


Panel Methods (6): Synthetic Control and Extensions

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).




- What can we do when the parallel trends assumption is violated?
- Solution 1: conditional on pre-treatment covariates → semi-parametric DID
- Solution 2: conditional on pre-treatment covariates & outcomes → synthetic control method
- This lecture
 - The synthetic control method
 - The latent factor approach
 - Doubly robust methods

Synthetic Control Grew Out of Applications




Public policy

California tobacco control



Political economy

Election-day registration



Development

Land titling in Benin

Source: Abadie, Diamond, and Hainmueller (2010); Xu (2017); Sanford (2019).

The Literature in One Map

Family	What changes	Representative work
Classic SCM	Convex donor weights	Abadie, Diamond, Hainmueller
Outcome models	Learn low rank structure	Xu; Athey et al.
Augmentation	Correct imperfect pre-fit	Ben-Michael, Feller, Rothstein
Synthetic DID	Weight units and time	Arkhangelsky et al.

Common estimand


Use outcomes from untreated units to estimate $Y_{1t}(0)$ after treatment.

Source: Arkhangelsky and Imbens (2024), survey of causal models for longitudinal and panel data.

The Synthetic Control Method (SCM)


SCM: Basic Idea

- $J + 1$ units in periods $1, 2, \dots, T$; one treated “1”, J controls
- Region “1” is exposed to the intervention after period T_0
- We aim to estimate the effect of the intervention on Region “1”




- [Athey and Imbens \(2017\)](#): “[a]rguably the most important innovation in the policy evaluation literature in the last 15 years.”
- A combo of many innovations
 - Take advantage of pre-treatment outcomes
 - Use cross-sectional instead of temporal correlations in data
 - Construct a convex combination of donors to construct a counterfactual
 - Reserve some pre-treatment periods for testing

DID vs. Synthetic Control




DAG Compatible with Additive Fixed Effects



- Two unobservables drive both D_t and Y_t : a fixed unit trait α and period shocks f_t (here, the time fixed effects)
- α shifts a unit by the same amount in every period; f_t shifts every unit by the same amount $\rightarrow$ [parallel trends](#); two-way demeaning removes both
- No arrow from past outcomes to later treatment: treatment is set in advance (strict exogeneity)

Source: Xu (2023), Figure 2. Gray: unobserved; dashed arrows: links involving unobservables.

What Synthetic Control Adds: The Same DAG, Augmented by Factors



- Now the unobservable is $U_t = \lambda' f_t$: the same fixed trait λ , but its effect moves with common factors f_t
- Units with different λ respond differently to the same shocks $\rightarrow$ trends need not be parallel; DID fails here
- λ is fixed and f_t is shared, so pre-treatment outcomes reveal them: **match on them** (synthetic control) or **estimate them** (factor models); TWFE is the special case $U_t = \alpha + f_t$
- Still no feedback from past outcomes to treatment: the same regime as DID; only how U_t enters differs

Source: Xu (2023), Figure 7.

$$Y_{it} = \tau_{it}D_{it} + \theta'_t Z_i + \xi_t + \lambda'_i f_t + \varepsilon_{it}$$

or

$$\begin{cases} Y_{it}(0) &= \theta'_t Z_i + \xi_t + \lambda'_i f_t + \varepsilon_{it} \\ Y_{it}(1) &= Y_{it}(0) + \tau_{it} \end{cases}$$

- Suppose there are R time-varying signals f_t out there
- Each unit (e.g. country, participant) picks up a fixed linear combination of these signals based on factor loadings λ_i
- Since these “confounders” are evidenced in the pre-treatment outcomes for both treated and controls, we can try to use this information to “balance on” these confounders
- We will discuss the model-based approach later

$$\begin{cases} Y_{it}(0) &= \theta'_t Z_i + \xi_t + \lambda'_i f_t + \varepsilon_{it} \\ Y_{it}(1) &= Y_{it}(0) + \tau_{it} \end{cases}$$

