



***The World's Largest Open Access Agricultural & Applied Economics Digital Library***

**This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.**

**Help ensure our sustainability.**

Give to AgEcon Search

AgEcon Search  
<http://ageconsearch.umn.edu>  
[aesearch@umn.edu](mailto:aesearch@umn.edu)

*Papers downloaded from AgEcon Search may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

*No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.*

# Implementing the panel event study

Damian Clarke  
 Department of Economics  
 University of Chile  
 Santiago, Chile  
 dclarke@fen.uchile.cl

Kathyia Tapia-Schythe  
 Department of Economics  
 University of California, Davis  
 Davis, CA  
 kattapia@ucdavis.edu

**Abstract.** Many studies estimate the impact of exposure to some quasiexperimental policy or event using a panel event study design. These models, as a generalized extension of “difference-in-differences” designs or two-way fixed-effects models, allow for dynamic leads and lags to the event of interest to be estimated, while also controlling for fixed factors (often) by area and time. In this article, we discuss the setup of the panel event study design in a range of situations and lay out several practical considerations for its estimation. We describe a command, `eventdd`, that allows for simple estimation, inference, and visualization of event study models in a range of circumstances. We then provide several examples to illustrate `eventdd`’s use and flexibility, as well as its interaction with various native Stata commands, and other relevant community-contributed commands such as `reghdfe` and `boottest`.

**Keywords:** st0655, `eventdd`, event studies, difference-in-differences, estimation, inference, visualization

## 1 Introduction

Recent developments in quasiexperimental methods have brought increasing attention to panel event study models. When one uses data covering a panel of observations (such as states) over time, the design seeks to estimate the impact of some event that occurs, or “switches on” in certain units and certain time periods.<sup>1</sup> These models seek to use as counterfactuals the areas in which the policy or event does not occur or has not yet occurred. By considering the variation in outcomes around the adoption of the event compared with a baseline reference period, one can estimate both event leads and lags, which allows for a clear visual representation of the event’s causal impact provided that key identifying assumptions are met.

These methods have been borne out of older difference-in-differences (DD) designs, or two-way fixed-effects models. These models often seek to examine the impact of natural experiments, where events are assigned to certain units due to some process

1. Throughout this article, for expositional purposes, we will refer to a geographical and a temporal dimension to these panel event study designs. However, we note that such models are suitable to other settings where one wishes to estimate the impact of an event that is assigned to certain units and where there is some temporal variation in outcomes to be studied. For example, rather than being based on different geographical units at different times, these models could be based on an event occurring in different age groups at different times. All results discussed in this article extend to settings such as these.

beyond the control of the analyst but owing to environmental or political factors (among others), and thus, generally do not assume that assignment is random. Indeed, as we lay out at more length in the following section, the key assumption underlying consistent estimation in event study models is that the occurrence of the event in a particular area is not systematically related to the changes in levels that would have occurred in the future in the absence of the event.

These models are widely used in empirical analyses in a range of contexts, having been applied to (among many others themes) automotive plant closures and opioid overdoses (Venkataramani et al. 2020), family planning access and childhood economic circumstance (Bailey, Malkova, and McLaren Forthcoming), healthcare reform and ambulatory care usage (Dimitrovová, Perelman, and Serrano-Alarcón 2020), and university reform and intergenerational mobility (Suhonen and Karhunen 2019). These cases suggest use across a range of fields, including social sciences, medicine and public health, and additional reviews of their frequency of use in several economic journals are provided in Abraham and Sun (2018); Roth (2019). A burgeoning literature has laid out several identification requirements in this setting (Freyaldenhoven, Hansen, and Shapiro 2019; Borusyak and Jaravel 2018; Abraham and Sun 2018; Athey and Imbens Forthcoming; Schmidheiny and Siegloch 2019). These methods can be used, with some restrictions, both in cases where events occur at the same time period in each unit and in cases where the adoption of events is staggered. Indeed, Athey and Imbens (Forthcoming) refer to these as “staggered adoption designs”, although here we follow the more common nomenclature of panel event studies.<sup>2</sup> Additionally, these methods are related to a much broader literature on staggered adoption of policies and the estimation of a single-coefficient model (de Chaisemartin and D’Haultfoeuille 2019; Callaway and Sant’Anna 2018; Goodman-Bacon 2018). While we briefly discuss these models in the methods section, our principal interest is on full panel-event study specifications that come with their own considerations.

In this article, we discuss these panel-event study models and practical issues related to their estimation and to inference in these settings. We also present the `eventdd` command, which allows for estimation and inference in event studies, as well as several postestimation procedures and the graphical presentation of estimates and confidence intervals (CIs).<sup>3</sup> This command can flexibly interact with both official Stata commands such as `regress` and `xtreg`, as well as the community-contributed regression command `reghdfe` (Correia 2014), which is highly convenient in two-way fixed-effects models such as those described in this article (Correia 2016). We discuss both estimation and inference in event study models. As well as standard inference procedures such as robust and cluster-robust inference, the `eventdd` command allows for wild bootstrap-based inference respecting the clustered nature of the occurrence of events and, specifically, the community-contributed `boottest` command (Roodman et al. 2019). After reviewing

2. Older literature, often related to the analyses of financial events and information shocks, refers to event studies. These event studies in finance are generally based on time-series observations and have quite different properties to the panel event studies used in policy analysis that we discuss in this article. A useful discussion of these finance-style event studies, and their application in Stata, is provided in Pacicco, Vena, and Venegoni (2018).

3. An earlier version of this command is available in Clarke and Tapia-Schythe (2020).

the theory behind panel event study models in section 2, we discuss the command syntax in section 3, before documenting the command's usage, applied to a particular empirical example, in section 4.

## 2 Methods

### 2.1 Estimation

Consider a panel covering a group, indexed as  $g$  and time periods  $t$ . We are interested in estimating the impact of the passage of an event that may occur at different times in different groups. We will denote as  $\text{Event}_g$  a variable recording the time period  $t$  in which the event is adopted in group  $g$ . Denoting the outcome of interest as  $y_{gt}$ , we can write the panel event study specification as<sup>4</sup>

$$y_{gt} = \alpha + \sum_{j=2}^J \beta_j (\text{Lead } j)_{gt} + \sum_{k=1}^K \gamma_k (\text{Lag } k)_{gt} + \mu_g + \lambda_t + \mathbf{X}'_{gt} \boldsymbol{\Gamma} + \varepsilon_{gt} \quad (1)$$

Here  $\mu_g$  and  $\lambda_t$  are group and time fixed effects,  $\mathbf{X}_{gt}$  are (optionally) time-varying controls, and  $\varepsilon_{gt}$  is an unobserved error term. In (1), leads and lags to the event of interest are defined as follows:

$$(\text{Lead } J)_{gt} = \mathbb{1}(t \leq \text{Event}_g - J) \quad (2)$$

$$(\text{Lead } j)_{gt} = \mathbb{1}(t = \text{Event}_g - j) \text{ for } j \in \{1, \dots, J-1\} \quad (3)$$

$$(\text{Lag } k)_{gt} = \mathbb{1}(t = \text{Event}_g + k) \text{ for } k \in \{1, \dots, K-1\} \quad (4)$$

$$(\text{Lag } K)_{gt} = \mathbb{1}(t \geq \text{Event}_g + K) \quad (5)$$

Leads and lags are thus binary variables indicating that the given group was a given number of periods away from the event of interest in the respective time period.  $J$

---

4. There are several ways to specify such a model. Slightly different notations are used by Schmidheiny and Siegloch (2019), who define the model as

$$y_{gt} = \sum_{j=\underline{j}}^{\bar{j}} \beta_j b_{gt}^j + \mu_g + \lambda_t + \varepsilon_{gt}$$

where

$$b_{gt}^j = \begin{cases} \mathbb{1}(t \leq \text{Event}_g + j) & \text{if } j = \underline{j} \\ \mathbb{1}(t = \text{Event}_g + j) & \text{if } \underline{j} < j < \bar{j} \\ \mathbb{1}(t \geq \text{Event}_g + j) & \text{if } j = \bar{j} \end{cases}$$

and where  $\underline{j}$  is equivalent to our definition of  $J$  and  $\bar{j}$  is equivalent to our  $K$ . Freyaldenhoven, Hansen, and Shapiro (2019) define a version of this model as

$$y_{gt} = \delta_{-K+} (1 - z_{g,t+(K-1)}) + \delta_{L+} z_{g,t-L} + \sum_{k=-(L-1)}^{K-1} \delta_{-k} \Delta z_{g,t+k} + \mu_g + \lambda_t + \varepsilon_{gt}$$

where  $z_{gt} \equiv$  Post event<sub>gt</sub> as defined in table 1,  $z_{g,t+k}$  and  $z_{g,t-k}$  refer to leads and lags of this variable, respectively, and  $\Delta$  refer to the first difference of these lead/lag terms. These models, and the one laid out in (1)–(5), are equivalent.

and  $K$  leads and lags are included, respectively, and, as indicated in (2) and (5), final leads and lags “accumulate” leads or lags beyond  $J$  and  $K$  periods. A single lead or lag variable is omitted to capture the baseline difference between groups where the event does and does not occur. In (1), as standard, this baseline omitted case is the first lead (one period prior to the reform), where  $j = 1$ .

A stylized example of such a setting is provided in table 1. We consider four groups forming a balanced panel of years from 2000–2009. The  $\text{Event}_g$  variable occurs at different times in different groups and, in the case of one group, does not occur. Here both four leads and four lags are included, such that  $J = K = 4$ . Lead and Lag 4 (exclusively) are switched on for periods in which the “Time to event” exceeds 4 leads or lags, respectively.

