

Package ‘TwoWayFEWeights’

May 7, 2026

Type Package

Title Estimation of the Weights Attached to the Two-Way Fixed Effects Regressions

Version 2.0.4

Maintainer Diego Ciccía <diego.ciccía@kellogg.northwestern.edu>

Description

Estimates the weights and measure of robustness to treatment effect heterogeneity attached to two-way fixed effects regressions.

Clément de Chaisemartin, Xavier D'Haultfœuille (2020) <DOI:10.1257/aer.20181169>.

License MIT + file LICENSE

Encoding UTF-8

Depends R(>= 3.4.0)

Imports dplyr, magrittr, cli, fixest, rlang, stats, utils, haven

Suggests tinytest, devtools

NeedsCompilation no

Author Diego Ciccía [aut, cre],
Felix Knau [aut],
Mélitine Malezieux [aut],
Doulo Sow [aut],
Shuo Zhang [aut],
Clément de Chaisemartin [aut],
Grant McDermott [ctb]

URL <https://github.com/chaisemartinPackages>

BugReports <https://github.com/chaisemartinPackages/twowayfeweights/issues>

RoxygenNote 7.3.2

Repository CRAN

Date/Publication 2024-07-26 09:20:09 UTC

Contents

print.twowayfeweights	2
twowayfeweights	3

Index	8
--------------	----------

print.twowayfeweights *A print method for twowayfeweights objects*

Description

Printed display of twowayfeweights objects.

Usage

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

Arguments

x	A twowayfeweights object.
...	Currently ignored.

Examples

```
# The following example is based on data from F. Vella and M. Verbeek (1998),
# "Whose Wages Do Unions Raise? A Dynamic Model of Unionism and Wage Rate
# Determination for Young Men".
# Results and further details about the estimation below can be found in
# Section V.C of de Chaisemartin & D'Haultfoeuille (2020a).
# Run the following lines to download the dataset in your local working
# directory and load it to your R environment:
```

```
repo = "chaisemartinPackages/twowayfeweights/main"
file = "wagepan_twfeweights.dta"
url = paste("https://raw.githubusercontent.com", repo, file, sep = "/")
wagepan = haven::read_dta(url)
```

```
# The default `type = "feTR"` estimation strategy uses a fixed-effects
# strategy under the assumption that parallel trends holds.
```

```
twowayfeweights(
  wagepan, # input data
  "lwage", "nr", "year", "union", # Y, G, T, & D
  type = "feTR", # estimation type ("feTR" is the default)
  summary_measures = TRUE, # show summary measures (optional)
  test_random_weights = "educ" # check randomness of weights (optional)
)
```

```
# The next line performs the same exercise using first differences of outcome
```

```

# and treatment:
twowayfeweights(
  wagepan,
  "diff_lwage", "nr", "year", "diff_union", # Y & D changed to differenced versions
  type          = "fdTR",                  # changed
  D0            = "union",                 # added (D0 is req'd for type="fdTR")
  summary_measures = TRUE,
  test_random_weights = "educ"
)

# Please note that the number of negative weights could be different from Section
# V.C. of de Chaisemartin and D'Haultfoeuille (2020a) due to rounding errors that
# affected older versions of the commands.

```

twowayfeweights	<i>Estimation of the weights attached to the two-way fixed effects regressions.</i>
-----------------	---

Description

Estimates the weights and measure of robustness to treatment effect heterogeneity attached to two-way fixed effects regressions.

Usage

```

twowayfeweights(
  data,
  Y,
  G,
  T,
  D,
  type = c("feTR", "feS", "fdTR", "fdS"),
  D0 = NULL,
  summary_measures = FALSE,
  controls = NULL,
  weights = NULL,
  other_treatments = NULL,
  test_random_weights = NULL,
  path = NULL
)

```

Arguments

data	A data frame or data matrix.
Y	Character string. The dependent variable in the regression. Y is the level of the outcome if one wants to estimate the weights attached to the fixed-effects regression, and Y is the first difference of the outcome if one wants to estimate the weights attached to the first-difference regression. Required.

