

NBER WORKING PAPER SERIES

TWO-WAY FIXED EFFECTS AND DIFFERENCES-IN-DIFFERENCES
WITH HETEROGENEOUS TREATMENT EFFECTS:
A SURVEY

Clément de Chaisemartin
Xavier D'Haultfoeuille

Working Paper 29691
<http://www.nber.org/papers/w29691>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2022

We are very grateful to Francesco Armillei, Kirill Borusyak, Bruno Ferman, Xavier Jaravel, Jonathan Roth, Jann Spiess, Gonzalo Vazquez-Bare, and Kaspar Wuthrich for their helpful comments. We have no sources of funding to disclose. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Clément de Chaisemartin and Xavier D'Haultfoeuille. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects:
A Survey
Clément de Chaisemartin and Xavier D'Haultfoeuille
NBER Working Paper No. 29691
January 2022
JEL No. C21,C23

ABSTRACT

Linear regressions with period and group fixed effects are widely used to estimate policies' effects: 26 of the 100 most cited papers published by the American Economic Review from 2015 to 2019 estimate such regressions. It has recently been shown that those regressions may produce misleading estimates, if the policy's effect is heterogeneous between groups or over time, as is often the case. This survey reviews a fast-growing literature that documents this issue, and that proposes alternative estimators robust to heterogeneous effects.

Clément de Chaisemartin
Department of Economics
University of California at Santa Barbara
Santa Barbara, CA 93106
and NBER
clementdechaisemartin@ucsb.edu

Xavier D'Haultfoeuille
CREST
5 avenue Henry Le Chatelier
91764 Palaiseau cedex
FRANCE
xavier.dhaultfoeuille@ensae.fr

Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey*

Clément de Chaisemartin[†] Xavier D’Haultfoeuille[‡]

Abstract

Linear regressions with period and group fixed effects are widely used to estimate policies’ effects: 26 of the 100 most cited papers published by the American Economic Review from 2015 to 2019 estimate such regressions. It has recently been shown that those regressions may produce misleading estimates, if the policy’s effect is heterogeneous between groups or over time, as is often the case. This survey reviews a fast-growing literature that documents this issue, and that proposes alternative estimators robust to heterogeneous effects.

Keywords: two-way fixed effects regressions, differences-in-differences, parallel trends, heterogeneous treatment effects, panel data, repeated-cross section data, policy evaluation.

JEL Codes: C21, C23

1 Introduction

A popular method to estimate the effect of a policy, or treatment, on an outcome is to compare over time groups experiencing different evolutions of their exposure to treatment. In practice, this idea is implemented by regressing $Y_{g,t}$, the outcome in group g and at period t , on group fixed effects, period fixed effects, and $D_{g,t}$, the treatment of group g at period t . For instance, to measure the effect of the minimum wage on employment in the US, researchers have often regressed employment in county g and year t on county fixed effects, year fixed effects, and the minimum wage in county g and year t .

Such two-way fixed effects (TWFE) regressions are probably the most-commonly used technique in economics to measure the effect of a treatment on an outcome. de Chaisemartin and

*We are very grateful to Francesco Armillei, Kirill Borusyak, Bruno Ferman, Xavier Jaravel, Jonathan Roth, Jann Spiess, Gonzalo Vazquez-Bare, and Kaspar Wüthrich for their helpful comments.

[†]Economics Department, Sciences Po, clement.dechaisemartin@sciencespo.fr

[‡]CREST-ENSAE, xavier.dhaultfoeuille@ensae.fr.

D’Haultfœuille (2021a) conducted a survey of the 20 papers with the most Google Scholar citations published by the American Economic Review in 2015, and of the similarly selected papers in 2016, 2017, 2018, and 2019. Of those 100 papers, 26 have estimated at least one TWFE regression to estimate the effect of a treatment on an outcome. TWFE regressions are also very commonly used in political science, sociology, and environmental sciences.

Researchers have long thought that TWFE estimators are equivalent to differences-in-differences (DID) estimators. With two groups and two periods, a DID estimator compares the outcome evolution from period 1 to 2 between a treatment group s that switches from untreated to treated, and a control group n that is untreated at both dates:

$$\text{DID} = Y_{s,2} - Y_{s,1} - (Y_{n,2} - Y_{n,1}). \quad (1)$$

DID relies on a parallel trends assumption: in the absence of the treatment, both groups would have experienced the same outcome evolution. Specifically, for every $g \in \{s, n\}$ and $t \in \{1, 2\}$, let $Y_{g,t}(0)$ and $Y_{g,t}(1)$ denote the potential outcomes in group g at period t without and with the treatment, respectively.¹ Parallel trends requires that the expected evolution of the untreated outcome be the same in both groups:

$$E[Y_{s,2}(0) - Y_{s,1}(0)] = E[Y_{n,2}(0) - Y_{n,1}(0)].$$

Under that assumption, DID is unbiased for the average treatment effect (ATE) in group s at period 2 (see, e.g., Abadie (2005)):

$$\begin{aligned} E[\text{DID}] &= E[Y_{s,2} - Y_{s,1} - (Y_{n,2} - Y_{n,1})] \\ &= E[Y_{s,2}(1) - Y_{s,1}(0) - (Y_{n,2}(0) - Y_{n,1}(0))] \\ &= E[Y_{s,2}(1) - Y_{s,2}(0)] + E[Y_{s,2}(0) - Y_{s,1}(0)] - E[Y_{n,2}(0) - Y_{n,1}(0)] \\ &= E[Y_{s,2}(1) - Y_{s,2}(0)], \end{aligned} \quad (2)$$

where the last equality follows from the parallel trends assumption. Parallel trends is partly testable, by comparing the outcome trends of groups s and n , before group s received the treatment. In practice, such pre-trends tests sometimes fail, but other times they indicate that the two groups were indeed on parallel paths before s got treated.²

¹Implicitly, this notation rules out dynamic treatment effects, and assumes that groups’ potential outcomes only depend on their current treatment, not on their past treatments. This restriction is not of essence to derive Equation (2) below, but it is of essence for some of the other results we cover, as noted later in the paper. We relax it in Section 3.2.

²Pre-trends tests come with caveats unveiled by a recent literature, see Kahn-Lang and Lang (2020), Bilinski and Hatfield (2018), and Roth (2021). Similarly, recent papers have proposed relaxations of the parallel trends assumption (see, e.g., Manski and Pepper, 2018; Rambachan and Roth, 2019; Freyaldenhoven et al., 2019). Though we allude to it in Section 3.2.1, this literature is mostly beyond the scope of this survey. See Roth et al. (2022) for a review.

Motivated by the fact that in the two-groups and two-periods design described above, DID is equal to the treatment coefficient in a TWFE regression with group and period fixed effects, researchers have also estimated TWFE regressions in more complicated designs with many groups and periods, variation in treatment timing, treatments switching on and off, and/or non-binary treatments. Recent research has shown that in those more complicated designs, TWFE estimators are unbiased for an ATE if parallel trends holds, and if another assumption is satisfied: the treatment effect should be constant, between groups and over time. Unlike parallel trends, this assumption is unlikely to hold, even approximately, in most of the applications where TWFE regressions have been used. For instance, the effect of the minimum wage on employment is likely to differ in counties with highly educated workers, and in counties with less educated workers.

The realization that one of the most commonly used empirical method in social science relies on an often-implausible assumption has spurred a flurry of methodological papers diagnosing the seriousness of the issue, and proposing alternative estimators. This review aims to provide an overview of this recent literature, which has developed in such a quick and dynamic manner that some practitioners may have gotten lost in the whirlwind of new working papers. We start by giving an overview of the papers that have identified TWFE’s regressions lack of robustness to heterogeneous treatment effects, and that have proposed diagnostic tools practitioners may use to assess the seriousness of this issue in their own application. We then give an overview of the papers that have proposed alternative estimators robust to heterogeneous treatment effects. Table 1 in the conclusion summarizes the heterogeneity-robust estimators available to applied researchers, depending on their research design. When available, the Stata and R commands implementing the diagnostics tools and alternative estimators discussed in this review are referenced, and the basic syntax of the Stata command is provided. We refer the reader to the commands’ help files for further details on their syntax.

2 TWFE regressions with heterogeneous treatment effects

2.1 TWFE regressions may not identify a convex combination of treatment effects

We consider observations that can be divided into G groups and T periods, respectively indexed by the placeholders g and t , which can refer to any group or time period. The data may be an individual-level panel, or repeated cross-sections where groups are, say, individuals’ county of birth. The data could also be a cross-section where cohort of birth plays the role of time.