- Let $W = (w_2, \dots, w_{J+1})'$ with $w_j \geq 0$ and $w_2 + \dots + w_{J+1} = 1$.
- Let $\bar{Y}_i^{K_1}, \dots, \bar{Y}_i^{K_M}$ be $M > R$ linear functions of pre-intervention outcomes
- Suppose that we can choose W^* such that:

$$Z_1 = \sum_{j=2}^{J+1} w_j^* Z_j, \quad \bar{Y}_1^k = \sum_{j=2}^{J+1} w_j^* \bar{Y}_j^k, \quad k \in \{K_1, \dots, K_M\}$$


- When T_0 is large, an approximately unbiased estimator of τ_{1t} is:

$$\hat{\tau}_{1t} = Y_{1t} - \sum_{j=2}^{J+1} w_j^* Y_{jt}, \quad t \in \{T_0 + 1, \dots, T\}$$

- Let $X_1 = (Z_1, \bar{Y}_1^{K_1}, \dots, \bar{Y}_1^{K_M})'$ be a $(k \times 1)$ vector of pre-intervention characteristics for the treated and X_0 , a $(k \times J)$ matrix, for the controls.
- The vector W^* is chosen to minimize $\|X_1 - X_0 W\|$, subject to our weight constraints.
 - We consider $\|X_1 - X_0 W\|_V = \sqrt{(X_1 - X_0 W)' V (X_1 - X_0 W)}$, where V is some $(k \times k)$ symmetric and positive semidefinite matrix.
 - Various ways to choose V (subjective assessment of predictive power of X , regression, minimize MSPE, cross-validation, etc.).

Example: Proposition 99 on Cigarette Consumption


- In 1988, California first passed comprehensive tobacco control legislation (cigarette tax, media campaign etc.)
- Using 38 states that had never passed such programs as controls



Cigarette Consumption: CA and the Rest of the U.S.


Example: Proposition 99 on Cigarette Consumption

- In 1988, California first passed comprehensive tobacco control legislation (cigarette tax, media campaign etc.)
- Using 38 states that had never passed such programs as controls




Cigarette Consumption: CA and Synthetic CA

Proposition 99: Gap, Placebos, and Inference




1. The gap

California minus synthetic
California



2. Placebos

Same gap computed for all 38
donor states



3. Inference

Ratio of post- to pre-treatment
MSPE

Source: Abadie, Diamond, and Hainmueller (2010).

1. Fit

Can donors reproduce the treated unit before treatment?

2. Gap

How does the treated-synthetic difference evolve after treatment?

3. Placebo

Would the method find similar gaps for untreated units?

4. Stress test

Do results survive donor and specification choices?

For every application, show pre-treatment fit, the post-treatment gap, and the placebo distribution.

Source: Cunningham, *Causal Inference: The Mixtape*, synthetic control chapter.

- Algorithmic
 - Deal with one treated unit at a time
 - Deal with one outcome at a time
 - Slow to implement and sometimes difficult to find a solution
 - Allow too much user discretion, e.g. cherry-picking $\bar{Y}_i^k$ results in over-rejection ([Ferman et al. 2020](#))
- Inference
 - Permutation inference and sensitivity analysis (e.g. [Hahn and Shi 2017](#); [Firpo and Possebom 2018](#); [Chernozhukov et al. 2021](#))
 - Inflated precision with nonstationary data ([Cattaneo et al. 2021](#))
- Identification
 - My opinion: intrinsically, a method based on strict exogeneity (fixed timing)

The Latent Factor Approach

- Recall that ADH (2010) use a factor-augmented model to motivate the synthetic control method:

$$Y_{it}(0) = \theta'_t Z_i + \xi_t + \lambda'_i f_t + \varepsilon_{it}$$


- What if we actually estimate the model using observations under the control condition only?
- Xu (2017) imports the so-called interactive fixed effects (IFE) model into a DID setting

$$Y_{it}(0) = X'_{it}\beta + \alpha_i + \xi_t + \lambda'_i f_t + \varepsilon_{it}$$

- Athey et al. (2021) extend it and introduce the matrix completion method
- Liu, Wang, and Xu (2024) put these methods in a general framework — “the counterfactual estimators”
- No negative weighting!
- Limitations: needs large T and N ; model dependency

