Impulse Response Heterogeneity

Òscar Jordà and Karel Mertens

AEA Continuing education 2023

Last updated: December 22, 2022

See also:
https://sites.google.com/site/oscarjorda/home/local-projections

Warmup observations

Where do the response dynamics come from?

suppose:

$$\begin{cases} \Delta y_t &= \beta \Delta s_t + \rho \Delta y_{t-1} + u_t^y \\ \Delta s_t &= \qquad\qquad \theta \Delta s_{t-1} + u_t^s \end{cases} ; \ \boldsymbol{u}_t \sim D(\boldsymbol{0}, I)$$

easy to see that:

$$\mathcal{R}_{ys}(h) = \underbrace{\beta\theta^h + \beta\theta^{h-1} + \ldots + \beta\theta\rho^{h-1}}_{\text{due to policy persistence}} + \underbrace{\beta\rho^h}_{\mathcal{R}^*_{ys}(h)}$$

$$\mathcal{R}_{ss}(h) = \theta^h$$

setting $\theta = 0$:

$\mathcal{R}^*_{ys}(h) = \beta\rho^h$ internal propagation

Actual response versus response conditional on future treatments

Simulated data from `counterfactual.do`

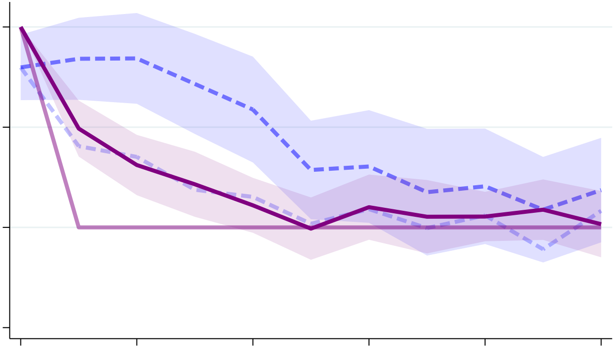

$\mathcal{R}_{ys}$: dashed blue line; $\mathcal{R}_{ss}$: purple. Lighter colors are $\mathcal{R}^*_{ys}$ and $\mathcal{R}^*_{ss}$

Back to the example

Suppose that instead of $\mathcal{R}_{ss}(h) = \theta^h$, you feed counterfactual $\mathcal{R}^c_{ss}(h)$
Can show that:

$$\begin{aligned}
\mathcal{R}^c_{ys}(0) &= \mathcal{R}^*_{ys}(0)\mathcal{R}^c_{ss}(0) \\
\mathcal{R}^c_{ys}(1) &= \mathcal{R}^*_{ys}(0)\mathcal{R}^c_{ss}(1) + \mathcal{R}^*_{ys}(1)\mathcal{R}^c_{ss}(0) \\
\mathcal{R}^c_{ys}(2) &= \mathcal{R}^*_{ys}(0)\mathcal{R}^c_{ss}(2) + \mathcal{R}^*_{ys}(1)\mathcal{R}^c_{ss}(1) + \mathcal{R}^*_{ys}(2)\mathcal{R}^c_{ss}(0) \\
\mathcal{R}^c_{ys}(3) &= \mathcal{R}^*_{ys}(0)\mathcal{R}^c_{ss}(3) + \mathcal{R}^*_{ys}(1)\mathcal{R}^c_{ss}(2) + \mathcal{R}^*_{ys}(2)\mathcal{R}^c_{ss}(1) + \mathcal{R}^*_{ys}(3)\mathcal{R}^c_{ss}(0) \\
\vdots &= \vdots
\end{aligned}$$

Example: `couterfactual.do`

Illustration: recovering the original response

Using original treatment path

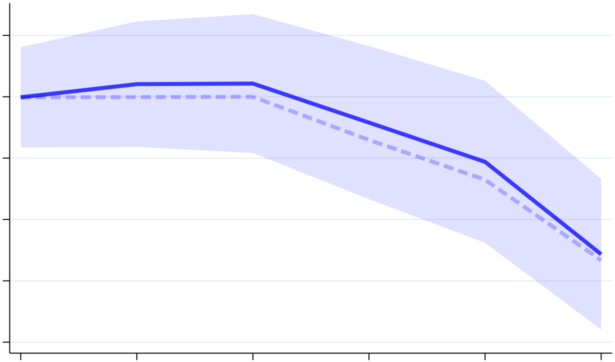

$\mathcal{R}_{ys}$: solid blue line; $\mathcal{R}^*_{ys}\mathcal{R}_{ss}$: dashed line

