,  
47, 59

hypotheses, 19, 19

inferences, 25  
insight::get\_data(), 5, 41, 53  
insight::get\_df, 9, 45, 56

lme4::predict.merMod, 11, 31, 34, 39, 47, 59

mgcv::predict.bam, 11, 31, 35, 39, 47, 59

plot\_comparisons, 28  
plot\_predictions, 32  
plot\_slopes, 35  
posterior\_draws, 39  
predictions, 40  
print.marginaleffects, 51

robustlmm::predict.rlmerMod, 11, 31, 35,  
39, 47, 59

slopes, 52  
specify\_hypothesis, 63  
stats::p.adjust, 9, 45, 56  
subset(), 5, 42, 54

weights(), 26