Let $\hat{\beta}_{fe}$ denote the coefficient of $D_{g,t}$, the treatment in group g at period t , in an OLS regression of $Y_{i,g,t}$, the outcome of individual i in group g at period t , on group fixed effects, period fixed

effects, and $D_{g,t}$:

$$Y_{i,g,t} = \hat{\alpha}_g + \hat{\gamma}_t + \hat{\beta}_{fe} D_{g,t} + \epsilon_{g,t}.$$

de Chaisemartin and D’Haultfoeuille (2020) show that under a parallel trends assumption on the potential outcome without treatment $Y_{g,t}(0)$,

$$E[\hat{\beta}_{fe}] = E\left[\sum_{(g,t):D_{g,t}\neq 0} W_{g,t} TE_{g,t}\right]. \quad (3)$$

If the treatment is binary, $TE_{g,t} = Y_{g,t}(1) - Y_{g,t}(0)$, the ATE in group g at time t . If the treatment is discrete or continuous, $TE_{g,t} = (Y_{g,t}(D_{g,t}) - Y_{g,t}(0))/D_{g,t}$, the effect of moving the treatment from 0 to $D_{g,t}$ scaled by $D_{g,t}$.³ The $W_{g,t}$ are weights summing to 1, that are proportional to and of the same sign as

$$N_{g,t}(D_{g,t} - D_{g,\cdot} - D_{\cdot,t} + D_{\cdot,\cdot}), \quad (4)$$

where $N_{g,t}$ is the number of observations in cell (g, t) , $D_{g,\cdot}$ is the average treatment of group g across periods, $D_{\cdot,t}$ is the average treatment at period t across groups, and $D_{\cdot,\cdot}$ is the average treatment across groups and periods.⁴

Equations (3) and (4) have two important consequences. First, they imply that $\hat{\beta}_{fe}$ may be biased for the average treatment effect across all treated (g, t) cells, the ATT. Indeed, Equation (4) shows that $W_{g,t}$ is not proportional to $N_{g,t}$, so $E[\hat{\beta}_{fe}]$ does not assign to $TE_{g,t}$ a weight proportional to the population of cell (g, t) . A special case where $\hat{\beta}_{fe}$ is unbiased for the ATT under the parallel trends assumption alone is when (i) the design is staggered, meaning that groups’ treatment can only increase over time and can change at most once;⁵ (ii) the treatment is binary; (iii) there is no variation in treatment timing: all treated groups start receiving the treatment at the same date. Then, one can show that $D_{g,t} - D_{g,\cdot} - D_{\cdot,t} + D_{\cdot,\cdot}$ is constant across the treated (g, t) s, so it follows from Equation (4) that $W_{g,t}$ is proportional to $N_{g,t}$, and Equation (3) then implies that $\hat{\beta}_{fe}$ is unbiased for the ATT. However, conditions (i)-(iii) are seldom met in practice. $\hat{\beta}_{fe}$ can also be unbiased for the ATT if one is ready to make more assumptions than just parallel trends. For instance, if one is also ready to assume that $D_{g,t} - D_{g,\cdot} - D_{\cdot,t} + D_{\cdot,\cdot}$ is uncorrelated with $TE_{g,t}$, the treatment effects that are up- and down-weighted by $\hat{\beta}_{fe}$ do not systematically differ, and one can then show that $\hat{\beta}_{fe}$ is unbiased for the ATT (see Corollary 2

³de Chaisemartin and D’Haultfoeuille (2020) derive Equation (3) assuming that groups’ potential outcomes only depend on their current treatment, not on their past treatments. With dynamic effects, Equation (3) still holds if the treatment is binary and staggered, except that some of the $TE_{g,t}$ s become effects of having been treated for more than one period.

⁴Equation (4) holds under the assumption that the number of observations follows the same evolution over time in every group.

⁵Together, (i) and (ii) imply that groups can only switch from untreated to treated, and may do so at different points in time. This is probably the definition of a staggered design many people have in mind. i) extends the definition of a staggered design to non-binary treatments.

in de Chaisemartin and D’Haultfoeuille, 2020).⁶ Unfortunately, this no-correlation condition is often implausible. To see this, note that $D_{g,t} - D_{g,\cdot} - D_{\cdot,t} + D_{\cdot,\cdot}$ is decreasing in $D_{g,\cdot}$, meaning that $\hat{\beta}_{fe}$ downweights the treatment effect of groups with the highest average treatment from period 1 to T . However, groups with the largest and lowest average treatment may have systematically different treatment effects. Similarly, $D_{g,t} - D_{g,\cdot} - D_{\cdot,t} + D_{\cdot,\cdot}$ is decreasing in $D_{\cdot,t}$, and the treatment effects at time periods with the highest average treatment across groups may also systematically differ from the treatment effects at time periods where the average treatment is lower.

Second, and perhaps more worryingly, Equation (4) implies that some of the weights $W_{g,t}$ may be negative. This means that in the minimum wage example, $\hat{\beta}_{fe}$ could be estimating something like 3 times the effect of the minimum wage on employment in Santa Clara county, minus 2 times the effect in Wayne county. Then, if raising the minimum wage by one dollar decreases employment by 5% in Santa Clara county and by 20% in Wayne county, one would have $E[\hat{\beta}_{fe}] = 3 \times -0.05 - (2 \times -0.2) = 0.25$. $E[\hat{\beta}_{fe}]$ would be positive, while the minimum wage’s effect on employment is negative both in Santa Clara and in Wayne county. This example shows that $\hat{\beta}_{fe}$ may not satisfy the “no-sign reversal property”: $E[\hat{\beta}_{fe}]$ could for instance be positive, even if the treatment effect is strictly negative in every (g, t) . This phenomenon can only arise when some of the weights $W_{g,t}$ are negative: when all those weights are positive, $\hat{\beta}_{fe}$ does satisfy the no-sign reversal property. Note that despite its intuitive appeal and its popularity among applied researchers, the no-sign reversal property is not grounded in statistical decision theory, unlike other commonly-used criteria to discriminate estimators such as the mean-squared error. Still, it is connected to the economic concept of Pareto efficiency. If an estimator satisfies “no-sign-reversal”, the estimand attached to it can only be positive if the treatment is not Pareto-dominated by the absence of treatment, meaning that not everybody is hurt by the treatment. Conversely, the estimand can only be negative if the treatment does not Pareto-dominate the absence of treatment. On the other hand, if an estimator does not satisfy “no-sign-reversal”, the estimand attached to it could for instance be positive, even if the treatment is Pareto-dominated.

Inasmuch as “no-sign-reversal” is a desirable property, it becomes interesting to understand when $\hat{\beta}_{fe}$ may satisfy it. Equation (4) shows that with a binary treatment, the weights attached to $\hat{\beta}_{fe}$ could all be positive. With a binary treatment, all the (g, t) s entering the summation in (3) must have $D_{g,t} = 1$, so for a weight $W_{g,t}$ to be strictly negative, one must have $1 + D_{\cdot,\cdot} < D_{g,\cdot} + D_{\cdot,t}$. This cannot happen if $D_{g,\cdot} + D_{\cdot,t} \leq 1$ for every (g, t) . Accordingly, all the weights are likely to be positive when there is no group that is treated most of the time, and no time periods where most groups are treated. In staggered designs, this has led Jakiela (2021) to propose to drop

⁶A special case of this “no-correlation” condition is if the treatment effect is constant, i.e. $TE_{g,t} = \delta$ for all (g, t) . Then, it directly follows from Equation (3) that $E[\hat{\beta}_{fe}] = \delta$. However, constant effect is most often an implausible assumption.

the last periods of the data, those when $D_{.,t}$ is the highest, to mitigate or eliminate the negative weights. One could also drop the always-treated groups, if there are any.

On the other hand, Equation (4) shows that with a non-binary treatment, it becomes more likely that some of the weights $W_{g,t}$ are negative. Gentzkow et al. (2011) study the effect of the number of newspapers in county g and year t on turnout in presidential elections. Assume that in year t , county g has 1 newspaper ($D_{g,t} = 1$), which is below its average number of newspapers across years, equal, say, to 2 ($D_{g,.} = 2$). At the same time, the average number of newspapers across counties in year t is equal to 2 ($D_{.,t} = 2$), which is above the average number of newspapers across all counties and years, equal, say, to 1 ($D_{.,.} = 1$). Then, it follows from (4) that the weight assigned to the effect of newspapers in county g and year t is strictly negative. More generally, a necessary condition to have that all weights are positive is that in every period where the population's treatment is higher than its average across periods ($D_{.,t} \geq D_{.,.}$), the treatment of each treated group must also be larger than its average across periods ($D_{g,t} \geq D_{g,.}$ for all gs such that $D_{g,t} \neq 0$). In practice, it seems that with a non-binary treatment, $\hat{\beta}_{fe}$ often has many large negative weights attached to it. de Chaisemartin and D'Haultfœuille (2020) review Gentzkow et al. (2011) and Enikolopov et al. (2011), two papers that have estimated TWFE regressions with a non-binary treatment. In both cases, more than half of the weights attached to $\hat{\beta}_{fe}$ are negative, and the sum of the negative weights is large (-0.53 and -2.26, respectively).⁷

The `twowayfeweights` Stata (see de Chaisemartin, D'Haultfœuille and Deeb, 2019) and R (see Zhang and de Chaisemartin, 2021) commands compute the weights $W_{g,t}$ in (3). The basic syntax of the Stata command is:

```
twowayfeweights outcome groupid timeid treatment, type(feTR)
```

The decomposition in (3) is the main result in de Chaisemartin and D'Haultfœuille (2020). Related results have appeared earlier in Theorems S1 and S2 of the Supplementary Material of de Chaisemartin and D'Haultfœuille (2015). Borusyak and Jaravel (2017) consider the case with a binary and staggered treatment. In their Lemma 1 and Proposition 1, they assume that the treatment effect varies with the duration elapsed since one has started receiving the treatment but does not vary across groups and over time. Then, they show that $\hat{\beta}_{fe}$ estimates a weighted sum of effects, that may assign negative weights to long-run treatment effects. Their Appendix C also contains another result related to that in Equation (3).⁸

⁷ Gentzkow et al. (2011) do not estimate the TWFE regression defined above, but a first-difference version of that regression. Many weights attached to their regression are negative, and the sum of negative weights is even larger than for the TWFE regression.

⁸ Prior to that, Chernozhukov et al. (2013) had shown that one-way FE regressions may be biased for the average treatment effect, though unlike TWFE regressions they always estimate a convex combination of effects.

2.2 The origin of the problem: forbidden comparisons

2.2.1 Forbidden comparisons when the treatment is binary and the design is staggered

Goodman-Bacon (2021) shows that when the treatment is binary and the design is staggered, meaning that groups can switch in but not out of treatment, we have

$$\widehat{\beta}_{fe} = \sum_{g \neq g', t < t'} v_{g,g',t,t'} DID_{g,g',t,t'}, \quad (5)$$

where $DID_{g,g',t,t'}$ is a DID comparing the outcome evolution of two groups g and g' from a pre period t to a post period t' , and where $v_{g,g',t,t'}$ are non-negative weights summing to one, with $v_{g,g',t,t'} > 0$ if and only if g switches treatment between t and t' while g' does not.⁹ Some of the $DID_{g,g',t,t'}$ s in Equation (5) compare a group switching treatment from t to t' to a group untreated at both dates, while other $DID_{g,g',t,t'}$ s compare a switching group to a group treated at both dates. The negative weights in (3) originate from this second type of DIDs.

