

Package ‘birdie’

September 9, 2026

Title Bayesian Instrumental Regression for Disparity Estimation

Version 0.8.0

Description Bayesian models for accurately estimating conditional distributions by race, using Bayesian Improved Surname Geocoding (BISG) probability estimates of individual race. Implements the methods described in McCartan, Fisher, Goldin, Ho and Imai (2025) <[doi:10.1080/01621459.2025.2526695](https://doi.org/10.1080/01621459.2025.2526695)>.

Depends R (>= 3.5.0)

Imports rlang (>= 0.1.2), Rcpp (>= 0.12.0), cli, vctrs, generics, dplyr, methods, stringi, stringr, RcppParallel (>= 5.0.1), SQUAREM

Suggests daarem, easycensus, wru, knitr, roxygen2, rmarkdown, rstan (>= 2.26.0), testthat (>= 3.0.0)

LinkingTo Rcpp (>= 0.12.0), cli, BH (>= 1.66.0), RcppEigen (>= 0.3.3.3.0), RcppParallel (>= 5.0.1), RcppThread, StanHeaders (>= 2.39.1)

License GPL (>= 3)

Encoding UTF-8

VignetteBuilder knitr

Config/testthat/edition 3

SystemRequirements GNU make

Biarch true

LazyData true

URL <https://corymccartan.com/birdie/>,
<https://github.com/CoryMcCartan/birdie>

BugReports <https://github.com/CoryMcCartan/birdie/issues>

RoxyenNote 7.3.2

NeedsCompilation yes

Author Cory McCartan [aut, cre],
Kosuke Imai [ctb],
Daniel Ho [ctb],
Jacob Goldin [ctb],
Robin Fisher [ctb],
The Stan Development Team [cph] (include/rstan)

Maintainer Cory McCartan <mccartan@psu.edu>

Repository CRAN

Date/Publication 2026-09-09 08:30:18 UTC

Contents

birdie	2
birdie-class	6
birdie-family	9
birdie.ctrl	9
bisg	11
census_race_geo_table	14
disparities	16
est_weighted	16
preproc	18
pseudo_vf	19
p_r_natl	20
Index	21

birdie	<i>Fit BIRDIE Models</i>
--------	--------------------------

Description

Fits one of three possible Bayesian Instrumental Regression for Disparity Estimation (BIRDIE) models to BISG probabilities and covariates. The simplest Categorical-Dirichlet model ([cat_dir\(\)](#)) is appropriate when there are no covariates or when all covariates are discrete and fully interacted with another. The more general Categorical mixed-effects model ([cat_mixed\(\)](#)) is a supports any number of fixed effects and up to one random intercept. For continuous outcomes a Normal linear model is available ([gaussian\(\)](#)).

Usage

```
birdie(  
  r_probs,  
  formula,  
  data,  
  family = cat_dir(),  
  prior = NULL,
```

```

weights = NULL,
algorithm = c("em", "gibbs", "em_boot"),
iter = 400,
warmup = 50,
prefix = "pr_",
ctrl = birdie.ctrl()
)

```

Arguments

<code>r_probs</code>	A data frame or matrix of BISG probabilities, with one row per individual. The output of <code>bisg()</code> can be used directly here.
<code>formula</code>	A two-sided formula object describing the model structure. The left-hand side is the outcome variable, which must be discrete. A single random intercept term, denoted with a vertical bar (" $1 \mid \text{<term>}$ "), is supported on the right-hand side.
<code>data</code>	An optional data frame containing the variables named in <code>formula</code> .
<code>family</code>	A description of the complete-data model type to fit. Options are: • <code>cat_dir()</code>: Categorical-Dirichlet model. All covariates must be fully interacted. • <code>cat_mixed()</code>: Categorical mixed-effects model. Up to one random effect is supported. • <code>gaussian()</code>: Linear model. See the Details section below for more information on the various models.
<code>prior</code>	A list with entries specifying the model prior. • For the <code>cat_dir</code> model, the only entry is <code>alpha</code>, which should be a matrix of Dirichlet hyperparameters. The matrix should have one row for every level of the outcome variable and one column for every racial group. The default prior (used when <code>prior=NULL</code>) is an empirical Bayes prior equal to the weighted-mean estimate of the outcome-race table. A fully noninformative prior with all entries set to ϵ can be obtained by setting <code>prior=NA</code>. When <code>prior=NULL</code> and <code>algorithm="em"</code> or <code>"em_boot"</code>, 1 is added to the prior so that the posterior mode, rather than the mean, is shrunk toward these values. • For the <code>cat_mixed</code> model, the prior list should contain three scalar entries: <code>scale_int</code>, the standard deviation on the Normal prior for the intercepts (which control the global estimates of $Y \mid R$), <code>scale_beta</code>, the standard deviation on the Normal prior for the fixed effects, and <code>scale_sigma</code>, the prior mean of the standard deviation of the random intercepts. These can be a single scalar or a vector with an entry for each racial group. • For the <code>gaussian</code> model, the prior list should contain two entries: <code>scale_int</code>, controlling the standard deviation on the Normal prior for the intercepts (which control the global estimates of $Y \mid R$), and <code>scale_beta</code>, controlling the standard deviation on the Normal prior for the fixed effects. These must be a single scalar. Each is expressed in terms of the estimated residual standard deviation (i.e., they are multiplied together to form the "true" prior).

