Brief Summary: Synthetic Difference-in-Differences (EMERG Reading Group)

Damian Clarke

March 13, 2023

Consider a setting with repeated units over time (a panel). Some units receive some type of treatment or policy, some do not, and we wish to estimate the impact of treatment/policy receipt. Two common settings are often used in these cases:

- Difference-in-difference (DiD) methods: generally a considerable number of treatment units, and a willingness to believe in parallel trend assumptions.
- Synthetic control (SC) methods (eg Abadie et al. [2010]): generally one or few treated units, and optimal reweighting used to proceed despite lack of parallel trends

Both methods are popular for analysis in observational settings with repeated cross sections or panel data. However, both have draw backs: DiD requires parallel trend assumption which may be too strong; SC requires convex hull assumption, and there is no clearly accepted preference on how to combine multiple treated units.

Arkhangelsky et al. [2021] propose an estimator which draws on the strength of both methods, while avoiding common pitfalls. This is the synthetic difference-in-differences estimator (SDID). They show that it is generally comparable or better than DiD or SC alone, and prove consistency and asymptotic normality under a series of assumptions which are weaker than the assumptions of DiD and SC.

The SDID estimator can be defined as follows:

$$\left(\widehat{\tau}^{sdid}, \widehat{\mu}, \widehat{\alpha}, \widehat{\beta}\right) = \arg\min_{\tau,\mu,\alpha,\beta} \left\{ \sum_{i=1}^{N}\sum_{t=1}^{T}(Y_{it} - \mu - \alpha_i - \beta_t - W_{it}\tau)^2 \widehat{\omega}_i^{sdid}\widehat{\lambda}_t^{sdid}. \right\} \tag{1}$$

Here, unit specific weights ω are optimally chosen to generate the control group, and time-specific weights λ are chosen to generate the pre-exposure period considered. These weights are chosen as:

$$\left(\widehat{\omega}_0, \widehat{\omega}^{sdid}\right) = \arg\min_{\omega_0\in\mathbb{R},\omega\in\Omega} \sum_{t=1}^{T_{pre}} \left(\omega_0 + \sum_{i=1}^{N_{co}}\omega_i Y_{it} - \frac{1}{N_{tr}}\sum_{i=N_{co}+1}^{N} Y_{it}\right)^2 + \zeta^2 T_{pre}||\omega||_2^2 \tag{2}$$

$$\text{where } \Omega = \left\{\omega \in \mathbb{R}_+^N, \text{ with } \sum_{i=1}^{N_{co}}\omega_i = 1 \text{ and } \omega_i = \frac{1}{N_{tr}} \text{ for all } i = N_{co}+1,\dots,N\right\},$$

$||\omega||_2$ refers to the Euclidean norm and ζ is a regularization parameter. A virtually identical procedure is followed for time-specific weights. In practice, unit specific weights seek to construct a control with parallel trends to the treated unit in the pre-control period, though α_i allows trends to be parallel, but *not* equal in levels.

To see how this is a generalization of SC and DiD, note that SC can be written as follows:

$$\left(\widehat{\tau}^{sc}, \widehat{\mu}, \widehat{\beta}\right) = \arg\min_{\tau,\mu,\beta} \left\{ \sum_{i=1}^{N}\sum_{t=1}^{T}(Y_{it} - \mu - \beta_t - W_{it}\tau)^2 \widehat{\omega}_i^{sc}. \right\} \tag{3}$$

Principally, this imposes that the synthetic control unit must match the treated unit *in levels*, given the omission of unit-specific fixed effects α_i which are allowed (but not required) in (1). And note that DiD can be written as:

$$\left(\widehat{\tau}^{did}, \widehat{\mu}, \widehat{\alpha}, \widehat{\beta}\right) = \arg\min_{\tau,\mu,\alpha,\beta} \left\{ \sum_{i=1}^{N}\sum_{t=1}^{T}(Y_{it} - \mu - \alpha_i - \beta_t - W_{it}\tau)^2 \right\}. \tag{4}$$

where the omission of re-weighting elements implies that we require parallel trends to hold generically. In practice, an example of what this implies can be seen in a simple single-treatment setting in the figure overleaf.