To see that, let us consider a simple example, first introduced by Borusyak and Jaravel (2017),¹⁰ with two groups and three periods. Group e , the early-treated group, is untreated at period 1 and treated at periods 2 and 3. Group ℓ , the late-treated group, is untreated at periods 1 and 2 and treated at period 3. In this simple example, Equation (5) reduces to

$$\widehat{\beta}_{fe} = (DID_{e,\ell,1,2} + DID_{\ell,e,2,3})/2, \quad (6)$$

with

$$\begin{aligned} DID_{e,\ell,1,2} &= Y_{e,2} - Y_{e,1} - (Y_{\ell,2} - Y_{\ell,1}), \\ DID_{\ell,e,2,3} &= Y_{\ell,3} - Y_{\ell,2} - (Y_{e,3} - Y_{e,2}). \end{aligned}$$

$DID_{e,\ell,1,2}$ compares the period-1-to-2 outcome evolution of group e , that switches from untreated to treated from period 1 to 2, to the outcome evolution of group ℓ that is untreated at both periods. Thus, $DID_{e,\ell,1,2}$ is similar to the DID estimator in Equation (1), and under parallel trends it is unbiased for the treatment effect in group e at period 2:

$$E[DID_{e,\ell,1,2}] = E[TE_{e,2}]. \quad (7)$$

$DID_{\ell,e,2,3}$, on the other hand, compares the period-2-to-3 outcome evolution of group ℓ , that switches from untreated to treated from period 2 to 3, to the outcome evolution of group e that

⁹Goodman-Bacon (2021) actually decomposes $\widehat{\beta}_{fe}$ as a weighted average of DIDs between cohorts of groups becoming treated at the same date, and between periods of time where their treatment remains constant. One can then further decompose his decomposition, as we do here.

¹⁰Borusyak and Jaravel (2017) have also coined the “forbidden comparisons” expression that we borrow here.

is treated at both dates. At both periods, e 's outcome is its treated potential outcome, which is equal to the sum of its untreated outcome and its treatment effect. Accordingly,

$$Y_{e,3} - Y_{e,2} = Y_{e,3}(0) + TE_{e,3} - (Y_{e,2}(0) + TE_{e,2}).$$

On the other hand, group ℓ is only treated at period 3, so

$$Y_{\ell,3} - Y_{\ell,2} = Y_{\ell,3}(0) + TE_{\ell,3} - Y_{\ell,2}(0).$$

Taking the expectation of the difference between the two previous equations,

$$E[\text{DID}_{\ell,e,2,3}] = E[TE_{\ell,3} - TE_{e,3} + TE_{e,2}], \quad (8)$$

where $E[Y_{e,3}(0) - Y_{e,2}(0)]$ and $E[Y_{\ell,3}(0) - Y_{\ell,2}(0)]$ cancel out under the parallel trends assumption. Finally, it follows from Equations (6), (7), and (8) that

$$E[\widehat{\beta}_{fe}] = E[1/2TE_{\ell,3} + TE_{e,2} - 1/2TE_{e,3}]. \quad (9)$$

In this simple example, Equation (3) reduces to (9). The right-hand side of Equation (9) is a weighted sum of three ATEs where one ATE receives a negative weight. As the previous derivation shows, this negative weight comes from the fact $\widehat{\beta}_{fe}$ leverages $\text{DID}_{\ell,e,2,3}$, a DID comparing a group switching from untreated to treated to a group treated at both periods.

If one is ready to assume that the treatment effect does not change over time, $TE_{e,3} = TE_{e,2}$, and (8) simplifies to

$$E[\text{DID}_{\ell,e,2,3}] = E[TE_{\ell,3}]. \quad (10)$$

Then, the negative weight in (8) disappears, and $\widehat{\beta}_{fe}$ estimates a weighted average of treatment effects. This extends beyond this simple example: Theorem S2 of the Web Appendix of de Chaisemartin and D'Haultfœuille (2020) and Equation (16) of Goodman-Bacon (2021) show that in staggered adoption designs with a binary treatment, $\widehat{\beta}_{fe}$ estimates a convex combination of effects, if the treatment effect does not change over time but may still vary across groups. This conclusion, however, no longer holds if the treatment is not binary or the design is not staggered. Moreover, assuming constant treatment effects over time is often implausible as this rules out both dynamic treatment effects and calendar time effects.

The decomposition in Equation (5) is key to understand why $\widehat{\beta}_{fe}$ may not identify a convex combination of treatment effects. On the other hand, it cannot be used to assess if $\widehat{\beta}_{fe}$ does indeed estimate a convex combination of effects in a given application. Consider an example similar to that above, but with a third group n that remains untreated from period 1 to 3. In this second example, the decomposition in (5) now indicates that $\widehat{\beta}_{fe}$ assigns a weight equal to $1/6$ to DIDs comparing a switcher to a group treated at both periods. On the other hand, all

the weights in (3) are positive in this second example. This phenomenon can also arise in real data sets. In the data of Stevenson and Wolfers (2006) used by Goodman-Bacon (2021) in his empirical application, if one restricts the sample to states that are not always treated and to the first ten years of the panel, all the weights in (3) are positive, but the sum of the weights in (5) on DID_s comparing a switcher to a group treated at both periods is equal to 0.06. Beyond these examples, one can show that having DID_s comparing a switcher to a group treated at both periods in (5) is necessary but not sufficient to have negative weights in (3). Similarly, the sum of the weights on DID_s comparing a switcher to a group treated at both periods in (5) is always larger than the absolute value of the sum of the negative weights in (3). The reason why Equation (5) “overestimates” the negative weights in (3) is that as soon as there are three distinct treatment dates, there is not a unique way of decomposing $\widehat{\beta}_{fe}$ as a weighted average of DID_s, and there exists other decompositions than Equation (5) putting less weight on DID_s using a group treated at both periods as the control group.¹¹

The `bacondecomp` Stata (see Goodman-Bacon et al., 2019) and R (see Flack and Edward, 2020) commands compute the DID _{g,g',t,t'} _s entering in (5), the weights assigned to them, as well as the sum of the weights on DID _{g,g',t,t'} _s using a group treated at both periods as the control group. The basic syntax of the `bacondecomp` Stata command is:

```
bacondecomp outcome treatment, ddetail
```

2.2.2 More forbidden comparisons when the design is not staggered or treatment is not binary

When the treatment does not follow a staggered design, or when it is not binary, $\widehat{\beta}_{fe}$ may leverage another type of comparison: it may compare the outcome evolution of a group m whose treatment increases more to the outcome evolution of a group ℓ whose treatment increases less.

¹¹To see that, let $t_0 < t_1 < t_2$ be three dates, let e be an early-treated group becoming treated at t_1 , let ℓ be a late-treated group becoming treated at t_2 , and let n be a group untreated yet at t_2 . Let $\underline{v} = \min(v_{\ell,e,t_1,t_2}, v_{e,n,t_0,t_2}) > 0$. One has

$$\text{DID}_{\ell,e,t_1,t_2} = \text{DID}_{\ell,n,t_0,t_2} - \text{DID}_{e,n,t_0,t_2} + \text{DID}_{e,\ell,t_0,t_1}. \quad (11)$$

Then, it follows from Equation (11) that

$$\begin{aligned} & v_{\ell,e,t_1,t_2} \text{DID}_{\ell,e,t_1,t_2} + v_{e,n,t_0,t_2} \text{DID}_{e,n,t_0,t_2} \\ &= (v_{\ell,e,t_1,t_2} - \underline{v}) \text{DID}_{\ell,e,t_1,t_2} + \underline{v} \text{DID}_{\ell,n,t_0,t_2} + \underline{v} \text{DID}_{e,\ell,t_0,t_1} + (v_{e,n,t_0,t_2} - \underline{v}) \text{DID}_{e,n,t_0,t_2}. \end{aligned} \quad (12)$$

Plugging Equation (12) into Equation (5) will yield a different decomposition of $\widehat{\beta}_{fe}$ as a weighted average of DID_s. But the weight on DID_s using a group treated at both periods as the control group is equal to v_{ℓ,e,t_1,t_2} in the left-hand-side of Equation (12), and to $(v_{\ell,e,t_1,t_2} - \underline{v})$ in its right-hand side. Accordingly, this new decomposition puts strictly less weight than Equation (5) on DID_s using a group treated at both periods as the control group.

In fact, with two groups m and ℓ and two periods, one can show that

$$\widehat{\beta}_{fe} = \frac{Y_{m,2} - Y_{m,1} - (Y_{\ell,2} - Y_{\ell,1})}{D_{m,2} - D_{m,1} - (D_{\ell,2} - D_{\ell,1})}, \quad (13)$$

where the right hand side of the previous display is the Wald-DID estimator studied by de Chaisemartin and D'Haultfœuille (2018). The Wald-DID compares the outcome evolution of groups m and ℓ , and scales that comparison by the differential evolution of m 's and ℓ 's treatments. de Chaisemartin and D'Haultfœuille (2018) show that the Wald-DID may not estimate a convex combination of effects, unless the treatment effect is constant over time and is the same in groups m and ℓ . This second requirement was not present in the binary and staggered case. In that case, we have seen before that if the treatment effect is constant over time, $\widehat{\beta}_{fe}$ estimates a convex combination of effects, even if the treatment effect varies between groups.