Counterfactuals in practice

Estimate usual LP but control for future treatments:

$$y_{t+h} - y_{t-1} = a_h + b_h \Delta s_t + \underbrace{\sum_{j=1}^{h} c_{jh} \Delta s_{t+j}}_{\text{new}} + d_h \Delta y_{t-1} + e_h \Delta s_{t-1} + v_{t+h}$$

Then b_h is an estimate of $\mathcal{R}^*_{ys}(h)$
$\mathcal{R}^c_{ss}(h)$ is supplied by user given particular counterfactual

Counterfactual response versus original response

Simulated data from `counterfactual.do`

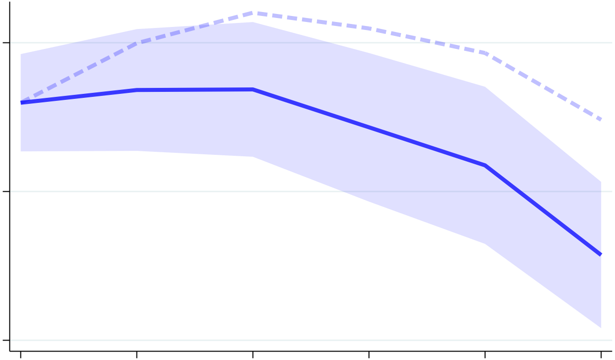

$\mathcal{R}_{ys}$: solid blue line; $\mathcal{R}^c_{ys} = \mathcal{R}_{ss} + 0.25$: dashed blue

Are these operations valid?

Mechanically: yes; causally: ?

- **Conditioning on future treatments:** are they randomly assigned?
- **Counterfactual treatment path:** how different from $\mathcal{R}_{ss}(h)$?
  Note:

$$\left(\mathcal{R}^*_{ss} - \hat{\mathcal{R}}_{ss}\right)' \Sigma_{ss}^{-1} \left(\mathcal{R}^*_{ss} - \hat{\mathcal{R}}_{ss}\right) \to \chi^2_H$$

- Worth reading: Viviano, Davide and Jelena Bradic. 2021.
  Dynamic covariate balancing: estimating treatment effects over time.

Impulse response heterogeneity: Kitagawa-Oaxaca-Blinder decompositions

Cloyne, Jordà, and Taylor (2020). Decomposing the fiscal multiplier

Back to index

Potential outcomes: A static setup first

Borrowing from applied micro

Think of observed y as coming from a latent mixture:

$$y = (1-s)\,y_0 + s\,y_1 = y_0 + s\,(y_1 - y_0); \quad s = 0,1$$

Assumption:

$$y_i \sim f(\mu_j; \sigma_j); \quad j = 0,1 \quad \text{unobservable random variables}$$

we would like: $E(y_1 - y_0)$ *average treatment effect*

Assume linear model for latent variables: $y_j;\ j = 0, 1$

let $y_j = \mu_j + v_j$, $E(v_j) = 0, j = 0, 1$. v_j captures heterogeneity

let $v_j = (\boldsymbol{x} - \boldsymbol{\mu}_x)\,\boldsymbol{\gamma}_j + \epsilon_j$ with $E(\epsilon_j) = 0$ and $E(\epsilon_j|\boldsymbol{x}) = 0$

then:

$$\underbrace{E_x[E(y_1|s=1;\boldsymbol{x}) - E(y_0|s=0;\boldsymbol{x})]}_{ATE} = [\mu_1 + E_x[E(\boldsymbol{x} - \boldsymbol{\mu}_x|s=1)]\,\boldsymbol{\gamma}_1]$$
$$- [\mu_0 + E_x[E(\boldsymbol{x} - \boldsymbol{\mu}_x|s=0)]\,\boldsymbol{\gamma}_0]$$

add/subtract counterfactual: $E_x[E(\boldsymbol{x} - \boldsymbol{\mu}_x|s=1)]\,\boldsymbol{\gamma}_0$

$$ATE = (\mu_1 - \mu_0) + E_x[E(\boldsymbol{x} - \boldsymbol{\mu}_x|s=1)]\,(\boldsymbol{\gamma}_1 - \boldsymbol{\gamma}_0) +$$
$$E_x\left[E(\boldsymbol{x} - \boldsymbol{\mu}_x|s=1) - E(\boldsymbol{x} - \boldsymbol{\mu}_x|s=0)\right]\,\boldsymbol{\gamma}_0$$

