

# Causal Inference Final Project: Replication and Extension

Sánchez-García, Rodon & Delgado-García, *Where has everyone gone? Depopulation and voting behaviour in Spain*, *European Journal of Political Research*

Giorgio Coppola

2026-05-28

This project replicates and extends Álvaro Sánchez-García, Toni Rodon, and Maria Delgado-García’s **Where has everyone gone? Depopulation and voting behaviour in Spain**, published in the *European Journal of Political Research* (vol. 64, pp. 296–319). DOI: [10.1111/1475-6765.12702](https://doi.org/10.1111/1475-6765.12702).

## 1 Paper summary

### 1.1 Research question

Sánchez-García, Rodon and Delgado-García (henceforth SRD) ask whether depopulation has an effect on voting behaviour. This factor has been understudied, while most academic attention has focused on the rural-urban divide. By the paper’s definition, depopulation means sustained municipal population loss over a relatively long period. The authors measure it over roughly two legislatures and classify municipalities as losing net population, gaining net population, or remaining broadly stable using year-specific standard-deviation cutoffs.

Spain is a useful case because population loss is common, politically salient, and unevenly distributed. The paper reports that 77 percent of Spanish municipalities have lost residents since 2001. The paper also considers whether depopulation produces support for parties that claim to defend neglected territories. In Spain, this includes local “Empty Spain” lists, such as Teruel Existe, which campaign for provinces affected by population loss and declining public services. These lists matter because they allow the authors to test whether depopulated municipalities turn toward parties built around this issue. The authors’ second explanation is that depopulation may change vote shares because younger residents are more likely to leave, so the remaining electorate becomes older and relatively more conservative.

The paper studies seven political parties or entities: PSOE, PP, Unidas Podemos (UP), Ciudadanos (Cs), non-state-wide parties (regional parties that compete only within a particular autonomous community, PANES in Spanish, NSWP in English), Empty Spain lists (ES), and Vox. It also asks whether the relationship between depopulation and voting changes with municipal size, age composition, public services, and private amenities.

### 1.2 Theoretical and empirical estimand

The **theoretical estimand** is an average causal effect over the units in the authors’ data. A unit is a municipality observed in a given election year. For each party or list, the estimand compares two potential outcomes for that unit: the party’s vote share if the municipality had lost population over roughly the previous two legislatures, and the party’s vote share if its population had remained broadly stable. The authors also consider the analogous comparison between population increase and no change.

The **empirical estimand** is the quantity the authors compute from the observed data to approximate that causal effect. In the replicated specification, it is the coefficient on **Decrease in**

a two-way fixed-effects regression of party vote share on depopulation status, municipality fixed effects, election-year fixed effects, and time-varying controls. In plain language, this coefficient compares changes in a party’s vote share within the same municipality when that municipality is classified as losing population rather than stable, after accounting for national election-year shocks and observed local changes. The coefficient on **Increase** is the same type of comparison for municipalities classified as gaining population. These quantities have a causal interpretation only if the assumptions below hold.

The authors also estimate a second treatment for municipalities classified as being at risk of disappearing. This risk category applies to municipalities with negative population growth, negative natural balance, and population density below 12.5 inhabitants per square kilometre. Its empirical target is an ATT for municipalities that enter the risk category, so it is different from the fixed-effects coefficients replicated here.

### 1.3 Dependent variable and treatment

Element	Definition
<b>Outcome</b> ( $Y$ )	Party vote share in Spanish parliamentary elections. The 2011–2019 analysis covers PSOE, PP, UP, Cs, PANES, ES, and Vox.
<b>Treatment</b> ( $D$ )	<code>depo_cat_te</code> , a three-category measure: no change, decrease, or increase in municipal population over roughly two legislatures. SRD compute the population-change rate and classify each municipality relative to that year’s cross-municipal distribution: above +0.5 SD as decrease, below $-0.5$ SD as increase, otherwise no change.
<b>Unit / panel</b>	A unit is a municipality observed in a given election year. The long panel covers 1989–2019 for the older party families; the shorter 2011–2019 panel allows the authors to include time-varying controls and newer parties.

### 1.4 Identification challenges

Depopulation is not randomly assigned. Municipalities lose population because of job opportunities elsewhere, ageing, weak local services, and broader urbanisation. Those same forces can also affect voting. A shrinking municipality may become more conservative because residents change their preferences, but it may also look more conservative because younger and more left-leaning residents leave while older residents remain.