Table 1. A stylized example

Group ( $g$ )	Year ( $t$ )	Event	Post event	Time to event	Lead 4	Lead 3	...	Lag 0	Lag 1	...	Lag 4
Group A	2000	2004	0	-4	1	0	...	0	0	...	0
Group A	2001	2004	0	-3	0	1	...	0	0	...	0
Group A	2002	2004	0	-2	0	0	...	0	0	...	0
Group A	2003	2004	0	-1	0	0	...	0	0	...	0
Group A	2004	2004	1	0	0	0	...	1	0	...	0
Group A	2005	2004	1	1	0	0	...	0	1	...	0
Group A	2006	2004	1	2	0	0	...	0	0	...	0
Group A	2007	2004	1	3	0	0	...	0	0	...	0
Group A	2008	2004	1	4	0	0	...	0	0	...	1
Group A	2009	2004	1	5	0	0	...	0	0	...	1
Group B	2000	2005	0	-5	1	0	...	0	0	...	0
Group B	2001	2005	0	-4	1	0	...	0	0	...	0
Group B	2002	2005	0	-3	0	1	...	0	0	...	0
Group B	2003	2005	0	-2	0	0	...	0	0	...	0
Group B	2004	2005	0	-1	0	0	...	0	0	...	0
Group B	2005	2005	1	0	0	0	...	1	0	...	0
Group B	2006	2005	1	1	0	0	...	0	1	...	0
Group B	2007	2005	1	2	0	0	...	0	0	...	0
Group B	2008	2005	1	3	0	0	...	0	0	...	0
Group B	2009	2005	1	4	0	0	...	0	0	...	1

Continued on next page

Table 1 (continued)

Group ( <i>g</i> )	Year ( <i>t</i> )	Event	Post event	Time to event	Lead 4	Lead 3	...	Lag 0	Lag 1	...	Lag 4
Group C	2000	.	0	.	0	0	...	0	0	...	0
Group C	2001	.	0	.	0	0	...	0	0	...	0
Group C	2002	.	0	.	0	0	...	0	0	...	0
Group C	2003	.	0	.	0	0	...	0	0	...	0
Group C	2004	.	0	.	0	0	...	0	0	...	0
Group C	2005	.	0	.	0	0	...	0	0	...	0
Group C	2006	.	0	.	0	0	...	0	0	...	0
Group C	2007	.	0	.	0	0	...	0	0	...	0
Group C	2008	.	0	.	0	0	...	0	0	...	0
Group C	2009	.	0	.	0	0	...	0	0	...	0
Group D	2000	2007	0	-7	1	0	...	0	0	...	0
Group D	2001	2007	0	-6	1	0	...	0	0	...	0
Group D	2002	2007	0	-5	1	0	...	0	0	...	0
Group D	2003	2007	0	-4	1	0	...	0	0	...	0
Group D	2004	2007	0	-3	0	1	...	0	0	...	0
Group D	2005	2007	0	-2	0	0	...	0	0	...	0
Group D	2006	2007	0	-1	0	0	...	0	0	...	0
Group D	2007	2007	1	0	0	0	...	1	0	...	0
Group D	2008	2007	1	1	0	0	...	0	1	...	0
Group D	2009	2007	1	2	0	0	...	0	0	...	0

Groups in which the event never occurs (such as Group C in table 1) act as pure controls. These units have 0s in all lead and lag terms and act as the counterfactual on which the estimation of impacts is based. Differences between these pure control groups and groups which adopt the event of interest are anchored at 0 in the omitted base period, that is, the first lead in (1). Hence, leads and lags capture the difference between treated and control groups, compared with the prevailing difference in the omitted base period. Unbiased estimation of postevent treatment effects thus relies fundamentally on the so called parallel trends assumption. In the absence of treatment, it is assumed that treated and control groups would have maintained similar differences as in the baseline period. Thus, these models have been demonstrated to be underidentified, or identified only up to a linear trend, when all units adopt treatment at some point in time (Schmidheiny and Siegloch 2019; Borusyak and Jaravel 2018). Schmidheiny and Siegloch (2019) show that in this case, it is necessary to bin leads and lags beyond certain maximum lead (*J*) and lag (*K*) periods.

The panel event study is an extension of the standard two-way fixed-effects (sometimes called DD) model, where a single “Post event” indicator is included for all periods posterior to the occurrence of the event in treated groups. This is simply

$$y_{gt} = \alpha + \beta \text{Post event}_{gt} + \mu_g + \lambda_t + \mathbf{X}'_{gt} \boldsymbol{\Gamma} + \varepsilon_{gt} \quad (6)$$

where following the notation from (2)–(5), Post event<sub>gt</sub> =  $\mathbb{1}(t \geq \text{Event}_g)$ . Estimation of event specification (1) provides two key pieces of information not observable in this single-coefficient model. First, the full set of event leads allows for the inspection of parallel trends in the pretreatment period. While this does not provide evidence that the units in which the event was adopted and not adopted would have necessarily followed similar trends in the postreform period (Kahn-Lang and Lang 2020) (which is the identifying assumption of these models), if trends in treated and untreated areas were not parallel even preevent, it is unlikely that they would be parallel postevent. Second, the policy lags allow for inspection of the temporal nature of treatment effects, noting any dynamics in the appearance of effects, for example, increasing or decreasing effects over time, and whether effects are transitory or permanent.

A developing literature, including articles by de Chaisemartin and D’Haultfoeuille (2019), Callaway and Sant’Anna (2018), and Goodman-Bacon (2018), point to challenges in interpreting the estimated  $\hat{\beta}$  from two-way fixed-effects models when treatment effects are heterogeneous (across either groups or time periods). Goodman-Bacon (2018), for example, demonstrates that treatment effects that are heterogeneous in time since treatment in contexts where treatments are adopted in different time periods in different groups can result in estimates that are biased away from a weighed average of the average treatment effect on the treated, a problem that is resolved in the panel event study design. However, results from Abraham and Sun (2018) suggest that specific types of heterogeneity concerns remain even in panel event study models examined here. In particular, they note undesired weighting of treatment effects if there is heterogeneity across treatment groups in particular lead and lag terms. Other concerns exist in event study designs, such as possible inferential problems related to selective survival of models based on pretrend tests (Roth 2019). The `eventdd` command will not account for corrections raised in these particular settings, because these are inherent to empirical estimation of panel event study designs. We do note, however, that there are several alternative estimators that are complementary to panel event study designs and that should be considered as part of a complete estimation and testing procedure, such as the stacked DD procedure of Abraham and Sun (2018), sensitivity tests described in Roth (2019) and Rambachan and Roth (2020), and alternative models to account for dynamic paths of treatment effects, such as those described in de Chaisemartin and D’Haultfoeuille (2019) and Callaway and Sant’Anna (2018). As many of these have existing estimation libraries in some languages, when discussing the command syntax of `eventdd` in section 3 and examples of use in section 4, we discuss ways in which `eventdd` and its returned objects have been designed to facilitate interaction with these other commands.

## 2.2 Inference

A standard inference concern where policies are assigned by some group such as a state, and outcomes are followed over time within these groups, is related to potential serial-correlation in the outcome variable over time (Bertrand, Duflo, and Mullainathan 2004). While the derivations from Bertrand, Duflo, and Mullainathan (2004) are based

on single-coefficient models of the form of (6), the crux of the concern relates to high serial correlation in the outcome variable of interest, and relatively little change in the independent variables of interest. This setting is replicated in event study models described in (1)–(5). It is thus fundamental to account for this within-cluster correlation when conducting inference in such models.

The standard solution is to allow for within-cluster auto-correlation by using a cluster-robust variance-covariance estimator (CRVE) to estimate standard errors and CIs on regression parameters. Such an estimator is provided as standard in Stata by specifying the `vce(cluster clustvar)` option in e(class) models.<sup>5</sup> However, as has been extensively documented, standard CRVEs are only asymptotically valid, where the asymptotic behavior depends on the number of clusters (or groups)  $G \rightarrow \infty$  (see, for example, the comprehensive review in Cameron and Miller [2015]). When standard clustering is used based on “too few” clusters, the CRVE is generally downward biased, resulting in overrejection of null hypotheses. This bias can be severe (Cameron and Miller 2015; MacKinnon and Webb 2018).

In practice, knowing how many clusters is “too few” depends on several factors. While rules of thumb such as the rule of 42 are laid out in Angrist and Pischke (2009), who suggest that standard clustering provides a good approximation if  $G \geq 42$  clusters, the performance of these methods under simulation has been shown to depend also on the relative size of clusters (MacKinnon and Webb 2017). A range of results surveyed in Cameron and Miller (2015) leads to their suggestion that if one is analyzing data with fewer than 50 clusters in a group-year panel (such as the case with panel event studies), alternative inference methods should be considered.

In this case where the quasiexperimental setup is based on fewer than about 50 clusters, the wild cluster bootstrap has been documented to be a successful resampling-based method to account for autocorrelation in variables underlying panel event studies, even in cases with fewer clusters (see, for example, Cameron, Gelbach, and Miller [2008]; Cameron and Miller [2015]; Roodman et al. [2019]). This has been efficiently implemented in Stata as described in Roodman et al. (2019) and programmed for Stata as `boottest` (Roodman 2015). Finally, note that in the case of very few clusters, and in particular few clusters where an event occurs, inference is further complicated. In cases such as this, several potential solutions have been proposed, such as those described in MacKinnon and Webb (2018) and Conley and Taber (2011). As we lay out in the following sections, the `eventdd` command allows simple access to various inference options depending on the context of interest, including standard clustering, bootstrap, and wild

5. Denoting the matrix of  $k$  independent variables as  $\mathbf{X}$ , the dependent variable vector as  $\mathbf{y}$ , and the vector of regression coefficients as  $\beta$ , this CRVE in Stata for a linear regression model with  $N$  observations in  $G$  clusters is estimated using

$$\hat{V}_{\text{CR}}(\hat{\beta}) = (\mathbf{X}'\mathbf{X})^{-1} \left( \sum_{g=1}^G \mathbf{X}_g \tilde{\mathbf{u}}_g \tilde{\mathbf{u}}_g' \mathbf{X}_g' \right) (\mathbf{X}'\mathbf{X})^{-1}$$

Here  $\tilde{\mathbf{u}}_g = \sqrt{c} \hat{\mathbf{u}}_g$ , where  $c$  is a small sample correction  $c = \{G/(G-1)\} \times \{(N-1)/(N-k)\}$  and  $\tilde{\mathbf{u}} = \mathbf{y} - \mathbf{X}\hat{\beta}$  are standard regression residuals (Cameron, Gelbach, and Miller 2008).

cluster bootstrap in various guises based on both Stata's native CRVE procedures, as well as the community-contributed `boottest` command.