Kitagawa-Oaxaca-Blinder decomposition components

recall:

$$ATE = \underbrace{(\mu_1 - \mu_0)}_{direct} + \underbrace{E_x[E(\boldsymbol{x} - \boldsymbol{\mu}_x|s = 1)]\,(\boldsymbol{\gamma}_1 - \boldsymbol{\gamma}_0)}_{indirect} +$$
$$\underbrace{E_x\left[E(\boldsymbol{x} - \boldsymbol{\mu}_x|s = 1) - E(\boldsymbol{x} - \boldsymbol{\mu}_x|s = 0)\right]\,\boldsymbol{\gamma}_0}_{composition}$$

direct: ATE under random assignment

indirect: treatment spillovers on covariates

composition: failure of random assignment? small sample bias

interesting null hypotheses

linear case, still working through applied micro motivation

$$y_i = \mu_0 + (\mathbf{x}_i - \bar{\mathbf{x}}_0)\,\boldsymbol{\gamma}_0 + s_i\left[\beta + (\mathbf{x}_i - \bar{\mathbf{x}}_1)\,\boldsymbol{\theta}\right] + \omega_i$$

note: $\beta = \mu_1 - \mu_0$; $\boldsymbol{\theta} = \boldsymbol{\gamma}_1 - \boldsymbol{\gamma}_0$; and $\omega_i = \epsilon_{0,i} + s_i\left(\epsilon_{1,i} - \epsilon_{0,i}\right)$

hence:

- $H_0 : \beta = 0$ null of no *direct* treatment effect
- $H_0 : \boldsymbol{\theta} = 0$ null of no *indirect* effect
- $H_0 : E(\mathbf{x}|s = 1) - E(\mathbf{x}|s = 0) = 0$ null of no *composition* effect
- $H_0 : \boldsymbol{\gamma}_0 = 0$ null of *random assignment* (hence no composition effect possible)

What does this mean for local projections?

let $\boldsymbol{y}_t = (y_t, y_{t+1}, \ldots, y_{t+H})$ and $\boldsymbol{y}$ denote the associated r.v.

assume conditional mean independence
let $E(\boldsymbol{y}_s) = \mu_s$ for $s \in \{0, 1\}$, wlog $\boldsymbol{y}_s = \boldsymbol{\mu}_s + \boldsymbol{v}_s$
under linearity $\boldsymbol{v}_s = (\boldsymbol{x} - \boldsymbol{\mu}_x)\, \boldsymbol{B}_s + \boldsymbol{\epsilon}_s$, then:

$$E(\boldsymbol{y}_s|\boldsymbol{x}) = \boldsymbol{\mu}_s; \quad E(\boldsymbol{v}_s) = 0; \quad E(\boldsymbol{\epsilon}_s|\boldsymbol{x}) = 0; \quad s \in \{0, 1\}$$

note: Angrist et al. (2017) assume stronger conditional ignorability

hence:

$$y_{t+h} = \underbrace{\mu_0^h + (\boldsymbol{x}_t - \bar{\boldsymbol{x}})\boldsymbol{\gamma}_0^h + s_t\, \beta^h}_{\text{usual local projection}} + \underbrace{s_t\,(\boldsymbol{x}_t - \bar{\boldsymbol{x}})\boldsymbol{\theta}^h}_{\text{Kitagawa term}} + \omega_{t+h};$$

$$h = 0, 1, \ldots, H; t = h, ..., T.$$

Kitagawa decomposition components

recall:

$$y_{t+h} = \underbrace{\mu_0^h + (\boldsymbol{x}_t - \overline{\boldsymbol{x}})\boldsymbol{\gamma}_0^h + s_t\,\beta^h}_{\text{usual local projection}} + \underbrace{s_t\,(\boldsymbol{x}_t - \overline{\boldsymbol{x}})\boldsymbol{\theta}^h}_{\text{Kitagawa terms}} + \omega_{t+h};$$

$$h = 0, 1, \ldots, H; t = h, ..., T.$$

direct effect: $\hat{\mu}_1^h - \hat{\mu}_0^h = \hat{\beta}^h$

indirect effect: $(\overline{\boldsymbol{x}}_1 - \overline{\boldsymbol{x}})(\hat{\boldsymbol{\gamma}}_1^h - \hat{\boldsymbol{\gamma}}_0^h) = (\overline{\boldsymbol{x}}_1 - \overline{\boldsymbol{x}})\hat{\boldsymbol{\theta}}^h$

composition effect: $(\overline{\boldsymbol{x}}_1 - \overline{\boldsymbol{x}}_0)\hat{\boldsymbol{\gamma}}_0^h$

ergodicity: needed to ensure $\overline{\boldsymbol{x}} \to \boldsymbol{\mu}_x$

Implications: state-dependence

note: suppose $\boldsymbol{x} = \boldsymbol{x}^*$ then total response is:

$$
\begin{aligned}
E(\boldsymbol{y}_1|\boldsymbol{x}^*, s = \delta) - E(\boldsymbol{y}_0|\boldsymbol{x}^*, s = 0) \\
&= \delta\mu_1 + \delta[\boldsymbol{x}^* - E(\boldsymbol{x})]\boldsymbol{\gamma}_1 - \{\mu_0 + [\boldsymbol{x}^* - E(\boldsymbol{x})]\boldsymbol{\gamma}_0\} \\
&= \delta\beta + \delta[\boldsymbol{x}^* - E(\boldsymbol{x})]\boldsymbol{\theta},
\end{aligned}
$$

remarks:

- dependence on $\boldsymbol{x}^*$ is only partial equilibrium
- need identification (instruments) for $\boldsymbol{x}$
- usual single variable stratification omits other terms in $\boldsymbol{x} \rightarrow$ bias

Example from a previous experiment: two episodes

how effective was monetary policy in …

1. November 1987 (post-stock market crash)
   - stocks 23% lower by end of October
   - Fed lowered funds rate 50bps
2. February 1996 (middle of a long expansion)
   - middle of stable funds rate

idea: two different scenarios, but similar policy paths $\rightarrow$
differences not due to different policy

funds rate path nearly identical and to baseline

baseline is average over the sample

Federal funds rate

(a) November 1987

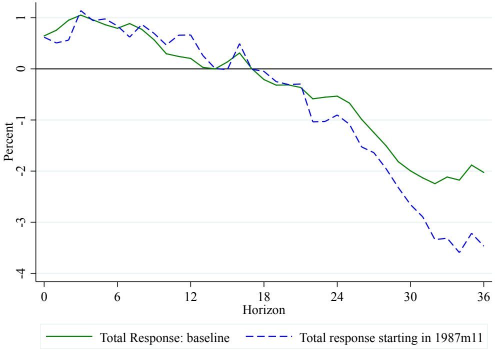

(b) February 1996

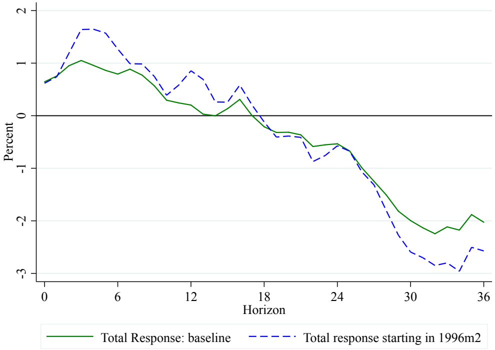

policy unable to boost activity post-1987crash

policy as usual February 1996

Industrial production

(a) November 1987

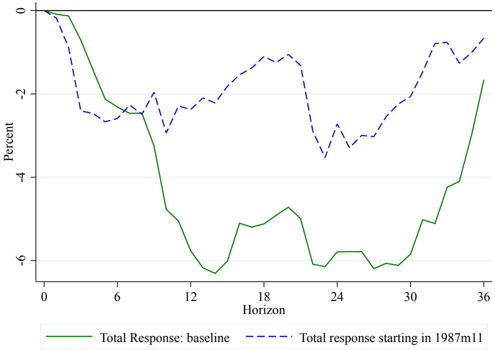

(b) February 1996

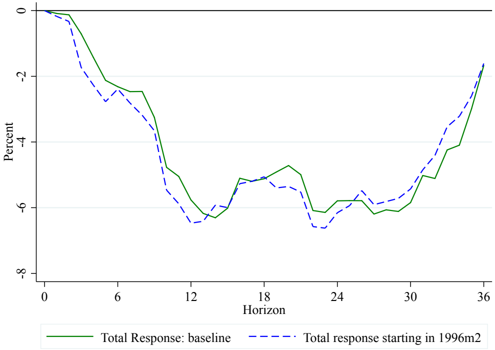

Example: GDP response to fiscal policy varies with monetary stance

Cloyne, Jordà, and Taylor 2023

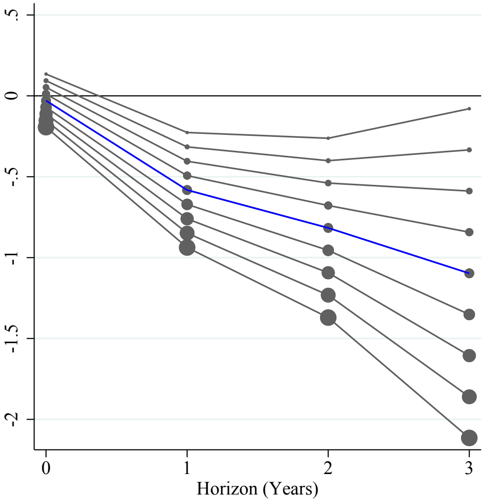

