


Panel Methods (3): Factorial Difference-in-Differences

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).



- Lecture 2 covered the canonical DID design: a treated group, an untreated comparison group, and parallel trends
- But many events leave **no one untreated**: a famine, a financial crisis, a pandemic, a nationwide reform
- Applied work still runs “DID” in these settings, comparing units that differ in a **pre-event characteristic**
- This lecture: what does that comparison identify, and under which assumptions? (Xu, Zhao, and Ding 2026, *JASA*)
- Lecture 4 then returns to TWFE regressions and what breaks under staggered timing

Motivation: How Social Capital Saved Lives

- Cao, Xu, and Zhang (2022): “How social capital saved lives during China’s Great Famine”
- Data structure:
 - A baseline factor G : time-invariant measure of social capital
 - Event time: {Pre-Famine, Famine, Post-Famine}
- Estimation: difference-in-differences (DID), or equivalently, two-way fixed effects (TWFE)
- Interpretation:
 - **Descriptively**: “the rise in the mortality rate during the famine years is significantly smaller in counties with higher social capital”
 - **Causally**: “we interpret these differences as the effects of social capital on famine relief”



Estimates from TWFE with a dynamic specification.

American Political Science Review (2019) 113, 2, 405–422

doi:[10.1017/S0003055419000017](https://doi.org/10.1017/S0003055419000017)

© American Political Science Association 2019

How Do Immigrants Respond to Discrimination? The Case of Germans in the US During World War I

VASILIKI FOUKA *Stanford University*

I study the effect of taste-based discrimination on the assimilation decisions of immigrant minorities. Do discriminated minority groups increase their assimilation efforts in order to avoid discrimination and public harassment or do they become alienated and retreat in their own communities? I exploit an exogenous shock to native attitudes, anti-Germanism in the United States during World War I, to empirically identify the reactions of German immigrants to increased native hostility. I use two measures of assimilation efforts: naming patterns and petitions for naturalization. In the face of increased discrimination, Germans increase their assimilation investments by Americanizing their own and their children's names and filing more petitions for US citizenship. These responses are stronger in states that registered higher levels of anti-German hostility, as measured by voting patterns and incidents of violence against Germans.

Discrimination × WWI

Source: Fouka (2019).

American Economic Review 2020, 110(11): 3454–3491
<https://doi.org/10.1257/aer.20191054>

Devotion and Development: Religiosity, Education, and Economic Progress in Nineteenth-Century France[†]

By MARA P. SQUICCIARINI*

This paper studies when religion can hamper diffusion of knowledge and economic development, and through which mechanism. I examine Catholicism in France during the Second Industrial Revolution (1870–1914). In this period, technology became skill-intensive, leading to the introduction of technical education in primary schools. I find that more religious locations had lower economic development after 1870. Schooling appears to be the key mechanism: more religious areas saw a slower adoption of the technical curriculum and a push for religious education. In turn, religious education was negatively associated with industrial development 10 to 15 years later, when schoolchildren entered the labor market. (JEL D83, I21, I26, N33, Z12)

Catholicism × Industrial Revolution

Source: Squicciarini (2020).



Elite strongholds × state-building reform

Agents of the Regime? Traditional Leaders and Electoral Politics in South Africa

The Journal of
Politics

Daniel de Kadt, University of California
Horacio A. Larreguy, Harvard University

Traditional leaders are often given subnational authority in developing democracies. Although ubiquitous, the electoral consequences of their power have received little attention, often due to empirical challenges. We argue that traditional leaders, whose power and resources largely depend on the state, have incentives to support incumbent political parties who can guarantee their survival and provide them with rents. We study this quid pro quo in the former Bantustans of South Africa, showing that an alignment between the incumbent African National Congress and chiefs maps to increased electoral support for the party. Our results suggest that chiefs boost African National Congress vote share by between 6.6 and 8.2 percentage points in the Bantustans, translating into a change in the national vote of between 2 and 2.5 percentage points. Our empirical strategy rules out that our results are driven by differences in ethnicity, race, or economic circumstances, which commonly confound the presence of traditional authorities.

