Package: wdsmatch (via r-universe)

October 20, 2026

Type Package

Title Weighted Double Score Matching for Survey-Weighted Causal Inference

Version 0.2.2

Description Implements weighted double score matching (WDSM) for estimating population-level causal effects from complex survey data. Combines propensity scores and prognostic scores with survey design weights for matching, survey-weighted imputation within match sets, and Hajek normalization to target the population average treatment effect (PATE) and the population average treatment effect on the treated (PATT). Supports both retrospective (treatment-dependent) and prospective (treatment-independent) sampling designs. Uses propensity probabilities and arm-specific prognostic scores for matching, with a complete quadratic bias correction in each arm's double score. Provides linearization-based multinomial replication variance estimates and centered normal Wald confidence intervals, retaining the original matching reuse coefficients without re-matching. Supplied scores can be held fixed for inference conditional on those scores. This weight-only interface does not encode survey strata, clusters, or design-specific replicate weights. The method is described by Zeng et al. (2026) <doi:10.48550/arXiv.2610.25870>.

License GPL-3

Encoding UTF-8

LazyData true

RoxygenNote 7.3.2

Depends R (>= 3.6.0)

Imports stats

Suggests testthat (>= 3.0.0), knitr, rmarkdown

Config/testthat/edition 3

URL https://cran.r-project.org/package=wdsmatch, https://arxiv.org/abs/2610.25870

Repository https://ykzeng-yale.r-universe.dev

Date/Publication 2026-10-20 04:00:54 UTC

RemoteUrl https://github.com/ykzeng-yale/wdsmatch

RemoteRef HEAD

RemoteSha e9352f344d10318c851664819dfe773d02c9e4d3

Contents

print.wdsmatch 2
summary.wdsmatch 3
survey_obs 3
wdsmatchATE 4
wdsmatchATT 8

Index **11**

`print.wdsmatch` *Print method for wdsmatch objects*

Description

Print method for wdsmatch objects

Usage

```
## S3 method for class 'wdsmatch'
print(x, digits = 4, ...)
```

Arguments

`x`	A `wdsmatch` object returned by `wdsmatchATE` or `wdsmatchATT`.
`digits`	Number of decimal places.
`...`	Additional arguments (ignored).

Value

Invisibly returns the input object `x`. Called for its side effect of printing a formatted summary to the console, including the point estimate, standard error, confidence interval, number of matches, and sample sizes.

summary.wdsmatch *Summary method for wdsmatch objects*

Description

Summary method for wdsmatch objects

Usage

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

Arguments

object	A wdsmatch object returned by wdsmatchATE or wdsmatchATT.
...	Additional arguments (ignored).

Value

Invisibly returns the input object. Called for its side effect of printing a formatted summary to the console.

survey_obs *Simulated Survey Observational Data*

Description

A simulated dataset drawn from a survey-weighted observational study with treatment-dependent (retrospective) sampling. Contains 6 covariates, a binary treatment indicator, observed outcome, and survey design weights. The true PATE is approximately 0.8 and the true PATT approximately 1.0.

Usage

```
survey_obs
```

Format

A data frame with approximately 120 rows and 9 variables:

Y Observed outcome (continuous).

Z Binary treatment indicator (1 = treated, 0 = control).

X1, X2, X3, X4, X5, X6 Pre-treatment covariates (continuous). The true propensity and prognostic models include an X1:X2 interaction.

survey_weight Survey design weight (inverse selection probability).

Details

Generated by a simulation where:

- Treatment assignment: $P(Z = 1|X) = \text{logit}^{-1}(0.3 + 0.6X_1 + 0.4X_2 - 0.3X_3 + 0.2X_1X_2)$.
- Outcome model: $Y(0) = 1 + X_1 + 0.5X_2 - 0.3X_3 + 0.2X_4 + 0.3X_1X_2 + \varepsilon$, with treatment effect $\tau(X) = 0.8 + 0.2X_1$.
- Survey selection: treatment-dependent (retrospective) with $P(S = 1|Z, X) = \text{logit}^{-1}(-2 + 0.3Z + 0.2X_1 + 0.15X_2)$.

Source

Simulated data; see `data-raw/make_survey_data.R`.