To see that with a non-binary or non-staggered treatment $\widehat{\beta}_{fe}$ may not estimate a convex combination of effects even if the treatment effect is constant over time, let us consider a simple example. Assume that group m goes from 0 to 2 units of treatment from period 1 to 2, while group ℓ goes from 0 to 1 unit. Then, the denominator of the Wald-DID is equal to $2 - 0 - (1 - 0) = 1$, so

$$\widehat{\beta}_{fe} = Y_{m,2} - Y_{m,1} - (Y_{\ell,2} - Y_{\ell,1}).$$

To simplify, let us also assume that in both groups, potential outcomes are linear in the number of treatment units, with slopes that are constant over time but may differ for groups m and ℓ :

$$\begin{aligned} Y_{m,t}(d) &= Y_{m,t}(0) + \delta_m d \\ Y_{\ell,t}(d) &= Y_{m,t}(0) + \delta_\ell d. \end{aligned}$$

Then, under parallel trends,

$$\begin{aligned} E[\widehat{\beta}_{fe}] &= E[Y_{m,2}(2) - Y_{m,1}(0) - (Y_{\ell,2}(1) - Y_{\ell,1}(0))] \\ &= E[Y_{m,2}(0) + 2\delta_m - Y_{m,1}(0) - (Y_{\ell,2}(0) + \delta_\ell - Y_{\ell,1}(0))] \\ &= E[Y_{m,2}(0) - Y_{m,1}(0)] - E[Y_{\ell,2}(0) - Y_{\ell,1}(0)] + 2\delta_m - \delta_\ell \\ &= 2\delta_m - \delta_\ell, \end{aligned}$$

a weighted sum of m and ℓ 's treatment effects, where group ℓ 's effect is weighted negatively. Intuitively, group ℓ is also treated at period two, and $\widehat{\beta}_{fe}$, which uses ℓ as a control group, subtracts its treatment effect out. This example also shows that $\widehat{\beta}_{fe}$ may fail to identify a convex combination of effects, even without variation in treatment timing: here, both m and ℓ start getting treated at period 2.

2.3 Decomposition results for other TWFE regression coefficients

2.3.1 Extensions of the decomposition in Equation (3)

A decomposition similar to (3) can be obtained for TWFE regressions with control variables, and for first-difference regressions where the outcome’s first difference is regressed on the treatment’s first difference and period fixed effects. de Chaisemartin and D’Haultfœuille (2020) also derive decompositions similar to (3), for $\widehat{\beta}_{fe}$ and for the first-difference coefficient, under common trends and under the assumption that the treatment effect does not change over time. The weights in all those decompositions are also computed by the `twowayfeweights` Stata and R commands.

2.3.2 Dynamic TWFE regressions

In staggered designs with a binary treatment, Sun and Abraham (2021) study dynamic TWFE regressions, also called event-study regressions:

$$Y_{g,t} = \widehat{\gamma}_g + \widehat{\lambda}_t + \sum_{\ell=-K, \ell \neq -1}^L \widehat{\beta}_\ell 1\{F_g = t - \ell\} + \varepsilon_{g,t}, \quad (14)$$

where F_g is the first period at which group g is treated. In words, the outcome is regressed on group and period fixed effects, and relative-time indicators $1\{F_g = t - \ell\}$ equal to 1 if group g started receiving the treatment ℓ periods ago. For $\ell \geq 0$, $\widehat{\beta}_\ell$ is supposed to estimate the cumulative effect of $\ell + 1$ treatment periods. For $\ell \leq -2$, $\widehat{\beta}_\ell$ is supposed to be a placebo coefficient testing the parallel trends assumption, by comparing the outcome trends of groups that will and will not start receiving the treatment in $|\ell|$ periods.

Actually, Sun and Abraham (2021) show that under parallel trends, for $\ell \geq 0$,

$$E[\widehat{\beta}_\ell] = E\left[\sum_g w_{g,\ell} TE_g(\ell) + \sum_{\ell' \neq \ell} \sum_g w_{g,\ell'} TE_g(\ell')\right], \quad (15)$$

where $TE_g(\ell)$ is the cumulative effect of $\ell + 1$ treatment periods in group g , and $w_{g,\ell}$ and $w_{g,\ell'}$ are weights such that $\sum_g w_{g,\ell} = 1$ and $\sum_g w_{g,\ell'} = 0$ for every ℓ' .¹² The first summation in the right-hand side of Equation (15) is a weighted sum across groups of the cumulative effect of $\ell + 1$ treatment periods, with weights summing to 1 but that may be negative. This first summation resembles that in the decomposition of the “static” TWFE coefficient in (3), and it implies that $\widehat{\beta}_\ell$ may be biased if the cumulative effect of $\ell + 1$ treatment periods varies across groups. The second summation is a weighted sum, across $\ell' \neq \ell$ and groups, of the cumulative effect of $\ell' + 1$

¹²Equation (15) follows from Proposition 3 in Sun and Abraham (2021), assuming no binning and that the treatment does not have an effect after $L + 1$ periods of exposure. A slight difference is that the decomposition in Sun and Abraham (2021) gathers groups that started receiving the treatment at the same period into cohorts. Their decomposition can then be further decomposed, as we do here.

treatment periods in group g , with weights summing to 0. This second summation was not present in the decomposition of the static TWFE coefficient. Importantly, its presence implies that $\widehat{\beta}_\ell$, which is supposed to estimate the cumulative effect of $\ell + 1$ treatment periods, may in fact be contaminated by the effects of $\ell' + 1$ treatment periods. As $\sum_g w_{g,\ell'} = 0$ for every ℓ' , this second summation disappears if $TE_g(\ell')$ does not vary across groups, but it is often implausible that the treatment effect does not vary across groups.

For $\ell \leq -2$, and without assuming parallel trends, Sun and Abraham (2021) show that $\widehat{\beta}_\ell$ estimates the sum of two terms. As intended, the first term measures deviations from parallel trends between groups that will and will not start receiving the treatment in $|\ell|$ periods. But the second term is similar to the second summation in the right-hand side of Equation (15): a weighted sum, across $\ell' \geq 0$ and groups, of the cumulative effect of $\ell' + 1$ treatment periods in group g , with weights summing to zero. Due to the presence of this second term, the expectation of $\widehat{\beta}_\ell$ may differ from zero even if parallel trends holds, and it may be equal to zero even if parallel trends fails. Thus, an important consequence of the results in Sun and Abraham (2021) is that in the presence of heterogeneous treatment effects, (14) cannot be used to test for parallel trends.

The `eventstudyweights` Stata command (see Sun, 2020) computes the weights attached to event-study regressions. Its basic syntax is:

```
eventstudyweights {rel_time_list}, absorb(i.groupid i.timeid)
cohort(first_treatment) rel_time(ry),
```

where `rel_time_list` is the list of relative-time indicators $1\{F_g = t - \ell\}$ included in (14), `first_treatment` is a variable equal to the period when group g got treated for the first time, and `ry` is a variable equal to `timeid` minus `first_treatment`, the number of periods elapsed since group g started receiving the treatment.

2.3.3 TWFE regressions with more than one treatment

Another case of interest is TWFE regressions with several treatments. For instance, to estimate separately the effect of medical and recreational marijuana laws on consumption, one may regress marijuana consumption in state g and year t on state and year fixed effects, on whether state g has a recreational marijuana law in year t , and on whether state g has a medical law in year t . de Chaisemartin and D'Haultfœuille (2021b) show that in those regressions, the coefficient on a given treatment identifies a weighted sum of that treatment's effect across (g, t) s, with weights summing to 1 but that may be negative, plus weighted sums of the effects of the other treatments in the regression, with weights summing to 0. In the example above, the coefficient on recreational laws may be contaminated by the effect of medical laws. The weights attached to TWFE regressions with several treatments are also computed by the `twowayfeweights` Stata and R commands.

3 Alternative heterogeneity-robust DID estimators

In this section, we review several recently-proposed alternatives to TWFE regressions. We restrict our attention to estimators relying on parallel trends assumptions, like TWFE regressions, but that do not restrict treatment effect heterogeneity between groups and over time, unlike TWFE regressions. This excludes papers that have assumed randomized treatment timing (see, e.g., Athey and Imbens, 2022; Roth and Sant’Anna, 2021) or sequential treatment randomization (see, e.g., Bojinov et al., 2021), rather than parallel trends. Intuitively, all the estimators below carefully choose valid control groups, to avoid making the “forbidden comparisons” that render TWFE estimators non-robust to heterogeneous treatment effects. We start by reviewing estimators ruling out dynamic effects, i.e. that assume that a group’s current outcome only depends on its current treatment, before reviewing estimators that allow dynamic effects. In complicated designs, say with a continuous treatment that changes often, allowing for dynamic effects comes with a number of costs: it may result in imprecise estimators, and may complicate the interpretation of the estimated effects. Then, one may want to carefully evaluate if past treatments are indeed likely to affect the current outcome.

3.1 Estimators ruling out dynamic effects

With a binary treatment de Chaisemartin and D’Haultfœuille (2020) propose to use the DID_M estimator, a weighted average, across t , of two types of DIDs:

1. a DID comparing the $t - 1$ to t outcome evolution of groups going from untreated to treated from $t - 1$ to t , the “switchers in”, and of groups untreated at both dates.
2. a DID comparing the $t - 1$ to t outcome evolution of groups treated at both dates, and of groups going from treated to untreated from $t - 1$ to t , the “switchers out”.

The first DID compares the outcome evolution of groups switching from untreated to treated, and groups untreated at both dates. It is therefore similar to the DID estimator in Equation (1), and it is unbiased for the treatment effect of the switching-in groups at period t , under a parallel trends assumption on the untreated outcome $Y_{g,t}(0)$. The second DID compares the outcome evolution of groups treated at both dates, and groups switching from treated to untreated. It is again similar to the DID estimator in Equation (1), switching “treatment” and “non-treatment”. Then, one can show that this second DID is unbiased for the treatment effect of the switching-out groups at period t , under a parallel trends assumption on the treated outcome $Y_{g,t}(1)$. Accordingly, DID_M relies on parallel trends assumptions on both $Y_{g,t}(0)$ and $Y_{g,t}(1)$, which implies parallel trends on the treatment effect $Y_{g,t}(1) - Y_{g,t}(0)$. de Chaisemartin and D’Haultfœuille (2020) propose placebo estimators that one can use to test those two assumptions. The placebos

compare the outcome trends of switchers and non-switchers, before the switchers switch. The multi-period DID estimator in Imai and Kim (2021) is related to the DID_M estimator. It can be used with a binary treatment, to estimate the switchers-in’s treatment effect.

