

Panel Methods (1): The Parametric Approach

Lectures on Causal Panel Analysis


Yiqing Xu
Stanford University

September 2026


Feel free to use or adapt the slides with attribution (CC BY 4.0).

Motivations


- Causal inference is challenging with observational data
- Panel data often come to the rescue (**false hope?**)
- In the social sciences, difference-in-differences (DiD) and two-way fixed effects (TWFE) are the most popular



What's a Panel Dataset?




What's a Panel Dataset?



What's Causal Inference with Panel Data?

What would happen if a group of units experienced a set of counterfactual treatment histories?



What's Special about Panel Data?

- The fundamental problem of causal inference (Holland 1986)

$$\tau_i = Y_{1i} - Y_{i0}$$

- A statistical solution makes use of others' information


$$\text{e.g. } ATE = E[Y_1] - E[Y_0]$$

- A scientific solution exploits homogeneity or invariance assumptions

e.g. The long-run growth rate of the US economy is 2.5%.


- Panel data allow us to construct treated counterfactuals using information from both **the past** and **the others** with the caveat that treatment assignment mechanism may be complicated
- Causal inference with panel data is also challenging because of
 - potentially (very) high dimensional treatment
 - limited knowledge about treatment assignment mechanisms
 - both cross-sectional and temporal interference (SUTVA violations)

Panel Models Impute Different Counterfactuals




- Each fit uses all control outcomes and the treated unit through period 6.

Source: Author's simulation.



Lectures 1 and 6 build counterfactuals from outcome models; the rest start from treatment-timing comparisons, with assumptions tailored to those comparisons.




Main question


Which observations does each estimator compare, and under what conditions does that comparison recover a causal effect?

Application 1: Direct Democracy and Naturalization

Treatment timing



Naturalization rates in six municipalities




Question

Do elected municipal councils raise naturalization rates relative to citizen referendums?


Source: Hainmueller and Hangartner (2019).

Application 2: Beer Taxes and Fatalities

Pooled comparison



After subtracting each state's mean



Question

Do higher beer taxes reduce traffic fatalities?

Source: Ruhm (1996); Stock and Watson (2007).

Panel Setup & Unobserved Confounding

Fixed Effects Regression

Modeling Time

Modeling Dynamic & Heterogeneous Effects

Single unit:

$$\mathbf{y}_i = \begin{pmatrix} y_{i1} \\ \vdots \\ y_{it} \\ \vdots \\ y_{iT} \end{pmatrix}_{T \times 1} \quad \mathbf{X}_i = \begin{pmatrix} x_{i,1,1} & x_{i,1,2} & x_{i,1,j} & \cdots & x_{i,1,K} \\ \vdots & \vdots & \vdots & & \vdots \\ x_{i,t,1} & x_{i,t,2} & x_{i,t,j} & \cdots & x_{i,t,K} \\ \vdots & \vdots & \vdots & & \vdots \\ x_{i,T,1} & x_{i,T,2} & x_{i,T,j} & \cdots & x_{i,T,K} \end{pmatrix}_{T \times K}$$

Panel with all units:

$$\mathbf{y} = \begin{pmatrix} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_i \\ \vdots \\ \mathbf{y}_N \end{pmatrix}_{NT \times 1} \quad \mathbf{X} = \begin{pmatrix} \mathbf{X}_1 \\ \vdots \\ \mathbf{X}_i \\ \vdots \\ \mathbf{X}_N \end{pmatrix}_{NT \times K}$$

Unobserved Effects Model: Farm Output

- For a randomly drawn cross-sectional unit i , the model is given by

$$y_{it} = \mathbf{x}_{it}\boldsymbol{\beta} + c_i + \varepsilon_{it}, \quad t = 1, 2, \dots, T$$

- y_{it} : output of farm i in year t
- $\mathbf{x}_{it}$: $1 \times K$ vector of variable inputs for farm i in year t , such as labor, fertilizer, etc. plus an intercept
- $\boldsymbol{\beta}$: $K \times 1$ vector of marginal effects of variable inputs, the **causal quantity of interest**
- c_i : farm effect, i.e. the sum of all time-invariant inputs known to farmer i (but unobserved for the researcher), such as soil quality, managerial ability, etc.
 - often called: **unobserved effect**, **unobserved heterogeneity**, capturing **unobserved confounding**
- ε_{it} : time-varying unobserved inputs, such as rainfall, unknown to the farmer at the time the decision on the variable inputs $\mathbf{x}_{it}$ is made
 - often called: **idiosyncratic error**
- What happens when we regress y_{it} on $\mathbf{x}_{it}$?