G	Character string. The variable identifying each group. Required.
T	Character string. The variable identifying each period. Required.
D	Character string. The treatment variable in the regression. D is the level of the treatment if one wants to estimate the weights attached to the fixed-effects regression, and D is the first difference of the treatment if one wants to estimate the weights attached to the first-difference regression. Required.
type	Character string. The type of estimation strategy. Can take one of four different values, each defining a unique combination of regression strategy (either fixed-effects or first-difference) and inference assumption (either common trends on its own or common trends plus an additional assumption about treatment stability over time). <ul style="list-style-type: none"> • "feTR" (the default). Fixed-effects regression under the common trends assumption. • "feS". Fixed-effects regression under common trends and the assumption that groups' treatment effect does not change over time. • "fdTR". first-difference regression under the common trends assumption. • "fdS": first-difference regression under common trends and the assumption that groups' treatment effect does not change over time.
D0	Character string. If type = "fdTR" is specified above, then the function requires a fifth argument, D0. D0 is the mean of the treatment in group g and at period t. It should be non-missing at the first period when a group appears in the data (e.g. at t=1 for the groups that are in the data from the beginning), and for all observations for which the first-difference of the group-level mean outcome and treatment are non missing.
summary_measures	Logical. Should the complementary results from the computation of the weights be displayed? Specifically, the option outputs: (i) the point estimate of the coefficient on the D variable from a TWFE regression, (ii) the minimum value of the standard deviation of the ATEs compatible with the coefficient from the TWFE regression and ATE across all treated (g,t) cells being equal to zero, (iii) the minimum value of the standard deviation of the ATEs compatible with the coefficient from the TWFE regression and ATE across all treated (g,t) cells having different signs (this is computed only if the sum of negative weights is different from 0). See the FAQ section for other details.
controls	Character string(s). An optional vector of control variables that are included in the regression. Controls should not vary within each <i>groupperiod cell</i> , because the results in <i>de Chaisemartin & D'Haultfoeuille (2020a)</i> apply to two-way fixed effects regressions with <i>group×period</i> level controls. If a control does vary within a <i>group×period cell</i> , the command will replace it by its average value within each <i>groupperiod cell</i> .
weights	Character string. Specifies a column name in the input data that replaces the default weighting scheme. If the regression is weighted, the weight variable can be specified in weight. If type="fdTR" is specified, then the weight variable should be non-missing at the first period when a group appears in the data (e.g. at t=1 for the groups that are in the data from the beginning), and for all observations for which the first-difference of the group-level mean outcome and treatment are non missing.

other_treatments	Character string(s). An optional vector of other treatment variables that are included in the regression. Note that this option can only be used when type = "feTR" is specified above. While the results in de Chaisemartin & D'Haultfoeuille (2020a) do not cover two-way fixed effects regressions with several treatments, those in de Chaisemartin & D'Haultfoeuille(2020b) do, so the command follows results from that second paper when other_treatments is specified. When it is specified, the command reports the number and sum of positive and negative weights attached to the treatment, but it does not report the summary measures of the regression's robustness to heterogeneous treatment effects, as these summary measures are no longer applicable when the regression has several treatment variables. The command also reports the weights attached to the other treatments. The weights reported by the command are those in Corollary 1 in de Chaisemartin & D'Haultfoeuille (2020b). See de Chaisemartin & D'Haultfoeuille (2020b) for further details.
test_random_weights	Character string(s). An optional vector that, when specified, will cause the function to estimate the correlation between each variable in the vector and the weights. Testing if those correlations significantly differ from zero is a way to assess whether the weights are as good as randomly assigned to groups and time periods.
path	File path for saving the results in a valid csv file that containing 3 variables (Group, Time, Weight). This option allows the user to see the weight attached to each group*time cell. If the other_treatments option is specified, the weights attached to the other treatments are also saved.

Details

This function estimates the weights attached to the two-way fixed effects regressions studied in de Chaisemartin & D'Haultfoeuille (2020a), as well as summary measures of these regressions' robustness to heterogeneous treatment effects.

Value

A list object with an additional "twowayfeweights" class attribute to enable bespoke print (and other) methods. Included among the slots of the returned list object is a data frame containing the weights attached to each group*time cell.

FAQ

Q. How can one interpret the summary measures of the regression's robustness to heterogeneous treatment effects?

When the two-way fixed effects regression has only one treatment variable, the command reports two summary measures of the robustness of the treatment coefficient beta to treatment heterogeneity across groups and over time. The first one is defined in point (i) of Corollary 1 in de Chaisemartin & D'Haultfoeuille (2020a). It corresponds to the minimal value of the standard deviation of the treatment effect across the treated groups and time periods under which beta and the average treatment effect on the treated (ATT) could be of opposite signs. When that number is large, this means that beta and the ATT can only be of opposite signs if there is a lot of treatment effect heterogeneity

across groups and time periods. When that number is low, this means that beta and the ATT can be of opposite signs even if there is not a lot of treatment effect heterogeneity across groups and time periods. The second summary measure is defined in point (ii) of Corollary 1 in de Chaisemartin & D'Haultfoeuille (2020a). It corresponds to the minimal value of the standard deviation of the treatment effect across the treated groups and time periods under which beta could be of a different sign than the treatment effect in all the treated group and time periods.