	The prior is stored after model fitting in the <code>\$prior</code> element of the fitted model object.
<code>weights</code>	An optional numeric vector specifying likelihood weights.
<code>algorithm</code>	The inference algorithm to use. One of 3 options: • <code>"em"</code>: An expectation-maximization algorithm which will perform inference for the maximum a posteriori (MAP) parameter values. Computationally efficient and supported by all the model families. No uncertainty quantification. • <code>"gibbs"</code>: A Gibbs sampler for performing full Bayesian inference. Generally more computationally demanding than the EM algorithm, but provides uncertainty quantification. Currently supported for <code>cat_dir()</code> and <code>gaussian()</code> model families. Computation-reliability tradeoff can be controlled with <code>iter</code> argument. • <code>"em_boot"</code>: Bootstrapped version of EM algorithm. Number of bootstrap replicates controlled by <code>iter</code> parameter. Provides approximate uncertainty quantification. Currently supported for <code>cat_dir()</code> and <code>gaussian()</code> model families.
<code>iter</code>	The number of post-warmup Gibbs samples, or the number of bootstrap replicates to use to compute approximate standard errors for the main model estimates. Only available when <code>family=cat_dir()</code> or <code>gaussian()</code> . Ignored if <code>algorithm="em"</code> . For bootstrapping, when there are fewer than 1,000 individuals or 100 or fewer replicates, a Bayesian bootstrap is used instead (i.e., weights are drawn from a $\text{Dirichlet}(1, 1, \dots, 1)$ distribution, which produces more reliable estimates.
<code>warmup</code>	Number of warmup iterations for Gibbs sampling. Ignored unless <code>algorithm="gibbs"</code> .
<code>prefix</code>	If <code>r_probs</code> is a data frame, the columns containing racial probabilities will be selected as those with names starting with <code>prefix</code> . The default will work with the output of <code>bisg()</code> .
<code>ctrl</code>	A list containing control parameters for the EM algorithm and optimization routines. A list in the proper format can be made using <code>birdie.ctrl()</code> .

Details

By default, `birdie()` uses an expectation-maximization (EM) routine to find the maximum *a posteriori* (MAP) estimate for the specified model. Asymptotic variance-covariance matrices for the MAP estimate are available for the Categorical-Dirichlet and Normal linear models via bootstrapping. Full Bayesian inference is supported via Gibbs sampling for the Categorical-Dirichlet and Normal linear models as well.

Whatever model or method is used, a finite-population estimate of the outcome-given-race distribution for the entire observed sample is always calculated and stored as `$est` in the returned object, which can be accessed with `coef.birdie()` as well.

The Categorical-Dirichlet model is specified as follows:

$$Y_i \mid R_i, X_i, \Theta \sim \text{Categorical}(\theta_{R_i X_i})$$

$$\theta_{r,x} \sim \text{Dirichlet}(\alpha_r),$$

where Y is the outcome variable, R is race, X are covariates (fixed effects), and θ_{rx} and α_r are vectors with length matching the number of levels of the outcome variable. There is one vector θ_{rx} for every combination of race and covariates, hence the need for formula to either have no covariates or a fully interacted structure.

The Categorical mixed-effects model is specified as follows:

$$Y_i \mid R_i, X_i, \Theta \sim \text{Categorical}(g^{-1}(\mu_{R_i X_i}))$$

$$\mu_{rxy} = W\beta_{ry} + Zu_{ry}$$

$$u_r \mid \vec{\sigma}_r, L_r \sim \mathcal{N}(0, \text{diag}(\vec{\sigma}_r)C_r\text{diag}(\vec{\sigma}_r))$$

$$\beta_{ry} \sim \mathcal{N}(0, s_{r\beta}^2)$$

$$\sigma_{ry} \sim \text{Inv-Gamma}(4, 3s_{r\sigma})C_r \sim \text{LKJ}(2),$$

where β_{ry} are the fixed effects, u_{ry} is the random intercept, and g is a softmax link function. Estimates for β_{ry} and σ_{ry} are stored in the `$beta` and `$sigma` elements of the fitted model object.