Traditional leaders × rise of Zuma

- In each study: a baseline factor G interacted with an event Z that hits every unit

Source: Chen, Wang, and Zhang (2025); de Kadt and Larreguy (2018).

Same DID estimator, a different research design:
“Factorial Difference-in-Differences”

- **Factorial**


- A classic topic in statistics
- Factorial experiments pioneered by R. A. Fisher and Frank Yates
- ≥ 2 treatment factors: their main effects and interaction are of interest

- **Difference-in-differences**

- Popular in economics and related fields
- A “research design” for causal inference with observational data
- Leverages panel data (units $\times$ time) to identify causal effects, usually the ATT

This Lecture

- Factorial DID is a different research design from a canonical DID
- Under no anticipation & parallel trends, the DID estimator identifies **treatment effect heterogeneity** of exposure to the event (with nuances)
- Identifying G 's **causal effect** requires stronger assumptions
- Factorial DID includes canonical DID as a special case with an additional assumption
- Extensions & practical recommendations



Two groups, two periods; no covariates

- Study population: $i = 1, 2, \dots, n$
- Timing of the event is fixed
- Time periods: $t = \text{pre, post}$
- Baseline factor: $G_i \in \{0, 1\}$
- Data: $\{G_i, Y_{i,\text{pre}}, Y_{i,\text{post}} : i = 1, 2, \dots, n\}$
- **Universal exposure** to the event in the post period:

$$Z_i = 1, \quad i = 1, 2, \dots, n$$

	Pre-Event	Post-Event
$G_i = 1$		$Z_i = 1$
$G_i = 0$		$Z_i = 1$

Factorial DID

	Pre-Event	Post-Event
$G_i = 1$		$Z_i = 1$
$G_i = 0$		$Z_i = 0$

Canonical DID

The Difference-in-Differences Estimator

	Pre-Event	Post-Event
$G_i = 1$	$\frac{1}{n_1} \sum_{i:G_i=1} Y_{i,\text{pre}}$	$\frac{1}{n_1} \sum_{i:G_i=1} Y_{i,\text{post}}$
$G_i = 0$	$\frac{1}{n_0} \sum_{i:G_i=0} Y_{i,\text{pre}}$	$\frac{1}{n_0} \sum_{i:G_i=0} Y_{i,\text{post}}$

Define

- $\hat{\tau}_{\text{DID}} = \frac{1}{n_1} \sum_{i:G_i=1} (Y_{i,\text{post}} - Y_{i,\text{pre}}) - \frac{1}{n_0} \sum_{i:G_i=0} (Y_{i,\text{post}} - Y_{i,\text{pre}})$
- Focus on: $\text{plim } \hat{\tau}_{\text{DID}} = \tau_{\text{DID}} = \mathbb{E}[\Delta Y_i \mid G_i = 1] - \mathbb{E}[\Delta Y_i \mid G_i = 0]$

Potential Outcomes


- Frame factorial DID as a factorial design with two factors: g and z

- Potential outcomes indexed by both: $Y_{i,t}(g, z)$

- Observed outcomes:

$$Y_{i,t} = \sum_{g,z \in \{0,1\}} \mathbf{1}\{G_i = g, Z_i = z\} \cdot Y_{i,t}(g, z) = Y_{i,t}(G_i, Z_i)$$

- Recall: in observed data, $Z_i = 1$ for all units $\Rightarrow$ **half of the potential outcomes are never observed**



Red: potential outcomes never observed under universal exposure.