The DID_M estimator can easily be extended to non-binary treatments taking a finite number of values. Then, it is a weighted average, across d and t , of DID’s comparing the $t - 1$ to t outcome evolution of groups whose treatment goes from d to some other value from $t - 1$ to t , and of groups with a treatment equal to d at both dates, normalized by the intensity of the treatment change experienced by the switchers. For instance, in Gentzkow et al. (2011), a county going from 2 to 4 newspapers is compared to a county with 2 newspapers at both dates.

The DID_M estimator is computed by the `did_multiplt` Stata (see de Chaisemartin, D’Haultfoeuille and Guyonvarch, 2019) and R (see Zhang and de Chaisemartin, 2020) commands. The basic syntax of the Stata command is:

```
did_multiplt outcome groupid timeid treatment
```

de Chaisemartin et al. (2022) extend the DID_M estimator to continuous treatments. To simplify, we present their estimators in the case with two time periods, though they readily extend to the case with more periods. de Chaisemartin et al. (2022) assume that from period one to two, the treatment of some units, hereafter referred to as the movers, changes. They also assume that the treatment of other units, hereafter referred to as the stayers, does not change. This assumption is likely to be met when the treatment is say, trade tariffs: tariffs’ reforms rarely apply to all products, so it is likely that tariffs of at least some products stay constant over time. On the other hand, this assumption is unlikely to be met when the treatment is say, precipitations: geographical units never experience the exact same precipitations over two consecutive years.

Under the assumption that there are some stayers, the estimator proposed by de Chaisemartin et al. (2022) compares the outcome evolution of movers and stayers, with the same period-one treatment. With a continuous treatment, such comparisons can either be achieved by reweighting stayers by propensity score weights, or by adjusting movers’ outcome change using a nonparametric regression of the outcome change on the period-one treatment among the stayers. Under parallel trends assumptions on all potential outcomes, the corresponding estimands identify a weighted average of the effect, across all movers, of moving their treatment from its period-one to its period-two value, scaled by the difference between these two values. This effect can be interpreted as some average of the effect of increasing the treatment by one unit on the outcome.

The estimators in de Chaisemartin et al. (2022) can be extended to the case where there are no stayers, provided there are quasi-stayers, meaning units whose treatment barely changes from period one to two. Alternatively, one could also use the estimator proposed by Graham and Powell (2012), which compares the outcome evolution of movers and quasi stayers whose treatment changes by less than some bandwidth, but without conditioning on units’ period-one treatment.

Their estimator relies on a linear treatment effect assumption, unlike those in de Chaisemartin et al. (2022). Both estimators require choosing a bandwidth, namely the lowest treatment change below which a unit can be considered as a quasi-stayer. Neither de Chaisemartin et al. (2022) or Graham and Powell (2012) derive an “optimal” bandwidth, so for now bandwidth choice is left to the discretion of the researcher. If the data has at least three periods, one could also use the correlated-random-coefficient estimator proposed by Chamberlain (1992). While it allows for some treatment effect heterogeneity, that estimator relies on a linear treatment effect assumption.

de Chaisemartin et al. (2022) show that after some relabelling, some of their estimators are equivalent or nearly equivalent to estimators that had been previously proposed by de Chaisemartin and D’Haultfoeuille (2018), Abadie (2005), and Callaway and Sant’Anna (2021). This implies that their estimators can be computed, up to small tweaks, by the companion software for those papers. We refer the reader to de Chaisemartin et al. (2022) for a precise description of how their estimators can be computed using existing software.

3.2 Estimators allowing for dynamic effects...

3.2.1 ... When the treatment is binary and the design is staggered.

With dynamic effects, group g ’s outcome at time t is allowed to depend on her past treatments. Then, researchers have proposed to replace the parallel trends assumption on the potential outcome without treatment $Y_{g,t}(0)$ by a parallel trends assumption on the outcome without having ever been treated $Y_{g,t}(\mathbf{0}_t)$, where $\mathbf{0}_t$ is a vector of t zeros (see Callaway and Sant’Anna, 2021; Sun and Abraham, 2021): for all $g \neq g'$,

$$E [Y_{g,t}(\mathbf{0}_t) - Y_{g,t-1}(\mathbf{0}_{t-1})] = E [Y_{g',t}(\mathbf{0}_t) - Y_{g',t-1}(\mathbf{0}_{t-1})]. \quad (16)$$

We now review the estimators that have been proposed by Callaway and Sant’Anna (2021), Sun and Abraham (2021), and Borusyak et al. (2021) for binary and staggered treatments, under the parallel trends assumption in Equation (16).

The estimators proposed by Callaway and Sant’Anna (2021)

In a staggered adoption design, groups can be aggregated into cohorts that start receiving the treatment at the same period. Callaway and Sant’Anna (2021) define their parameters of interest as $TE_{c,c+\ell}$, the average treatment effect at period $c + \ell$ of the cohort that started receiving the treatment at period c , for every $c \geq 2$ and $\ell \geq 0$ such that $\ell + c \leq T$. To estimate, say, $TE_{c,c}$, the treatment effect at time c of the cohort that started receiving the treatment at c , they propose to compare c ’s outcome evolution from $c - 1$ to c to the same outcome evolution

among all never-treated groups n , assuming for now that such groups exist. This DID estimator is unbiased:

$$\begin{aligned}
& E [Y_{c,c} - Y_{c,c-1} - (Y_{n,c} - Y_{n,c-1})] \\
&= E [Y_{c,c}(\mathbf{0}_{c-1}, 1) - Y_{c,c-1}(\mathbf{0}_{c-1}) - (Y_{n,c}(\mathbf{0}_c) - Y_{n,c-1}(\mathbf{0}_{c-1}))] \\
&= E [Y_{c,c}(\mathbf{0}_{c-1}, 1) - Y_{c,c}(\mathbf{0}_c)] + E [Y_{c,c}(\mathbf{0}_c) - Y_{c,c-1}(\mathbf{0}_{c-1}) - (Y_{n,c}(\mathbf{0}_c) - Y_{n,c-1}(\mathbf{0}_{c-1}))] \\
&= E [Y_{c,c}(\mathbf{0}_{c-1}, 1) - Y_{c,c}(\mathbf{0}_c)],
\end{aligned}$$

where the last equality follows from Equation (16). More generally, to estimate c 's effect at period $c + \ell$, Callaway and Sant'Anna (2021) propose to compare its $c - 1$ to $c + \ell$ outcome evolution to that of the never treated groups.

Callaway and Sant'Anna (2021) also propose estimators of more aggregated treatment effects, such as the average effect of having been treated for $\ell + 1$ periods across all groups reaching that treatment horizon, or the average treatment effect across all treated units that do not belong to a group treated throughout the panel. When there is no group that is always treated, this latter parameter is in fact the ATT.

Callaway and Sant'Anna (2021) also propose estimators similar to those above, but that use the not-yet-treated instead of the never-treated as controls. This is very useful when there is no never-treated group: in that case, the cohort-and-period specific effects $TE_{c,c+\ell}$ can still be estimated, for every $c \geq 2$ and $\ell \geq 0$ such that $\ell + c \leq U$, where U is the last period when at least one group is still untreated. Even when there are never-treated groups, the not-yet-treated is a larger control group, so this second set of estimators may be more precise than the first. Note that in staggered adoption designs with a binary treatment, the DID_M estimator proposed by de Chaisemartin and D'Haultfœuille (2020) also compares switchers to not-yet-treated, and is unbiased for switchers' instantaneous treatment effect.

Callaway and Sant'Anna (2021) also propose estimators relying on a conditional parallel trends assumption. Rather than assuming that all groups have the same evolution of their never-treated outcome, conditional parallel trends requires that groups with the same value of some time-invariant covariates have the same evolution of their never-treated outcome, an often weaker assumption. In the minimum wage example, if the proportion of college graduates at the start of the panel is conditioned upon, conditional parallel trends assumes that in the absence of a minimum wage, counties with the same proportion of college graduates would have experienced the same employment evolution, rather than assuming that all counties would have experienced the same employment evolution.

Finally, Callaway and Sant'Anna (2021) propose placebo estimators to test the parallel trends assumption underlying their estimators. Their placebo estimators are robust to heterogeneous effects, unlike the coefficients $\hat{\beta}_\ell$ for $\ell \leq -2$ from the event-study regression in (14).

The estimators proposed by Callaway and Sant’Anna (2021) are computed by the `csdid` Stata command (see Rios-Avila et al., 2021), and by the `did` R command (see Sant’Anna and Callaway, 2021). The basic syntax of the Stata command is

```
csdid outcome, time(timeid) gvar(cohort)
```

where `cohort` is a variable equal to the period when a group starts receiving the treatment.

The estimators proposed by Sun and Abraham (2021)

Sun and Abraham (2021) also propose DID estimators of the cohort-and-period specific effects $TE_{c,c+\ell}$ that only rely on the parallel trends assumption in Equation (16), and that are robust to heterogeneous treatment effects. Their estimators either use the never-treated groups as controls, or the last-treated groups if there are no never-treated. With the former control group, their estimators of the $TE_{c,c+\ell}$ parameters are identical to those proposed by Callaway and Sant’Anna (2021) with the same control group. Operationally, they show that their estimators can be computed via a simple linear regression, which may reduce computing time. Unlike Callaway and Sant’Anna (2021), they do not propose estimators relying on a conditional parallel trends assumption, and they also do not propose estimators using the not-yet-treated as controls.