Examples

```
data(survey_obs)
head(survey_obs)

# Estimate PATE
fit <- wdsmatchATE(Y = survey_obs$Y, X = survey_obs[, 3:8],
                   Z = survey_obs$Z, weights = survey_obs$survey_weight,
                   M = 3, varest = FALSE)
fit
```

`wdsmatchATE` *Weighted Double Score Matching Estimator for Population Average Treatment Effect*

Description

Estimates the population average treatment effect (PATE) using weighted double score matching with survey design weights. Matching uses propensity probabilities and arm-specific prognostic scores. Missing potential outcomes are imputed by survey-weighted averaging within match sets and aggregated using Hajek normalization. Optional polynomial bias correction adjusts for matching discrepancies. Inference uses fixed-reuse multinomial replication with centered normal Wald intervals.

Usage

```
wdsmatchATE(
  Y,
  X,
  Z,
  weights,
  M = 5,
  ps = NULL,
```

```
  pg = NULL,
  model.ps = NULL,
  model.pg = NULL,
  sampling = c("retrospective", "prospective"),
  use.bias.correction = TRUE,
  varest = TRUE,
  boots = 200,
  alpha = 0.05,
  tie.seed = 20260917L,
  tie.tolerance = 64 * .Machine$double.eps
)
```

Arguments

`Y`	Finite numeric vector of observed outcomes, with no missing values.
`X`	Numeric matrix or data frame with finite numeric covariate columns, one row per unit, and unique names other than `Y` or `Z`.
`Z`	Binary treatment indicator (1 = treated, 0 = control), with one entry per unit and both groups represented.
`weights`	Finite, strictly positive numeric vector of survey design weights, with one entry per unit. Required.
`M`	Positive integer number of nearest neighbors (default 5). Each required donor arm must contain at least `M` units.
`ps`	Optional finite numeric vector of pre-estimated propensity probabilities strictly between zero and one. If `NULL` (default), estimate internally using `model.ps`. Supplied scores remain fixed during replication; see Details.
`pg`	Optional matrix of pre-estimated prognostic scores, with columns `psi0` (control) and `psi1` (treated) in that order and one row per unit. PATT also accepts a one-column matrix containing only the control-side score; with two columns, only the first is used. If `NULL` (default), estimate the required arm-specific scores using `model.pg`. Supplied scores remain fixed during replication; see Details.
`model.ps`	Formula for the propensity model, such as `Z ~ X1 + X2`. If `NULL`, use all columns of `X`. Used only when `ps` is `NULL`.
`model.pg`	Formula for the prognostic regression, such as `Y ~ X1 + X2`. If `NULL`, use all columns of `X`. Used only when `pg` is `NULL`.
`sampling`	Character: `"retrospective"` (default) for treatment-dependent sampling with survey-weighted propensity fitting, or `"prospective"` for treatment-independent sampling with unweighted propensity fitting. This choice also applies to replication refits.
`use.bias.correction`	Logical: apply the arm-specific complete quadratic bias correction (default `TRUE`); see Details for the interpretation of `FALSE`.
`varest`	Logical: compute replication variance and a centered normal Wald interval (default `TRUE`).
`boots`	Integer number of multinomial replicates, at least 2 when variance is requested (default 200).

alpha	Significance level strictly between zero and one (default 0.05). The interval confidence level is 1 - alpha.
tie.seed	Nonnegative integer seed for recipient-specific random tie selection (default 20260917), independent of the bootstrap seed. Use set.seed() separately to reproduce bootstrap draws.
tie.tolerance	Nonnegative relative tolerance on squared distances. The default is 64 * .Machine$double.eps, scaled by the larger of one and the Mth squared distance. Set zero to randomize exact ties only.

Details

Under retrospective sampling, the propensity score is fitted with survey weights; under prospective sampling it is fitted without survey weights. Prognostic regressions are unweighted within each treatment arm in both designs. These modeling choices require the corresponding identification and sampling assumptions; weighting alone does not establish those assumptions.

Matching uses unweighted pooled standardization of propensity probabilities and the relevant prognostic score. The bias-correction regression uses survey weights and a complete quadratic basis in that arm's own double score: an intercept, both coordinates, their squares, and their interaction. Distances are Euclidean and matching is with replacement. At the Mth-distance boundary, the required donors are sampled uniformly without replacement, separately for every recipient, while retaining all strictly closer donors. A separate, reproducible random-number stream controls this selection and leaves the caller's random-number state unchanged. Thus tie randomization does not change the bootstrap count stream. Original matches remain fixed in replication. Numerical boundary equivalence is controlled by tie.tolerance; use zero for exact floating-point equality only. This rule avoids reusing one fixed first-M subset throughout a tied score cell; it does not establish inference guarantees for discrete scores. Propensity fitting starts from zero coefficients and checks convergence and finite, nonsaturated fitted probabilities.