- Individual conditional effect of the event at factor level g :

$$\tau_{i,Z|G=g} = Y_{i,\text{post}}(g, 1) - Y_{i,\text{post}}(g, 0)$$

- Individual interaction effect:

$$\begin{aligned}\tau_{i,\text{cm}} &= Y_{i,\text{post}}(1, 1) - Y_{i,\text{post}}(1, 0) \\ &\quad - Y_{i,\text{post}}(0, 1) + Y_{i,\text{post}}(0, 0) \\ &= \tau_{i,Z|G=1} - \tau_{i,Z|G=0}\end{aligned}$$

the **causal moderation** of G on Z



- **Average causal interaction** (VanderWeele 2009; Bansak 2021):

$$\tau_{\text{cm}} = \mathbb{E}[Y_{i,\text{post}}(1, 1) - Y_{i,\text{post}}(1, 0) - Y_{i,\text{post}}(0, 1) + Y_{i,\text{post}}(0, 0)] = \mathbb{E}[\tau_{i,Z|G=1}] - \mathbb{E}[\tau_{i,Z|G=0}]$$

“Social capital reduced the mortality increase caused by the famine.”


- **Effect modification** — associative:

$$\tau_{\text{em}} = \mathbb{E}[\tau_{i,Z|G=1} \mid G_i = 1] - \mathbb{E}[\tau_{i,Z|G=0} \mid G_i = 0]$$

“Places with higher social capital experienced a smaller mortality increase caused by the famine.”

- The verbs differ because the estimands differ: τ_{cm} is causal, τ_{em} compares **different units**

Identification: A Roadmap



Identification under Canonical DID Assumptions

No Anticipation

$$Y_{i,\text{pre}}(g, 0) = Y_{i,\text{pre}}(g, 1) \text{ for all } i \text{ and } g = 0, 1$$

Parallel Trends


$$\mathbb{E}[\Delta Y_i(1, 0) \mid G_i = 1] = \mathbb{E}[\Delta Y_i(0, 0) \mid G_i = 0]$$

Proposition


Under no anticipation and parallel trends,

$$\tau_{\text{DID}} = \tau_{\text{em}} = \mathbb{E}[\tau_{i,Z|G=1} \mid G_i = 1] - \mathbb{E}[\tau_{i,Z|G=0} \mid G_i = 0]$$

Two different **conditional effects**, across two different groups



Why τ_{DID} May Not Identify Causal Moderation under Parallel Trends



- Imagine an unobservable U that determines how units respond to the event
 - If both groups share the same U , the event moves them equally and $\tau_{\text{DID}} = 0$, whether U is small or large
- U may be **correlated with G** : here, high- U units ($U_i = 3/4$) are concentrated in $G = 1$
- Then $\tau_{\text{DID}} = \tau_{\text{em}}$ mixes the moderation of G with **compositional differences** in U
- Therefore τ_{DID} cannot be interpreted as the causal moderation of G

Identifying Causal Moderation

- Under no anticipation and parallel trends:

$$\tau_{\text{DID}} = \tau_{\text{em}} = \mathbb{E}[Y_{i,\text{post}}(1, 1) - Y_{i,\text{post}}(1, 0) \mid G_i = 1] - \mathbb{E}[Y_{i,\text{post}}(0, 1) - Y_{i,\text{post}}(0, 0) \mid G_i = 0]$$


- When does τ_{em} equal τ_{cm} , which drops the conditioning on G_i ?

Factorial Parallel Trends (sufficient condition)

$\Delta Y_i(g, z) \perp\!\!\!\perp G_i$ (orthogonality between G and first-differenced potential outcomes)

- Weaker than (quasi-)random assignment of G : $Y_{i,t}(g, z) \perp\!\!\!\perp G_i$
- Under no anticipation, parallel trends, and $\Delta Y_i(g, z) \perp\!\!\!\perp G_i$: $\tau_{\text{DID}} = \tau_{\text{em}} = \tau_{\text{cm}}$

Reconciling Canonical and Factorial DID



Exclusion Restriction on Z for Group $G_i = 0$

$$\mathbb{E}[Y_{i,\text{post}}(0, 1) \mid G_i = 0] = \mathbb{E}[Y_{i,\text{post}}(0, 0) \mid G_i = 0]$$

the event does not move $G = 0$ units: their Δ is zero

Proposition

Under no anticipation, parallel trends, and the exclusion restriction,

$$\tau_{DID} = \tau_{att} = \mathbb{E}[Y_{i,\text{post}}(1, 1) - Y_{i,\text{post}}(1, 0) \mid G_i = 1]$$