## 3 The `eventdd` command

### 3.1 Syntax

Panel event studies can be implemented in Stata using the following command syntax:

```
eventdd depvar [indepvars] [if] [in] [weight], timevar(timevar)
    [ci(type[, ...]) baseline(#) level(#) accum leads(#) lags(#) noend
    keepbal(varname) method(type[, absorb(absvars) * ...]) wboot
    wboot_op(string) balanced inrange noline graph_op(string) coef_op(string)
    endpoints_op(string) keepdummies]
```

The required `depvar` should specify the dependent variable of interest, and then `indepvars` should specify (where relevant) the optional controls, including fixed effects to be included in the panel event study model (1) but not including leads and lags, that should be entered in the regression. `pweights`, `aweights`, `fweights`, and `iweights` are allowed; see [U] 11.1.6 weight. The `method()` option specifies the estimation procedure for the underlying model and can be `ols` (ordinary least squares), `fe` (fixed effects), or `hdfe` (absorbing multiple levels of fixed effects with the community-contributed `reghdfe` command). If no estimation method is specified, `ols` is used by default. In the case of fixed-effect (`fe`) or high-dimension fixed-effect (`hdfe`) models, fixed effects can be absorbed (as discussed in the options below) and thus need not be entered in the standard `varlist` syntax. In the case of `fe` specifications, data must first be `xtset` in Stata. Based on this syntax, `eventdd` takes care of the generation of all lead and lag terms, estimation and inference, and the production of an event study plot. The `eventdd` command requires previous installation of the `matsort` (Millar 2005) command from the Statistical Software Components Archive. Examples of usage of `eventdd` are provided in section 4 of this article.

### 3.2 Options

`timevar(timevar)` is a required option. The time variable specified should contain a standardized value, where 0 corresponds to the time period in which the event of interest occurs for a given unit,  $-1$  refers to one year prior to the event, 1 refers to one year following the event, and so forth. For any units in which the event does not occur (pure controls), this variable should contain missing values.

`ci(type[, ...])` specifies the type of graph for the CIs. The types available are `rarea` for an interval with area shading (`twoway rarea`), `rcap` for an interval with capped spikes (`twoway rcap`), and `rline` for an interval with lines (`twoway rline`). Only

one type can be specified, and all intervals will be the same type. The appearance can be modified with the inclusion of any graphing option for the CIs permitted in `rarea`, `rcap`, or `rline` depending on the type of CI indicated, including area options, line options, and connect options, respectively. This does not allow the use of the general options such as titles and legends, which should be specified in the `graph_option` option. By default, standard `rcap` graphical output will be provided.

`baseline(#)` specifies the reference period for the event study, which is a baseline omitted category to which all other periods should be compared in the event study output. The default is `baseline(-1)` as in (1).

`level(#)` specifies the confidence level, as a percentage, for the CIs. The default is `level(95)` or as set by `set level`. This sets the levels for CIs in regression output, as well as the event study plot and returned matrices. This will also be passed to `boottest` if wild clustered CIs are requested.

`accum` specifies that all periods beyond some specified values should be accumulated into final points, indicated as  $J$  and  $K$  in (1). For example, if `accum` is specified and `leads(#)` and `lags(#)` are both set equal to 10, a single coefficient will be displayed in regressions and graphical output capturing 10 or more periods pre- or postreform. By default, all possible leads and lags will be included in models and graphical output.

`leads(#)` indicates the maximum amount of preevent periods to consider in the event study. This can be specified (and must be specified) only if `accum`, `keepbal()`, or `inrange` is also specified. Only integer values are permitted.

`lags(#)` indicates the maximum amount of postevent periods to consider in the event study. This can be specified (and must be specified) only if either `accum`, `keepbal()`, or `inrange` is also specified. Only integer values are permitted.

`noend` requests that accumulative endpoints not be shown on graphical output when the `accum` option is specified.

`keepbal(varname)` specifies that only units that are balanced in the panel should be kept for estimation. Here `varname` indicates the panel variable (for example, `state`) that indicates units. In this case, “balance” refers to balance over calendar time. An alternative option (`balanced`), discussed below, allows for only balanced leads and lags relative to treatment to be considered in graphical output.

`method(type[, absorb(absvars) * ...])` specifies the method of estimation for the event study model underlying graphical output. `ols` requests that the model be fit by ordinary least squares using Stata’s `regress` command, `fe` requests that the model be fit by fixed-effects (within) estimation using Stata’s `xtreg`, `fe` command, and `hdfe` requests that the model be fit using the community-contributed `reghdfe` command (if installed). `*` represents any other estimation options included and permitted by `regress`, `xtreg`, or `reghdfe` that will be passed to the specified estimation command. This allows for the inclusion of clustered standard errors or other variance estimators (see [R] `vce_option`) and allows for alternative levels for CIs to

be used (see `level()`). For `ols`, unit-specific fixed effects and time-specific fixed effects must be included in the `indepvars` indicated in the command syntax. For `fe`, unit-specific fixed effects should not be included in the `indepvars` indicated, but time-specific fixed effects still need to be. Finally, for `hdfe`, the `absorb(abvars)` option should also be specified to indicate which fixed effects should be controlled in the regression (refer to `reghdfe`, if installed, for additional details), and any fixed effects indicated in `absorb(abvars)` should not be included in the `indepvars` indicated. `hdfe` cannot be used in combination with the `wboot` option. The default is `method(ols)`.

`wboot` indicates that inference in the event study plot produced by the command should be based on wild cluster bootstrapped CIs. When indicated, CIs for each lead and lag term will be calculated using a wild cluster bootstrap. This requires the community-contributed `boottest` command of Roodman (2015) (if installed). This option may not be combined with the `hdfe` estimation option.

`wboot_op(string)` allows for the inclusion of any other wild bootstrap option permitted in `boottest`, including `seed(#)` to set the seed for simulation-based calculations allowing replication of the CIs and `bootclust(varname)` to specify which variables to cluster the wild bootstrap upon. Setting the level (which is 95 by default) should be indicated in the `level()` option of the command, and this will be passed to `wboot_op()`. The `nograph` option is specified automatically when the `wboot` option is used.

`balanced` requests that only “balanced” leads and lags be plotted. This will produce a graph showing only leads and lags for which each treated unit has data, and thus, all coefficients plotted will be based on all units in the data. While only balanced leads and lags will be plotted, all units and time periods will be included in the estimation of the event study.

`inrange` requests that only the specified leads and lags be plotted. While only leads and lags indicated in `leads(#)` and `lags(#)` will be plotted, all units and time periods will be included in the estimation of the event study.

`noline` requests that the line before the event on the *x* axis not be shown on graphical output.

`graph_op(string)` allows for the inclusion of any other graphing options permitted in `twoway_options`, including `title_options`, `added_lines_options`, and `axis_label_options`. This also allows for the use of alternative labels for graph axes. By default, standard graphical output will be provided.

`coef_op(string)` allows for the inclusion of any graphing option for the coefficients permitted in `scatter`, including `marker_options` and `marker_label_options`. This does not allow for the use of the general options of `graph_op()`. By default, standard graphical output will be provided.

`endpoints_op(string)` allows for the inclusion of any graphing option for the endpoint coefficients permitted in `scatter`, including `marker_options` and `marker_label_`

*options*. This is available only if specifying the `accum` option and does not allow for the use of the general options of `graph_op()`. By default, standard graphical output will be provided.

`keepdummies` requests that the dummy variables of all leads and lags used in the estimation be included in the database. One must save the data before running the command with the `keepdummies` option (the first time this option is used), or otherwise data in memory will be lost. This option is necessary to perform joint significance tests using a wild or score bootstrap with the postestimation commands (see discussion below).

### 3.3 Stored results

`eventdd` stores the following in `e()`:

Scalars	
<code>e(baseline)</code>	baseline period specified
<code>e(level)</code>	confidence level
Macros	
<code>e(cmd)</code>	<code>eventdd</code>
<code>e(cmdline)</code>	command as typed
<code>e(depvar)</code>	name of dependent variable
<code>e(wtype)</code>	weight type
<code>e(clustvar)</code>	name of cluster variable
<code>e(vce)</code>	<i>vctype</i> specified in <code>vce()</code>
<code>e(vcetype)</code>	title used to label Std. err.
Matrices	
<code>e(b)</code>	coefficient vector
<code>e(V)</code>	variance–covariance matrix of the estimators
<code>e(leads)</code>	all event leads, their lower bound, the point estimate, and their upper bound
<code>e(lags)</code>	all event lags, their lower bound, the point estimate, and their upper bound
<code>e(V_leads_lags)</code>	variance–covariance matrix of the leads and lags estimators

Note that methods related to event study models such as that described by Rambachan and Roth (2020) rely on access to point estimates and standard errors of lead and lag terms, which are available through the matrices returned here.

### 3.4 Postestimation commands

Several postestimation commands are available after using the `eventdd` command. These are available for joint tests of leads and lags or the joint significance of all lead and lag parameters. Specifically, the below-listed postestimation commands are of special interest after `eventdd`.

Command	Description
<code>estat leads</code>	joint significance test for leads
<code>estat lags</code>	joint significance test for lags
<code>estat eventdd</code>	joint significance test for leads and lags

Unless otherwise requested, these postestimation commands conduct  $F$  tests of the joint significance of parameters. However, wild clustered bootstrap versions of the joint tests can be conducted with the following options:

Options	Description
<code>wboot</code>	joint significance test using <code>boottest</code> command; requires specifying the <code>keepdummies</code> option in <code>eventdd</code> ; <code>nograph</code> option is already specified in <code>boottest</code>
*	specify any additional options that should be passed to the joint significance test; options should be permitted by <code>test</code> or <code>boottest</code> (if specifying the <code>wboot</code> option)

`boottest` does not work after `reghdfe` with more than one set of fixed effects.

## 4 Examples based on an empirical application

We now provide several illustrations of the performance of `eventdd` to estimate the panel event study in empirical applications. We use data from Stevenson and Wolfers (2006) of the no-fault divorce reforms and female suicide in United States. These data have been used in other articles to demonstrate the functionality of recent advances in two-way fixed-effect models (see Goodman-Bacon [2018]) and are drawn from examples used in documenting such methods when used in Stata (Goodman-Bacon, Goldring, and Nichols 2019).<sup>6</sup> The data consist of a balanced panel with 49 states observed from 1964 to 1996 with different timing of unilateral divorce reforms across the states.

