

On Event Studies and Distributed-Lags in Two-Way Fixed Effects Models: Identification, Equivalence, and Generalization*

KURT SCHMIDHEINY

SEBASTIAN SIEGLOCH

February 2, 2023

Abstract

We discuss properties and pitfalls of panel-data event study designs. We derive three main results. First, assuming constant treatment effects before and/or after some event time, also known as binning, is a natural restriction imposed on theoretically infinite effect windows. Binning identifies dynamic treatment effects in the absence of never-treated units and is particularly suitable in case of multiple events. Second, event study designs with binned endpoints and distributed-lag models are numerically identical leading to the same parameter estimates after correct reparametrization. Third, classic dummy variable event study designs can be generalized to models that account for multiple treatments of different signs and varying intensities, which are common in public and labor economics. We demonstrate the practical relevance of our methodological points in an application studying the effects of unemployment benefit duration on job search effort.

Keywords: event study, distributed-lag, applied microeconomics, credibility revolution

JEL codes: C23, C51, H00, J08

* Kurt Schmidheiny (kurt.schmidheiny@unibas.ch) is affiliated with the University of Basel, CEPR and CESifo; Sebastian Siegloch (siegloch@wiso.uni-koeln.de) is affiliated with the University of Cologne, CEPR, ZEW, IZA, and CESifo. We thank Samara Gunter, Justin McCrary, Jesse Shapiro, Divya Singh, Juan Carlos Suárez Serrato, Tony Strittmatter, and seminar participants at IEB Barcelona, Basel, Calgary, Mannheim, St. Gallen, Würzburg, IHS Vienna for providing valuable comments. We also thank Jakob Miethe (LMU Munich) for collaborating on programming the R code. Mona Köhler and Tim Bayer provided excellent research assistance. Both authors are grateful to the University of California at Berkeley – in particular Enrico Moretti and Emmanuel Saez – for their hospitality during the Academic Year 2017/2018 when this project was started. Siegloch is thankful to the German Research Foundation DFG for financing the research stay under the Research Fellowship program (# 361846460). He is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC 2126/1-390838866. An early version of this paper circulated as “On Event Study Designs and Distributed-Lag Models: Equivalence, Generalization and Practical Implications”.

1 Introduction

The credibility revolution in empirical economics has led to more transparent (quasi-) experimental research designs. This shift has increased the policy relevance and the scientific impact of empirical work (Angrist and Pischke, 2010). Differences-in-differences (DD) models in general, and event study designs in particular are very popular in this respect as the underlying identifying assumption of parallel post-treatment trends is intuitive and its plausibility can be assessed (though not verified) visually by inspecting whether pre-treatment trends are parallel. Moreover, the intuition of the empirical model is straightforward. The specification usually boils down to the well-known two-way fixed effects (TWFE) panel data model where the regressors of interest are a set of event indicators that are defined relative to the event.


Originating from the finance literature¹, event study designs are now widely used in applied economics, mostly public and labor economics, where an event is typically defined as a policy change whose effects are investigated. Figure 1 plots the use of event study designs in the Top Five economics journals over time.² While we see a gradual increase since 1990, there is a surge since 2010 driven by the three journals focusing on applied microeconomic work among the Top-Five, i.e. the American Economic Review (AER), the Quarterly Journal of Economics (QJE), and the Journal of Political Economy (JPE).

However, the intuitive appeal of event studies entails a risk of modeling them rather loosely. In more than one-third of the event study papers published in the AER, QJE, or JPE since 2010, no regression equation is specified. In addition, among the studies that specify an empirical model, only 15% provide some information on the choices that have been made at the ends of the effect window – in many cases only verbally and not formally. We show below that these details of the empirical model often imply identifying assumptions, which lead to misspecification if they are not warranted. The lack of detail also impedes reproducibility and easily leads to mistakes when the methodology is applied in other contexts.

¹ Dating back to Dolley (1933), see also MacKinlay (1997) for a survey on the financial literature.

² We proxy the use by the share of studies mentioning the term “event study”. More than 80% of the studies mentioning event study designs implement one.