- This is exactly the canonical DID of Lecture 2: “untreated” units are units for which the event is assumed not to matter

Identifying G 's Conditional Effect

- One more question the interaction might answer: the **effect of G itself** during the event

Exclusion Restriction on G Absent the Event

$$\mathbb{E}[Y_{i,\text{post}}(1, 0)] = \mathbb{E}[Y_{i,\text{post}}(0, 0)]$$


without the event, the baseline factor would not move the outcome

- Under no anticipation, parallel trends, factorial parallel trends, and this exclusion restriction:

$$\tau_{\text{DID}} = \tau_{\text{cm}} = \mathbb{E}[Y_{i,\text{post}}(1, 1) - Y_{i,\text{post}}(0, 1)] = \tau_{G|Z=1}$$

- In the famine example: the average causal effect of social capital on mortality **during the famine years**

Summary of Identification Results



Canonical vs Factorial DID: Two Research Designs

Canonical DID (reframed)

Define the canonical DID research design as the combination of:

- 2×2 data structure & universal exposure
- Identification: under no anticipation & parallel trends & exclusion restriction on Z , $\hat{\tau}_{DID}$ identifies τ_{att}

Factorial DID

Define the factorial DID research design as the combination of:

- 2×2 data structure & universal exposure
- Identification:
 - ① Under no anticipation & parallel trends, $\hat{\tau}_{DID}$ identifies τ_{em}
 - ② Adding $\Delta Y_i(g, z) \perp\!\!\!\perp G_i$, $\hat{\tau}_{DID}$ identifies τ_{cm}
 - ③ Adding also the exclusion restriction on G , $\hat{\tau}_{DID}$ identifies $\tau_{G|Z=1}$

- Researchers often run the following TWFE regression:

$$Y_{it} = \alpha_i + \xi_t + \tau G_i \cdot \text{Post}_t + \beta^T X_i \cdot \text{Post}_t + \epsilon_{it}$$

Table 1

Clans and mortality rate during the great famine.

Cao, Xu & Zhang (2022)

Outcome variable mean	Outcome variable: Mortality rate (%)					
	(1)	(2)	(3)	(4)	(5)	(6)
	13.997	13.997	13.997	13.997	13.997	13.997
High clan density x Famine period	-1.456*** (0.596)	-2.325*** (0.635)	-2.254*** (0.698)			
Log (#Genealogies/pop) x Famine period				-3.385*** (0.398)	-4.108*** (0.597)	-4.296*** (0.661)
Grain output (PC) x Famine period		0.006* (0.003)	0.005* (0.003)		0.005* (0.003)	0.005* (0.003)
Non-farming land ratio x Famine period		0.006 (0.031)	0.002 (0.033)		0.013 (0.031)	0.009 (0.033)
Urbanization rate x Famine period		-0.075** (0.031)	-0.076** (0.034)		-0.088** (0.030)	-0.087** (0.033)
Distance from Beijing x Famine period		0.004*** (0.001)	0.004*** (0.001)		0.004*** (0.001)	0.004*** (0.001)
Distance from provincial capital x Famine period		0.002 (0.003)	0.003 (0.003)		0.001 (0.003)	0.002 (0.003)
Historical migrants * Famine period		-0.078*** (0.011)	-0.081*** (0.012)		-0.053*** (0.011)	-0.055*** (0.013)
Crop suitability index for rice * Famine period		0.359*** (0.160)	0.392*** (0.171)		0.460*** (0.164)	0.470*** (0.175)
Share of minority * Famine period		-0.049*** (0.017)	-0.050*** (0.018)		-0.067*** (0.016)	-0.067*** (0.018)
Averaged years of schooling * Famine period		-0.822** (0.406)	-0.899** (0.431)		-0.505 (0.408)	-0.372 (0.434)
Observations	17,342	17,342	17,342	17,342	17,342	17,342
Number of counties	1448	1448	1448	1448	1448	1448
R-squared	0.388	0.407	0.428	0.390	0.409	0.430
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
County-specific time trends	No	No	Yes	No	No	Yes