The authors address part of this problem with municipality fixed effects, year fixed effects, and time-varying controls. Municipality fixed effects remove stable differences such as geography, historical party alignment, and long-run local culture. Year fixed effects remove national shocks common to all municipalities. The controls account for observed changes in population size, age structure, unemployment, and sectoral employment. The main remaining concern is local change that is not observed in the data but affects both population loss and voting, such as a school closure, a local economic shock, or a mobilisation campaign.

### 1.5 Method

SRD first estimate two-way fixed-effects models. Their longer 1989–2019 analysis focuses on parties observed across the full period. Their shorter 2011–2019 analysis adds controls and includes newer parties. My replication targets this controlled specification for the seven parties or entities:

$$Y_{i,t}^j = \alpha_i + \delta_t + \beta_1 \mathbf{1}[D_{i,t} = \text{dec.}] + \beta_2 \mathbf{1}[D_{i,t} = \text{inc.}] + X'_{i,t} \gamma + \varepsilon_{i,t}^j,$$

The authors estimate the model with `fixest::feols`, municipality and year fixed effects, and standard errors clustered by region, municipality, and year. They then study heterogeneity by interacting depopulation with municipal size, age-structure change, public services, and amenities. Separately, they use `fect` to estimate the ATT for entering the depopulation-risk category.

## 1.6 DAG

```

dag <- ggdag::dagify(
  Y ~ D + M + alpha + delta + X + U,
  M ~ D,
  D ~ alpha + delta + X + U,
  exposure = "D",
  outcome = "Y",
  coords = list(
    x = c(D = 0, Y = 4, M = 2, alpha = -1.2, delta = 5.2, X = 2, U = 2),
    y = c(D = 0, Y = 0, M = -0.7, alpha = 1.6, delta = 1.6, X = -4.0, U = 2.4))
)

label_df <- tibble::tribble(
  ~x, ~y, ~label,
  -1.10, 0, "D: depopulation",
  5.10, 0, "Y: vote share",
  2, -1.55, "M: composition of population",
  -1.2, 2.35, "alpha_i: mun. FE",
  5.2, 2.35, "delta_t: year FE",
  2, -4.70, "X_it: controls",
  2, 3.15, "U_it: unobserved")

ggdag::ggdag_status(dag,
  text = FALSE,
  use_labels = NULL,
  node_size = 18,
  stylized = TRUE) +
ggplot2::geom_point(data = data.frame(x = 2, y = -0.7),
  mapping = ggplot2::aes(x = x, y = y),
  shape = 21,
  size = 18,
  fill = "#f4c842",
  colour = "grey60",
  stroke = 0.5,
  inherit.aes = FALSE) +
ggplot2::geom_label(data = label_df,
  mapping = ggplot2::aes(x = x, y = y, label = label),
  size = 2.8,
  label.size = 0.25,
  label.r = ggplot2::unit(0.15, "lines"),
  colour = "grey25",
  fill = "white",

```

```

    alpha      = 0.95,
    inherit.aes = FALSE) +
ggplot2::scale_color_manual(
  values  = c(exposure = "#2c7fb8", outcome = "#d7301f", latent = "grey75"),
  na.value = "grey60",
  guide   = "none") +
ggplot2::scale_fill_manual(
  values  = c(exposure = "#2c7fb8", outcome = "#d7301f", latent = "grey75"),
  na.value = "grey75",
  guide   = "none") +
ggplot2::expand_limits(x = c(-2.8, 6.8), y = c(-5.1, 3.5)) +
ggdag::theme_dag()