- When we ignore the panel structure and regress y_{it} on $\mathbf{x}_{it}$ we get

$$y_{it} = \mathbf{x}_{it}\beta + v_{it}, \quad t = 1, 2, \dots, T$$

with **composite error** $v_{it} \equiv c_i + \varepsilon_{it}$

- Main assumption to obtain consistent estimates for β is:
 - $E[v_{it} | \mathbf{x}_{i1}, \mathbf{x}_{i2}, \dots, \mathbf{x}_{iT}] = E[v_{it} | \mathbf{x}_{it}] = 0$ for $t = 1, 2, \dots, T$
 - $\mathbf{x}_{it}$ are **strictly exogenous**: the composite error v_{it} in each time period is uncorrelated with the past, current, and future regressors
 - But: labour input $\mathbf{x}_{it}$ likely depends on soil quality c_i and so we have omitted variable bias and $\hat{\beta}$ is not consistent
 - No correlation between $\mathbf{x}_{it}$ and $v_{it} \Rightarrow$ no correlation between unobserved effect c_i and $\mathbf{x}_{it}$ for all t
 - Violations are common: whenever we omit a time-constant variable that is correlated with the regressors (**unobserved confounding**)

Unobserved Effects Model: Program Evaluation

- Program evaluation model:

$$y_{it} = \text{prog}_{it} \beta + c_i + \varepsilon_{it}, \quad t = 1, 2, \dots, T$$

- y_{it} : log wage of individual i in year t
- prog_{it} : indicator coded 1 if individual i participates in training program at t and 0 otherwise
- β : effect of program
- c_i : sum of all time-invariant unobserved characteristics that affect wages, such as ability, etc.
- What happens when we regress y_{it} on prog_{it} ? $\hat{\beta}$ not consistent since prog_{it} is likely correlated with c_i (e.g. ability)
- Always ask: Is there a time-constant unobserved variable (c_i) that is correlated with the regressors? If yes, pooled OLS is problematic
- Additional problem: $v_{it} \equiv c_i + \varepsilon_{it}$ are serially correlated for same i since c_i is present in each t and thus pooled OLS standard errors are invalid


Panel Setup & Unobserved Confounding

Fixed Effects Regression

Modeling Time

Modeling Dynamic & Heterogeneous Effects


Simpson's Paradox: Pooled and Within-State Slopes



- Each dot is a state-year observation.
- Subtracting each state's means changes the slope from $+0.365$ to -0.656 .

Source: U.S. traffic-fatality panel; Ruhm (1996) and Stock and Watson (2007).

Between-State and Within-State Variation



- Colored dots are annual observations; colored lines are state-specific fits.
- Gray dots show all other state-years; diamonds show state means.
- The right panel subtracts each state's mean on both axes.

Source: U.S. traffic-fatality panel; Ruhm (1996) and Stock and Watson (2007).

- Our unobserved effects model is:

$$y_{it} = \mathbf{x}_{it}\boldsymbol{\beta} + c_i + \varepsilon_{it}, \quad t = 1, 2, \dots, T$$

- If we have data on multiple time periods, we can think of c_i as **fixed effects** or “nuisance parameters” to be estimated
- OLS estimation with fixed effects yields:

$$(\hat{\boldsymbol{\beta}}, \hat{c}_1, \dots, \hat{c}_N) = \underset{\mathbf{b}, m_1, \dots, m_N}{\operatorname{argmin}} \sum_{i=1}^N \sum_{t=1}^T (y_{it} - \mathbf{x}_{it}\mathbf{b} - m_i)^2$$

this amounts to including N farm dummies in regression of y_{it} on $\mathbf{x}_{it}$

$$(\hat{\beta}, \hat{c}_1, \dots, \hat{c}_N) = \underset{\mathbf{b}, m_1, \dots, m_N}{\operatorname{argmin}} \sum_{i=1}^N \sum_{t=1}^T (y_{it} - \mathbf{x}_{it}\mathbf{b} - m_i)^2$$

The first-order conditions (FOC) for this minimization problem are:

$$\sum_{i=1}^N \sum_{t=1}^T \mathbf{x}'_{it} (y_{it} - \mathbf{x}_{it}\hat{\beta} - \hat{c}_i) = 0$$

and

$$\sum_{t=1}^T (y_{it} - \mathbf{x}_{it}\hat{\beta} - \hat{c}_i) = 0$$

for $i = 1, \dots, N$.