Basic Idea

- In a panel setting, treat the treated observations' $Y(0)$ as **missing data**
- Predict $Y(0)$ based on an **outcome model**
(Use pre-treatment data for model selection)
- Estimate ATT by averaging differences between $Y(1)$ and $\hat{Y}(0)$



A model-based counterfactual estimator proceeds in the following steps:


- Step 1. Train the model using observations under the control condition ($D_{it} = 0$).
- Step 2. Predict the counterfactual outcome $\hat{Y}_{it}(0)$ for each observation under the treatment condition ($D_{it} = 1$) and obtain an estimate of the individual treatment effect: $\hat{\tau}_{it} = Y_{it} - \hat{Y}_{it}(0)$.
- Step 3. Generate estimates for the causal quantities of interest

$$ATT = \mathbb{E}[\tau_{it} | D_{it} = 1, \forall i \in \mathcal{I}, \forall t], \quad \text{or}$$

$$ATT_s = \mathbb{E}[\tau_{it} | D_{i,t-s} = 0, \underbrace{D_{i,t-s+1} = D_{i,t-s+2} = \dots = D_{it} = 1}_{s \text{ periods}}, \forall i \in \mathcal{I}].$$


Indexing: s counts periods under treatment, so $s = 1$ is the first treated period and $s = 0$ the last untreated one (the **fect** convention). Lecture 4's event-study specification indexed the same two periods as $s = 0$ and $s = -1$.

Relative Time: Re-Indexing the Panel




- Each treated observation is re-indexed by its distance to the adoption date, with $s = 1$ the first treated period and $s = 0$ the last untreated one; $\widehat{ATT}_s$ averages within a diagonal

Reading a Dynamic Treatment Effects Plot



- Pre-treatment averages are close to zero; post-treatment averages step up: no visible pre-trend
- The gray bars at the bottom show how many observations each average uses; distant leads and lags rest on few observations

The Warning Sign: A Pre-Trend



- Pre-treatment averages drift away from zero: the comparison units were already on a different path, and the post-treatment estimates inherit that drift

We review three estimation strategies:

- FEct:

$$\hat{Y}_{it}(0) = X_{it}\hat{\beta} + \hat{\alpha}_i + \hat{\xi}_t$$

- IFect (Gobillon and Magnac 2016; Xu 2017):

$$\hat{Y}_{it}(0) = X_{it}\hat{\beta} + \hat{\lambda}'_i \hat{f}_t$$

- Matrix Completion (MC) (Athey et al. 2021):

$$\hat{Y}_{it}(0) = X_{it}\hat{\beta} + \hat{L}_{it},$$

where matrix $\{L_{it}\}_{N \times T}$ is a lower-rank matrix approximation of $\{Y(0)\}_{N \times T}$ with missing values


Remarks:

- DID is a special case of FEct
- Both IFect and MC are estimated via iterative algorithms
- Cross-validation to choose the tuning parameter

Xu (2017) proposes a three-step approach based on a latent factor model:


$$\begin{aligned}
 \text{Control} \quad Y_{it}(0) &= X'_{it}\beta + \alpha_i + \xi_t + \lambda'_i f_t + \varepsilon_{it} \\
 \text{Treated} \quad Y_{it}(0) &= X'_{it}\beta + \alpha_i + \xi_t + \lambda'_i f_t + \varepsilon_{it} && (\text{pre}) \\
 Y_{it}(1) &= X'_{it}\beta + \alpha_i + \xi_t + \lambda'_i f_t + \varepsilon_{it} + \tau_{it} && (\text{post})
 \end{aligned}$$

- 1 Estimate using controls
- 2 Predict on treated $\hat{Y}_{it}(0)$
- 3 Calculate the ATT
- 4 Expectation-Maximization
(Gobillon and Magnac 2016)





Election Day Registration (EDR) and Voter Turnout



Causal inference is a missing data problem.