Multinomial replication draws counts for the original sample units and keeps the original match sets and weighted reuse coefficients fixed. With bias correction enabled, internally estimated nuisance scores are refitted using the counts, with additional survey weighting for retrospective propensity fitting and outcome bias correction. Supplied scores remain fixed. Thus, when scores are supplied, the resulting inference is conditional on them and excludes uncertainty from their external estimation. The corresponding model formula is not used to replace a supplied score during replication.

With B = boots successful replicates, the variance estimate is mean((boot.estimates - mean(boot.estimates))^2), using divisor B. The interval is estimate + c(-1, 1) * qnorm(1 - alpha/2) * se. All requested replicates must succeed. An invalid fit, degenerate score, unsuccessful replicate, or nonfinite calculation raises an error; failed draws are not omitted or replaced by zeros. A successful numerical check does not establish adequate statistical overlap.

With use.bias.correction = FALSE, point estimation uses direct matched imputation without polynomial outcome adjustment, and replication uses the fixed-reuse expression without a bias-correction regression. Unused nuisance scores are not refitted for these raw fixed-match replicates. Matching discrepancies remain, so the asymptotic justification for the bias-corrected estimator does not automatically apply to this option.

Double robustness concerns consistency under the stated identification, sampling, positivity, and regularity assumptions when a required score model is correctly specified; it does not guarantee negligible finite-sample bias or nominal interval coverage in every setting. This interface accepts

unit-level weights, not survey strata, clusters, or design-specific replicate weights. The individual-unit multinomial procedure is not a general variance estimator for arbitrary complex survey designs.

Value

A list of class `wdsmatch` with components:

`estimate`	Point estimate of PATE.
`variance`	Replication variance with divisor B, or `NA` when `varest = FALSE`.
`se`	Replication standard error, or `NA` without inference.
`ci`	Centered normal Wald interval as `c(lower, upper)`, or two `NA` values without inference.
`boot.estimates`	All requested finite replication estimates, or `NULL` without inference.
`alpha`	Requested significance level.
`interval.type`	Interval-construction metadata.
`variance.divisor`	Replication-variance divisor metadata.
`n.boot`	Number of retained replicates.
`estimand`	Target estimand.
`sampling`	Sampling-design option.
`settings`	Estimation and replication settings.
`M`	Number of matches used.
`n`	Sample size.
`n.treated`	Number of treated units.
`n.control`	Number of control units.
`call`	The matched call.
`diagnostics`	Point-representation agreement and numerical propensity-fit diagnostics for the point estimate and replicates, plus tie counts, seed, tolerance, and donor-reuse diagnostics.

References

Zeng Y, Tong G, Tong J, Lu H, Mukherjee B, Li F. Where to weight? Estimating population causal effects with weighted double score matching in complex surveys. Preprint. arXiv. Posted online October 19, 2026. doi:10.48550/arXiv.2610.25870.

Examples

```
data(survey_obs)
fit <- wdsmatchATE(
  Y = survey_obs$Y,
  X = survey_obs[, c("X1", "X2", "X3", "X4", "X5", "X6")],
  Z = survey_obs$Z,
  weights = survey_obs$survey_weight,
  M = 3,
```

```
  model.ps = Z ~ X1 + X2 + X3 + X4 + X5 + X6 + X1:X2,
  model.pg = Y ~ X1 + X2 + X3 + X4 + X5 + X6 + X1:X2,
  sampling = "retrospective",
  varest = FALSE
)
fit
```

---

wdsmatchATT *Weighted Double Score Matching Estimator for Population Average Treatment Effect on the Treated*

Description

Estimates the population average treatment effect on the treated (PATT) using survey-weighted double score matching. Treated units are matched to controls using the control-side double score consisting of propensity probability and control prognostic score. Only the missing control outcome requires imputation; treated outcomes remain observed. Aggregation uses treated-side survey weights and Hajek normalization.

Usage

```
wdsmatchATT(
  Y,
  X,
  Z,
  weights,
  M = 5,
  ps = NULL,
  pg = NULL,
  model.ps = NULL,
  model.pg = NULL,
  sampling = c("retrospective", "prospective"),
  use.bias.correction = TRUE,
  varest = TRUE,
  boots = 200,
  alpha = 0.05,
  tie.seed = 20260917L,
  tie.tolerance = 64 * .Machine$double.eps
)
```

Arguments

Y	Finite numeric vector of observed outcomes, with no missing values.
X	Numeric matrix or data frame with finite numeric covariate columns, one row per unit, and unique names other than Y or Z.