```

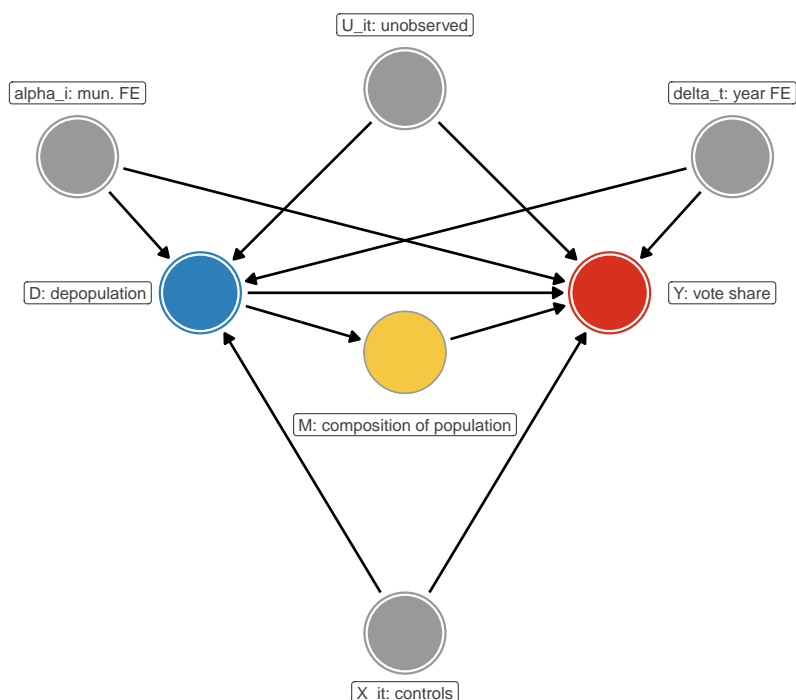


Figure 1: DAG for the identification analysis. Municipality fixed effects represent stable local traits; year fixed effects represent national shocks; controls represent observed local changes. The remaining concern is unobserved local change,  $U_{i,t}$ .

The DAG represents the data-generating process. Municipality fixed effects handle stable local differences; year fixed effects handle national shocks; observed controls  $X_{i,t}$  handle observed local changes. The remaining confounder is  $U_{i,t}$ : local time-varying factors that affect both depopulation and voting but are not observed. The mediator  $M_{i,t}$ , the political composition of the remaining voters, makes the sorting channel explicit: depopulation can change the vote share either through preference change among the people who stay (the direct path  $D \rightarrow Y$ ) or because younger, more left-leaning voters disproportionately leave (the indirect path  $D \rightarrow M \rightarrow Y$ ), even if no remaining voter changes their mind. The paper's empirical strategy does not separately

identify the two channels: without individual-level data and without exogenous variation in  $M$ , the regression coefficient on `Decrease` returns their sum.

### 1.7 Potential outcomes notation

For municipality  $i$ , election year  $t$ , and party  $j$ , let  $D_{i,t} \in \{\text{no change, decrease, increase}\}$  and let  $Y_{i,t}^j(d)$  be the potential vote share under depopulation status  $d$ . Taking the expectation over the units in the relevant panel, the main comparisons are

$$\tau_1^j \equiv \mathbb{E}[Y_{i,t}^j(\text{dec.}) - Y_{i,t}^j(\text{no ch.})], \quad \tau_2^j \equiv \mathbb{E}[Y_{i,t}^j(\text{inc.}) - Y_{i,t}^j(\text{no ch.})].$$

Consistency requires that the observed vote share equals the potential outcome under the municipality’s observed depopulation status. If effects differ across municipalities or years, the TWFE coefficient should be read as an average within-municipality comparison, not as the effect for every place.

### 1.8 Assumptions: empirical $\rightarrow$ theoretical estimand

1. **Conditional parallel trends.** In the absence of depopulation, municipalities in different depopulation categories would have followed similar vote-share trends after conditioning on fixed effects and controls.
2. **No anticipation.** Voters should not change behaviour because they expect future population loss.
3. **No spillovers.** One municipality’s vote share should not depend on another municipality’s depopulation status. This is demanding because migration links sending and receiving places.
4. **Comparable treatment.** The “decrease” category should capture a similar type of demographic process across municipalities, even though the same population loss can mean different things in a city and in a small village.

### 1.9 Assumptions: estimate $\rightarrow$ empirical estimand

1. **Correct model specification.** The regression needs to represent the relevant controls and fixed effects well enough for the coefficient on depopulation to be interpretable.
2. **Valid inference.** The clustered standard errors need to behave well with the number of clusters and the sample used for each party.
3. **Reliable treatment measurement.** Municipalities near the  $\pm 0.5$  SD cutoffs should not be frequently misclassified.

### 1.10 Challenges

- **Sorting.** The paper’s preferred interpretation is partly compositional: young people leave, older residents remain, and party vote shares shift. That is substantively important, but it limits a simple “depopulation changes preferences” interpretation.
- **Spatial spillovers.** Emigrants usually move somewhere else in Spain, so depopulation in one municipality can change the electorate in another.
- **Time-varying local shocks.** A plant closure, school closure, or local political campaign could affect both population and voting.
- **Endogenous moderators.** Public services and amenities may be consequences of depopulation, not purely pre-treatment moderators.

## 2 Replication of the main controlled regression specification

I replicate **Table 2** of the paper, “The effect of depopulation on electoral results (2011–2019)” (p. 306), which contains the paper’s main controlled fixed-effects estimates linking depopulation to party vote shares across the seven outcomes (PSOE, PP, UP, Cs, PANES, ES, Vox). This is the analytical table the rest of the paper builds on: the moderator analyses (Figures 2–5) use the same specification with interactions added, and the depopulation-risk results (Figure 7)

extend it with a different identification strategy.