Main Results



Factors and Factor Loadings




Placebo Test




Example: Property Rights and Land Improvement

Sanford (2019)


- Do property rights lead to improved land quality?
- “Experiment”: giving peasants in Borgou, Benin land titles
- Use satellite (remote sensing) data to measure land improvement, i.e., switch from annual crops to perennial crops (bushes and trees)
- Use IFeCt to construct counterfactuals




Original Satellite Images




Treated and Counterfactual Averages





Matrix Completion Methods



- Recall that our main goal is to predict treated counterfactuals
- Taking advantage of the matrix structure, matrix completion methods use non-treated data to achieve this goal
- The basic idea to find a lower-rank representation of the matrix to impute the “missing data”
- Xu (2017) is a special case of this approach

- Non-parametric block bootstrap
 - sample with replacement across units
 - valid when N is large, $\frac{N_{tr}}{N}$ is fixed
- A permutation-based test for Sharp Nulls ([Chernozhukov et al. 2021](#))
 - e.g. $Y_{it}(1) = Y_{it}(0), \forall i \in \mathcal{T}, t > T_{0i}$
 - randomization over time (by blocks) instead of across units
 - valid if T is large, errors are stationary, weakly dependent, and estimators are consistent or stable
 - exact if errors are i.i.d.
- Liu, Wang, and Xu (2024) propose a set of diagnostic tests

Doubly and Triply Robust Methods

Doubly and Triply Robust Methods

- So far, we've surveyed two groups of methods: (1) those constructing balancing weights; (2) those modeling the conditional outcomes
- Combining the two approaches will likely produce doubly robust estimators
- Some methods we discussed, including semi-parametric DID and panel matching, are already doing a simple version of it (balancing plus regression)
- We review three methods that formally adopt this idea
 - Augmented synthetic control ([Ben-Michael et al. 2021](#)): modeling first
 - Synthetic DID ([Arkhangelsky et al. 2021](#)): weighting first
 - Triply robust panel estimator ([Athey, Imbens, Qu, and Viviano 2026](#)): unit weights, time weights, and the factor model of the previous section, tuned together


Augmented Synthetic Control (Ben-Michael et al. 2021)

- Assuming unit 1 being treated ($D_1 = 1; D_{-1} = 0$), pretreatment covariates X_i
- Basic Idea**
 - Run an outcome model (e.g. Ridge, FEct, IFEct, MC, etc.) and obtain model fit $\hat{m}(X_i)$
 - Balance on the residual averages, obtaining weights $\hat{\gamma}_i$ for the controls
 - Treated average is constructed using:

$$\begin{aligned}\hat{Y}_1^{aug}(0) &= \underbrace{\sum_{i \in \mathcal{C}} \hat{\gamma}_i Y_i}_{SCM} + \underbrace{\hat{m}(X_1) - \sum_{i \in \mathcal{C}} \hat{\gamma}_i \hat{m}(X_i)}_{debias} \\ &= \hat{m}(X_1) + \sum_{i \in \mathcal{C}} \hat{\gamma}_i (Y_i - \hat{m}(X_i))\end{aligned}$$

- The balancing weights take care of the remaining biases from the outcome model; the estimator is thus **doubly robust**
- Inference via jackknife

Revisiting California Prop 99




Source: Own estimates with the `augsynth` package on the Abadie, Diamond, and Hainmueller (2010) data; bands are jackknife+ intervals (`augsynth_prop99_2026.R`).

Synthetic DID (Arkhangelsky et al. 2021)

- Key insight: weight not only control units, but also pretreatment periods


Notation: the time weights λ_t in what follows are unrelated to the factor loadings λ_i of the previous section.



Source: Figures on the SDID frames are from the authors' slides; see Arkhangelsky et al. (2021).

Unit and Time Weights


- SC:
 - Using **pre-treatment data**, we learn an average of controls that's predictive of California.
 - Assuming this relationship remains valid post-treatment, we use the **same average of controls** to impute treatment-free observations for California
- Forecasting:
 - Using **controls**, we learn an average of periods forecasting what we see post-treatment
 - Assuming this relationship remains valid for the treated, we use the **same average of periods** to impute treatment-free observations for California



Source: The authors' slides; see Arkhangelsky et al. (2021).