The Normal linear model is specified as follows:

$$Y_i \mid R_i, \vec{X}_i, \Theta \sim \mathcal{N}(\vec{X}_i^\top \vec{\theta}, \sigma^2)$$

$$\sigma^2 \sim \text{Inv-Gamma}(n_\sigma/2, l_\sigma^2 n_\sigma/2)$$

$$\beta_{\text{intercept}} \sim \mathcal{N}(0, s_{\text{int}}^2)$$

$$\beta_k \sim \mathcal{N}(0, s_\beta^2),$$

where $\vec{\theta}$ is a vector of linear model coefficients. Estimates for θ and σ are stored in the `$beta` and `$sigma` elements of the fitted model object.

More details on the models and their properties may be found in the paper referenced below.

Value

An object of class `birdie`, for which many methods are available. The model estimates may be accessed with `coef.birdie()`, and updated BISG probabilities (conditioning on the outcome) may be accessed with `fitted.birdie()`. Uncertainty estimates, if available, can be accessed with `$se` and `vcov.birdie()`.

References

McCartan, C., Fisher, R., Goldin, J., Ho, D.E., & Imai, K. (2025). Estimating Racial Disparities when Race is Not Observed. *Journal of the American Statistical Association*. Available at [doi:10.1080/01621459.2025.2526695](https://doi.org/10.1080/01621459.2025.2526695).

Examples

```
data(pseudo_vf)

r_probs = bisg(~ nm(last_name) + zip(zip), data=pseudo_vf)

# Process zip codes to remove missing values
pseudo_vf$zip = proc_zip(pseudo_vf$zip)

fit = birdie(r_probs, turnout ~ 1, data=pseudo_vf)
print(fit)
fit$se # uncertainty quantification

fit = birdie(r_probs, turnout ~ zip, data=pseudo_vf, algorithm="gibbs")

fit = birdie(r_probs, turnout ~ (1 | zip), data=pseudo_vf,
             family=cat_mixed(), ctrl=birdie.ctrl(abstol=1e-3))

summary(fit)
coef(fit)
fitted(fit)
```

birdie-class	<i>Class "birdie" of BIRDiE Models</i>
--------------	--

Description

The output of `birdie()` is an object of class `birdie`, which supports many generic functions. Notably `coef.birdie()` returns the main model estimates of outcome given race, and `fitted.birdie()` returns a table analogous to the output of `bisg()` with updated race probabilities.

Usage

```
## S3 method for class 'birdie'
coef(object, subgroup = FALSE, ...)

## S3 method for class 'birdie'
fitted(object, ...)

## S3 method for class 'birdie'
residuals(object, x_only = FALSE, ...)

## S3 method for class 'birdie'
predict(object, adj = NULL, ...)

## S3 method for class 'birdie'
simulate(object, nsim = 1, seed = NULL, ...)
```

```

## S3 method for class 'birdie'
plot(x, log = FALSE, ...)

## S3 method for class 'birdie'
tidy(x, subgroup = FALSE, ...)

## S3 method for class 'birdie'
glance(x, ...)

## S3 method for class 'birdie'
augment(x, data, ...)

## S3 method for class 'birdie'
formula(x, ...)

## S3 method for class 'birdie'
family(object, ...)

## S3 method for class 'birdie'
nobs(object, ...)

## S3 method for class 'birdie'
vcov(object, ...)

## S3 method for class 'birdie'
print(x, ...)

## S3 method for class 'birdie'
summary(object, ...)

```

Arguments

<code>object, x</code>	A birdie model object
<code>subgroup</code>	If TRUE, return subgroup-level (rather than marginal) coefficient estimates as a 3D array.
<code>...</code>	Potentially further arguments passed from other methods
<code>x_only</code>	if TRUE, calculate fitted values using covariates only (i.e., without using sur-names).
<code>adj</code>	A point in the simplex that describes how BISG probabilities will be thresholded to produce point predictions. The probabilities are divided by <code>adj</code> , then the racial category with the highest probability is predicted. Can be used to trade off types of prediction error. Must be nonnegative but will be normalized to sum to 1. The default is to make no adjustment.
<code>nsim</code>	The number of vectors to simulate. Defaults to 1.
<code>seed</code>	Used to seed the random number generator. See <code>stats::simulate()</code> .
<code>log</code>	If TRUE, plot estimated probabilities on a log scale.

data A data frame to augment with $\Pr(R \mid Y, X, S)$ probabilities

Details

The internal structure of `birdie` objects is not designed to be accessed directly. The generics listed here should be used instead.

Value

Varies, depending on the method. See generic functions' documentation for details.