## 2.1 Code

```
df <- load_rdata("replication/database/short_time.RData")

rhs <- "depo_cat_te + agriculture_workers + services_workers + unemployed +
      var_older_te + var_younger_te + population_log | mun_code + year"

fit_one <- function(y) {
  fml <- as.formula(paste(y, "~", rhs))
  tryCatch(
    feols(fml, data = df, cluster = c("region", "mun_code", "year"), notes = FALSE),
    error = \(e) feols(fml, data = df, cluster = c("mun_code", "year"), notes = FALSE)
  )
}

models <- c("psoe", "pp", "podemos", "cs", "nswp", "es", "vox") |>
  set_names() |>
  map(fit_one)
```

## 2.2 Results: replicated versus published

```
coef_map <- c(
  "depo_cat_te2_decrease" = "Decrease",
  "depo_cat_te3_increase" = "Increase",
  "agriculture_workers"   = "% workers agriculture",
  "services_workers"     = "% workers services",
  "unemployed"           = "% unemployed",
  "var_older_te"         = "D % 60 y/o or older",
  "var_younger_te"       = "D % 16 y/o or younger",
  "population_log"       = "(Log) Population")

modelssummary::modelssummary(
  models,
  coef_map      = coef_map,
  gof_omit     = "AIC|BIC|RMSE|Log.Lik|Std.Errors|R2 Adj.|R2 Within Adj.|FE: region",
  stars        = c("*" = 0.10, "***" = 0.05, "****" = 0.01),
  title        = "Replicated controlled fixed-effects estimates, 2011-2019.",
  output       = "kableExtra") |>
  kableExtra::kable_styling(bootstrap_options = c("striped", "hover", "condensed"),
                             full_width      = FALSE,
                             position        = "center",
                             latex_options   = c("striped", "scale_down", "HOLD_position"))
```

Table 1: Replicated controlled fixed-effects estimates, 2011-2019.

	psoe	pp	podemos	cs	nswp	es	vox
Decrease	-0.233*	0.381***	-0.419**	0.006	0.249	-0.012	0.082
	(0.106)	(0.057)	(0.129)	(0.242)	(0.208)	(0.008)	(0.342)
Increase	0.817*	0.690	-0.983	0.967*	-0.512	0.009	-0.384
	(0.334)	(0.858)	(0.756)	(0.314)	(0.262)	(0.011)	(0.588)
% workers agriculture	0.006	0.014	0.010*	-0.002	0.003	0.000	-0.031
	(0.019)	(0.015)	(0.004)	(0.014)	(0.011)	(0.001)	(0.015)
% workers services	-0.005	0.003	0.025**	0.007	-0.002	0.000	-0.036*
	(0.007)	(0.018)	(0.006)	(0.010)	(0.005)	(0.001)	(0.015)
% unemployed	0.022	0.064	-0.064	-0.017	0.048	-0.001	-0.005
	(0.030)	(0.059)	(0.035)	(0.014)	(0.030)	(0.001)	(0.038)
D % 60 y/o or older	-0.002	-0.008	-0.007	-0.030	-0.024	-0.001	0.082*
	(0.023)	(0.020)	(0.023)	(0.025)	(0.016)	(0.001)	(0.027)
D % 16 y/o or younger	0.013	-0.072	-0.017	0.079	-0.108	-0.001	0.193
	(0.041)	(0.076)	(0.047)	(0.040)	(0.077)	(0.001)	(0.110)
(Log) Population	-4.824	2.514	4.540	-6.855***	3.011	-0.062	-0.806
	(2.338)	(5.504)	(3.967)	(0.205)	(4.517)	(0.088)	(6.446)
Num.Obs.	40 396	40 396	40 396	32 342	33 507	6732	32 342
R2	0.927	0.939	0.830	0.790	0.984	0.334	0.771
R2 Within	0.008	0.003	0.018	0.014	0.008	0.001	0.006
FE: mun_code	X	X	X	X	X	X	X
FE: year	X	X	X	X	X	X	X

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

The replication reproduces the two depopulation coefficients in each column to four decimal places. For PSOE, PP, UP, Cs, PANES, and Vox, the standard errors also match the published estimates after rounding. The exception is the ES column, where the published article reports standard errors around  $10^{-5}$  and highly significant coefficients, while the current rerun gives standard errors near 0.008 (Decrease) and 0.011 (Increase) with no stars. Sample sizes are slightly smaller in the rerun for every column. The side-by-side table below summarizes the comparison.