Therefore, for $i = 1, \dots, N$,

$$\hat{c}_i = \frac{1}{T} \sum_{t=1}^T (y_{it} - \mathbf{x}_{it}'\hat{\beta}) = \bar{y}_i - \bar{\mathbf{x}}_i'\hat{\beta},$$

where

$$\bar{\mathbf{x}}_i \equiv \frac{1}{T} \sum_{t=1}^T \mathbf{x}_{it}, \quad \bar{y}_i \equiv \frac{1}{T} \sum_{t=1}^T y_{it}.$$

Plug this result into the first FOC to obtain:

$$\begin{aligned} \hat{\beta} &= \left(\sum_{i=1}^N \sum_{t=1}^T (\mathbf{x}_{it} - \bar{\mathbf{x}}_i)' (\mathbf{x}_{it} - \bar{\mathbf{x}}_i) \right)^{-1} \left(\sum_{i=1}^N \sum_{t=1}^T (\mathbf{x}_{it} - \bar{\mathbf{x}}_i)' (y_{it} - \bar{y}_i) \right) \\ \hat{\beta} &= \left(\sum_{i=1}^N \sum_{t=1}^T \ddot{\mathbf{x}}_{it}' \ddot{\mathbf{x}}_{it} \right)^{-1} \left(\sum_{i=1}^N \sum_{t=1}^T \ddot{\mathbf{x}}_{it}' \ddot{y}_{it} \right) \end{aligned}$$

with time-demeaned variables $\ddot{\mathbf{x}}_{it} \equiv \mathbf{x}_{it} - \bar{\mathbf{x}}_i$, $\ddot{y}_{it} \equiv y_{it} - \bar{y}_i$.

Running a regression with the time-demeaned variables $\check{y}_{it} \equiv y_{it} - \bar{y}_i$ and $\check{\mathbf{x}}_{it} \equiv \mathbf{x}_{it} - \bar{\mathbf{x}}_i$ is numerically equivalent to a regression of y_{it} on $\mathbf{x}_{it}$ and unit specific dummy variables.

Fixed effects estimator is often called the **within estimator** because it only uses the time variation within each cross-sectional unit.

The regression with the time-demeaned variables is consistent for β even when $\text{Cov}[\mathbf{x}_{it}, c_i] \neq 0$, because time-demeaning eliminates the unobserved effects:

$$y_{it} = \mathbf{x}_{it}\beta + c_i + \varepsilon_{it}$$


$$\bar{y}_i = \bar{\mathbf{x}}_i\beta + c_i + \bar{\varepsilon}_i$$

$$(y_{it} - \bar{y}_i) = (\mathbf{x}_{it} - \bar{\mathbf{x}}_i)\beta + (c_i - c_i) + (\varepsilon_{it} - \bar{\varepsilon}_i)$$

$$\check{y}_{it} = \check{\mathbf{x}}_{it}\beta + \check{\varepsilon}_{it}$$

Panel Connectedness: Municipalities and Years

Panel Structure: Bipartite Graph



- A line marks a usable municipality–year observation.
- Red rings mark a municipality or year represented once.
- Rumendingen appears only in 2003, so it adds no within-municipality variation.

Source: Hainmueller and Hangartner (2019) replication data; selected municipalities, 1998–2009; `panelView(type = "network")`.

Naturalization Panel: Variables and Timing

- The panel contains 1,400 Swiss municipalities observed annually from 1991 to 2009.
- The outcome is the naturalization rate:

$$y_{it} = 100 \times \frac{\text{naturalizations}_{it}}{\text{eligible foreign population}_{i,t-1}}.$$

- $x_{it} = 1$ when an elected municipal council decides applications and 0 when citizens vote in referendums.
- Some municipalities changed their decision rules over time.