Functions

- `coef(birdie)`: Return estimated outcome-given-race distributions. When `subgroup=FALSE` this always returns a finite-population estimate of the outcome-given-race distribution for the observed sample.
- `fitted(birdie)`: Return an updated race probability table. `bisg()` estimates $\Pr(R \mid G, X, S)$; this table is $\Pr(R \mid Y, G, X, S, \text{Theta-hat})$.
- `residuals(birdie)`: Return the residuals for the outcome variable as a matrix. Useful in sensitivity analyses and to get an idea of how well race, location, names, etc. predict the outcome.
- `predict(birdie)`: Create point predictions of individual race. Returns factor vector of individual race labels. Strongly not recommended for any kind of inferential purpose, as biases may be extreme and in unpredictable directions.
- `simulate(birdie)`: Simulate race from the posterior distribution $\Pr(R \mid Y, G, X, S, \text{Theta-hat})$. Does not account for uncertainty in model parameters.
- `plot(birdie)`: Visualize the estimated conditional distributions for a BIRDiE model. If available, marginal standard error estimates (`$se`) will be visualized with 95% confidence-level error bars.
- `tidy(birdie)`: Put BIRDiE model coefficients in a tidy format.
- `glance(birdie)`: Glance at a BIRDiE model.
- `augment(birdie)`: Augment data with individual race predictions from a BIRDiE model.
- `formula(birdie)`: Extract the formula used to specify a BIRDiE model.
- `family(birdie)`: Return the BIRDiE complete-data model family.
- `nobs(birdie)`: Return the number of observations used to fit a BIRDiE model.
- `vcov(birdie)`: Return the estimated variance-covariance matrix for the BIRDiE model estimates, if available.
- `print(birdie)`: Print a summary of the model fit.
- `summary(birdie)`: Print a more detailed summary of the model fit.

Examples

```
methods(class="birdie")
```

birdie-family

*BIRDIE Complete-Data Model Families***Description**

BIRDIE supports a number of complete-data outcome models, including categorical regression models. Models specific to BIRDIE are listed here. See the Details section of `birdie()` for more information about each model.

Usage

```
cat_dir(link = "identity")
```

```
cat_mixed(link = "softmax")
```

Arguments

`link` The link function. Only one option available for categorical regression models.

Value

A list of class family containing the specification.

Examples

```
cat_dir()
cat_mixed()
```

birdie.ctrl

*Control of BIRDIE Model Fitting***Description**

Constructs control parameters for BIRDIE model fitting. All arguments have defaults.

Usage

```
birdie.ctrl(
  abstol = 1e-06,
  reltol = 1e-06,
  max_iter = 1000,
  fix_sigma = FALSE,
  accel = c("squarem", "anderson", "daarem", "none"),
  order = switch(match.arg(accel), none = 0L, anderson = -1L, daarem = -1L, squarem = 1L),
  anderson_restart = TRUE
)
```

Arguments

<code>abstol</code>	The absolute tolerance used in checking convergence or in estimating linear model coefficients.
<code>reltol</code>	The relative tolerance used in checking convergence. Ignored if <code>accel = "squarem"</code> or <code>"daarem"</code> .
<code>max_iter</code>	The maximum number of EM iterations.
<code>fix_sigma</code>	If TRUE when <code>model=gaussian()</code> , fix sigma to an initial estimate, in order to avoid estimation collapse when the outcomes are discrete.
<code>accel</code>	The acceleration algorithm to use in doing EM. The default <code>"squarem"</code> is good for most purposes, though <code>"anderson"</code> may be faster when there are few parameters or very tight tolerances. <code>"daarem"</code> is an excellent choice as well that works across a range of problems, though it requires installing the small <code>daarem</code> package. <code>"none"</code> is not recommended unless other algorithms are running into numerical issues. See the references below for details on these schemes.
<code>order</code>	The order to use in the acceleration algorithm. Interpretation varies by algorithm. Can range from 1 (default) to 3 for SQUAREM and from 1 to the number of parameters for Anderson and DAAREM (default -1 allows the order to be determined by problem size).
<code>anderson_restart</code>	Whether to use restarts in Anderson acceleration.

Value

A list containing the control parameters.

References

- Varadhan, R., & Roland, C. (2004). Squared extrapolation methods (SQUAREM): A new class of simple and efficient numerical schemes for accelerating the convergence of the EM algorithm.
- Walker, H. F., & Ni, P. (2011). Anderson acceleration for fixed-point iterations. *SIAM Journal on Numerical Analysis*, 49(4), 1715-1735.
- Henderson, N. C., & Varadhan, R. (2019). Damped Anderson acceleration with restarts and monotonicity control for accelerating EM and EM-like algorithms. *Journal of Computational and Graphical Statistics*, 28(4), 834-846.

Examples

```
str(birdie.ctrl(max_iter=100))
```

bisg

*Bayesian Improved Surname Geocoding (BISG)***Description**

Calculates individual probabilities of belonging to racial groups given last name, location, and other covariates (optional). The standard function `bisg()` treats the input tables as fixed. An alternative function `bisg_me()`, assumes that the input tables are subject to measurement error, and uses a Gibbs sampler to impute the individual race probabilities, using the model of Imai et al. (2022).