Unit and Time Weights

- Synthetic Diff-in-Diff (SDID)
 - Do both synthetic control and forecasting and combine via diff-in-diff
 - Only one of these relationships has to remain valid (double robustness)
 - Our synthetic control can be parallel to California



1. Estimate unit weights $\hat{\omega}$ from pre-treatment data $\Rightarrow$ a synthetic control unit: $\hat{\omega}_0 + \hat{\omega}^T Y_{co,pre} \approx Y_{\overline{tr},pre}$.
2. Estimate time weights $\hat{\lambda}$ from control data $\Rightarrow$ a synthetic pre-treatment period: $\hat{\lambda}_0 + Y_{co,pre} \hat{\lambda} \approx Y_{co,\overline{post}}$.
3. Apply diff-in-diff to the resulting synthetic 2×2 panel by minimizing:

$$\sum_{i=1}^N \sum_{t=1}^T (Y_{it} - \mu - \alpha_i - \beta_t - X_{it}\gamma - D_{it}\tau)^2 \hat{\omega}_i \hat{\lambda}_t$$

	Synthetic Pre-Treatment	Average Post-treatment
Synthetic Control	$\hat{\omega}^T Y_{co,pre} \hat{\lambda}$	$\hat{\omega}^T Y_{co,\overline{post}}$
Average Treated	$Y_{\overline{tr},pre} \hat{\lambda}$	$Y_{\overline{tr},\overline{post}}$

* Inference via jackknife

Source: The authors' slides; see Arkhangelsky et al. (2021).

Triply Robust Panel Estimator (Athey, Imbens, Qu, and Viviano 2026)

- SDID weights units and periods around a two-way fixed effects model. TROP adds the third ingredient from the previous section, a **low-rank factor component**, and lets the data tune all three
- Working model for the untreated outcome: $Y_{it}(0) = \alpha_i + \beta_t + L_{it} + \varepsilon_{it}$, with L low rank
- Three components, one tuning parameter each:
 - ➊ **Regression adjustment**: estimate L by nuclear-norm penalized least squares (penalty λ_{nn} ; $\lambda_{nn} = \infty$ switches the factor component off)
 - ➋ **Unit weights** $\omega_j = \exp(-\lambda_{unit} d_{unit}(j, i))$: decay with the mean squared outcome gap between control j and the treated unit over their shared untreated periods
 - ➌ **Time weights** $\theta_s = \exp(-\lambda_{time} |s - t|)$: decay with the distance to the treated period, so the recent past counts most
- Fit the weighted, penalized model on the untreated cells, impute $\hat{Y}_{it}(0) = \hat{\alpha}_i + \hat{\beta}_t + \hat{L}_{it}$, and take $\hat{\tau}_{it} = Y_{it} - \hat{Y}_{it}(0)$; the weights are specific to the treated cell, so any assignment pattern works

TROP: Tuning by Leave-One-Out, and Why “Triply” Robust

- Choosing $(\lambda_{unit}, \lambda_{time}, \lambda_{nn})$: pretend each untreated cell is treated, impute it from the other untreated cells, and pick the triple whose pseudo-effects are closest to zero,

$$Q(\lambda) = \sum_{i,t} (1 - W_{it}) \hat{\tau}_{it}(\lambda)^2,$$

the placebo idea from the synthetic control section, turned into a loss function

- The grid nests the menu: flat weights and no factor is DID/TWFE; flat weights with a factor is matrix completion; no factor with particular unit and time weights is synthetic control or SDID
- Triple robustness**: with fixed weights the bias is bounded by a **product**,

$$|\mathbb{E}[\hat{\tau} - \tau \mid L]| \leq \underbrace{\|\Delta_u(\omega)\|}_{\text{unit imbalance}} \times \underbrace{\|\Delta_t(\theta)\|}_{\text{time imbalance}} \times \underbrace{\|B\|}_{\text{bias of } \hat{L}}$$

so it vanishes if the weighted controls match the treated unit's loadings, or the weighted pre-periods match the treated period's factors, or the factor fit is unbiased