Source: Hainmueller and Hangartner (2019).

```
library(haven)
library(dplyr)
library(fixest)

d <- read_dta("figs/Swiss_Panel_long.dta") |>
  mutate(year = as.integer(year))

d |> select(munID, muni_name, year,
           nat_rate, repdem) |> slice(30:38)
```

Data structure

4,655 municipality-years; 245 municipalities; 1991–2009.

```
m_pool <- feols(nat_rate ~ repdem, data = d)
etable(m_pool)
```

	Estimate	Standard error
Representative democracy	2.503	0.129

Question

Does the cross-sectional difference reflect a causal effect, or stable differences across municipalities?

```
d <- d |>
  group_by(muniID) |>
  mutate(across(c(nat_rate, repdem),
    ~ .x - mean(.x), .names = "{.col}_dm")) |>
  ungroup()

m_within <- feols(nat_rate_dm ~ repdem_dm, data = d)
```

Within comparison

The coefficient uses only changes within the same municipality.

```
m_fe <- feols(nat_rate ~ repdem | muniID,  
              data = d, cluster = ~muniID)
```

	Estimate	Clustered SE
Representative democracy	3.023	0.186

- `| muniID` absorbs municipality fixed effects.
- `cluster = muniID` allows serial correlation within municipalities.

```
m_twfe <- feols(nat_rate ~ repdem | munID + year ,  
               data = d, cluster = ~munID)  
etable(m_fe , m_twfe)
```

	Unit FE	Unit + year FE
Representative democracy	3.023	1.204
Clustered SE	0.186	0.303

The coefficient falls sharply once year fixed effects remove changes common to all municipalities.

```
feols(nat_rate_dm ~ repdem_dm, data = d)
feols(nat_rate ~ repdem + i(muniID), data = d)
feols(nat_rate ~ repdem | muniID, data = d)
```

Same coefficient, different computation

Demeaning makes the within-municipality comparison explicit. For large panels, absorbing fixed effects is more convenient.

$$y_{it} - \lambda \bar{y}_i = (\mathbf{x}_{it} - \lambda \bar{\mathbf{x}}_i) \boldsymbol{\beta} + (v_{it} - \lambda \bar{v}_i), \quad 0 \leq \lambda \leq 1$$

- The random-effects estimator uses changes within units and differences in average levels across units.
- The parameter λ determines how much of each unit mean is subtracted.
- It requires $E[c_i | \mathbf{x}_i] = 0$: persistent unobserved unit differences are unrelated to all regressors.
- It can estimate coefficients on time-invariant regressors.

Pooled OLS, Fixed Effects, and Random Effects

	Pooled OLS and random effects	Fixed effects
Comparisons used	Within and between	Within each unit
Relation between c_i and $\mathbf{x}_i$	Assume no correlation	Allows correlation
Time-invariant regressors	Estimated	Not separately identified
Main risk	Time-invariant and time-varying confounding	Time-varying confounding; x_{it} responds to earlier outcome shocks

Model choice

Use fixed effects when stable unit differences may be correlated with the regressors. The random-effects model requires the design to make $E[c_i | \mathbf{x}_i] = 0$ credible. Comparing the fixed- and random-effects estimates does not establish this assumption.

R Demo: Beer Taxes and Traffic Fatalities

```
library(plm)
data("Fatalities", package = "AER")
Fatalities$fatal_rate <- with(Fatalities, fatal / pop * 10000)
p <- pdata.frame(Fatalities, index = c("state", "year"))

pool <- plm(fatal_rate ~ beertax, p, model = "pooling")
fe <- plm(fatal_rate ~ beertax, p, model = "within")
re <- plm(fatal_rate ~ beertax, p, model = "random")
twfe <- plm(fatal_rate ~ beertax, p,
            model = "within", effect = "twoways")

state_se <- function(m) sqrt(diag(vcovHC(
  m, method = "arellano", type = "HC1", cluster = "group")))
```

Data

The four models use 336 observations from 48 states over seven years.

Beer-Tax Results Across Models

Model	Beer-tax coefficient
Pooled OLS	0.365 (0.119)
Random effects	-0.052 (0.109)
State FE	-0.656 (0.289)
State + year FE	-0.640 (0.350)

- Parentheses contain standard errors clustered by state using `state_se()`.
- Pooled OLS and state fixed effects give estimates with opposite signs.
- The random-effects estimate lies between the pooled and state fixed-effects estimates.
- A causal interpretation requires beer-tax changes to be unrelated to unobserved fatality shocks in every year.

Source: Author's calculations using `AER::Fatalities`; Ruhm (1996); Stock and Watson (2007).