Table 2: Published versus replicated estimates. Coefficients are reported with standard errors in parentheses.

Party	Decrease coefficient		Increase coefficient		Sample size		
	Published (Dec.)	Replicated (Dec.)	Published (Inc.)	Replicated (Inc.)	N (Pub.)	N (Repl.)	Diff.
PSOE	-0.2329 (0.1064)	-0.2329 (0.1064)	0.8166 (0.3345)	0.8166 (0.3345)	40,400	40,396	4
PP	0.3811 (0.0567)	0.3811 (0.0567)	0.6904 (0.8576)	0.6904 (0.8576)	40,400	40,396	4
UP	-0.4188 (0.1294)	-0.4188 (0.1294)	-0.9834 (0.7563)	-0.9834 (0.7563)	40,400	40,396	4
Cs	0.0056 (0.2425)	0.0056 (0.2425)	0.9673 (0.3138)	0.9673 (0.3138)	32,343	32,342	1
PANES	0.2488 (0.2078)	0.2488 (0.2078)	-0.5124 (0.2618)	-0.5124 (0.2618)	34,424	33,507	917
ES	-0.0118 (0.0000)	-0.0118 (0.0076)	0.0092 (0.0000)	0.0092 (0.0109)	7,460	6,732	728
VOX	0.0822 (0.3423)	0.0822 (0.3423)	-0.3842 (0.5883)	-0.3842 (0.5883)	32,343	32,342	1

### 2.3 Discrepancies

1. **Sample size.** The rerun uses slightly fewer observations than the published estimates. The gap is 4 observations for PSOE, PP, and UP; 1 observation for Cs and Vox; 917 for PANES; and 728 for ES. I do not force the published sample size by hand. The current `fixest` run reports dropped observations from missing values and fixed-effect singletons,

and the coefficient estimates remain unchanged at the reported precision.

2. **ES standard errors.** The ES models are restricted to Aragón and Castilla y León in 2015, 2016, and November 2019. With the current package versions, the three-way clustered ES model fails when `fixest` tries to compute the variance-covariance matrix. The script therefore falls back to municipality-year clustering for this column only. The ES coefficients match the paper, but the uncertainty does not: the published ES standard errors are effectively zero, while the rerun gives about 0.008 for `Decrease` and 0.011 for `Increase`. I therefore treat the ES significance stars as not replicated.

### 3 Extension: heterogeneous treatment effects via causal forest

#### 3.1 Motivation

SRD already show that the relationship between depopulation and voting differs across municipalities. Their moderator analyses consider one factor at a time: population size, age composition, public services, and amenities. My extension asks whether those patterns remain when several moderators are considered together.

I use a causal forest for this extension. There are two reasons. First, TWFE with a non-binary, fluctuating treatment is known to recover a weighted average of within-unit comparisons whose weights can be perverse under heterogeneous treatment effects (de Chaisemartin & d'Haultfoeuille, 2020). Whether that matters here depends on whether the depopulation–vote relationship is in fact heterogeneous. The paper studies heterogeneity through one-moderator-at-a-time interactions and provides no formal test that effect heterogeneity is present at all. Second, a causal forest estimates how the association between depopulation and Vox vote share varies across municipalities considered jointly across observables, rather than one interaction at a time. The forest does not fix identification, since it inherits the paper’s within-municipality assumptions, but it documents a property (heterogeneity) that is directly relevant to how the paper’s headline coefficients should be read.

I focus on **Vox**. The controlled fixed-effects estimates show no average relationship between ordinary depopulation and Vox vote share, but the paper’s depopulation-risk analysis suggests that Vox gains support when municipalities become at risk of disappearing. Vox is therefore a useful case for checking whether a null average hides differences across types of municipalities.

#### 3.2 Design

- **Sample:** I follow the paper’s moderator analyses and compare only `no change` and `decrease`; municipalities classified as `increase` are dropped.
- **Treatment:**  $W_{i,t} = 1$  when the municipality is in the `decrease` category and 0 when it is in `no change`.
- **Outcome:** Vox vote share.
- **Fixed effects:** I remove municipality and year fixed effects from both the outcome and treatment before fitting the forest.
- **Covariates:** the six controls from the fixed-effects specification plus `public_services` and `degurba`, the EU degree-of-urbanisation category. I exclude `amenities` because it is missing for more than 20,000 observations in the short panel.
- **Missing data:** Vox is missing before the party entered the election data, so those observations drop out together with rows missing the selected covariates.
- **Algorithm:** `grf::causal_forest` with 2000 trees and default tuning.