Notes: In columns (1)-(3), high clan density is a dummy variable indicating whether the number of genealogies in a county divided by the county's population size is above the national mean. In columns (4)-(6), we use the continuous measure. Famine period is defined 1958-1961. Grain output per capita, non-farming land ratio, and urbanization rate are measured in 1957; distance from Beijing and distance from provincial capital are spherical distances; crop suitability index for rice is from GAEZ (Global Agro-Ecological Zones); historical migrants denote the inflow of prominent families in the late Song Dynasty; the share of ethnic minorities and averaged years of schools are calculated based on people older than 40 years old in 1982 using the 1982 population census data. See Appendix B for details on data source. Standard errors clustered at the county level are reported in parentheses. ***p < 0.01, **p < 0.05, *p < 0.1. For easier comparisons of the coefficients, all regressions

- Believe $\Delta Y_i(g, z) \perp\!\!\!\perp G_i \mid X_i$ is more plausible than $\Delta Y_i(g, z) \perp\!\!\!\perp G_i$
- Simple improvements to the TWFE regression: (a) demean X_i ; (b) add interaction terms

$$Y_{it} = \alpha_i + \xi_t + \tau G_i \cdot \text{Post}_t + \beta^T X_i \cdot \text{Post}_t + \gamma^T G_i \cdot X_i \cdot \text{Post}_t + \epsilon_{it}$$

- $\hat{\tau}$ identifies $\mathbb{E}[\tau_{\text{em}}(x)]$, or τ_{cm} under linearity
- Requires overlap: $0 < \mathbb{P}(G_i = 1 \mid X_i = x) < 1$
- Can leverage more flexible models of ΔY_i on X_i for subgroups $G_i = 1, 0$:
 - Transform data into wide form; replace Y_i with ΔY_i
 - Apply estimators developed for selection-on-observables designs (stratification, matching, balancing, IPW, AIPW, outcome modeling, double machine learning, ...)


Extension to Multiple Pre- and Post-Periods



- Lesson 1: pretrend tests can help assess parallel trends, but **not** $\Delta Y(g, z) \perp\!\!\!\perp G$
- Lesson 2: using post-periods as non-event periods requires an additional “no carryover effect” assumption


Example: Clans and Calamity

- Event: the Great Famine (1958–1961)
- G: social capital, proxied by genealogy books
 - No genealogies: 412 counties
 - Have pre-PRC genealogies: 509 counties
- Outcome: county mortality rate




Source: Cao, Xu, and Zhang (2022); replication data in the `fdid` R package.

Raw Data: Group Means




Conditional on Covariates

- Pre-event covariates:
 - Log population size
 - Per capita grain production
 - Ratio of non-farming land
 - Urbanization ratio
 - Distance from Beijing
 - Distance from the provincial capital
 - Share of ethnic minorities
 - Rice suitability
 - Average years of education




- Estimates are stable across estimators, from TWFE to AIPW

Dynamic Estimates: DID




Dynamic Estimates: TWFE with Covariates


TWFE⁺: covariates $\times$ Post



TWFE*: additional interaction terms




Dynamic Estimates: AIPW





- No visible pretrend; the mortality gap opens exactly in the famine years

Multi-valued G



- $G = 0$: 412 counties (no genealogies)
- $G = 1$: 254 counties (some genealogies)
- $G = 2$: 255 counties (many genealogies)

Multi-valued G: AIPW Estimates



- The mortality gap concentrates where social capital is strongest

How to Correctly Interpret Your Findings?