The specification of the baseline two-way fixed-effect DD style model of female suicide on no-fault divorce reforms used is

$$\text{asmrs}_{st} = \gamma_s + \lambda_t + \tau \text{post}_{st} + \mathbf{X}'_{st} \boldsymbol{\Gamma} + \varepsilon_{st}$$

This is the analogue of (6) applied to this case in particular. Here `asmrs` refers to the female suicide rate for all women in state  $s$  at time  $t$ ,  $\gamma_s$  is a fixed effect by state,  $\lambda_t$  is a temporal (year) fixed effect, `post` takes the value of 1 after the implementation of a no-fault divorce reform, and  $\varepsilon_{st}$  is a stochastic error. The controls ( $\mathbf{X}_{st}$ ) include per capita income (`pcinc`), homicide mortality (`asmrh`), and the Aid to Families with Dependent Children rate for a family of four (`cases`). Here  $\tau$  is the parameter that captures the average impact of unilateral divorce on suicide rate assuming standard DD parallel trends.<sup>7</sup>

6. These data are available online at [http://www.damianclarke.net/stata/bacon\\_example.dta](http://www.damianclarke.net/stata/bacon_example.dta), and can be imported directly into Stata by typing `webuse set www.damianclarke.net/stata/` and then `webuse bacon_example`.

7. In reality, as laid out in section 2, when this effect is heterogeneous over time, this will not capture the average treatment effects but rather a weighted average of pairwise comparisons (Goodman-Bacon 2018; de Chaisemartin and D'Haultfoeuille 2019; Callaway and Sant'Anna 2018).

## 4.1 Estimation of the panel event study

To estimate a panel event study specification corresponding to the no-fault divorce reform, one first creates the standardized version of the time-to-reform variable, presuming such a variable is not already available in the data. In this case in particular, the creation of the variable in Stata simply requires subtracting the reform period,  $\text{Event}_s$ , called  $\text{Event}_g$  in section 2 (and  $\_nfd$ , for “no fault divorce”, in the data), from the time period  $t$  (called `year` in the data):

```
. webuse set www.damianclarke.net/stata/
(prefix now "http://www.damianclarke.net/stata")
. webuse bacon_example
(Stevenson and Wolfers (2006) divorce example, provided by Goldring et al.)
. generate timeToTreat = year - _nfd
(429 missing values generated)
```

Note that as expected, missing values are generated for states in which the reform is not adopted at any point in this period and that act as pure controls in the panel event study. Below, you can see how the data are set for the first 10 observations, documenting the relationship between the absolute time period (`year`), the time the reform was implemented (`_nfd`), and the relative time to the reform’s implementation (`timeToTreat`):

```
. sort stfips year
. list stfips year _nfd timeToTreat in 1/10, noobs sepby(stfips) abbreviate(11)
```

stfips	year	_nfd	timeToTreat
1	1964	1971	-7
1	1965	1971	-6
1	1966	1971	-5
1	1967	1971	-4
1	1968	1971	-3
1	1969	1971	-2
1	1970	1971	-1
1	1971	1971	0
1	1972	1971	1
1	1973	1971	2

The second step is to estimate the event study, as per (1)–(5). In this example, the general form of the event study model, including all leads and lags available, is

$$\begin{aligned} \text{asmrs}_{st} = & \alpha + \beta_{21}(\text{Lead 21})_{st} + \cdots + \beta_2(\text{Lead 2})_{st} \\ & + \gamma_0(\text{Lag 0})_{st} + \cdots + \gamma_{27}(\text{Lag 27})_{st} \\ & + \mathbf{X}'_{st} \mathbf{\Gamma} + \mu_s + \lambda_t + \varepsilon_{st} \end{aligned} \quad (7)$$

where as above  $\text{asmrs}$  is the female suicide rate for all women and a series of  $J = 21$  leads and  $K = 27$  lags are considered relative to the event of interest (fully saturating the model). As is generally standard, the reference period is set as  $-1$ : the period immediately preceding the adoption of the event in each state. Fixed effects for state and time are included as  $\mu_s$  and  $\lambda_t$ , respectively.

The `eventdd` command provides a simple syntax to generate all necessary leads and lags for (7), fit the event study model, and plot point estimates and CIs. The command requires the `timevar(timeToTreat)` option to indicate the standardized “time to treatment” variable generated previously. Below, we request that the command run quietly (`quietly`); however, later in this section, we document an example where full regression output is displayed. In the following syntax, the `method(,)` option is used to pass specific options to the underlying regression command.

```
. #delimit ;
delimiter now ;
. quietly
> eventdd asmrh pcinc cases i.year i.stfips, timevar(timeToTreat)
> method( , cluster(stfips)) graph_op(ytitle("Suicides per 1m women")
> xlabel(-20(5)25));
. #delimit cr
delimiter now cr
```

The command stores all event leads, their lower bound, the point estimate, and their upper bound. For example, if we wish to visualize the estimates on the full set of leads, as well as their upper and lower CIs, we can simply examine the returned `leads` matrix:

```
. matrix list e(leads)
e(leads)[21,4]
      Lead          LB          Est          UB
r21      1          0          0          0
r20      2  -5.6533713  -.51573968  4.621892
r19      3  -8.3800554  -2.3125875  3.7548807
r18      4  -4.6691194  .22835743  5.1258345
r17      5  -8.129283  -2.7754233  2.5784359
r16      6  -6.8603821  -.75055814  5.3592658
r15      7 -10.123232  -1.2564343  7.6103635
r14      8 -10.710829  -2.7376504  5.2355275
r13      9 -12.330478  -5.0007019  2.3290732
r12     10 -11.226175  -1.1506662  8.9248428
r11     11 -17.51306  -9.3819475 -1.2508357
r10     12 -14.165311  -.04345011  14.078411
r9      13 -10.856544  -1.3885684  8.0794067
r8      14 -6.3317404  4.3279953  14.987731
r7      15 -7.7231503  .84775668  9.418664
r6      16 -8.3638954  -1.022577  6.3187418
r5      17 -17.12311  -4.4348736  8.2533636
r4      18 -10.074435  -.51595128  9.0425329
r3      19 -3.328351  8.8427277  21.013805
r2      20 -34.54834  -12.084179  10.379983
r1      21 -31.114813  -22.920727  -14.72664
```

Because we do not specify the estimation method in the `method()` option, `eventdd` uses Stata's `regress` command to fit the model by ordinary least-squares regression (if we were to specify `method(ols, cluster(stfips))`, the same result would be obtained). We can also request other estimators for the underlying event study model; if we specify the `fe` option, the model would be fit with the fixed-effects estimator.<sup>8</sup>

In the same way, we can estimate the results efficiently absorbing multiple levels of fixed effects via the `reghdfe` command by indicating `hdfe` in the `method()` option, which is quite useful when we have to control for many fixed effects. Note that in this case, the fixed effects of interest must be indicated using the `absorb()` option which is passed to the `reghdfe` command. For instance, if we wish to absorb the temporal and geographic fixed effects, the necessary syntax is as follows:

```
. #delimit ;
delimiter now ;
.eventdd asmrh pcinc cases, timevar(timeToTreat)
> method(hdfe, absorb(i.year) cluster(stfips))
> graph_op(ytitle("Suicides per 1m women") xlabel(-20(5)25));
(MWFE estimator converged in 2 iterations)
warning: missing F statistic; dropped variables due to collinearity or too few
> clusters
HDFE Linear regression
Absorbing 2 HDFE groups
Statistics robust to heteroskedasticity
Number of obs      =      1,617
F( 51,      48) =      .
Prob > F          =      .
R-squared          =      0.7212
Adj R-squared      =      0.6964
Within R-sq.       =      0.0731
Number of clusters (stfips) =      49      Root MSE      =      10.8116
(Std. err. adjusted for 49 clusters in stfips)
```

---

8. Assuming the data have been `xtset`, this would simply be

```
. #delimit ;
delimiter now ;
. quietly
> eventdd asmrh pcinc cases i.year, timevar(timeToTreat)
> method(fe, cluster(stfips)) graph_op(ytitle("Suicides per 1m women")
> xlabel(-20(5)25));
. #delimit cr
delimiter now cr
```

where note that we no longer include the state fixed effects, because these will be taken care of given the panel fixed-effect estimator (`xtreg`) used to estimate the regression.

asmrs	Coefficient	Robust std. err.	t	P> t	[95% conf. interval]
pcinc	-.0011046	.0004071	-2.71	0.009	-.0019232 -.0002861
asmrh	1.08064	.5968879	1.81	0.076	-.119483 2.280764
cases	-190.3716	134.4991	-1.42	0.163	-460.8002 80.05698
lead21	-22.92073	4.011063	-5.71	0.000	-30.98551 -14.85594
lead20	-12.08418	10.99637	-1.10	0.277	-34.19385 10.0255
lead19	8.842727	5.957829	1.48	0.144	-3.13629 20.82174
lead18	-.5159513	4.678946	-0.11	0.913	-9.923602 8.891699
lead17	-4.434874	6.210982	-0.71	0.479	-16.92289 8.053142
lead16	-1.022577	3.593627	-0.28	0.777	-8.248049 6.202895
lead15	.8477567	4.19552	0.20	0.841	-7.587901 9.283414
lead14	4.327995	5.218016	0.83	0.411	-6.163529 14.81952
lead13	-1.388568	4.634641	-0.30	0.766	-10.70714 7.930001
lead12	-.0434501	6.912751	-0.01	0.995	-13.94247 13.85557
lead11	-9.381948	3.980237	-2.36	0.023	-17.38475 -1.379145
lead10	-1.150666	4.932033	-0.23	0.817	-11.06718 8.765851
lead9	-5.000702	3.587977	-1.39	0.170	-12.21481 2.213409
lead8	-2.73765	3.902927	-0.70	0.486	-10.58501 5.10971
lead7	-1.256434	4.34036	-0.29	0.773	-9.983313 7.470445
lead6	-.7505582	2.990802	-0.25	0.803	-6.763968 5.262852
lead5	-2.775423	2.620752	-1.06	0.295	-8.044798 2.493952
lead4	.2283574	2.39735	0.10	0.925	-4.591837 5.048552
lead3	-2.312587	2.970068	-0.78	0.440	-8.28431 3.659136
lead2	-5.157397	2.514907	-0.21	0.838	-5.572299 4.54082
lag0	.2507466	2.722144	0.09	0.927	-5.222491 5.723984
lag1	-1.619351	2.941537	-0.55	0.585	-7.533709 4.295006
lag2	-1.687107	3.898178	-0.43	0.667	-9.524919 6.150706
lag3	-.7444709	2.862572	-0.26	0.796	-6.500058 5.011116
lag4	-2.956354	2.832628	-1.04	0.302	-8.651735 2.739027
lag5	-2.377841	2.75474	-0.86	0.392	-7.916617 3.160935
lag6	-3.311888	3.568157	-0.93	0.358	-10.48615 3.862372
lag7	-5.136502	3.401946	-1.51	0.138	-11.97657 1.70357
lag8	-6.991146	3.086374	-2.27	0.028	-13.19672 -.785575
lag9	-4.82321	3.089481	-1.56	0.125	-11.03503 1.388607
lag10	-8.814158	3.6746	-2.40	0.020	-16.20244 -1.42588
lag11	-7.27331	3.631759	-2.00	0.051	-14.57545 .0288304
lag12	-6.151559	4.089512	-1.50	0.139	-14.37407 2.070957
lag13	-8.276837	3.946249	-2.10	0.041	-16.2113 -.3423707
lag14	-6.593221	3.867273	-1.70	0.095	-14.3689 1.182453
lag15	-7.850839	4.070836	-1.93	0.060	-16.0358 .3341258
lag16	-7.234422	4.270836	-1.69	0.097	-15.82151 1.35267
lag17	-8.516898	4.344278	-1.96	0.056	-17.25166 .2178592
lag18	-9.991582	3.758781	-2.66	0.011	-17.54912 -2.434046
lag19	-11.53613	3.861769	-2.99	0.004	-19.30074 -3.771526
lag20	-9.219165	4.501869	-2.05	0.046	-18.27078 -.167551
lag21	-10.79088	4.417864	-2.44	0.018	-19.67359 -1.908173
lag22	-10.65478	4.608349	-2.31	0.025	-19.92049 -1.389076
lag23	-12.08658	5.29214	-2.28	0.027	-22.72714 -1.446016
lag24	-10.67796	6.147523	-1.74	0.089	-23.03838 1.682466
lag25	-10.26777	7.459044	-1.38	0.175	-25.26518 4.729644
lag26	-16.69255	10.54234	-1.58	0.120	-37.88934 4.504239
lag27	-4.4344752	8.147106	-0.05	0.958	-16.81533 15.94638
_cons	85.59069	10.90519	7.85	0.000	63.66433 107.5171