Variation in the multiplier by horizon and stance

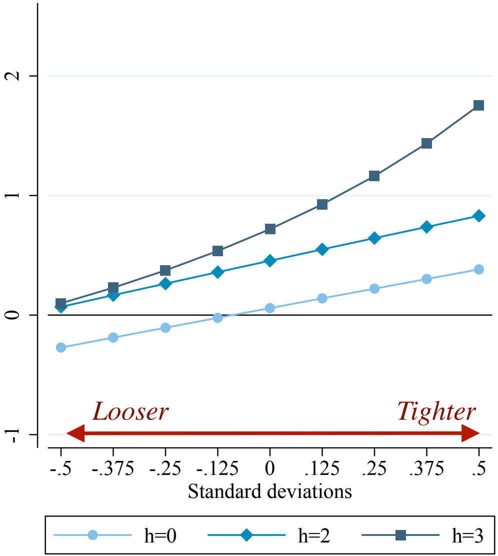

Time varying estimates of the multiplier

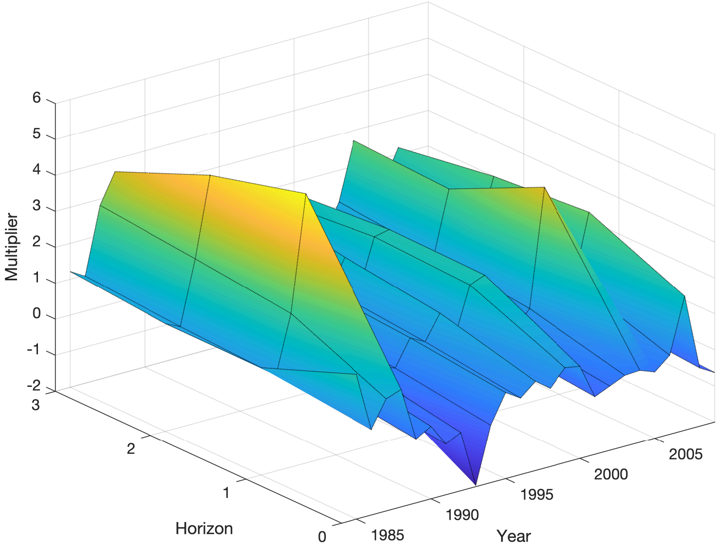

STATA example: Usual LPIV

`kob_example.do`

Response of real GDP to 1pp fiscal consolidation

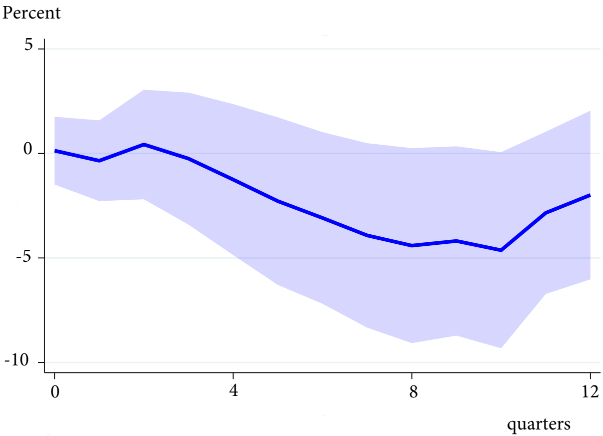

90% error bands

STATA example: Choosing two dates

`kob_example.do`

Response of real GDP to 1pp fiscal consolidation

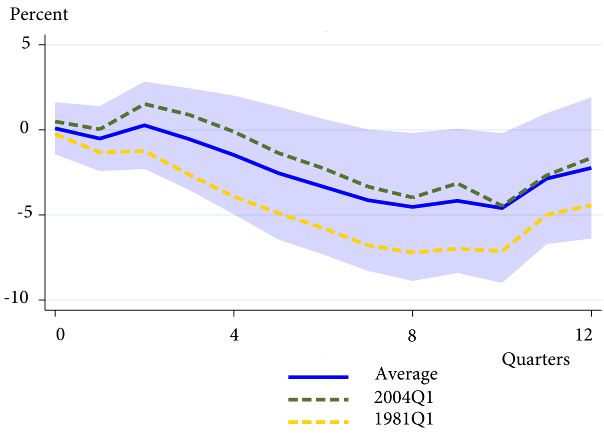

90% error bands

STATA example: Monetary offset

`kob_example.do`

Response of real GDP to 1pp fiscal consolidation

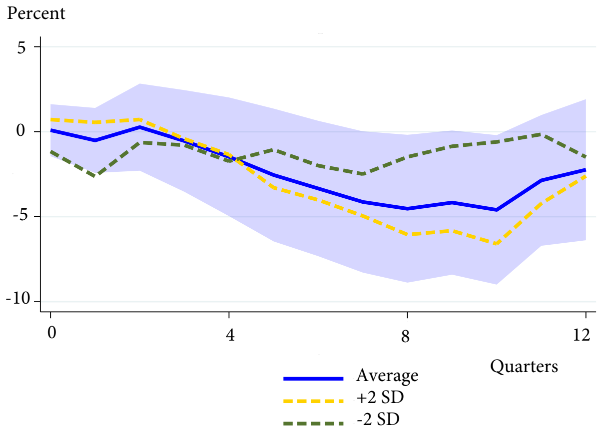

90% error bands

Panel data applications

Back to index

Differences-in-differences with LPs

Dube, Girardi, Jordà and Taylor

D-i-D with multiple treated groups & treatment periods

- TWFE implementation of DiD (static or distributed lags) can be severely biased.
  - Estimate is an average with possibly negative weights. Bad!
- LP-DiD = local projections + clean controls (Cengiz et al 2019)
  - No negative weights. Good!
  - Simple reweighting to recover ATT

Difference-in-Differences (DiD)

2x2 Setting

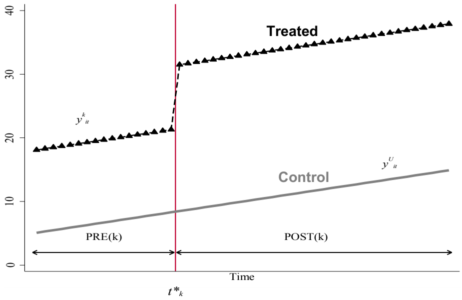

Staggered Setting