#### 3.3 Code and fit

```
covars <- c("agriculture_workers", "services_workers", "unemployed",
           "var_older_te", "var_younger_te", "population_log",
           "public_services", "degurba")
```

```

d_clean <- df |>
  filter(depo_cat_te %in% c("1_no_change", "2_decrease")) |>
  mutate(W = as.integer(depo_cat_te == "2_decrease"),
         public_services = as.numeric(public_services),
         degurba = as.numeric(degurba)) |>
  filter(!is.na(vox),
         if_all(all_of(covars), \ (x) !is.na(x)))

demeaned <- fixest::demean(
  X = d_clean |> select(vox, W) |> as.matrix(),
  f = d_clean |> select(mun_code, year) |> as.data.frame())

X <- d_clean |> select(all_of(covars)) |> as.matrix()
Y_resid <- demeaned[, "vox"]
W_resid <- demeaned[, "W"]

set.seed(2026)
cf <- grf::causal_forest(X, Y_resid, W_resid, num.trees = 2000, seed = 2026)

ate_vec <- grf::average_treatment_effect(cf)
ate <- tibble(estimate = ate_vec[["estimate"]],
             std_error = ate_vec[["std.err"]])
calibration <- coefstest_to_tibble(grf::test_calibration(cf))
blp <- coefstest_to_tibble(grf::best_linear_projection(cf, X))
cates <- predict(cf)$predictions

vi <- tibble(variable = covars,
             importance = grf::variable_importance(cf)[, 1]) |>
  arrange(importance)

pop_bins <- d_clean |>
  transmute(
    pop_quintile = cut(population_log,
                      quantile(population_log, probs = seq(0, 1, .2)),
                      include.lowest = TRUE,
                      labels = c("Q1 (smallest)", "Q2", "Q3", "Q4", "Q5 (largest)")),
    population = exp(population_log),
    cate = cates) |>
  group_by(pop_quintile) |>
  summarise(n = n(),
            median_population = median(population),
            mean_CATE = mean(cate),
            median_CATE = median(cate),
            .groups = "drop")

```

### 3.4 Is there heterogeneity at all?

```

calibration |>
  knitr::kable(
    digits = c(0, 4, 4, 2, 4),
    col.names = c("Term", "Estimate", "Std. error", "t value", "p value"),

```

```
caption = "Calibration test for the causal forest.")
```

Table 3: Calibration test for the causal forest.

Term	Estimate	Std. error	t value	p value
mean.forest.prediction	1.3816	0.3148	4.39	0
differential.forest.prediction	1.7239	0.1327	12.99	0

`test_calibration` checks whether the forest separates observations with higher and lower estimated effects. A `differential.forest.prediction` coefficient significantly above zero is evidence that within-municipality treatment effects vary systematically with observables. Under such heterogeneity the published TWFE coefficient on `Decrease` is not the effect for a particular municipality but a weighted average across switchers, which is the aggregation problem flagged in de Chaisemartin & d'Haultfœuille (2020).

### 3.5 Distribution of CATEs

```
ggplot(tibble(cate = cates), aes(cate)) +
  geom_histogram(bins = 60, fill = "#2c7fb8", colour = "white", alpha = 0.9) +
  geom_vline(xintercept = ate$estimate, colour = "#d7301f", linewidth = 0.7) +
  annotate("text", x = ate$estimate, y = Inf, vjust = 2,
         hjust = -0.05, label = "ATE", colour = "#d7301f") +
  labs(x = "Conditional ATE on Vox vote share",
       y = "Number of units") +
  theme_minimal(base_size = 12)
```