Absorbed degrees of freedom:

Absorbed FE	Categories	- Redundant	= Num.	Coefs
stfips	49	49	0	*
year	33	0	33	

\* = FE nested within cluster; treated as redundant for DoF computation

```
. #delimit cr
delimiter now cr
```

The standard command output consists of the regression output (all the above output, including the warning, comes directly from the regression estimated by `reghdfe`), and the event study lead and lag coefficients along with their CIs are plotted as in figure 1. As discussed in Stevenson and Wolfers (2006), the event study plot provides evidence of a reduction in rates of female suicide following the passage of no-fault divorce laws, with significant declines observed eight years following reform passage. We note that in this specification where all possible leads and lags are included (the default behavior of `eventdd`), we do observe several significant differences in the prereform period, in lead 11, and lead 21. Note, however, that these leads are sufficiently far from the time period of treatment that not all treated states are observed, and so these significant declines are likely due to compositional changes in these variables. We discuss this further below and limit analysis to balanced periods when discussing the `balanced` option of the command. Nevertheless, if desired, we can also formally test the joint significance of all the lead terms simultaneously with the hypothesis

$$H_0 : \beta_{21} = \beta_{20} = \cdots = \beta_2 = 0 \quad \text{versus} \quad H_1 : H_0 \text{ does not hold}$$

This can be simply assessed postestimation using one of the postestimation commands designed for use with `eventdd`:

```
. estat leads

Joint significance test for leads

F-stat:          32.1312
P-value:         0.0000

Degrees of freedom      (20,48)
```

Similar such postestimation commands exist to test the joint significance of the postimplementation coefficients (`estat lags`) or both the lead and lag terms in a single sequence (`estat eventdd`).

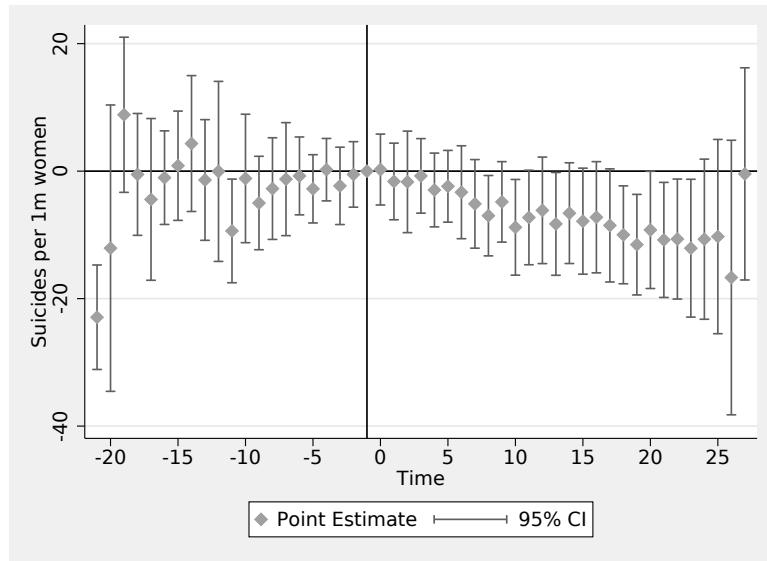


Figure 1. Event study example based on no-fault divorce reforms.

NOTES: Event study model follows the no-fault divorce analysis described in (Stevenson and Wolfers 2006), and replication/extension of Goodman-Bacon (2018). Point estimates are displayed along with their 95% CIs as described in (7). The baseline (omitted) base period is one year prior to the adoption of the reform in each reforming state, indicated by the vertical line in the plot.

This “fully saturated” model where all possible leads and lags are plotted is the default output in the `eventdd` command. However, many alternative estimation procedures are permitted and indeed are likely preferred, for example to avoid the behavior observed above where leads and lags far from treatment will not be balanced given that only states adopting in certain early or late time periods will be observed in these lead/lag terms. Here we discuss several such alternatives, documenting their syntax in the `eventdd` command. Graphical output in each case is summarized in figure 2.

**Limiting visualized leads and lags.** It may be a matter of interest to show only some lead/lag periods in the plot. For example, one such case discussed below relates to plotting only those lead/lag terms in which each treated state is observed. Generically, the `inrange` option allows for specifying that only certain coefficients and CIs should be included in the plot. We note here that in this case, the underlying regression model *will* include all periods as in the first case, and thus, these lead/lag terms will simply correspond to a restricted range from figure 1. For instance, if we want to show only the results between the time periods  $-10$  and  $10$ , the command will be

```

. #delimit ;
delimiter now ;
. quietly
> eventdd asmrh cases i.year, timevar(timeToTreat) inrange leads(10)
> lags(10) method(fe, cluster(stfips)) graph_op(ytitle("Suicides per 1m women"));
. #delimit cr
delimiter now cr

```

The output in this case is displayed in figure 2a. A special case of plotting limited leads/lags consists of the case in which one wishes to show only coefficients and CIs for which all states have a lead and lag term. We refer to this as a balanced plot, which can be produced quite simply using the `balanced` option. In this case, while all leads and lags are included in the underlying panel event study model, and only certain periods are plotted on the graph (like `inrange`), we do not need to know *a priori* which periods are balanced, because `eventdd` automatically identifies them. As figure 2b shows, in our case the balanced periods comprise periods between 5 years prereform and 11 years postreform.<sup>9</sup> In this case, the syntax simply requires indicating the `balanced` option:

```

. #delimit ;
delimiter now ;
. quietly
> eventdd asmrh cases i.year, timevar(timeToTreat) balanced
> method(fe, cluster(stfips)) graph_op(ytitle("Suicides per 1m women"));
. #delimit cr
delimiter now cr

```

**Restricting samples or accumulating leads/lags.** In contrast to simply focusing on particular coefficients in the unaltered baseline model, one may wish to work with particular subsamples that meet inclusion criteria, or accumulate leads and lags into periods that exceed some defined time, as an alternative way to avoid unbalanced leads and lags, as well as to avoid problems related to underidentification where all units are treated (Schmidheiny and Siegloch 2019). Consider the case where we wish to include 15 leads and 10 lags but to fit only the model with units that effectively have data for each of these periods. In the case of these data in particular that are yearly from 1964–1996, any units adopting no fault divorce reform between 1978 and 1996 will have (at least) 15 leads and 10 lags. Units adopting prior to 1978 will have fewer than 15 leads, and units adopting after 1996 will have fewer than 10 lags. To implement an estimation based on a balanced panel of observations with these lead/lag terms, one can use the `keepbal(varname)` option, where `varname` indicates the panel unit over which balance should be applied (`stfips` in this case where the treatment unit is states). It is additionally necessary to explicitly indicate the period of interest for plotting within the balanced panel, for instance, `leads(15)` and `lags(10)`. This is all implemented in the command below.

9. This can also be seen quite simply by tabulating (`tab`) the time variable used to indicate time to treatment. For example, in the case of the `timeToTreat` variable indicating relative time to the arrival of the divorce reform, we observe that for each of leads –5 up until lag 11, all periods are observed in the 36 states in which the event occurs. However, earlier leads and later lags are observed only in a subsample of states given that for others, these leads (lags) are earlier (later) than the first (last) year in the data sample.

```

. #delimit ;
delimiter now ;
. eventdd asmr pcinc asmrh cases, timevar(timeToTreat) keepbal(stfips)
> leads(15) lags(10) method(hdfe, absorb(i.stfips i.year) cluster(stfips))
> graph_op(ytitle("Suicides per 1m women"));
(MWFE estimator converged in 5 iterations)
warning: missing F statistic; dropped variables due to collinearity or too few
> clusters

HDFE Linear regression                               Number of obs      =      507
Absorbing 2 HDFE groups                           F(     28,      15) =      .
Statistics robust to heteroskedasticity           Prob > F        =
                                                R-squared        =      0.4758
                                                Adj R-squared   =      0.3831
                                                Within R-sq.   =      0.1416
Number of clusters (stfips) =          16          Root MSE        =      10.1003
                                                (Std. err. adjusted for 16 clusters in stfips)