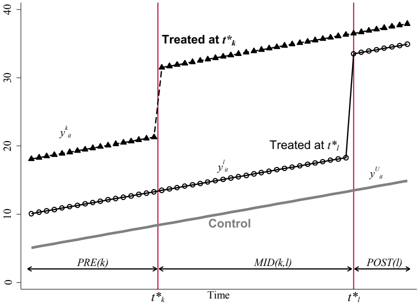

(Visual examples from Goodman-Bacon, 2021)

The conventional (until recently) DiD estimator: TWFE

- let $P_t = 1$ for post, 0 for pre; $A_i = 1$ for treated, 0 for control.
- Static TWFE

$$y_{it} = \alpha_i + \delta_t + \beta^{TWFE} D_{it} + \epsilon_{it}; \quad D_{it} = P_t \times A_i$$

- Event-study (distributed lags) TWFE

$$y_{it} = \alpha_i + \delta_t + \sum_{m=-Q}^{M} \beta_m^{TWFE} D_{it-m} + \epsilon_{it}$$

- OK in the 2x2 setting, or when treatment occurs at the same time.
- Biased even under parallel trends with staggered treatment, if treatment effects are dynamic and heterogeneous.

The problems with TWFE in the staggered setting

- TWFE as weighted-average of 2x2 comparisons (Goodman-Bacon 2021)
  1. Newly treated vs Never treated;
  2. Newly treated vs Not-yet treated;
  3. Newly treated vs Earlier treated.

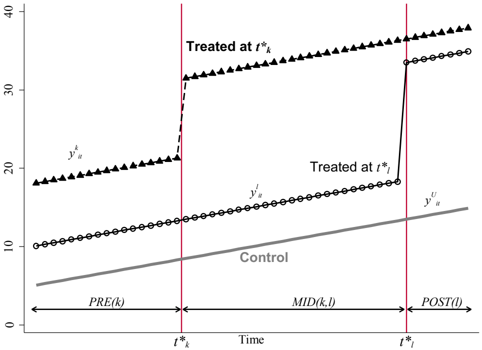

The problems with TWFE in the staggered setting

- TWFE as a weighted-average of cell-specific ATTs (de Chaisemartin & D'Haultfoeuille 2020)

$$E\left[\hat{\beta}^{TWFE}\right] = E\left[\sum_{(g,t):D_{gt}=1} \frac{N_{g,t}}{N_1} w_{g,t} \Delta_{g,t}\right]$$

$\rightarrow$ Weights can be negative!