Their estimators are computed by the `eventstudyinteract` Stata command (see Sun, 2021). Its basic syntax is

```
eventstudyinteract outcome {rel_time_list}, absorb(i.groupid i.timeid)
cohort(first_treatment) control_cohort(controlgroup)
```

where `rel_time_list` is the list of relative-time indicators $1\{F_g = t - \ell\}$ one would include in the dynamic TWFE regression in (14), `first_treatment` is a variable equal to the period when group g got treated for the first time, and `controlgroup` is an indicator for the control group observations (e.g.: the never treated).

The estimators proposed by Borusyak et al. (2021)

Borusyak et al. (2021) have proposed estimators that may be more efficient than those in Callaway and Sant’Anna (2021) and Sun and Abraham (2021), under some assumptions. Their estimators can be obtained by running a TWFE regression of the outcome on group and time fixed effects, and fixed effects for every treated (g, t) cell. To be concrete, if the data has 50 groups, 10 time periods, and 100 treated (g, t) cells, the regression has a constant and 158 fixed effects (49 for groups, 9 for time periods, and 100 for the treated (g, t) cells). Under the assumptions of the Gauss-Markov theorem, the coefficients from this regression are the linear estimators of the population coefficients with the lowest variance. But under parallel trends, the population coefficient on the fixed effect for treated cell (g, t) is actually equal to $TE_{g,t}$, the ATE in cell (g, t) , so the estimators in Borusyak et al. (2021) are the linear estimators of those ATEs with

the lowest variance. With estimators of $TE_{g,t}$ in hand, one can estimate $TE_{c,c+\ell}$ as the average of all the $TE_{g,t}$ s such that group g started receiving the treatment at period c and $t = c + \ell$. Again, Gauss-Markov ensures that this estimator is the best linear estimator of $TE_{c,c+\ell}$. As the estimators in Callaway and Sant’Anna (2021) and Sun and Abraham (2021) are also linear estimators, those in Borusyak et al. (2021) have a lower variance.

A second, numerically equivalent way of computing the estimators in Borusyak et al. (2021) amounts to fitting a regression of the outcome on group and time fixed effects in the sample of untreated observations, and using that regression to predict the counterfactual outcome of treated observations. Estimates of the treatment effect of those observations are then merely obtained by subtracting their counterfactual to their actual outcome. This second method is computationally faster than the first. It also readily generalizes to more complicated specifications, such as triple-differences, or models allowing for group-specific linear trends. Using this representation of their estimator, Borusyak et al. (2021) show that it can also be used to estimate the effect of a binary and non-staggered treatment, if that treatment does not have dynamic effects. This second method is the one used by the `did_imputation` Stata command (see Borusyak, 2021) and by the `didimputation` R command (see Butts, 2021) to compute the estimators proposed by Borusyak et al. (2021). The basic syntax of the Stata command is:

```
did_imputation outcome groupid timeid first_treatment,
```

where `first_treatment` is a variable equal to the period when group g first got treated.

Liu et al. (2021) and Gardner (2021) have independently proposed the same estimators as Borusyak et al. (2021),¹³ but the result showing that this estimator is efficient under the assumptions of the Gauss-Markov theorem only appears in Borusyak et al. (2021). Note that Wooldridge (2021) has also proposed an estimation strategy connected, and in some cases numerically equivalent, to that of Borusyak et al. (2021).

Understanding the differences between those estimators

Under parallel trends, the estimators in Borusyak et al. (2021) may offer precision gains with respect to those in Callaway and Sant’Anna (2021) or Sun and Abraham (2021), under the assumptions of the Gauss-Markov theorem. Those require, among other things, that the never treated potential outcomes $Y_{g,t}(\mathbf{0}_t)$ be independent of each other, both across groups and over time. It is, of course, often implausible that the potential outcomes of the same group are uncorrelated over time. With serial correlation, it is no longer guaranteed that the estimators in Borusyak et al. (2021) will always be more efficient than those in Callaway and Sant’Anna (2021) and Sun and Abraham (2021), but simulations in Borusyak et al. (2021) show that one can often still expect efficiency gains when serial correlation is moderate.

¹³Before that, Gobillon and Magnac (2016) have proposed a similar strategy to estimate treatment effects under a factor model.

If trends are not exactly parallel, the estimators in Borusyak et al. (2021) may be more or less biased than those in Callaway and Sant’Anna (2021) or Sun and Abraham (2021) depending on the nature of the violation of parallel trends. Borusyak et al. (2021) do not provide a closed-form of their estimators, but one can show that with only one treated group s , which starts to receive the treatment at period t_s , their estimator of that group’s treatment effect at $t_s + \ell$ is

$$Y_{s,t_s+\ell} - \frac{1}{t_s - 1} \sum_{k=1}^{t_s-1} Y_{s,k} - \frac{1}{G - 1} \sum_{g \neq s} \left(Y_{g,t_s+\ell} - \frac{1}{t_s - 1} \sum_{k=1}^{t_s-1} Y_{g,k} \right). \quad (17)$$

Instead, the estimator in Callaway and Sant’Anna (2021) and Sun and Abraham (2021) is

$$Y_{s,t_s+\ell} - Y_{s,t_s-1} - \frac{1}{G - 1} \sum_{g \neq s} (Y_{g,t_s+\ell} - Y_{g,t_s-1}). \quad (18)$$

Equations (18) and (17) show that the estimator in Callaway and Sant’Anna (2021) and Sun and Abraham (2021) use groups’ $t_s - 1$ outcome, the last period before s gets treated, as the baseline outcome, while the estimator in Borusyak et al. (2021) uses the average outcome from period 1 to $t_s - 1$. This is why the latter estimator is often more precise. However, it is also more biased, when parallel trends does not exactly hold and the discrepancy between groups’ trends gets larger over longer horizons, as would for instance happen when there are group-specific linear trends. In such instances, Roth (2021) notes that leveraging earlier pre-treatment periods increases the bias of a DID estimator, since one makes comparisons from earlier periods. If, on the other hand, parallel trends fails due to anticipation effects arising a few periods before t_s , Equations (17) and (18) imply that the estimators in Borusyak et al. (2021) is less biased than those in Callaway and Sant’Anna (2021) or Sun and Abraham (2021). Beyond the simple example we consider here, deriving a closed-form expression of the estimators in Borusyak et al. (2021) is not straightforward. Whether the conclusions we derive in this simple example carry through to more complicated designs is thus an open question.

In any case, if one views parallel trends as a reasonable first-order approximation rather than an assumption that holds exactly, it may make sense to investigate how sensitive one’s findings are to violations of parallel trends. To do so, one may for instance implement the partial identification approaches in Manski and Pepper (2018) or Rambachan and Roth (2019). The latter approach assumes that parallel trends do not hold exactly, and that the magnitude of placebo estimators is informative as to the magnitude of the bias in the actual estimators caused by differential trends. The estimators proposed by Callaway and Sant’Anna (2021) and Sun and Abraham (2021) may be more amenable to the approach in Rambachan and Roth (2019) than the estimators proposed by Borusyak et al. (2021). Consider again the same simple example as above. For any $\ell \leq t_s - 2$, one can construct the following placebo estimator:

$$Y_{s,t_s-1} - Y_{s,t_s-\ell-2} - \frac{1}{G - 1} \sum_{g \neq s} (Y_{g,t_s-1} - Y_{g,t_s-\ell-2}). \quad (19)$$

This placebo compares the treated and control groups' outcome evolution, from period $t_s - \ell - 2$ to $t_s - 1$, namely over $\ell + 1$ periods before group s got treated. It exactly mimicks the estimator of group s 's treatment effect at period $t_s + \ell$ proposed by Callaway and Sant'Anna (2021) and Sun and Abraham (2021), which compares the same groups, over the same number of periods. Accordingly, the magnitude of that placebo may indeed be informative as to the magnitude of the bias of the estimator in Equation (18), as requested by Rambachan and Roth (2019). Building a placebo that would similarly mimick the estimator proposed by Borusyak et al. (2021) is not feasible, precisely because that estimator leverages all pre-treatment periods to construct its baseline. See de Chaisemartin and D'Haultfœuille (2021a) for more discussion of the advantages of having placebos that mimick actual estimators.

Another difference between these approaches is that Borusyak et al. (2021) impose parallel trends for every group and between every pair of consecutive time periods.¹⁴ Callaway and Sant'Anna (2021), on the other hand, impose a weaker parallel trends assumption: from period c onwards, cohort c must be on the same trend as the never-treated groups, but before that cohort c may have been on a different trend. The assumption in Callaway and Sant'Anna (2021) is the minimal assumption ensuring that all the $TE_{c,c+\ell}$ can be unbiasedly estimated, but it is conditional on the design: which groups are required to be on parallel trends at which dates depends on groups' realized treatments. It is also not testable. We refer the reader to Marcus and Sant'Anna (2021) and Borusyak et al. (2021) for further discussion on the differences between parallel trends assumptions.

Overall, this discussion suggests that whether the estimators in Borusyak et al. (2021) should be preferred to those in Callaway and Sant'Anna (2021) and Sun and Abraham (2021) may depend on one's degree of confidence in the parallel trends assumption, on the type of violations of this assumption that seems more likely to arise in the application at hand, and on one's willingness to undertake a sensitivity analysis such as the one proposed by Rambachan and Roth (2019).

3.2.2 ... When the treatment is not binary or the design is not staggered.

de Chaisemartin and D'Haultfœuille (2021a) propose treatment effect estimators robust to heterogeneous and dynamic treatment effects and which can be used even if the treatment is not binary or the design is not staggered. In their survey of 26 highly cited 2015-2019 AER papers using a TWFE regression, they find that 19% have a binary treatment and a staggered design, so being able to accommodate more general designs is important.