Q. How can I tell if the first summary measure is high or low?

Assume that the first summary measure is equal to x . How can you tell if x is a low or a high amount of treatment effect heterogeneity? This is not an easy question to answer, but here is one possibility. Let us assume that the treatment effects of (g,t) cells are drawn from a uniform distribution. Then, to have that the mean of that distribution is 0 while its standard deviation is x , the treatment effects should be uniformly distributed on the $[-\sqrt{3}x, \sqrt{3}x]$ interval. Then, you can ask yourself: is it reasonable to assume that some (g,t) cells have a treatment effect as large as $\sqrt{3}x$, while other cells have a treatment effect as low as $-\sqrt{3}x$? If the answer is negative (you think that it is not reasonable to assume that the treatment effect will exceed the $\pm\sqrt{3}x$ bounds for some (g,t) cells), this means that the uniform distribution of treatment effects compatible with an ATT of 0 and a standard deviation of x seems implausible to you. Then, you can consider that the command's first summary measure is high, and that it is unlikely that beta and the ATT are of a different sign. Conversely, if the answer is positive (you believe that the treatment effect might exceed the bounds for some (g,t) cells), it may not be unlikely that beta and the ATT are of a different sign.

The previous sensitivity exercise assumes that treatment effects follow a uniform distribution. You may find it more reasonable to assume that they are, say, normally distributed. Then you can conduct the following, similar exercise. If the treatment effects of (g,t) cells are drawn from a normal distribution with mean 0 and standard deviation x normal distribution, 95\

Q. How can I tell if the second summary measure is high or low?

Assume that the second summary measure is equal to x . To fix ideas, let us assume that $\beta > 0$. Let us assume that the treatment effects of (g,t) cells are drawn from a uniform distribution. Then, one could have that those effects are all negative, with a standard deviation equal to x , for instance if they are uniformly drawn from the $[-2\sqrt{3}x, 0]$. Then, you can ask yourself: is it reasonable to assume that some (g,t) cells have a treatment effect as low as $-2\sqrt{3}x$? If the answer is negative (you are not willing to assume that some (g,t) cells have a treatment effect lower than $-2\sqrt{3}x$), this means that the uniform distribution of treatment effects compatible with sign reversal and a standard deviation of x seems implausible to you. Then, you can consider that the command's second summary measure is high, and that sign reversal is unlikely. If the treatment effects of (g,t) cells are all negative, they cannot follow a normal distribution, so we do not discuss that possibility here.

References

de Chaisemartin, C and D'Haultfoeuille, X (2020a). American Economic Review, vol. 110, no. 9. Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects.

de Chaisemartin, C and D'Haultfoeuille, X (2020b). Two-way fixed effects regressions with several treatments.

Examples

The following example is based on data from F. Vella and M. Verbeek (1998),

```
# "Whose Wages Do Unions Raise? A Dynamic Model of Unionism and Wage Rate
# Determination for Young Men".
# Results and further details about the estimation below can be found in
# Section V.C of de Chaisemartin & D'Haultfoeuille (2020a).
# Run the following lines to download the dataset in your local working
# directory and load it to your R environment:

repo = "chaisemartinPackages/twowayfeweights/main"
file = "wagepan_twfeweweights.dta"
url = paste("https://raw.githubusercontent.com", repo, file, sep = "/")
wagepan = haven::read_dta(url)

# The default `type = "feTR"` estimation strategy uses a fixed-effects
# strategy under the assumption that parallel trends holds.
twowayfeweights(
  wagepan,                # input data
  "lwage", "nr", "year", "union", # Y, G, T, & D
  type = "feTR",         # estimation type ("feTR" is the default)
  summary_measures = TRUE, # show summary measures (optional)
  test_random_weights = "educ" # check randomness of weights (optional)
)

# The next line performs the same exercise using first differences of outcome
# and treatment:
twowayfeweights(
  wagepan,
  "diff_lwage", "nr", "year", "diff_union", # Y & D changed to differenced versions
  type = "fdTR", # changed
  D0 = "union", # added (D0 is req'd for type="fdTR")
  summary_measures = TRUE,
  test_random_weights = "educ"
)

# Please note that the number of negative weights could be different from Section
# V.C. of de Chaisemartin and D'Haultfoeuille (2020a) due to rounding errors that
# affected older versions of the commands.
```

Index

`print.twowayfeweights`, 2

`twowayfeweights`, 3