- ASCM and SDID need one of two things to hold; TROP needs one of three

TROP against the Alternatives: RMSE Relative to the Best Estimator

	TROP	SDID	SC	DID	MC	DIFP		TROP	SDID	SC	DID	MC	DIFP
CPS log wage, min. wage	1.00	1.14	1.44	1.91	1.26	1.22	Germany GDP, treated unit	1.00	1.13	1.82	5.17	1.27	1.89
CPS unemployment, min. wage	1.00	1.05	1.11	1.89	1.10	1.09	Basque GDP, random	1.00	1.02	4.55	9.11	1.70	2.47
CPS hours, min. wage	1.00	1.19	1.22	1.25	1.11	1.20	Basque GDP, simulated	1.00	1.55	2.11	3.05	1.30	1.66
CPS log wage, gun law	1.00	1.15	1.13	2.04	1.33	1.32	Basque GDP, treated unit	1.00	1.35	2.79	1.25	3.02	2.94
CPS log wage, abortion	1.00	1.14	1.43	1.99	1.25	1.24	Prop 99 packs, random	1.00	1.22	1.48	2.16	1.14	1.45
CPS log wage, random	1.00	1.12	1.10	1.95	1.24	1.17	Prop 99 packs, simulated	1.00	1.28	1.93	2.88	1.44	1.41
PWT log GDP, democracy	1.00	1.44	1.59	7.85	1.76	1.54	Prop 99 packs, treated unit	1.00	1.04	1.72	2.18	1.27	1.30
PWT log GDP, education	1.00	1.51	2.10	6.90	1.60	1.57	Boatlift log wage, random	1.00	1.34	1.62	1.35	1.04	1.62
PWT log GDP, random	1.00	1.32	1.51	4.31	1.41	1.50	Boatlift log wage, simulated	1.00	1.50	1.31	1.46	1.12	1.58
Germany GDP, random	1.00	1.46	2.82	3.58	1.56	2.46	Boatlift log wage, treated unit	1.24	1.93	1.00	1.35	1.17	2.07
Germany GDP, simulated	1.00	1.79	5.82	7.58	1.93	4.72							



Out-of-sample RMSE divided by the smallest RMSE in the row; 21 designs calibrated to real panels, 10 treated units and 10 post periods.

DIFP: synthetic control with an intercept. Source: Athey, Imbens, Qu, and Viviano (2026), Table 1.

TROP: What the Calibrated Simulations Show

- Semi-synthetic designs: keep a real panel's factor structure and serial correlation, assign treatment as in the application, add a known effect. Panels: CPS wages and unemployment (minimum wage), Penn World Table (democracy), German reunification, the Basque Country, Prop 99, the Mariel boatlift; 21 designs with 10 treated units and 10 post periods
- TROP has the lowest RMSE in 20 of the 21. Relative to the best estimator in each design, DID's RMSE reaches 9 times the best (Basque), synthetic control 5.8 times, matrix completion 3 times, SDID 1.9 times
- Which ingredient matters: dropping the factor adjustment or the time weights raises RMSE by up to 85–91%; dropping the unit weights costs at most 10%; dropping time weights and the factor together can multiply it by five. **The recent past matters most**
- Their advice: simulate from your own panel and see which estimator wins there; failing that, prefer SDID or TROP to TWFE/DID

Conclusions

Concluding Remarks

- SCM makes the counterfactual construction visible: donor pool, weights, pre-fit, and post-treatment gap.
- All of these methods learn from the pre-treatment period. The learned relation between treated and donor outcomes must remain useful after treatment.
- Choose the method after examining the treatment pattern, the number of treated units, and the length of the pre-treatment period.
- Defend the donor pool, then report placebo and sensitivity checks.

Practical Recommendations

- Plot the raw paths and treatment timing before fitting a model.
- Defend the donor pool using subject knowledge; do not let software choose the comparison set by default.
- Report pre-fit, gaps, donor weights, and placebo results together.
- After fixing the research design, use alternative estimators to check sensitivity.