For simplicity, we describe those estimators when the treatment is binary but non-staggered, and refer the reader to de Chaisemartin and D'Haultfœuille (2021a) for details as to how the approach below can be adapted to non-binary treatments. When the treatment is binary and

¹⁴de Chaisemartin and D'Haultfœuille (2020) and Sun and Abraham (2021) also impose that assumption.

non-staggered, groups untreated at $t = 1$ can still be aggregated into cohorts, according to the date at which they first get treated. Then, for every $c \geq 2$ and $\ell \geq 0$ such that $\ell + c \leq T$, let N_c be the number of groups in cohort c and let

$$FTT_{c,c+\ell} = \frac{1}{N_c} \sum_{g \in c} (Y_{g,c+\ell}(\mathbf{0}_{c-1}, 1, D_{g,c+1}, \dots, D_{g,c+\ell}) - Y_{g,c+\ell}(\mathbf{0}_{c+\ell}))$$

be the average effect, across all groups in cohort c , of having been treated for the first time ℓ periods ago, relative to having never been treated. Notice that in the non-staggered case, a group treated at c may revert to being untreated after that, and $FTT_{c,c+\ell}$ is defined conditional on groups' treatment trajectories from $c + 1$ to $c + \ell$, $(D_{g,c+1}, \dots, D_{g,c+\ell})$. de Chaisemartin and D'Haultfœuille (2021a) show that under the parallel trends assumption in Equation (16), $FTT_{c,c+\ell}$ can be unbiasedly estimated by a DID comparing the $c - 1$ to $c + \ell$ outcome evolution of groups in cohort c and groups untreated from period 1 to $c + \ell$.

In a staggered design, one has

$$FTT_{c,c+\ell} = \frac{1}{N_c} \sum_{g \in c} (Y_{g,c+\ell}(\mathbf{0}_{c-1}, 1, 1, \dots, 1) - Y_{g,c+\ell}(\mathbf{0}_{c+\ell})),$$

so $FTT_{c,c+\ell}$ is just the cumulative effect of having been treated for $\ell + 1$ periods. Outside of staggered designs, $FTT_{c,c+\ell}$ may be harder to interpret, as it may bundle together many different treatment effects. For instance, if some of the groups that get treated for the first time at period c remain treated at $c + 1$ while others go back to being untreated at $c + 1$, $FTT_{c,c+1}$ aggregates together $E(Y_{g,c+1}(\mathbf{0}_{c-1}, 1, 1) - Y_{g,c+1}(\mathbf{0}_{c+1}))$, the effect of two treatment periods for the first set of groups, and $E(Y_{g,c+1}(\mathbf{0}_{c-1}, 1, 0) - Y_{g,c+1}(\mathbf{0}_{c+1}))$, the effect of having been treated for one period one period ago for the second set of groups. One could try to separately estimate the effects of all the treatment trajectories observed in the data, as Callaway and Sant'Anna (2021) successfully do in the staggered case. But in the staggered case, there can be at most $T + 1$ treatment trajectories, while in the non-staggered case there may be up to 2^T of them. Trying to separately estimate the effects of all trajectories may often yield noisy estimates.

Instead, de Chaisemartin and D'Haultfœuille (2021a) show that the $FTT_{c,c+\ell}$ parameters can be aggregated into an economically interpretable parameter, namely the cost-benefit ratio a planner would use to conduct a cost-benefit analysis comparing the actual treatments of groups that were initially untreated to the counterfactual *status quo* scenario where they would have remained untreated throughout the panel. In other words, the planner seeks to assess whether the treatment changes that occurred over the panel were beneficial. de Chaisemartin and D'Haultfœuille (2021a) show that this cost-benefit ratio is equal to a weighted sum, across c and ℓ , of the $FTT_{c,c+\ell}$ parameters. Accordingly, this cost-benefit ratio can simply be estimated by a weighted sum of DIDs. Importantly, this cost-benefit ratio can also be interpreted as some average of the effect produced by a one-unit increase of the treatment. By symmetry, de Chaisemartin

and D’Haultfoeulle (2021a) show that the cost-benefit ratio comparing the actual treatments of groups that were initially treated to the counterfactual *status quo* scenario where they would have remained treated throughout the panel can also be estimated by a weighted sum of DIDs.

The estimators proposed by de Chaisemartin and D’Haultfoeulle (2021a) are computed by the `did_multipligt` Stata and R commands. To compute those estimators rather than those proposed in de Chaisemartin and D’Haultfoeulle (2020), the Stata command’s basic syntax is:

```
did_multipligt outcome groupid timeid treatment, robust_dynamic dynamic(#)  
average_effect
```

where `dynamic(#)` specifies the horizon over which effects of a first treatment switch have to be estimated.

The estimators in de Chaisemartin and D’Haultfoeulle (2021a) can be used with a binary treatment switching on and off, with a discrete treatment, or with a continuous and staggered treatment (groups start getting treated at different dates, with differing intensities, but once a group gets treated its treatment intensity never changes). The estimators proposed by Callaway et al. (2021) can also accommodate continuous and staggered treatments. For continuous and non-staggered treatments, in their Section 4.3 de Chaisemartin et al. (2022) extend their baseline estimators to allow for dynamic effects. With respect to their baseline estimators, the main difference is that when allowing for dynamic effects, fewer units can be used as controls. Without dynamic effects, at period t , any unit whose treatment has not changed between $t - 1$ and t can be used as a valid control. With dynamic effects, only units whose treatments have not changed from period 1 to t can be used as valid controls. Therefore, the need for “stayers” becomes even stronger when allowing for dynamic effects: many units need to keep the same value of the treatment for a large number of time periods. Developing estimators robust to dynamic effects that can be used with a continuous treatment and no stayers has not been done yet and is another promising area for future research.

The estimators in de Chaisemartin and D’Haultfoeulle (2021a) can, of course, also be used with a binary and staggered treatment. Without covariates in the estimation, they are then equivalent to the estimators proposed by Callaway and Sant’Anna (2021) using the not-yet-treated as controls. With covariates, the estimators in Callaway and Sant’Anna (2021) and de Chaisemartin and D’Haultfoeulle (2021a) differ. Callaway and Sant’Anna (2021) consider time-invariant covariates, and assume that trends are parallel once we condition on them. de Chaisemartin and D’Haultfoeulle (2021a) instead consider time-varying covariates and assume that trends are parallel once the linear effect of those time-varying covariates on the outcome is accounted for. This for instance allows them to include group-specific linear trends in the estimation. With covariates, the parallel trends conditions in Callaway and Sant’Anna (2021) and de Chaisemartin and D’Haultfoeulle (2021a) are not nested, and in principle one could combine both.

Finally, it is worth noting that de Chaisemartin and D’Haultfoeuille (2021*b*) propose estimators for the case with several treatments. They propose both estimators that generalize the DID_M estimator in de Chaisemartin and D’Haultfoeuille (2020) and rule out dynamic effects, and estimators that generalize those in Callaway and Sant’Anna (2021) and allow for dynamic effects.

4 Conclusion, and avenues for future research

The literature reviewed in this survey has shown that TWFE regressions may not always estimate a convex combination of treatment effects. In such cases, it may be hard to give them a causal interpretation, as TWFE coefficients could for instance be of a different sign than every unit’s treatment effect. Table 1 below summarizes the alternative estimators available to applied researchers, depending on their research design and on whether they are ready or not to rule out dynamic effects. The table shows that the literature so far has mostly focused on providing alternative estimators for the case with a binary treatment and staggered adoption. Heterogeneity-robust DID estimators that can be used in more complicated designs are scarce, while many applications where TWFE regressions have been used either do not have a staggered design, or do not have a binary treatment. Developing more estimators that can be used in such designs is a promising avenue for future research. This can often be done by building upon the insights gained from studying the binary-and-staggered case. For instance, the estimators proposed by de Chaisemartin and D’Haultfoeuille (2021*a*) build upon those proposed by Callaway and Sant’Anna (2021) for the binary-and-staggered case. We hope that the whirlwind of DID working papers shall continue, till heterogeneity-robust DID estimators are as widely applicable as TWFE regressions. It is also important to stress that at this stage, it is still unclear whether researchers should systematically abandon TWFE estimators: those estimators sometimes estimate a convex combination of effects, and they often have a lower variance than the heterogeneity-robust estimators reviewed in the previous section. A comparison of the mean-squared error of TWFE and heterogeneity-robust DID estimators in a broad set of applications is another promising avenue for future research.

Table 1: A summary of available heterogeneity-robust DID estimators

Panel A: Estimators ruling out dynamic effects (outcome unaffected by past treatments)		
<i>Treatment</i>	<i>Estimators available</i>	<i>Stata¹⁵ commands</i>
Binary	de Chaisemartin and D’Haultfœuille (2020) Imai and Kim (2018) Borusyak et al. (2021)	<code>did_multipligt</code> <code>did_imputation</code>
Discrete	de Chaisemartin and D’Haultfœuille (2020)	<code>did_multipligt</code>
Continuous, with stayers	de Chaisemartin et al. (2022)	See paper
Continuous, w/o stayers	de Chaisemartin et al. (2022) Graham and Powell (2012) Chamberlain (1992)	See paper <code>gmm</code> <code>gmm</code>
Several treatments	de Chaisemartin and D’Haultfœuille (2021 <i>b</i>)	<code>did_multipligt</code>
Panel B: Estimators allowing dynamic effects (outcome affected by past treatments)		
<i>Treatment</i>	<i>Estimators available</i>	<i>Stata¹⁵ commands</i>
Binary and staggered	Callaway and Sant’Anna (2021) Sun and Abraham (2021) Borusyak et al. (2021) de Chaisemartin and D’Haultfœuille (2021 <i>a</i>)	<code>csdid</code> <code>eventstudyinteract</code> <code>did_imputation</code> <code>did_multipligt</code>
Binary or discrete, non-staggered	de Chaisemartin and D’Haultfœuille (2021 <i>a</i>)	<code>did_multipligt</code>
Continuous and staggered	de Chaisemartin and D’Haultfœuille (2021 <i>a</i>) Callaway et al. (2021)	<code>did_multipligt</code>
Continuous and non-stagg., with stayers	de Chaisemartin et al. (2022)	See paper
Continuous and non-stagg., w/o stayers	No estimator available yet	
Several treatments	de Chaisemartin and D’Haultfœuille (2021 <i>b</i>)	<code>did_multipligt</code>

¹⁵All the Stata commands have R equivalents with the same name, except `eventstudyinteract` that does not have an R equivalent, and `csdid` whose R equivalent is called `did`.