Problems that Fixed Effects Do Not Solve

$$y_{it} = \mathbf{x}_{it}\beta + c_i + \varepsilon_{it}, \quad t = 1, 2, \dots, T$$

- y_{it} is murder rate and x_{it} is police spending per capita
- What happens when we regress y on x and city fixed effects?
 - $\hat{\beta}_{FE}$ inconsistent unless **strict exogeneity** conditional on c_i holds
 - $E[\varepsilon_{it} | \mathbf{x}_{i1}, \mathbf{x}_{i2}, \dots, \mathbf{x}_{iT}, c_i] = 0, \quad t = 1, 2, \dots, T$
 - implies ε_{it} uncorrelated with past, current, and future regressors
- Most common violations:
 - ① **Time-varying omitted variables**
 - economic boom leads to more police spending and fewer murders
 - can include time-varying controls, but avoid post-treatment bias
 - ② **Simultaneity & feedback**
 - if city adjusts police based on past murder rate, then spending_{t+1} is correlated with ε_t (since higher ε_t leads to higher murder rate at t)
 - strictly exogenous x cannot react to what happens to y in the past or the future!
- Fixed effects do not obviate need for good research design...

Summary: Fixed Effects, Pooled OLS (and Random Effects)

Assm. 1 Regressors are strictly exogenous conditional on the time-invariant unobservables

Assm. 2 Regressors are uncorrelated with time-varying & time-invariant unobservables

- Results:
 - Fixed effects estimator is consistent given Assm. 1, but rules out time-invariant regressors
 - Pooled OLS (and random effects) are consistent under Assum. 2, and allow for time-invariant regressors
 - Specification tests cannot establish a credible research design.
- Assum. 2 is strong, so fixed effects models are typically more credible
 - One main advantage of panel data is to account for **time-invariant unobserved confounding**


Panel Setup & Unobserved Confounding

Fixed Effects Regression

Modeling Time

Modeling Dynamic & Heterogeneous Effects


Common Changes in the Swiss Panel



- From 1991 to 2009, both variables increase overall.
- A nationwide change can be mistaken for the effect of a local institutional change.

Source: Hainmueller and Hangartner (2019); author's calculations.

Common Linear Trend




$$y_{it} = \mathbf{x}_{it}\beta + c_i + \gamma t + \varepsilon_{it}$$

Common-trend assumption

The model assumes the same linear time trend for every unit: γ per period after accounting for $\mathbf{x}_{it}$ and c_i .

Year Fixed Effects




$$y_{it} = \mathbf{x}_{it}\beta + c_i + \delta_t + \varepsilon_{it}$$

Year-effect assumption

The model gives each year its own shift δ_t and assumes that shift is the same for every unit.

Unit-Specific Linear Trends




$$y_{it} = \mathbf{x}_{it}\beta + c_i + g_it + \delta_t + \varepsilon_{it}$$

Unit-trend assumption

After accounting for $\mathbf{x}_{it}$, c_i , and δ_t , unit i may follow its own linear trend g_it .

After removing a separate linear trend for every unit, little within-unit treatment variation may remain.

Naturalization Results Across Time Models



- Year fixed effects remove changes shared across municipalities.
- Municipality-specific trends remove each municipality's fitted linear change.

Source: Hainmueller and Hangartner (2019); 95% confidence intervals use municipality-clustered standard errors.

Panel Setup & Unobserved Confounding

Fixed Effects Regression

Modeling Time

Modeling Dynamic & Heterogeneous Effects

$$y_{it} = x_{it}\beta_0 + x_{it-1}\beta_1 + x_{it-2}\beta_2 + c_i + \varepsilon_{it}, \quad t = 1, 2, \dots, T$$

- Model recognizes that effect of change in x may occur with a lag
 - effect of new tax credit for children on fertility rate
- Interpretation of coefficients:
 - Consider **temporary increase** in x_{it} from level m to $m + 1$ at t , which lasts only one period

$$y_{t-1} = m\beta_0 + m\beta_1 + m\beta_2 + c_i$$

$$y_t = (m + 1)\beta_0 + m\beta_1 + m\beta_2 + c_i$$

$$y_{t+1} = m\beta_0 + (m + 1)\beta_1 + m\beta_2 + c_i$$

$$y_{t+2} = m\beta_0 + m\beta_1 + (m + 1)\beta_2 + c_i$$

$$y_{t+3} = m\beta_0 + m\beta_1 + m\beta_2 + c_i$$

- $\beta_0 = y_t - y_{t-1}$ immediate change in y due to temporary one-unit increase in x (impact propensity)
- $\beta_1 = y_{t+1} - y_{t-1}$ change in y one period after temporary one-unit increase in x
- $\beta_2 = y_{t+2} - y_{t-1}$ change in y two periods after temporary one-unit increase in x
- $y_{t+3} = y_{t-1}$ change in y is zero three periods after temporary one-unit increase in x

$$y_{it} = x_{it+1}\beta_{-1} + x_{it}\beta_0 + x_{it-1}\beta_1 + x_{it-2}\beta_2 + c_i + \varepsilon_{it}, \quad t = 1, 2, \dots, T$$