```

asmrs	Coefficient	Robust std. err.	t	P> t	[95% conf. interval]	
pcinc	-.0019977	.0007185	-2.78	0.014	-.0035291	-.0004663
asmrh	2.287467	1.540722	1.48	0.158	-.996504	5.571437
cases	-390.4487	180.673	-2.16	0.047	-.775.544	-5.353367
lead15	1.292386	6.839513	0.19	0.853	-13.28569	15.87046
lead14	4.959406	8.015248	0.62	0.545	-12.12469	22.0435
lead13	4.003705	9.322968	0.43	0.674	-15.86773	23.87514
lead12	-3.997656	5.732168	-0.70	0.496	-16.21548	8.22017
lead11	-2.629473	4.276906	-0.61	0.548	-11.74548	6.486536
lead10	-3.987131	5.759912	-0.69	0.499	-16.26409	8.289831
lead9	-7.773707	4.653677	-1.67	0.116	-17.69278	2.145371
lead8	-6.673529	6.393574	-1.04	0.313	-20.30111	6.954051
lead7	-3.861924	4.692105	-0.82	0.423	-13.86291	6.13906
lead6	-6.44329	3.524894	-1.83	0.088	-13.95642	1.069845
lead5	-7.389526	3.375333	-2.19	0.045	-14.58388	-1.951743
lead4	-11.18006	6.114788	-1.83	0.087	-24.21342	1.853305
lead3	-2.146437	5.282028	-0.41	0.690	-13.40481	9.111939
lead2	-2.254288	3.599554	-0.63	0.541	-9.926557	5.41798
lag0	-3.458685	2.757283	-1.25	0.229	-9.335696	2.418325
lag1	-2.050697	3.988558	-0.51	0.615	-10.55211	6.450712
lag2	3.8915	4.477937	0.87	0.399	-5.652997	13.436
lag3	-5.949	5.387559	-1.10	0.287	-17.43231	5.53431
lag4	1.151061	3.715023	0.31	0.761	-6.767323	9.069445
lag5	4.283462	6.478434	0.66	0.519	-9.524993	18.09192
lag6	-3.08539	2.849672	-1.08	0.296	-9.159321	2.988541
lag7	1.776614	7.27749	0.24	0.810	-13.73499	17.28822
lag8	.7996726	5.648138	0.14	0.889	-11.23905	12.83839
lag9	-1.215031	4.864573	-0.25	0.806	-11.58362	9.15356
lag10	-2.948988	4.776411	-0.62	0.546	-13.12967	7.231691
_cons	98.91732	14.70189	6.73	0.000	67.58098	130.2536

Absorbed degrees of freedom:

Absorbed FE	Categories	- Redundant	= Num. Coefs	
stfips	16	16	0	*
year	33	0	33	

\* = FE nested within cluster; treated as redundant for DoF computation

```
. #delimit cr
delimiter now cr
```

Given that we now restrict to only certain states based on their period of adoption (as well as nonadopting states), the lead and lag estimates will differ from those from the fully saturated model discussed previously. In the output of the above command, we observe that the estimation sample consists only of 507 observations for adopting states with balance in the indicated leads/lags, as well as states that do not adopt (versus 1,617 observations in the full-sample specification). The corresponding event study plot is presented in figure 2c, where we note that the considerable change in estimation sample (chosen simply for expositional reasons) produces quite different results.

An alternative way to work with the imbalance in standardized time periods is to stipulate that all periods beyond some specified values should be accumulated into final lead and lag points, as indicated in (2) and (5). This is implemented with the `accum` option. When this is specified, the panel event study is provided based on the number of leads and lags indicated in the `leads(#)` and `lags(#)` option, respectively, accumulating all periods beyond these periods into the final lead and lag term. For instance, if we specify `leads(15)` and `lags(10)`, a single coefficient will capture the period -15 and earlier and the period 10 and later. This is illustrated in the following syntax, with the resulting graphical output presented in figure 2d.

```
. #delimit ;
delimiter now ;
. quietly
> eventdd asmrh pcinc cases i.year, timevar(timeToTreat) accum leads(15)
> lags(10) method(fe, cluster(stfips)) graph_op(ytitle("Suicides per 1m women"));
. #delimit cr
delimiter now cr
```

Because these endpoints have a different interpretation to additional leads and lags, acting as an estimate of long-term impacts of the event for all periods beyond intermediate leads/lags, by default the endpoint estimates will be plotted in an alternative color. This behavior can be controlled fully using the `endpoints_op()` options, allowing for options such as marker styles and colors to be passed to the underlying scatterplot (additional discussion is provided in section 4.3 of this article). Alternatively, as documented below, the `noend` option can be invoked, which will omit these final accumulative endpoints from graphical output, as shown in figure 2e:

```

. #delimit ;
delimiter now ;
. quietly
> eventdd asmrh cases i.year, timevar(timeToTreat) accum leads(15)
> lags(10) noend method(fe, cluster(stfips))
> graph_op( ytitle("Suicides per 1m women"));
. #delimit cr
delimiter now cr

```

Finally, as discussed in section 2, the reference period for any estimated panel event study will be assumed to be the period immediately prior to the occurrence of the event in each state, unless otherwise indicated. This can be simply changed via the `baseline(#)` option. While the choice of `-1` as the baseline period is arbitrary, it is frequently adopted, and so alternative baseline periods should be based on some empirical or theoretical consideration, although both models will be equivalent up to a single constant shift. Below we provide the syntax setting an alternative baseline period, with all coefficients and standard errors referring to differences relative to 11 years prior to the event of interest. By default, the `eventdd` command places a vertical reference line at period `-1` to visually indicate the period immediately prior to the passage of the event. However, if this reference line is not desired, the `noline` option can be specified, as documented in figure 2f. If one wishes to provide alternative reference periods, these can be passed directly to the graphing command. For example, to add an alternative reference line in period 0, one should specify `graph_op(xline(0))`.

```

. #delimit ;
delimiter now ;
. quietly
> eventdd asmrh cases i.year, timevar(timeToTreat) ci(rarea)
> method(fe, cluster(stfips)) graph_op(ytitle("Suicides per 1m women")
> xlabel(-20(5)25));
. #delimit cr
delimiter now cr

```

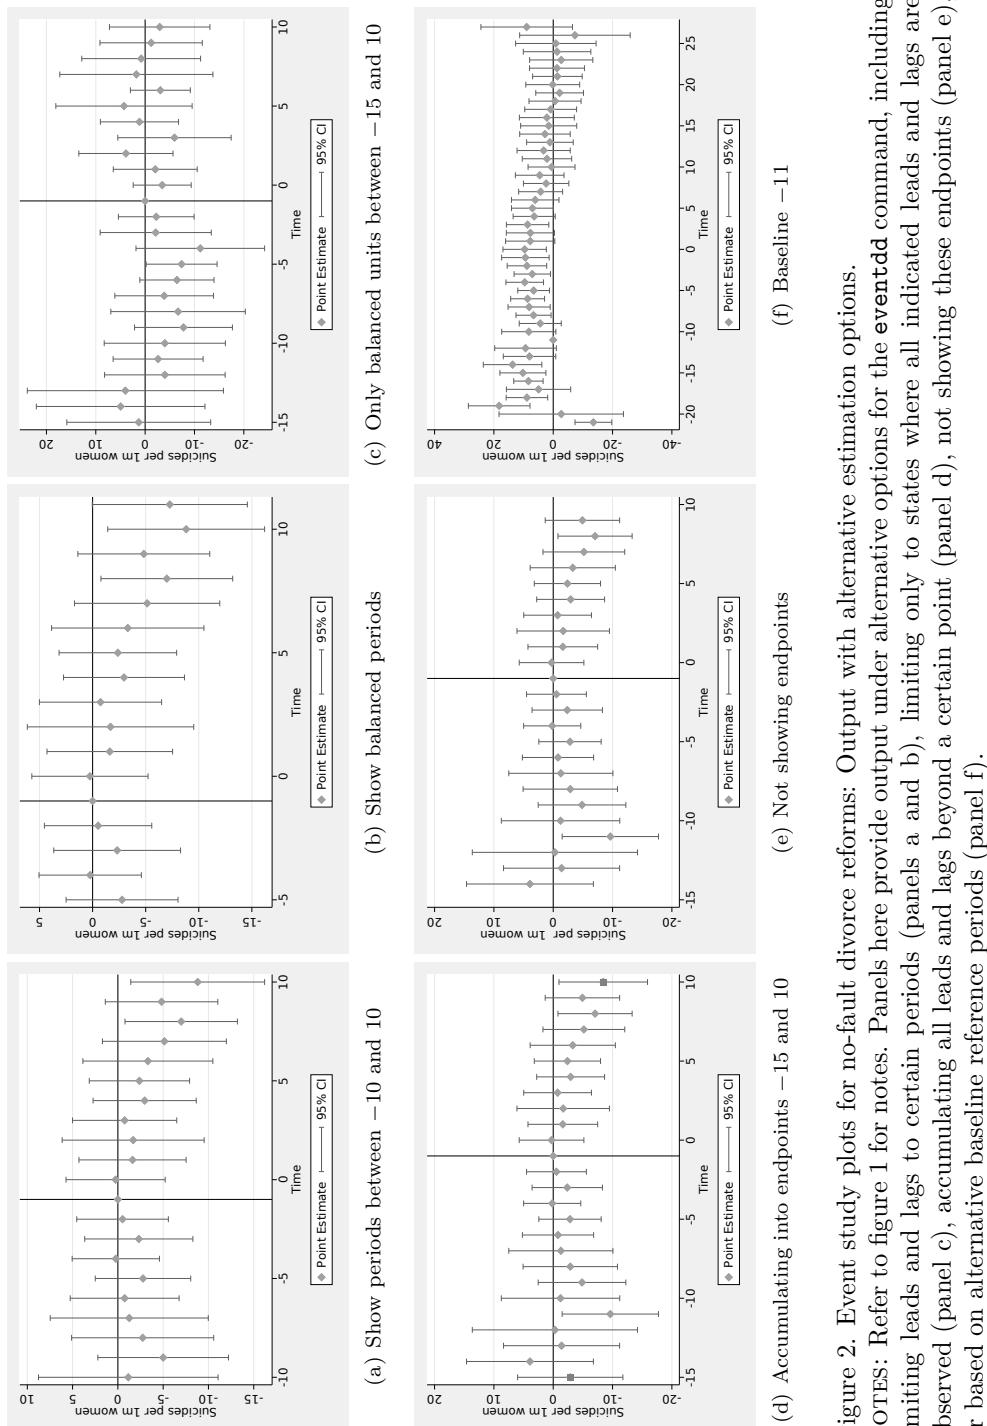


Figure 2. Event study plots for no-fault divorce reforms: Output with alternative estimation options.

NOTES: Refer to figure 1 for notes. Panels here provide output under alternative options for the `eventdd` command, including limiting leads and lags to certain periods (Panels a and b), limiting only to states where all indicated leads and lags are observed (panel c), accumulating all leads and lags beyond a certain point (panel d), not showing these endpoints (panel e), or based on alternative baseline reference periods (panel f).