References

- Abadie, A. (2005), ‘Semiparametric difference-in-differences estimators’, *Review of Economic Studies* **72**(1), 1–19.
- Athey, S. and Imbens, G. W. (2022), ‘Design-based analysis in difference-in-differences settings with staggered adoption’, *Journal of Econometrics* **226**, 62–79.
- Bilinski, A. and Hatfield, L. A. (2018), Nothing to see here? non-inferiority approaches to parallel trends and other model assumptions. arXiv preprint arXiv:1805.03273.
- Bojinov, I., Rambachan, A. and Shephard, N. (2021), ‘Panel experiments and dynamic causal effects: A finite population perspective’, *Quantitative Economics* **12**, 1171–1196.
- Borusyak, K. (2021), ‘DID_IMPUTATION: Stata module to perform treatment effect estimation and pre-trend testing in event studies’.
URL: <https://ideas.repec.org/c/boc/bocode/s458957.html>
- Borusyak, K. and Jaravel, X. (2017), Revisiting event study designs. Working Paper.
- Borusyak, K., Jaravel, X. and Spiess, J. (2021), Revisiting event study designs: Robust and efficient estimation. arXiv preprint arXiv:2108.12419.
- Butts, K. (2021), ‘didimputation: Imputation Estimator from Borusyak, Jaravel, and Spiess (2021) in R’.
URL: <https://cran.r-project.org/web/packages/didimputation/index.html>
- Callaway, B., Goodman-Bacon, A. and Sant’Anna, P. H. (2021), Difference-in-differences with a continuous treatment. arXiv preprint arXiv:2107.02637.
- Callaway, B. and Sant’Anna, P. H. (2021), ‘Difference-in-differences with multiple time periods’, *Journal of Econometrics* **225**, 200–230.
- Chamberlain, G. (1992), ‘Efficiency bounds for semiparametric regression’, *Econometrica* pp. 567–596.
- Chernozhukov, V., Fernández-Val, I., Hahn, J. and Newey, W. (2013), ‘Average and quantile effects in nonseparable panel models’, *Econometrica* **81**(2), 535–580.
- de Chaisemartin, C. and D’Haultfoeuille, X. (2015), Fuzzy differences-in-differences. ArXiv e-prints, eprint 1510.01757v2.
- de Chaisemartin, C. and D’Haultfoeuille, X. (2018), ‘Fuzzy differences-in-differences’, *The Review of Economic Studies* **85**(2), 999–1028.

- de Chaisemartin, C. and D’Haultfœuille, X. (2020), ‘Two-way fixed effects estimators with heterogeneous treatment effects’, *American Economic Review* **110**(9), 2964–2996.
- de Chaisemartin, C. and D’Haultfœuille, X. (2021a), Difference-in-differences estimators of intertemporal treatment effects. arXiv preprint arXiv:2007.04267.
- de Chaisemartin, C. and D’Haultfœuille, X. (2021b), Two-way fixed effects regressions with several treatments. arXiv preprint arXiv:2012.10077.
- de Chaisemartin, C., D’Haultfœuille, X. and Deeb, A. (2019), ‘twowayfeweights: Estimation of the Weights Attached to the Two-Way Fixed Effects Regressions in Stata’.
URL: <https://ideas.repec.org/c/boc/bocode/s458611.html>
- de Chaisemartin, C., D’Haultfœuille, X. and Guyonvarch, Y. (2019), ‘did_multipleqt: DID Estimation with Multiple Groups and Periods in Stata’.
URL: <https://ideas.repec.org/c/boc/bocode/s458643.html>
- de Chaisemartin, C., D’Haultfœuille, X., Pasquier, F. and Vazquez-Bare, G. (2022), Difference-in-differences estimators of the effect of a continuous treatment. arXiv preprint arXiv:2201.06898.
- Enikolopov, R., Petrova, M. and Zhuravskaya, E. (2011), ‘Media and political persuasion: Evidence from russia’, *American Economic Review* **101**(7), 3253–3285.
- Flack, E. and Edward (2020), ‘bacondecomp: Goodman-Bacon Decomposition in R’.
URL: <https://cran.r-project.org/web/packages/bacondecomp/index.html>
- Freyaldenhoven, S., Hansen, C. and Shapiro, J. M. (2019), ‘Pre-event trends in the panel event-study design’, *American Economic Review* **109**(9), 3307–38.
- Gardner, J. (2021), Two-stage differences in differences. Working paper.
- Gentzkow, M., Shapiro, J. M. and Sinkinson, M. (2011), ‘The effect of newspaper entry and exit on electoral politics’, *American Economic Review* **101**(7), 2980–3018.
- Gobillon, L. and Magnac, T. (2016), ‘Regional policy evaluation: Interactive fixed effects and synthetic controls’, *Review of Economics and Statistics* **98**(3), 535–551.
- Goodman-Bacon, A. (2021), ‘Difference-in-differences with variation in treatment timing’, *Journal of Econometrics* **225**, 254–277.
- Goodman-Bacon, A., Goldring, T. and Nichols, A. (2019), ‘BACONDECOMP: Stata module to perform a Bacon decomposition of difference-in-differences estimation’.
URL: <https://ideas.repec.org/c/boc/bocode/s458676.html>

- Graham, B. S. and Powell, J. L. (2012), ‘Identification and estimation of average partial effects in “irregular” correlated random coefficient panel data models’, *Econometrica* **80**(5), 2105–2152.
- Imai, K. and Kim, I. S. (2018), On the use of two-way fixed effects regression models for causal inference with panel data. Working Paper.
- Imai, K. and Kim, I. S. (2021), ‘On the use of two-way fixed effects regression models for causal inference with panel data’, *Political Analysis* **29**(3), 405–415.
- Jakiela, P. (2021), Simple diagnostics for two-way fixed effects. arXiv preprint arXiv:2103.13229.
- Kahn-Lang, A. and Lang, K. (2020), ‘The promise and pitfalls of differences-in-differences: Reflections on 16 and pregnant and other applications’, *Journal of Business & Economic Statistics* **38**(3), 613–620.
- Liu, L., Wang, Y. and Xu, Y. (2021), A practical guide to counterfactual estimators for causal inference with time-series cross-sectional data. arXiv preprint arXiv:2107.00856.
- Manski, C. F. and Pepper, J. V. (2018), ‘How do right-to-carry laws affect crime rates? coping with ambiguity using bounded-variation assumptions’, *Review of Economics and Statistics* **100**(2), 232–244.
- Marcus, M. and Sant’Anna, P. H. (2021), ‘The role of parallel trends in event study settings: An application to environmental economics.’, *Journal of the Association of Environmental and Resource Economists* **8**(2), 235–275.
- Rambachan, A. and Roth, J. (2019), An honest approach to parallel trends. Working paper.
- Rios-Avila, F., Sant’Anna, P. and Callaway, B. (2021), ‘Cstdid: Stata module for the estimation of difference-in-difference models with multiple time periods’.
URL: <https://EconPapers.repec.org/RePEc:boc:bocode:s458976>
- Roth, J. (2021), ‘Pre-test with caution: Event-study estimates after testing for parallel trends’, *American Economic Review: Insights* **forthcoming**.
- Roth, J. and Sant’Anna, P. H. (2021), Efficient estimation for staggered rollout designs. arXiv preprint arXiv:2102.01291.
- Roth, J., Sant’Anna, P. H., Bilinski, A. and Poe, J. (2022), What’s trending in difference-in-differences? a synthesis of the recent econometrics literature. arXiv preprint arXiv:2201.01194.
- Sant’Anna, P. and Callaway, B. (2021), ‘did: Treatment effects with multiple periods and groups in r’.
URL: <https://cran.r-project.org/web/packages/did/index.html>

- Stevenson, B. and Wolfers, J. (2006), ‘Bargaining in the shadow of the law: Divorce laws and family distress’, *The Quarterly Journal of Economics* **121**(1), 267–288.
- Sun, L. (2020), ‘EVENTSTUDYWEIGHTS: Stata module to estimate the implied weights on the cohort-specific average treatment effects on the treated (CATTs) (event study specifications)’.
URL: <https://ideas.repec.org/c/boc/bocode/s458833.html>
- Sun, L. (2021), ‘EVENTSTUDYINTERACT: Stata module to implement the interaction weighted estimator for an event study’.
URL: <https://ideas.repec.org/c/boc/bocode/s458978.html>
- Sun, L. and Abraham, S. (2021), ‘Estimating dynamic treatment effects in event studies with heterogeneous treatment effects’, *Journal of Econometrics* **225**, 175–199.
- Wooldridge, J. (2021), ‘Two-way fixed effects, the two-way mundlak regression, and difference-in-differences estimators’. Available at SSRN 3906345.
- Zhang, S. and de Chaisemartin, C. (2020), ‘did_multiplegt: DID Estimation with Multiple Groups and Periods in R’.
URL: <https://cran.r-project.org/web/packages/DIDmultiplegt/index.html>
- Zhang, S. and de Chaisemartin, C. (2021), ‘TwowayFEWeights: Estimation of the Weights Attached to the Two-Way Fixed Effects Regressions in R’.
URL: <https://cran.r-project.org/web/packages/TwoWayFEWeights/index.html>