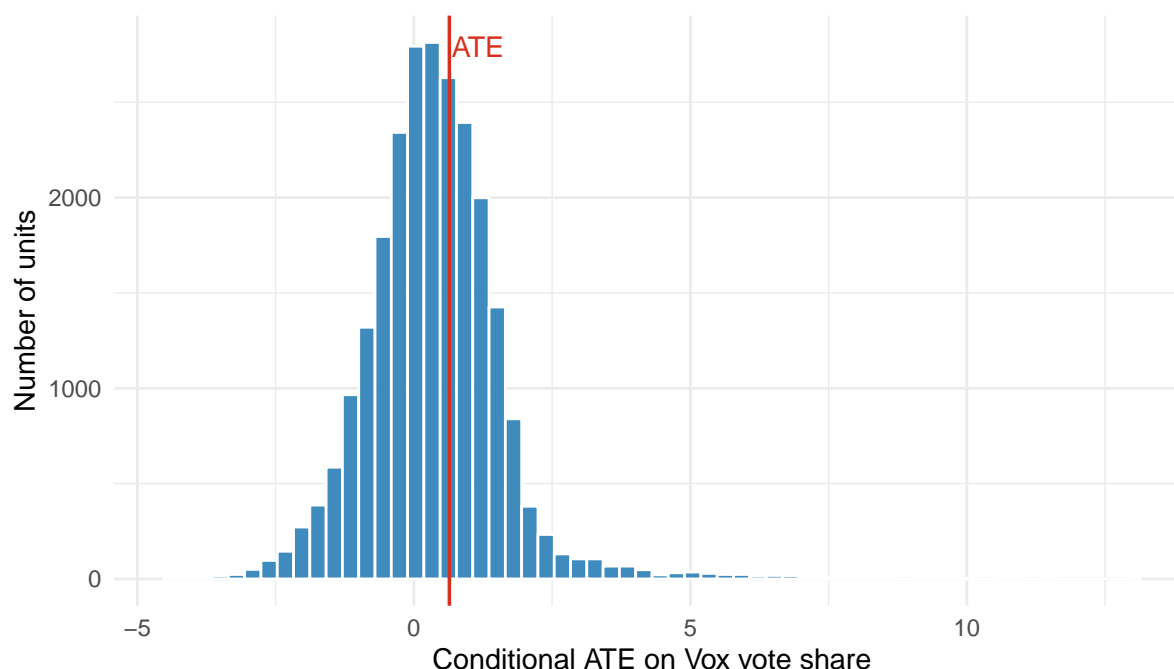


Figure 2: Distribution of estimated CATEs across the sample. The vertical line marks the average treatment effect from the forest; the histogram shows the spread of conditional effects.

### 3.6 Heterogeneity drivers

```
ggplot(vi, aes(importance, factor(variable, levels = variable))) +
  geom_col(fill = "#2c7fb8") +
  labs(x = "Variable importance (share of splits)", y = NULL) +
  theme_minimal(base_size = 12)
```

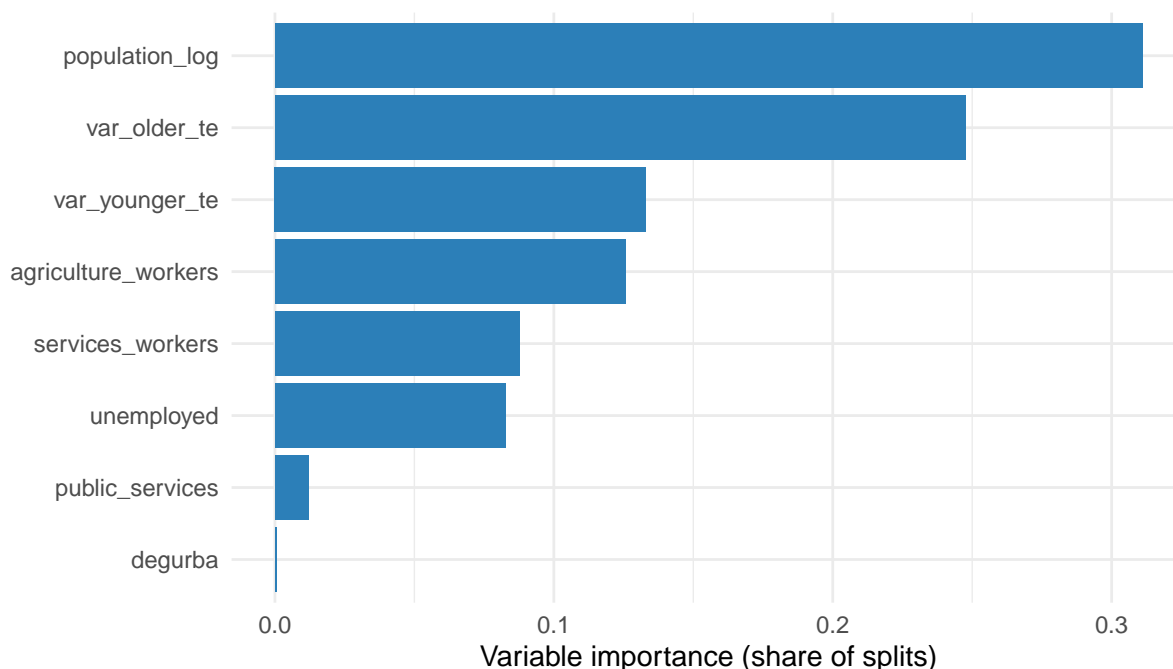


Figure 3: Variable importance: share of forest splits that use each covariate. Larger bars indicate covariates that drive more of the heterogeneity in treatment effects.

```
blp |>
  knitr::kable(
    digits = c(0, 4, 4, 2, 4),
    col.names = c("Term", "Estimate", "Std. error", "t value", "p value"),
    caption = "Best linear projection of the forest's CATEs onto the covariates.")
```

Table 4: Best linear projection of the forest's CATEs onto the covariates.