LP-DiD Estimator

No Covariates, Outcome Lags

$$
\begin{aligned}
y_{i,t+k} - y_{i,t-1} = \quad & \beta^{k\ LP-DiD} \Delta D_{it} && \} \text{ treatment indicator} \\
& + \delta_t^k && \} \text{ time effects} \\
& + e_{it}^k \ ; && \quad \text{for } k = 0, \ldots, K\,.
\end{aligned}
$$

restricting the sample to observations that are either:

$$
\begin{cases}
\text{treatment} & \Delta D_{it} = 1\,, \\
\text{clean control} & \Delta D_{i,t+h} = 0 \text{ for } h = -H, \ldots, k\,.
\end{cases}
$$

Key advantage of LP over distributed lags TWFE formulation of DiD:
differencing is in outcomes, not treatments.

LP-DiD Estimator

$$
\begin{aligned}
y_{i,t+k} - y_{i,t-1} = \quad & \beta^{k\ LP-DiD} \Delta D_{it} && \} \text{ treatment indicator} \\
& + \textstyle\sum_{p=1}^{P} \gamma_{0,p}^{k} \Delta y_{i,t-p} && \} \text{ outcome lags} \\
& + \textstyle\sum_{m=1}^{M} \sum_{p=0}^{P} \gamma_{m,p}^{k} \Delta x_{m,i,t-p} && \} \text{ covariates} \\
& + \delta_t^k && \} \text{ time effects} \\
& + e_{it}^k \; ; && \quad \text{for } k = 0, \ldots, K\,.
\end{aligned}
$$

restricting the sample to observations that are either:

$$
\begin{cases}
\text{treatment} & \Delta D_{it} = 1\,, \\
\text{clean control} & \Delta D_{i,t+h} = 0 \text{ for } h = -H, \ldots, k\,.
\end{cases}
$$

An equivalent specification to implement LP-DiD

- Instead can use dummies to rule out unclean controls

$$
\begin{aligned}
y_{i,t+k} - y_{i,t-1} = \quad & \beta^{k\ LP-DiD} \Delta D_{it} && \} \text{ treatment indicator} \\
& + \theta^k UC_{i,t} && \} \text{ UC indicator} \\
& + \textstyle\sum_{p=1}^{P} \gamma_{0,p}^k (1 + \rho_{0,p}^k UC_{i,t})\ \Delta y_{i,t-p} && \} \text{ outcome lags} \times \text{UC} \\
& + \textstyle\sum_{m=1}^{M} \sum_{p=0}^{P} \gamma_{m,p}^k (1 + \rho_{m,p}^k UC_{i,t}) \Delta x_{m,i,t-p} && \} \text{ covariates} \times \text{UC} \\
& + \delta_t^k (1 + \phi_t^k UC_{i,t}) && \} \text{ time effects} \times \text{UC} \\
& + e_{it}^k \ ; && \quad \text{for } k = 0, \ldots, K \, .
\end{aligned}
$$

- UC_{it} =1 if previously treated.
- With absorbing treatment, $UC_{it} = \sum_{j=-H (j \neq 0)}^{k} \Delta D_{i,t+j}$

Simulation Evidence

- N=500 units; T=50 time periods.
- DGP: $Y_{0it} = \rho Y_{0,i,t-1} + \lambda_i + \gamma_t + \epsilon_{it}; \quad -1 < \rho < 1; \quad \lambda_i, \gamma_t, \epsilon_{it} \sim N(0, 25)$
- Binary staggered treatment.
- TE grows in time for 20 periods, and is stronger for early adopters.

1. Exogenous treatment
   - Units randomly assigned to 10 groups of size $N/10$
   - One group never treated; others treated at $\tau = 11, 13, 15 \ldots, 27$.
2. Endogenous treatment
   - Probability of treatment depends on past outcome dynamics.
   - Negative shocks increase probability of treatment.
   - Parallel trends holds only conditional on outcome lag.

Simulation Evidence

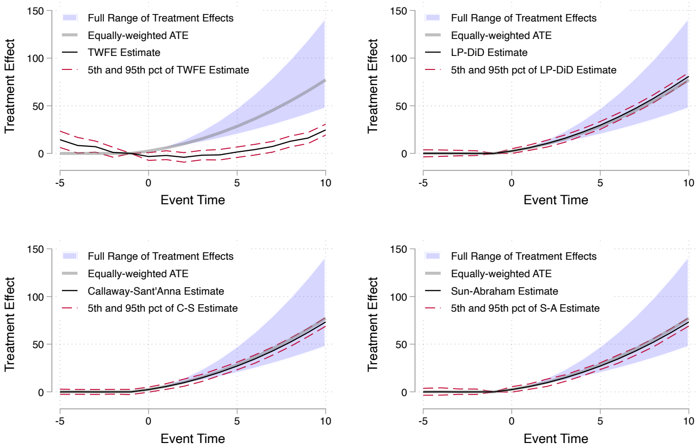

Average estimates and 95% and 5% percentiles from 200 replications.