Usage

```

bisg(
  formula,
  data = NULL,
  p_r = p_r_natl(),
  p_rgx = NULL,
  p_rs = NULL,
  save_rgx = TRUE,
  save_sr = FALSE
)

bisg_me(
  formula,
  data = NULL,
  p_r = p_r_natl(),
  p_rgx = NULL,
  p_rs = NULL,
  iter = 1000,
  warmup = 100,
  cores = 1L
)

## S3 method for class 'bisg'
summary(object, p_r = NULL, ...)

## S3 method for class 'bisg'
predict(object, adj = NULL, ...)

## S3 method for class 'bisg'
simulate(object, nsim = 1, seed = NULL, ...)

```

Arguments

<code>formula</code>	A formula specifying the BISG model. Must include the special term <code>nm()</code> to identify the surname variable. Certain geographic variables can be identified similarly: <code>zip()</code> for ZIP codes, and <code>state()</code> for states. If no other predictor
----------------------	---

variables are provided, then `bisg()` will automatically be able to build a table of census data to use in inference. If other predictor variables are included, or if other geographic identifiers are used, then the user must specify the `p_rgx` argument below. In this case, do *not* label the geographic variables with `zip()` or `state()`, just include them as normal variables in the formula. The left-hand side of the formula is ignored. See the examples section below for sample formulas.

<code>data</code>	The data frame containing the variables in formula.
<code>p_r</code>	The prior distribution of race in the sample, as a numeric vector. Defaults to U.S. demographics as provided by <code>p_r_natl()</code> . Can also set <code>p_r="est"</code> or <code>"estimate"</code> to estimate this from the observed surname and geographic distribution using an EM algorithm. Since the prior distribution on race strongly affects the calibration of the BISG probabilities and thus the accuracy of downstream estimates, users are encouraged to think carefully about an appropriate value for <code>p_r</code> . If no prior information on the racial makeup of the sample is available, and yet the sample is very different from the overall U.S. population, then <code>p_r="estimate"</code> will likely produce superior results.
<code>p_rgx</code>	The distribution of race given location (G) and other covariates (X) specified in formula. Should be provided as a data frame, with columns matching the predictors in formula, and additional columns for each racial group containing the conditional probability for that racial group given the predictors. For example, if Census tracts are the only predictors, <code>p_rgx</code> should be a data frame with a tract column and columns <code>white</code>, <code>black</code>, etc. containing the racial distribution of each tract. If formula contains only labeled terms (like <code>zip()</code>), then by default <code>p_rgx</code> will be constructed automatically from the most recent Census data. Thus, if a custom <code>p_rgx</code> is desired, do <i>not</i> use geographic labellers like <code>zip()</code> in the formula. The <code>p_rgx</code> table will be normalized by row, so it can be provided as population counts as well. Column totals are used to calculate the reference race distribution, so counts are preferred and are required for <code>bisg_me()</code>. The <code>census_race_geo_table()</code> function can be helpful to prepare tables, as can be the <code>build_dec()</code> and <code>build_acs()</code> functions in the <code>censable</code> package. Note that no missing values are allowed in the predictors used in <code>p_rgx</code>. One can include a level <code>"<none>"</code> in the predictor columns of <code>p_rgx</code>, which will be imputed with the overall <code>p_r</code>. Thus it is important to preprocess <code>p_rgx</code> and/or the input data to ensure a 100% match rate.
<code>p_rs</code>	The distribution of race given last name. As with <code>p_rgx</code> , should be provided as a data frame, with a column of names and additional columns for each racial group. Users should not have to specify this argument in most cases, as the table will be built from published Census surname tables automatically. Counts are required for <code>bisg_me()</code> . One of the last names can be <code>"<generic>"</code> , which, if included, will be used for any names not found in the table.
<code>save_rgx</code>	If TRUE, save the <code>p_rgx</code> table (matched to each individual) as the <code>"p_rgx"</code> and <code>"gx"</code> attributes of the output. Necessary for some sensitivity analyses.
<code>save_sr</code>	If TRUE, save the <code>p_sr</code> table (surname given race; matched to each individual) as the <code>"p_sr"</code> and <code>"s"</code> attributes of the output.

<code>iter</code>	How many sampling iterations in the Gibbs sampler
<code>warmup</code>	How many burn-in iterations in the Gibbs sampler
<code>cores</code>	How many parallel cores to use in computation. Around 4 seems to be optimal, even if more are available.
<code>object</code>	An object of class <code>bisg</code> , the result of running <code>bisg()</code> .
<code>...</code>	Additional arguments to generic methods (ignored).
<code>adj</code>	A point in the simplex that describes how BISG probabilities will be thresholded to produce point predictions. The probabilities are divided by <code>adj</code> , then the racial category with the highest probability is predicted. Can be used to trade off types of prediction error. Must be nonnegative but will be normalized to sum to 1. The default is to make no adjustment.
<code>nsim</code>	The number of vectors to simulate. Defaults to 1.
<code>seed</code>	Used to seed the random number generator. See <code>stats::simulate()</code> .