## 4.2 Inference options

The previous subsection describes several alternative estimation procedures that are potentially of relevance in the estimation of a panel event study design. However, as discussed in section 2 of this article, there are several inference considerations that must be weighed when implementing a panel event study model. Until now, the command has always been implemented with `cluster(stfips)`, indicating that a CRVE should be estimated, where clusters are based at the level of the state—the level at which the event is assigned in this case. As discussed in section 2.2, in this example based on 49 states, and hence 49 clusters, a CRVE is likely the appropriate inference mode for this model.

However, the `eventdd` command allows for inference using a wild clustered bootstrap as a postestimation procedure, via its interaction with the `boottest` command (provided this command is installed on the user's system). This is indicated by the `wboot` option, which by default assumes that a clustered wild bootstrap is desired, with the cluster variable indicated in the `cluster()` option. This is especially useful when there are few clusters in the panel. However, note that given that this procedure is based around bootstrap resampling, the inference procedure likely will take longer than inference based on Stata's native CRVE and, additionally, that the `wboot` option may not be combined with the `hdf` estimation option. However, `boottest` offers considerable other benefits, including the option to undertake inference with two-way clustering, which may exhibit preferable size properties in the case of very few clusters (MacKinnon and Webb 2018). Any option that should be passed directly to `boottest` can be indicated in the `wboot_op()` option, as illustrated with the `seed()` option below, ensuring replicability in pseudo-random bootstrap resamples if desired. Figure 3 contrasts the differences between the previous CRVE-based inference procedure with the wild cluster bootstrap inference procedure illustrated here.

```

. #delimit ;
delimiter now ;
. quietly
> eventdd asmr pcinc asmrh cases i.year, timevar(timeToTreatment) accum leads(10)
> lags(10) method(fe, cluster(stfips)) wboot wboot_op(seed(1303))
> graph_op(ytitle("Suicides per 1m women"));
. #delimit cr
delimiter now cr

```

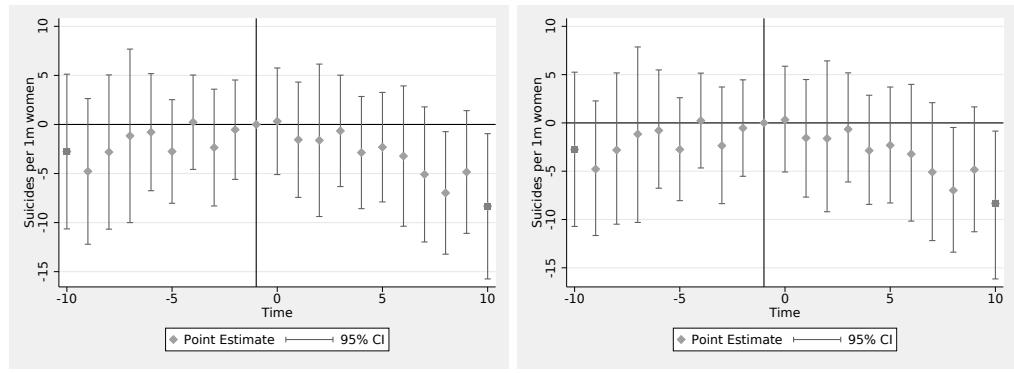


Figure 3. Visualizing alternative inference procedures for event study models

Finally, note that as standard, `eventdd` provides 95% CIs in the command's output, returned objects, and the resulting graph and legend. The `level()` option (which should be specified as a suboption to `method()`) allows for alternative levels to be indicated, where, for example, 90% CIs are requested below. Graphical output differs only in the CIs provided [figure 4a versus b].

```

. #delimit ;
delimiter now ;
. quietly
> eventdd asmr pcinc asmrh cases i.year, timevar(timeToTreat) accum leads(15)
> lags(10) method(fe, cluster(stfips) level(90))
> graph_op(ytitle("Suicides per 1m women"));
. #delimit cr
delimiter now cr

```

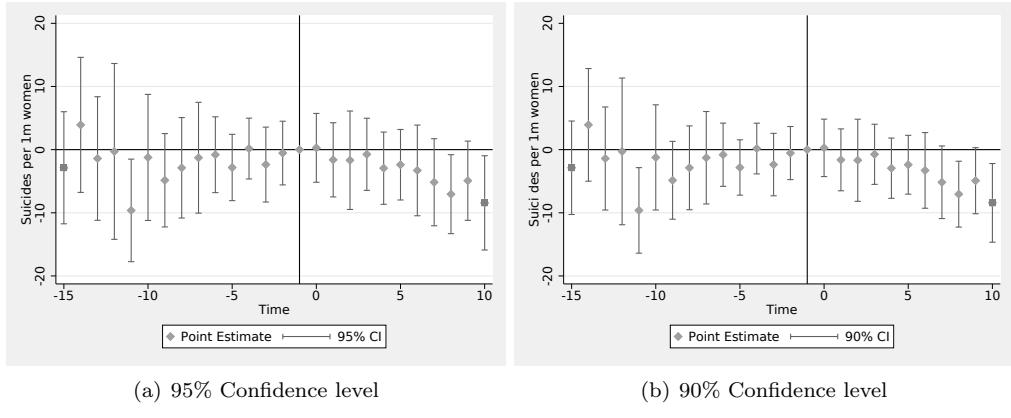


Figure 4. Default event study plots with alternative CIs

### 4.3 Altering standard appearance

`eventdd` allows for several ways to visualize the CIs using a range of Stata's standard `twoway` graph types. The command requires that the user specify one of the following types of CIs by specifying `ci(rarea)` for an interval with area shading, `ci(rcap)` for an interval with capped peaks, and `ci(rline)` for an interval with lines. Figure 5 shows the initial event study from figure 1, however now with the three alternative types of plots available. If a `ci()` type is specified, this will apply for all intervals displayed. By default, an `rcap` plot is provided.

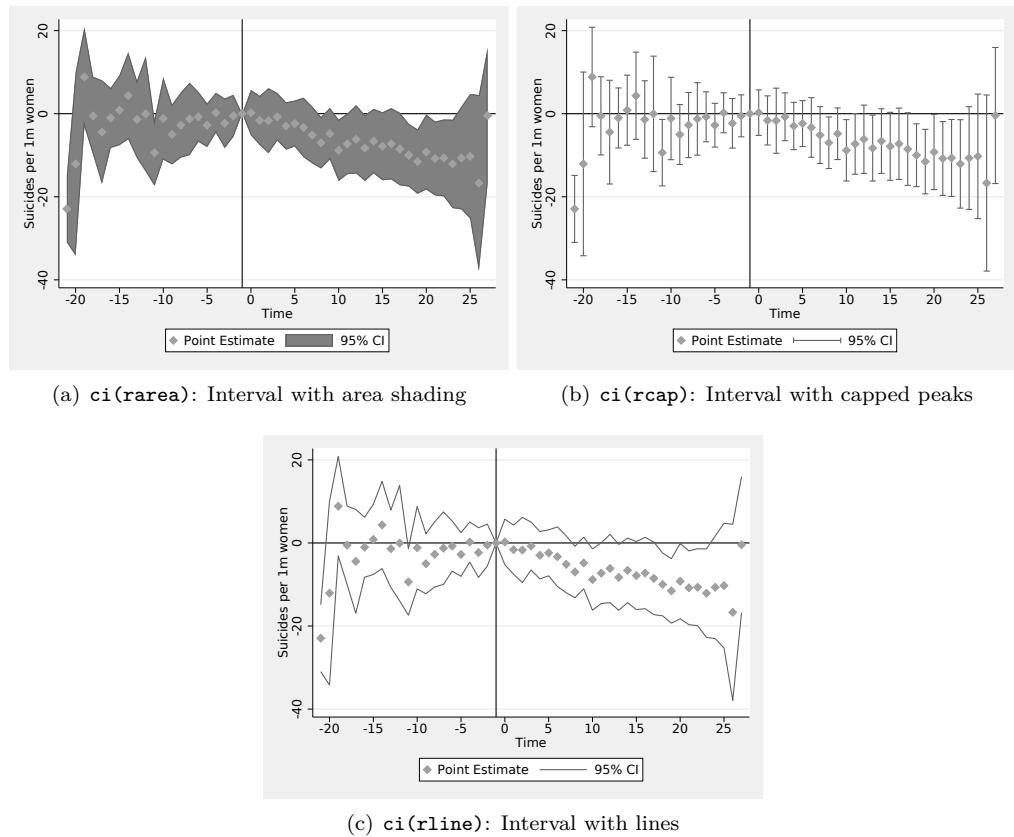


Figure 5. Alternative visualization options for event study CIs

These graph types can be fully controlled using suboptions within the `ci()` option (for example `ci(rline, lcolor(black))` to specify `lcolor()`), though the suboptions included must be compatible with the actual type of CI requested. The compatibility of options can be confirmed in Stata's help files for `twoway rcap`, `twoway rarea`, or `twoway rline` for each of the accepted `ci()` options. Similarly, we can specify any options desired for the graphing of the coefficients in the plot with the `coef_op()` option, and if we are accumulating periods into final points, we can specify graphing options for these points in `endpoints_op()`. In both cases, these accept any valid options for Stata's `twoway scatter` plot type. Finally, a `graph_op()` option allows for the inclusion of any general graphing options, such as alternative labeling schemes, graph schemes, or title options. In figure 6, we compare a standard output (left) with an alternative output (right), taking advantage of Stata's transparency options and alternative color schemes. The `eventdd` syntax used to generate figure 6(b) is provided below, followed by the resulting output.

```

. #delimit ;
delimiter now ;
. quietly
> eventdd asmrh pcinc asmrh cases i.year, timevar(timeToTreat) accum leads(15)
> lags(20) method(fe, cluster(stfips)) ci(rarea, fcolor(ltblue%45))
> graph_op(xlabel(-15 "{&le} -15" -10 "-10" -5 "-5" 0 "0" 5 "5" 10 "10" 15 "15"
> 20 "{&le} 20") scheme(s1mono) ytitle("Suicides per 1m women"))
> coef_op(msymbol(Oh)) endpoints_op(msymbol(0));
. #delimit cr
delimiter now cr