Simulation Evidence

endogenous treatment scenario

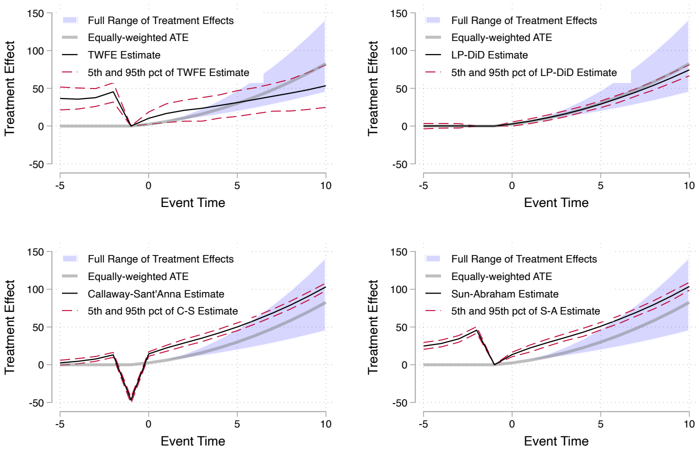

Average estimates and 95% and 5% percentiles from 200 replications.

Banking Deregulation and the Labor Share

1970-1996: staggered introduction of (inter-state & intra-state) banking deregulation in US states.

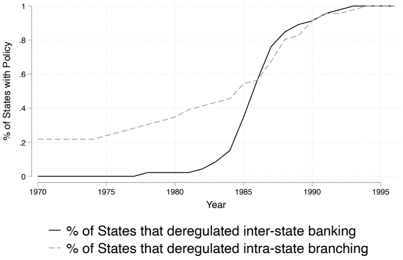

- Leblebicioglu & Weinberger (2020) use static & event-study TWFE to estimate effects on the labor share.
- Negative effect of *inter-state* banking deregulation (≈ -1p.p.).
- No effect of *intra-state* branching deregulation.

TWFE estimates

- negative effect from inter-state
- no effect from intra-state

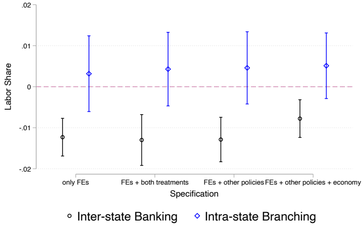

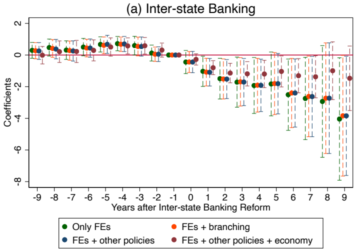

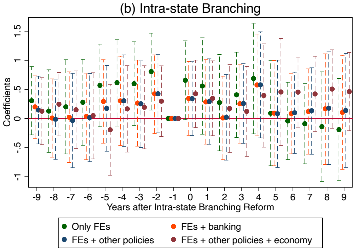

Forbidden comparisons in the TWFE specification

- TWFE uses 'forbidden' comparisons:
  earlier liberalizers are controls for later liberalizers.
- Use Goodman-Bacon (2021) decomposition to assess their influence.
- Contribution of unclean comparisons to TWFE estimates:
  - 36% for inter-state banking deregulation;
  - 70% for intra-state branching deregulation.

Effect of banking deregulation on the labor share

LP-DiD estimates

(a) Inter-state banking deregulation

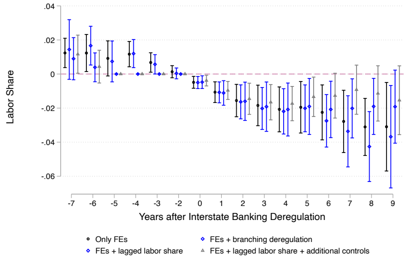

(b) Intra-state branching deregulation

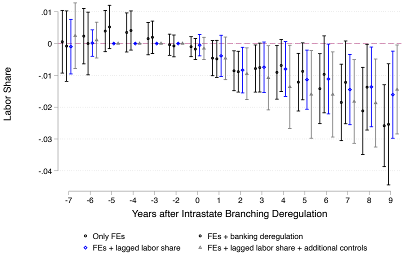

- LP-DiD avoids unclean comparisons & allows controlling for *y* lags.
- Negative effect of inter-state branking deregulation is confirmed.
- But also intra-state branching deregulation has negative effect.