Value

An object of class `bisg`, which is just a data frame with some additional attributes. The data frame has rows matching the input data and columns for the race probabilities. The target and reference race distributions are saved as the "p_r" and "q_r" attributes, respectively.

Methods (by generic)

- `summary(bisg)`: Summarize predicted race probabilities. Returns vector of individual entropies.
- `predict(bisg)`: Create point predictions of individual race. Returns factor vector of individual race labels. Strongly not recommended for any kind of inferential purpose, as biases may be extreme and in unpredictable directions.
- `simulate(bisg)`: Simulate race from the $\Pr(R \mid G, X, S)$ distribution.

Functions

- `bisg()`: The standard BISG model.
- `bisg_me()`: The measurement error BISG model.

Surname Matching

The Census surname table can be inspected with the following code:

```
readRDS(system.file("extdata", "names_2010_counts.rds", package="birdie"))
```

Surnames are processed with `proc_name()` before being matched to the table. The code also recognizes double-barrelled (hyphenated) surnames and attempts to match each part if the overall name is not found in the surname table. Specifying `save_rs=TRUE` will save the matched surname table and a lookup vector that matches each individual to their surname table row. The overall match rate is reported as part of the `summary()` output.

References

- Elliott, M. N., Fremont, A., Morrison, P. A., Pantoja, P., and Lurie, N. (2008). A new method for estimating race/ethnicity and associated disparities where administrative records lack self-reported race/ethnicity. *Health Services Research*, 43(5p1):1722–1736.
- Fiscella, K. and Fremont, A. M. (2006). Use of geocoding and surname analysis to estimate race and ethnicity. *Health Services Research*, 41(4p1):1482–1500.
- Imai, K., Olivella, S., & Rosenman, E. T. (2022). Addressing census data problems in race imputation via fully Bayesian Improved Surname Geocoding and name supplements. *Science Advances*, 8(49), eadc9824.

Examples

```
data(pseudo_vf)
bisg(~ nm(last_name), data=pseudo_vf)

r_probs = bisg(~ nm(last_name) + zip(zip), data=pseudo_vf)
summary(r_probs)
head(predict(r_probs))

# Using a custom p_rgx
# Requires a Census API key
## Not run:
d_census = census_race_geo_table(geo = "zcta", year = 2023, survey = "acs5")
# Manually process ZIPs & convert invalid to <none>; 'GEOID' matches d_census
pseudo_vf$GEOID = proc_zip(pseudo_vf$zip)
r_probs_acs = bisg(~ nm(last_name) + GEOID, data=pseudo_vf, p_rgx=d_census)
plot(r_probs$pr_black, r_probs_acs$pr_black)

## End(Not run)

data(pseudo_vf)
bisg_me(~ nm(last_name) + zip(zip), data=pseudo_vf)
```

census_race_geo_table [Download Census Race Data](#)

Description

Downloads and prepares race-by-geography tables from U.S. census data, using the [easycensus](#) package. Requires that an api key be set up through [easycensus::cens_auth\(\)](#) in that package, usually by storing it in the CENSUS_API_KEY environment variable. Supports data from the decennial census and the American Community Survey at a variety of levels of geographic detail. The output of this function can be used directly in [bisg\(\)](#).

Usage

```
census_race_geo_table(
  geo = c("us", "state", "county", "zcta", "tract"),
  ...,
  year = 2010,
  survey = c("dec", "acs1", "acs5"),
  GEOIDs = TRUE,
  counts = TRUE
)
```

Arguments

<code>geo</code>	The geographic level to return. Common options are listed in the function signature, but any of the geographies listed at easycensus::cens_geo() may be used.
<code>...</code>	Further subgeographies to return, as in easycensus::cens_geo() .
<code>year</code>	The year for the data
<code>survey</code>	The data product to use: either the decennial census ("dec"), or the the 1-year or 5-year ACS.
<code>GEOIDs</code>	If TRUE, return the GEOID column as the unique geographic identifier; if FALSE, return a human-readable name. For example, with <code>geo="state"</code> , setting <code>GEOIDs=FALSE</code> will return a column named <code>state</code> with entries like "Massachusetts".
<code>counts</code>	If TRUE, return the table as actual population counts; if FALSE, return table as percentages within each geography.

Value

A data frame with geographic identifier column(s) and six columns `white`, `black`, etc. containing the counts or proportion of residents in each racial group.

Examples

```
census_race_geo_table("zcta", year=2010)
## Not run:
# Census API key required
census_race_geo_table("us", year=2010)
census_race_geo_table("state", year=2021, survey="acs1")
census_race_geo_table("county", state="NH", year=2020, GEOIDs=FALSE)

## End(Not run)
```

disparities	<i>Compute Racial Disparities from Model Estimates</i>
-------------	--

Description

This function lets you easily compute differences in conditional expectations between all pairs of specified racial groups.