```

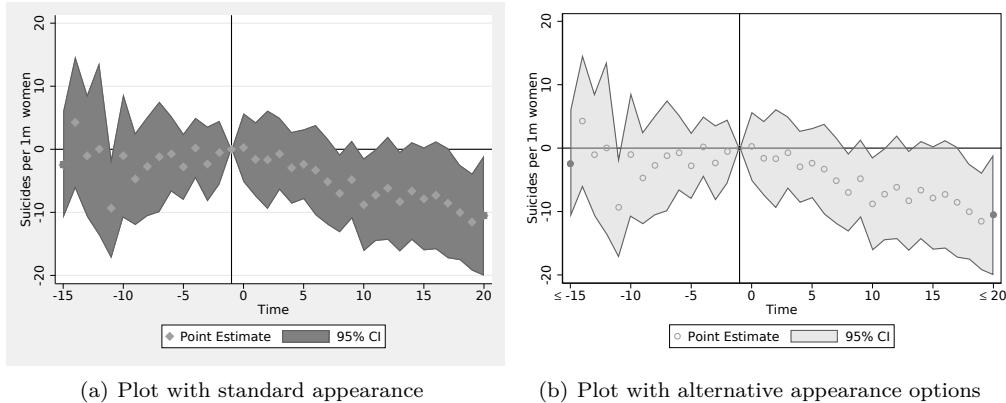


Figure 6. Event study plots no-fault divorce reforms: Appearance options

## 5 Conclusions

The panel event study is an increasingly frequently used tool in the applied analysts' toolbox. It allows for the clear presentation of estimated impacts in quasiexperimental (observational) contexts, when one wishes to consider the impact of some event that occurs at (potentially) different times in different geographical areas. What's more, while the discussion and examples provided in this article are structured around geographical clustering of events (such as the application of divorce reforms studied in Stevenson and Wolfers [2006]) and applied to demonstrate other two-way fixed-effects methods (Goodman-Bacon 2018), this setting can similarly be applied where there is the temporal arrival of some event of interest in other dimensions, such as by age or other demographic groups.

In this article, we discussed a growing literature laying out panel event study designs and introduced a flexible command, `eventdd`, that allows for their estimation and visualization in Stata. We introduced several estimation and inference concerns and showed how the command can simply deal with such concerns in an applied setting.

While `eventdd` can be based on Stata's native commands such as `regress` or `xtreg` and CRVES, it can also interact with several extremely powerful community-contributed commands, allowing for extensions such as the efficient estimation of high-dimensional fixed-effects equations, and the use of a wild cluster bootstrap for inference.

## 6 Acknowledgments

We are grateful to an anonymous referee for very useful suggestions related to command syntax and structure. The authors acknowledge the financial support of the Universidad de Santiago de Chile and the Millennium Nucleus for the Study of the Life Course and Vulnerability, funded by the Ministry of Economics of the Government of Chile.

## 7 Programs and supplemental materials

To install a snapshot of the corresponding software files as they existed at the time of publication of this article, type

```
. net sj 21-4
. net install st0655      (to install program files, if available)
. net get st0655         (to install ancillary files, if available)
```

## 8 References

Abraham, S., and L. Sun. 2018. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. ArXiv Working Paper No. arXiv:1804.05785v1. <http://arxiv.org/abs/1804.05785v1>.

Angrist, J. D., and J.-S. Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton, NJ: Princeton University Press.

Athey, S., and G. W. Imbens. Forthcoming. Design-based analysis in difference-in-differences settings with staggered adoption. *Journal of Econometrics*. <https://doi.org/10.1016/j.jeconom.2020.10.012>.

Bailey, M. J., O. Malkova, and Z. M. McLaren. Forthcoming. Does access to family planning increase children's opportunities? Evidence from the war on poverty and the early years of title X. *Journal of Human Resources*. <https://doi.org/10.3368/jhr.55.1.1216-8401R1>.

Bertrand, M., E. Duflo, and S. Mullainathan. 2004. How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics* 119: 249–275. <https://doi.org/10.1162/003355304772839588>.

Borusyak, K., and X. Jaravel. 2018. Revisiting event study designs, with an application to the estimation of the marginal propensity to consume. [https://scholar.harvard.edu/files/borusyak/files/event\\_studies\\_may8\\_website.pdf](https://scholar.harvard.edu/files/borusyak/files/event_studies_may8_website.pdf).

Callaway, B., and P. H. C. Sant'Anna. 2018. Difference-in-differences with multiple time periods and an application on the minimum wage and employment. DETU Working Papers 1804, Department of Economics, Temple University. <https://ideas.repec.org/p/tem/wpaper/1804.html>.

Cameron, A. C., J. B. Gelbach, and D. L. Miller. 2008. Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics* 90: 414–427. <https://doi.org/10.1162/rest.90.3.414>.

Cameron, A. C., and D. L. Miller. 2015. A practitioner's guide to cluster-robust inference. *Journal of Human Resources* 50: 317–372. <https://doi.org/10.3386/jhr.50.2.317>.

Clarke, D., and K. Tapia-Schythe. 2020. eventdd: Stata module to panel event study models and generate event study plots. Statistical Software Components S458737, Department of Economics, Boston College. <https://ideas.repec.org/c/boc/bocode/s458737.html>.

Conley, T. G., and C. R. Taber. 2011. Inference with “difference in differences” with a small number of policy changes. *Review of Economics and Statistics* 93: 113–125. [https://doi.org/10.1162/REST\\_a\\_00049](https://doi.org/10.1162/REST_a_00049).

Correia, S. 2014. reghdfe: Stata module to perform linear or instrumental-variable regression absorbing any number of high-dimensional fixed effects. Statistical Software Components S457874, Department of Economics, Boston College. <https://ideas.repec.org/c/boc/bocode/s457874.html>.

———. 2016. Linear models with high-dimensional fixed effects: An efficient and feasible estimator. <http://scorreia.com/research/hdfe.pdf>.

de Chaisemartin, C., and X. D'Haultfoeuille. 2019. Two-way fixed effects estimators with heterogeneous treatment effects. NBER Working Paper No. 25904, The National Bureau of Economic Research. <https://doi.org/10.3386/w25904>.

Dimitrovová, K., J. Perelman, and M. Serrano-Alarcón. 2020. Effect of a national primary care reform on avoidable hospital admissions (2000–2015): A difference-in-difference analysis. *Social Science & Medicine* 252: 112908. <https://doi.org/10.1016/j.socscimed.2020.112908>.

Freyaldenhoven, S., C. Hansen, and J. M. Shapiro. 2019. Pre-event trends in the panel event-study design. *American Economic Review* 109: 3307–3338. <https://doi.org/10.1257/aer.20180609>.

Goodman-Bacon, A. 2018. Difference-in-differences with variation in treatment timing. NBER Working Paper No. 25018, The National Bureau of Economic Research. <https://doi.org/10.3386/w25018>.

Goodman-Bacon, A., T. Goldring, and A. Nichols. 2019. bacondecomp: Stata module to perform a Bacon decomposition of difference-in-differences estimation. Statistical Software Components S458676, Department of Economics, Boston College. <https://ideas.repec.org/c/boc/bocode/s458676.html>.

Kahn-Lang, A., and K. Lang. 2020. The promise and pitfalls of differences-in-differences: Reflections on *16 and Pregnant* and other applications. *Journal of Business & Economic Statistics* 38: 613–620. <https://doi.org/10.1080/07350015.2018.1546591>.

MacKinnon, J. G., and M. D. Webb. 2017. Wild bootstrap inference for wildly different cluster sizes. *Journal of Applied Econometrics* 32: 233–254. <https://doi.org/10.1002/jae.2508>.

———. 2018. The wild bootstrap for few (treated) clusters. *Econometrics Journal* 21: 114–135. <https://doi.org/10.1111/ectj.12107>.

Millar, P. 2005. matsort: Stata module to sort a matrix by a given column. Statistical Software Components S449504, Department of Economics, Boston College. <https://ideas.repec.org/c/boc/bocode/s449504.html>.

Pacicco, F., L. Vena, and A. Venegoni. 2018. Event study estimations using Stata: The estudy command. *Stata Journal* 18: 461–476. <https://doi.org/10.1177/1536867X1801800211>.

Rambachan, A., and J. Roth. 2020. An honest approach to parallel trends. Working Paper, Harvard University. [https://scholar.harvard.edu/files/jroth/files/honestparalleltrends\\_main.pdf](https://scholar.harvard.edu/files/jroth/files/honestparalleltrends_main.pdf).

Roodman, D. 2015. boottest: Stata module to provide fast execution of the wild bootstrap with null imposed. Statistical Software Components S458121, Department of Economics, Boston College. <https://ideas.repec.org/c/boc/bocode/s458121.html>.

Roodman, D., M. Ø. Nielsen, J. G. MacKinnon, and M. D. Webb. 2019. Fast and wild: Bootstrap inference in Stata using boottest. *Stata Journal* 19: 4–60. <https://doi.org/10.1177/1536867X19830877>.

Roth, J. 2019. Pre-test with caution: Event-study estimates after testing for parallel trends. Working Paper, Harvard University. [https://scholar.harvard.edu/files/jroth/files/roth\\_pretrends\\_20190730.pdf](https://scholar.harvard.edu/files/jroth/files/roth_pretrends_20190730.pdf).

Schmidheiny, K., and S. Siegloch. 2019. On event study designs and distributed-lag models: Equivalence, generalization and practical implications. IZA Discussion Paper No. 12079, Institute of Labor Economics (IZA). <http://ftp.iza.org/dp12079.pdf>.

Stevenson, B., and J. Wolfers. 2006. Bargaining in the shadow of the law: Divorce laws and family distress. *Quarterly Journal of Economics* 121: 267–288. <https://doi.org/10.1093/qje/121.1.267>.

Suhonen, T., and H. Karhunen. 2019. The intergenerational effects of parental higher education: Evidence from changes in university accessibility. *Journal of Public Economics* 176: 195–217. <https://doi.org/10.1016/j.jpubeco.2019.07.001>.

Venkataramani, A. S., E. F. Bair, R. L. O'Brien, and A. C. Tsai. 2020. Association between automotive assembly plant closures and opioid overdose mortality in the United States: A difference-in-differences analysis. *JAMA Internal Medicine* 180: 254–262. <https://doi.org/10.1001/jamainternmed.2019.5686>.

**About the authors**

Damian Clarke is an associate professor at the Department of Economics of the Universidad de Chile, a research fellow at IZA, and affiliated with the Millennium Institute of Market Imperfections and Public Policy.

Kathya Tapia-Schythe is a doctoral student in economics at the University of California, Davis and associated researcher with the Millennium Nucleus for the Study of the Life Course and Vulnerability.