Term	Estimate	Std. error	t value	p value
(Intercept)	1.8373	1.6965	1.08	0.2788
agriculture_workers	-0.0336	0.0078	-4.30	0.0000
services_workers	-0.0071	0.0086	-0.82	0.4103
unemployed	-0.0546	0.0239	-2.29	0.0223
var_older_te	0.0248	0.0377	0.66	0.5115
var_younger_te	0.1470	0.0502	2.93	0.0034
population_log	0.4169	0.1078	3.87	0.0001
public_services	0.1401	0.0958	1.46	0.1438
degurba	-0.8308	0.3757	-2.21	0.0270

The best linear projection summarizes which covariates are associated with higher or lower estimated effects. It considers the candidate moderators together rather than one at a time.

### 3.7 Does the SRD finding hold up?

The calibration test suggests substantial heterogeneity: the `differential.forest.prediction` test statistic is 13.0. Estimated CATEs range from about -4.4 to 13.0 percentage points. This does not mean every municipality-level estimate should be interpreted on its own. It means the model finds systematic differences in the depopulation-Vox relationship across observed municipal characteristics.

The main substantive result is the population “gradient”. In the smallest quintile, where the median municipality has about 67 residents, depopulation is associated with a mean CATE of -0.71 percentage points for Vox. In the largest quintile, where the median municipality has about 7,870 residents, the mean CATE is 1.43 points. This is consistent with the paper’s moderator analysis by population size: Vox does not gain from ordinary depopulation in the smallest places, but it does better in larger municipalities that are losing population.

The variable-importance ranking points to size and composition. The top three moderators are `population_log`, `var_older_te`, `var_younger_te`. The best linear projection gives the same main message: the slope on `population_log` is positive and statistically significant.

Table 5: Estimated CATE on Vox vote share by quintile of log municipal population.

Population quintile	N	Median population	Mean CATE	Median CATE
Q1 (smallest)	4,883	67	-0.705	-0.717
Q2	4,845	188	-0.116	-0.097
Q3	4,860	479	0.346	0.308
Q4	4,863	1,455	0.879	0.868
Q5 (largest)	4,862	7,870	1.430	1.160

**What does the extension add?** It shows that the heterogeneity found in the paper is not driven only by one chosen interaction. Municipal size and age-structure change remain the most important moderators when the covariates are considered together.

**Note:** The forest’s average estimate (0.64 pp, SE 0.12) is not directly comparable to the published fixed-effects coefficient for Vox (0.082, not significant). The forest is fit after removing fixed effects from both  $Y$  and  $W$ , so I use it to study heterogeneity rather than to replace the paper’s average-effect estimate.

**Summary:** The extension supports the paper’s main substantive claim but qualifies how the published numbers should be read. The forest produces strong, formal evidence of effect heterogeneity (calibration  $t \approx 13$ ), which is the condition under which TWFE aggregation is most vulnerable. The direction of the Vox finding survives: depopulation is associated with lower Vox support in the smallest municipalities and higher Vox support in larger shrinking ones. The published null-on-average coefficient masks this gradient and should not be interpreted as depopulation does not influence Vox vote share, but rather as an average that obscures opposing effects across the municipal size distribution, negative in the smallest places and positive in larger shrinking ones.

## 4 References

- Sánchez-García, Á., Rodon, T., & Delgado-García, M. (2025). Where has everyone gone? Depopulation and voting behaviour in Spain. *European Journal of Political Research*, 64, 296–319. doi:[10.1111/1475-6765.12702](https://doi.org/10.1111/1475-6765.12702).
- Lundberg, I., Johnson, R., & Stewart, B. M. (2021). What is your estimand? Defining the target quantity connects statistical evidence to theory. *American Sociological Review*,

86(3), 532–565.

- Liu, L., Wang, Y., & Xu, Y. (2024). A practical guide to counterfactual estimators for causal inference with time-series cross-sectional data. *American Journal of Political Science*, 68(1), 160–176.
- de Chaisemartin, C., & d’Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9), 2964–2996.
- Tibshirani, J., Athey, S., Sverdrup, E., & Wager, S. (2026). *grf: Generalized Random Forests* (R package version 2.6.1).