Usage

```
disparities(x, subgroup = FALSE, races = TRUE)
```

Arguments

x	A birdie model object.
subgroup	If TRUE, return subgroup-level (rather than marginal) disparity estimates.
races	A character vector of racial groups to compute disparities for. The special value TRUE, the default, computes disparities for all racial groups.

Value

A data frame containing a row with every possible disparity for the specified races, which are identified by columns race_1 and race_2. The reported disparity is estimate_1 - estimate_2.

Examples

```
data(pseudo_vf)
r_probs = bisg(~ nm(last_name) + zip(zip), data=pseudo_vf)
fit = birdie(r_probs, turnout ~ 1, data=pseudo_vf)

disparities(fit)
disparities(fit, races=c("white", "black"))
```

est_weighted	<i>Calculate Weighted Estimate of (Discrete) Outcomes By Race</i>
--------------	---

Description

Calculates the "standard" weighted estimator of conditional distributions of an outcome variable Y by race R , using BISG probabilities. This estimator, while commonly used, is only appropriate if $Y \perp R \mid X, S$, where S and X are the last names and covariates (possibly including geography) used in making the BISG probabilities. In most cases this assumption is not plausible and [birdie\(\)](#) should be used instead. See the references below for more discussion as to selecting the right estimator.

Up to Monte Carlo error, the weighted estimate is equivalent to performing multiple imputations of the race vector from the BISG probabilities and then using them inside a weighted average or linear regression.

Usage

```
est_weighted(
  r_probs,
  formula,
  data = NULL,
  weights = NULL,
  prefix = "pr_",
  se_boot = 0
)

## S3 method for class 'est_weighted'
print(x, ...)

## S3 method for class 'est_weighted'
summary(object, ...)
```

Arguments

r_probs	A data frame or matrix of BISG probabilities, with one row per individual. The output of <code>bisg()</code> can be used directly here.
formula	A two-sided formula object describing the estimator structure. The left-hand side is the outcome variable, which must be discrete. Subgroups for which to calculate estimates may be specified by adding covariates on the right-hand side. Subgroup estimates are available with <code>coef(..., subgroup=TRUE)</code> and <code>tidy(..., subgroup=TRUE)</code> .
data	An optional data frame containing the variables named in formula.
weights	An optional numeric vector specifying weights.
prefix	If r_probs is a data frame, the columns containing racial probabilities will be selected as those with names starting with prefix. The default will work with the output of <code>bisg()</code> .
se_boot	The number of bootstrap replicates to use to compute an approximate covariance matrix for the estimator. If no bootstrapping is used, an analytical estimate of standard errors will be returned as \$se. For bootstrapping, when there are fewer than 1,000 individuals or 100 or fewer replicates, a Bayesian bootstrap is used instead (i.e., weights are drawn from a $\text{Dirichlet}(1, 1, \dots, 1)$ distribution, which produces more reliable estimates.
...	Additional arguments to generic methods (ignored).
object, x	An object of class est_weighted.

Value

An object of class `est_weighted`, inheriting from `birdie`, for which many methods are available. The model estimates may be accessed with `coef()`. Uncertainty estimates, if available, can be accessed with `$se` and `vcov.birdie()`.

Methods (by generic)

- `print(est_weighted)`: Print a summary of the model fit.
- `summary(est_weighted)`: Print a more detailed summary of the model fit.

References

McCartan, C., Fisher, R., Goldin, J., Ho, D.E., & Imai, K. (2025). Estimating Racial Disparities when Race is Not Observed. *Journal of the American Statistical Association*. Available at [doi:10.1080/01621459.2025.2526695](https://doi.org/10.1080/01621459.2025.2526695).

Examples

```
data(pseudo_vf)

r_probs = bisg(~ nm(last_name) + zip(zip), data=pseudo_vf)

# Process zip codes to remove missing values
pseudo_vf$zip = proc_zip(pseudo_vf$zip)

est_weighted(r_probs, turnout ~ 1, data=pseudo_vf)

est = est_weighted(r_probs, turnout ~ zip, data=pseudo_vf)
tidy(est, subgroup=TRUE)
```

```
preproc
```

```
Preprocess Last Names and Geographic Identifiers
```

Description

These functions are called automatically by `bisg()` but may be useful, especially when geographic variables are included in a `birdie()` model. `proc_zip()` and `proc_state()` preprocess their corresponding geographic identifiers. States are partially matched to state names and abbreviations and are returned as FIPS codes. ZIP codes are crosswalked to Census ZCTAs. Missing identifiers are replaced with "<none>". `proc_name()` processes last names in accordance with Census processing rules (<https://www2.census.gov/topics/genealogy/2010surnames/surnames.pdf>). Names are converted to Latin characters, capitalized, stripped of prefixes and suffixes, and otherwise standardized.