Z	Binary treatment indicator (1 = treated, 0 = control), with one entry per unit and both groups represented.
weights	Finite, strictly positive numeric vector of survey design weights, with one entry per unit. Required.
M	Positive integer number of nearest neighbors (default 5). Each required donor arm must contain at least M units.
ps	Optional finite numeric vector of pre-estimated propensity probabilities strictly between zero and one. If NULL (default), estimate internally using model.ps. Supplied scores remain fixed during replication; see Details.
pg	Optional matrix of pre-estimated prognostic scores, with columns psi0 (control) and psi1 (treated) in that order and one row per unit. PATT also accepts a one-column matrix containing only the control-side score; with two columns, only the first is used. If NULL (default), estimate the required arm-specific scores using model.pg. Supplied scores remain fixed during replication; see Details.
model.ps	Formula for the propensity model, such as Z ~ X1 + X2. If NULL, use all columns of X. Used only when ps is NULL.
model.pg	Formula for the prognostic regression, such as Y ~ X1 + X2. If NULL, use all columns of X. Used only when pg is NULL.
sampling	Character: "retrospective" (default) for treatment-dependent sampling with survey-weighted propensity fitting, or "prospective" for treatment-independent sampling with unweighted propensity fitting. This choice also applies to replication refits.
use.bias.correction	Logical: apply the arm-specific complete quadratic bias correction (default TRUE); see Details for the interpretation of FALSE.
varest	Logical: compute replication variance and a centered normal Wald interval (default TRUE).
boots	Integer number of multinomial replicates, at least 2 when variance is requested (default 200).
alpha	Significance level strictly between zero and one (default 0.05). The interval confidence level is 1 - alpha.
tie.seed	Nonnegative integer seed for recipient-specific random tie selection (default 20260917), independent of the bootstrap seed. Use set.seed() separately to reproduce bootstrap draws.
tie.tolerance	Nonnegative relative tolerance on squared distances. The default is 64 * .Machine$double.eps, scaled by the larger of one and the Mth squared distance. Set zero to randomize exact ties only.

Details

PATT uses only the control-side prognostic and bias-correction regressions; a treated-outcome regression is not required. Treatment identification uses one-sided unconfoundedness of the control potential outcome, together with the required survey identification, positivity, and regularity assumptions.

Matching, score fitting, arm-specific quadratic bias correction, strict numerical checks, and the behavior of supplied scores follow wdsmatchATE. The PATT replication expression retains the original weighted control reuse coefficients and normalizes by the replicate's treated-side survey-weight total. Variance uses divisor boots; intervals are centered normal Wald intervals. All requested replicates must succeed, otherwise an informative error is raised.

Supplied scores remain fixed during replication, so their external estimation uncertainty is excluded. With use.bias.correction = FALSE, matching discrepancies remain and the bias-corrected asymptotic justification does not automatically apply. The weight-only interface and individual-unit multinomial replication do not provide general variance estimation for arbitrary clustered or stratified survey designs. See wdsmatchATE for full details.

Value

A list of class wdsmatch with the same components as wdsmatchATE, with estimate targeting PATT and estimand identifying PATT.

References

Zeng Y, Tong G, Tong J, Lu H, Mukherjee B, Li F. Where to weight? Estimating population causal effects with weighted double score matching in complex surveys. Preprint. arXiv. Posted online October 19, 2026. doi:10.48550/arXiv.2610.25870.

Examples

```
data(survey_obs)
fit <- wdsmatchATT(
  Y = survey_obs$Y,
  X = survey_obs[, c("X1", "X2", "X3", "X4", "X5", "X6")],
  Z = survey_obs$Z,
  weights = survey_obs$survey_weight,
  M = 3,
  model.ps = Z ~ X1 + X2 + X3 + X4 + X5 + X6 + X1:X2,
  model.pg = Y ~ X1 + X2 + X3 + X4 + X5 + X6 + X1:X2,
  sampling = "retrospective",
  varest = FALSE
)
fit
```

Index

∗ **datasets**
survey_obs, 3

print.wdsmatch, 2

summary.wdsmatch, 3
survey_obs, 3

wdsmatchATE, *2, 3*, 4, *10*
wdsmatchATT, *2, 3*, 8