- Abadie (2021): synthetic controls in comparative case studies.
- Ben-Michael, Feller, and Rothstein (2021): augmented synthetic control.
- Arkhangelsky et al. (2021): synthetic difference-in-differences.
- Athey, Imbens, Qu, and Viviano (2026): the triply robust panel estimator, with simulations calibrated to real panels.
- Arkhangelsky and Imbens (2024): DID, fixed effects, factor models, and synthetic controls in a common framework.
- Cunningham (2021): worked Texas prison example with R code.

References

- Abadie, Alberto. 2021. "Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects." *Journal of Economic Literature* 59 (2): 391–425.
- Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2010. "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program." *Journal of the American Statistical Association* 105 (490): 493–505.
- Arkhangelsky, Dmitry, Susan Athey, David A. Hirshberg, Guido W. Imbens, and Stefan Wager. 2021. "Synthetic Difference-in-Differences." *American Economic Review* 111 (12): 4088–118.
- Arkhangelsky, Dmitry, and Guido W. Imbens. 2024. "Causal Models for Longitudinal and Panel Data: A Survey." *The Econometrics Journal* 27 (3): C1–C61.
- Athey, Susan, Mohsen Bayati, Nikolay Doudchenko, Guido Imbens, and Khashayar Khosravi. 2021. "Matrix Completion Methods for Causal Panel Data Models." *Journal of the American Statistical Association* 116 (536): 1716–30.
- Athey, Susan, and Guido Imbens. 2017. "The State of Applied Econometrics: Causality and Policy Evaluation." *Journal of Economic Perspectives* 31 (2): 3–32.
- Athey, Susan, Guido Imbens, Zhaonan Qu, and Davide Viviano. 2026. "Triply Robust Panel Estimators." Forthcoming, *Journal of Applied Econometrics*.
- Ben-Michael, Eli, Avi Feller, and Jesse Rothstein. 2021. "The Augmented Synthetic Control Method." *Journal of the American Statistical Association* 116 (536): 1789–803.
- Cattaneo, Matias D., Yingjie Feng, and Rocio Titiunik. 2021. "Prediction Intervals for Synthetic Control Methods." *Journal of the American Statistical Association* 116 (536): 1865–80.
- Chernozhukov, Victor, Kaspar Wüthrich, and Yinchu Zhu. 2021. "An Exact and Robust Conformal Inference Method for Counterfactual and Synthetic Controls." *Journal of the American Statistical Association* 116 (536): 1849–64.
- Cunningham, Scott. 2021. *Causal Inference: The Mixtape*. Yale University Press.

References (cont.)

- Ferman, Bruno, Cristine Pinto, and Vitor Possebom. 2020. "Cherry Picking with Synthetic Controls." *Journal of Policy Analysis and Management* 39 (2): 510–32.
- Firpo, Sergio, and Vitor Possebom. 2018. "Synthetic Control Method: Inference, Sensitivity Analysis and Confidence Sets." *Journal of Causal Inference* 6 (2).
- Gobillon, Laurent, and Thierry Magnac. 2016. "Regional Policy Evaluation: Interactive Fixed Effects and Synthetic Controls." *Review of Economics and Statistics* 98 (3): 535–51.
- Hahn, Jinyong, and Ruoyao Shi. 2017. "Synthetic Control and Inference." *Econometrics* 5 (4): 52.
- Liu, Licheng, Ye Wang, and Yiqing Xu. 2024. "A Practical Guide to Counterfactual Estimators for Causal Inference with Time-Series Cross-Sectional Data." *American Journal of Political Science* 68 (1): 160–76.
- Sanford, Luke. 2019. "Geospatial Synthetic Controls for Agricultural Impact Evaluation." Working paper.
- Xu, Yiqing. 2017. "Generalized Synthetic Control Method: Causal Inference with Interactive Fixed Effects Models." *Political Analysis* 25 (1): 57–76.
- Xu, Yiqing. 2023. "Causal Inference with Time-Series Cross-Sectional Data: A Reflection." In *The Oxford Handbook of Engaged Methodological Pluralism in Political Science*, 353–74. Oxford University Press.