Usage

```
proc_zip(x)

proc_state(x)

proc_name(x, to_latin = TRUE)
```

Arguments

<code>x</code>	A character vector of names or geographic identifiers to process
<code>to_latin</code>	If TRUE, convert names to Latin characters only. Strongly recommended if non-Latin characters are present, since these will not match Census tables. However, the conversion is slightly time-consuming and so can be disabled with this flag.

Value

A processed character vector

Functions

- `proc_zip()`: Match ZIP codes to ZCTAs and fill in missing values.
- `proc_state()`: Match state names and abbreviations and fill in missing values.
- `proc_name()`: Process names to a Census-standardized format.

Examples

```
proc_name("Smith Jr.")
proc_zip("00501")
proc_state("Washington")
```

pseudo_vf	<i>A pseudo-voterfile</i>
-----------	---------------------------

Description

A dataset containing 5,000 fake voter records. Created by randomizing a subset of the North Carolina voter file. Turnout records are completely randomly generated.

Usage

```
pseudo_vf
```

Format

A data frame with 5,000 rows and 4 records:

last_name Voter's last name

zip 5-digit ZIP code. May be NA

race One of "white", "black", "hisp", "asian", "aian", or "other"

turnout 1 if the voter voted in the most recent election, 0 otherwise

Source

<https://www.ncsbe.gov/results-data/voter-registration-data>

Examples

```
data(pseudo_vf)
print(pseudo_vf)
```

p_r_natl

*National Racial Demographics***Description**

Returns the proportion of the U.S. population in six racial groups in a given year. Group definitions necessarily follow those used by the Census Bureau in its surname tables:

- white: Non-Hispanic White alone
- black: Non-Hispanic Black alone
- hisp: Hispanic, any race
- asian: Non-Hispanic Asian, Native Hawaiian, or Pacific Islander alone
- aian: Non-Hispanic American Indian/Alaska Native
- other: Non-Hispanic, two or more races, or other race

Usage

```
p_r_natl(year = 2024, vap = FALSE)
```

Arguments

year	The year to return demographics for.
vap	If TRUE, return statistics for the voting-age population (18+) rather than the full U.S. population. Voting-age statistics are available from 2009 onward.

Details

Except for 2010, values come from the one-year American Community Survey (ACS). The 2010 value comes from the decennial census. Because the standard 2020 ACS one-year estimates were not released, the 2020 value is interpolated from 2019 and 2021. Voting-age values come from ACS Selected Population Profiles. To maintain a consistent ACS series, the 2010 and 2020 voting-age values are interpolated from the adjacent years.

Value

A named numeric vector of length 6.

Examples

```
p_r_natl(year=2010)
```

Index

- * **bisg**
 - bisg, 11
- * **datasets**
 - pseudo_vf, 19
- * **estimators**
 - birdie, 2
 - birdie-class, 6
 - birdie.ctrl, 9
 - disparities, 16
 - est_weighted, 16
- * **misc**
 - pseudo_vf, 19
- * **preproc**
 - census_race_geo_table, 14
 - preproc, 18
- augment.birdie (birdie-class), 6
- birdie, 2, 5, 17
- birdie(), 6, 9, 16, 18
- birdie-class, 6
- birdie-family, 9
- birdie.ctrl, 9
- birdie.ctrl(), 4
- bisg, 11
- bisg(), 3, 4, 6, 8, 13, 14, 17, 18
- bisg_me (bisg), 11
- cat_dir (birdie-family), 9
- cat_dir(), 2, 3
- cat_mixed (birdie-family), 9
- cat_mixed(), 2, 3
- census_race_geo_table, 14
- census_race_geo_table(), 12
- coef.birdie (birdie-class), 6
- coef.birdie(), 4, 5
- disparities, 16
- easycensus, 14
- easycensus::cens_auth(), 14
- easycensus::cens_geo(), 15
- est_weighted, 16
- family.birdie (birdie-class), 6
- fitted.birdie (birdie-class), 6
- fitted.birdie(), 5
- formula.birdie (birdie-class), 6
- gaussian(), 2, 3
- glance.birdie (birdie-class), 6
- nobs.birdie (birdie-class), 6
- p_r_natl, 20
- p_r_natl(), 12
- plot.birdie (birdie-class), 6
- predict.birdie (birdie-class), 6
- predict.bisg (bisg), 11
- preproc, 18
- print.birdie (birdie-class), 6
- print.est_weighted (est_weighted), 16
- proc_name (preproc), 18
- proc_name(), 13
- proc_state (preproc), 18
- proc_zip (preproc), 18
- pseudo_vf, 19
- residuals.birdie (birdie-class), 6
- simulate.birdie (birdie-class), 6
- simulate.bisg (bisg), 11
- stats::simulate(), 7, 13
- summary.birdie (birdie-class), 6
- summary.bisg (bisg), 11
- summary.est_weighted (est_weighted), 16
- tidy.birdie (birdie-class), 6
- vcov.birdie (birdie-class), 6
- vcov.birdie(), 5, 17