Estimand	Interpretation in the running example	Key identifying assumptions
τ_{em} , effect modification	How did the average effect of exposure to famine on mortality differ between counties with high and low social capital?	No anticipation; parallel trends
τ_{att} , effect of exposure for $G = 1$	What was the average effect of exposure to famine on mortality in counties with high social capital?	+ exclusion restriction on Z in group $G = 0$
τ_{cm} , causal moderation	To what extent did high (vs. low) social capital mitigate the impact of famine on mortality?	No anticipation; parallel trends; $\Delta Y(g, z) \perp\!\!\!\perp G \mid X$
$\tau_{G Z=1}$, G 's conditional effect	What was the average causal effect of social capital on mortality during the famine years?	+ exclusion restriction on G absent the event

```
library(fdid)
mortality$any_gen <- as.integer(
  mortality$pczupu > 0) # 412 vs 509 counties
s <- fdid_prepare(
  data = mortality,
  Y_label = "mortality",
  G_label = "any_gen",
  unit_label = "countyid",
  time_label = "year")

fit <- fdid(s, tr_period = 1958:1961,
  ref_period = 1957,
  method = "ols1")
```

Observed interaction

The famine $\times$ social-capital coefficient is an effect-modification contrast.

Moderation interpretation

Causal moderation requires factorial parallel trends.

- What is your estimand?! (Lundberg, Johnson, and Stewart 2021)
- Assess no anticipation and parallel trends (e.g., by checking pretrends)
- If the goal is causal moderation, justify $\Delta Y(g, z) \perp\!\!\!\perp G \mid X$ and show robustness of findings (e.g., by conducting sensitivity analysis)
- If using TWFE models, demean covariates and add interaction terms $G_i \cdot X_i \cdot \text{Post}_t$; check overlap
- Should not automatically assume no carryover effects

- When an event hits every unit, there is no untreated group; FDID compares before-after changes across levels of a **baseline factor**
- The arithmetic is canonical DID; the estimand is not: under no anticipation and parallel trends, τ_{DID} is **effect modification** — an associative quantity
- **Factorial parallel trends** ($\Delta Y(g, z) \perp\!\!\!\perp G$) upgrades it to **causal moderation**; an exclusion restriction on G upgrades it to G 's conditional effect; an exclusion restriction on Z recovers canonical DID
- The broader lesson of these lectures: **estimators are not research designs** — the same arithmetic answers different questions under different assumptions
- Lecture 4: back to TWFE regressions, and what breaks under staggered timing

- `fdid` (R): factorial DID with the famine application data
- `panelView` (R & Stata): treatment status and outcome visualization
- Tutorial: <https://yiqingxu.org/packages/fdid/>

References

- Bansak, Kirk. 2021. "Estimating Causal Moderation Effects with Randomized Treatments and Non-Randomized Moderators." *Journal of the Royal Statistical Society: Series A* 184 (1): 65–86.
- Cao, Jiarui, Yiqing Xu, and Chuanchuan Zhang. 2022. "Clans and Calamity: How Social Capital Saved Lives during China's Great Famine." *Journal of Development Economics* 157:102865.
- Chen, Joy, Erik H. Wang, and Xiaoming Zhang. 2025. "From Powerholders to Stakeholders: State-Building with Elite Compensation in Early Medieval China." *American Journal of Political Science* 69 (2): 607–23.
- de Kadt, Daniel, and Horacio A. Larreguy. 2018. "Agents of the Regime? Traditional Leaders and Electoral Politics in South Africa." *Journal of Politics* 80 (2): 382–99.
- Fouka, Vasiliki. 2019. "How Do Immigrants Respond to Discrimination? The Case of Germans in the US during World War I." *American Political Science Review* 113 (2): 405–22.
- Lundberg, Ian, Rebecca Johnson, and Brandon M. Stewart. 2021. "What Is Your Estimand? Defining the Target Quantity Connects Statistical Evidence to Theory." *American Sociological Review* 86 (3): 532–65.
- Squicciarini, Mara P. 2020. "Devotion and Development: Religiosity, Education, and Economic Progress in Nineteenth-Century France." *American Economic Review* 110 (11): 3454–91.
- VanderWeele, Tyler J. 2009. "On the Distinction between Interaction and Effect Modification." *Epidemiology* 20 (6): 863–71.
- Xu, Yiqing, Anqi Zhao, and Peng Ding. 2026. "Factorial Difference-in-Differences." *Journal of the American Statistical Association* 121 (554): 963–75.