In their paper, Arkhangelsky et al. [2021] also provide an applied discussion of how to implement these methods. They cover points such as:

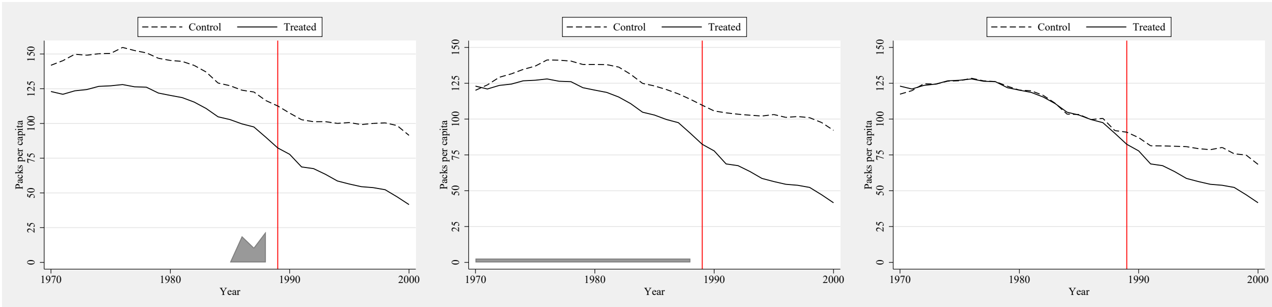

(a) SDID: Outcome Trends (b) DID: Outcome Trends (c) SC: Outcome Trends

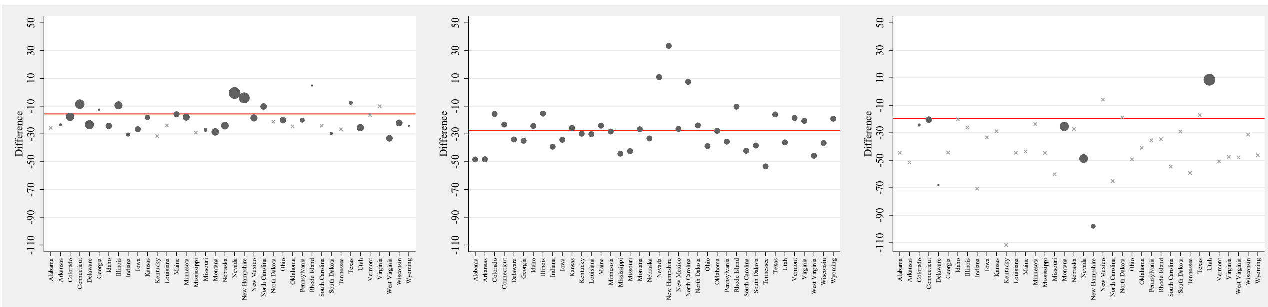

(d) SDID: Unit-Specific Weights (e) DID: Unit-Specific Weights (f) SC: Unit-Specific Weights

Figure 1: Comparison of estimators

(a) Inference, describing the following alternatives:
   - Bootstrap: works well when multiple units are treated, computationally costly.
   - Jack-knife: works well when multiple units are treated, somewhat conservative, computationally cheap.
   - Placebo: works well even if only one treated unit, requires homoscedasticity

(b) Multiple treated units I: seamlessly incorporated in cases of "block designs" where multple units adopt at a specific time period.

(c) Multiple treated units II: additionally propose what to do in cases with "staggered adoption" where units adopt at varying times. This consists of generating a weighted ATT from underlying adoption-specific SDID estimates: $\widehat{ATT} = \sum_a (T^a_{post}/T_post) \times \tau^{sdid}_a$, where a refers to adoption dates.

Clarke et al. [2023] discuss some additional computational and empirical details including:

- Defining inference procedures in the case of the staggered adoption design
- Various proposed manners to deal with controls
- How to map this into an "event study" style design

References

A. Abadie, A. Diamond, and J. Hainmueller. 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, 2010.

D. Arkhangelsky, S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager. Synthetic difference-in-differences. *American Economic Review*, 111(12):4088–4118, December 2021. doi: 10.1257/aer.20190159. URL `https://www.aeaweb.org/articles?id=10.1257/aer.20190159`.

D. Clarke, D. Pailañir, S. Athey, and G. Imbens. Synthetic Difference In Differences Estimation. econ.EM 2301.11859, arXiv, 2023. URL `https://doi.org/10.48550/arXiv.2301.11859`.