- Can use estimate of β_{-1} to test for anticipation effects
 - Consider **temporary increase** in x_{it} from level m to $m+1$ at t
 - $y_{t-2} = \beta_{-1}m + m\beta_0 + m\beta_1 + m\beta_2 + c_i$
 - $y_{t-1} = \beta_{-1}(m+1) + m\beta_0 + m\beta_1 + m\beta_2 + c_i$
- Anticipation effect: $\beta_{-1} = y_{t-1} - y_{t-2}$ change in y in period $t-1$, the period before the temporary one-unit increase in x
- Placebo test: if x causes y , but y does not cause x , then β_{-1} should be close to zero

- Let D_{it} be a binary indicator coded 1 if unit i switched from control to treatment between t and $t - 1$; 0 otherwise
 - Lags: D_{it-1} : unit switched between $t - 1$ and $t - 2$
 - Leads: D_{it+1} : unit switches between $t + 1$ and t
- Include lags and leads into the fixed effects model:

$$y_{it} = D_{it+2}\beta_{-2} + D_{it+1}\beta_{-1} + D_{it}\beta_0 + D_{it-1}\beta_1 + D_{it-2}\beta_2 + c_i + \varepsilon_{it}$$

- Interpretation of coefficients:
 - Leads β_{-1} , β_{-2} , etc. test for anticipation effects
 - Switch β_0 tests for immediate effect
 - Lags β_1 , β_2 , etc. test for long-run effects
 - highest lag dummy can be coded 1 for all post-switch years


Dynamic Effect of Switching to Representative Democracy



Lagged Dependent Variables (LDVs)

$$y_{it} = \alpha y_{it-1} + c_i + \varepsilon_{it}, \quad t = 1, 2, \dots, T$$

- y_{it} could be capital stock of firm i at time t , and α the capital depreciation rate
- Models with unit fixed effects and lagged y do not produce consistent estimators (**Nickell bias**)
 - after taking first differences to eliminate c_i , the differenced residual $\Delta\varepsilon_{it}$ is correlated with the lagged dependent variable Δy_{it-1} by construction
- We might use past levels y_{it-2} as an instrument for Δy_{it-1} and apply a general method of moment (GMM) estimator, but this requires strong assumptions (e.g. no serial correlation in ε_{it})
- When $T \rightarrow \infty$, the Nickell bias goes to 0; controlling for LDVs is fine as long as they are not correlated with treatment assignment
- However, many existing studies do not survive adding one LDV



- So far we have assumed that the treatment effect is constant across units
- We can allow for heterogeneous treatment effects by including interaction of treatment with other regressors


$$y_{it} = \text{treat}_{it}\alpha_0 + (\text{treat}_{it} \cdot x_{it})\alpha_1 + x_{it}\beta + c_i + t + \varepsilon_{it}, \quad t = 1, 2, \dots, T$$

- Often the treatment is interacted with a time-invariant regressor:

$$y_{it} = \text{treat}_{it}\alpha_0 + (\text{treat}_{it} \cdot x_i)\alpha_1 + c_i + t + \varepsilon_{it}, \quad t = 1, 2, \dots, T$$

- Note: The lower order term on the time-invariant x_i is collinear with the fixed effects and drops out

Heterogeneous Effect of Direct Democracy



Source: Own estimates on the Hainmueller and Hangartner (2019) data: two-way fixed effects with a linear interaction (line, 95% band) and by quartile of the SVP share (points); municipality-clustered standard errors (syp heterogeneity 2026.R)

References

- Hainmueller, Jens, and Dominik Hangartner. 2019. "Does Direct Democracy Hurt Immigrant Minorities? Evidence from Naturalization Decisions in Switzerland." *American Journal of Political Science* 63 (3): 530–47.
- Holland, Paul W. 1986. "Statistics and Causal Inference." *Journal of the American Statistical Association* 81 (396): 945–60.
- Ruhm, Christopher J. 1996. "Alcohol Policies and Highway Vehicle Fatalities." *Journal of Health Economics* 15 (4): 435–54.
- Stock, James H., and Mark W. Watson. 2007. *Introduction to Econometrics*. 2nd ed. Pearson Addison Wesley.

Parts of this lecture build on teaching slides by Jens Hainmueller.