Figure 1: The rise of event studies in economics



Notes: This graph plots the three-year moving average of the share of studies mentioning event study designs in top economics journals. We use a 3-year moving average to control for mean reversion. The Top 5 journals are the American Economic Review (AER), the Quarterly Journal of Economics (QJE), the Journal of Political Economy (JPE), *Econometrica*, and the *Review of Economic Studies*. We report results for AER, QJE, and JPE separately as these journals are known to publish many applied microeconomic studies.

This paper’s primary purpose is to clarify the understanding of event study designs in methodological and practical terms. We show below that the inherent imprecision in event study specifications can easily give rise to a misspecified model or a model that relies on strong implicit assumptions. We derive three main methodological contributions that are important for applied researchers when setting up event studies. While we derive these three points formally, we choose an intuitive, rather non-technical way of presenting them to make the insights widely accessible to applied researchers. For the same reason, we also discuss various practical implications and pitfalls along the way, making use of simple numerical examples to visualize our main results.

The three contributions are: First, researchers need to define a range for which the dynamic treatment effects are studied. We call this range the *effect window*. While this choice is a practical necessity due to limited data availability, it is not innocuous. Two approaches are commonly used in the literature, and we systematically discuss the underlying assumptions, advantages, and caveats. In the first approach, all unit-period observations outside of the effect window are excluded from the estimation sample. This approach requires the presence of never-treated units to identify dynamic treatment effects. The ultimate question is whether never-treated units can be regarded as a suitable

control group. Moreover, dropping observations outside of the effect window is less straightforward in settings where units can receive multiple treatments, which are quite common in public and labor economics. The second, in our view preferable approach is to define the last lag (lead) as open intervals capturing all known events that (will) have happened in the past (future). We refer to this practice as *binning*. Interestingly, an early application of event studies in economics formally discusses and implements the binning approach, but few studies refer to it (McCrary, 2007, pp. 334). We show that binning imposes implicit assumptions that enable researchers to identify dynamic treatment effects in the absence of never-treated units. Without binning it would be impossible to separate dynamic policy effects from secular time effects. We show in an application that a failure to bin endpoints leads to biased estimates, while the empirical event study model looks innocuous at first sight. Binning is also readily applicable to settings where units might receive multiple treatments (see the third contribution below).

Second, we show that event study (ES) designs *with binned endpoints* and distributed-lag (DL) models with a limited number of leads and lags are equivalent. In a DL model, an outcome is regressed on leads and lags of the treatment variable, i.e. the tax rate instead of an indicator of a tax rate change. We show that the DL model is a reparametrization of the ES model. Binning of endpoints in the ES model thereby corresponds to limiting the number of leads and lags in the DL model. Event study estimates can be recovered from DL estimates by properly normalizing the DL model and cumulating the post-treatment and pre-treatment effects away from zero. We argue that this isomorphism provides a more transparent and intuitive way to understand the role of parameter restrictions for identification: a distributed lag model can be simply spelled-out, whereas a complete ES specification needs to specify how the endpoints of the effect window are treated (see above). The isomorphism also offers the distributed-lag model as an alternative implementation in statistical software which is less error-prone as the binning of endpoints can be tricky – especially in the case of multiple treatments.

Third, we show that the simple event study design can be generalized to account for multiple treatments and/or treatments of different signs and varying intensities. Such institutional settings

are common in public and labor economics; consider, for instance, a sequence of state-level minimum wage changes or tax reforms of different signs and sizes. Initially, event studies in economics were predominantly used in situations, where a single event, measured by an indicator variable, hits a unit at a unique point in time. In settings with multiple treatments of different signs and varying intensities, distributed lag models were frequently used – sometimes in addition to classic event study models (see, e.g., Suárez Serrato and Zidar, 2016, Fuest et al., 2018). We show that the equivalence result between DL and ES models also holds in the general case if and only if endpoints of the effect window are binned. Moreover, we discuss important modeling assumptions that are necessary when applying event studies in environments with multiple, heterogeneous events.

In the final part of the paper, we demonstrate the practical relevance of our three contributions, replicating and expanding the study by Baker and Fradkin (2017), henceforth BF2017. In the original paper, the authors suggest a neat way to measure worker search intensity based on Google search data. BF2017 apply their new measure to test whether search intensity responds negatively to increases in the potential benefit duration (PBD) induced by state-level reforms following the Great Recession. While their difference-in-difference estimates clearly show the expected negative relationship, the original event study results are inconclusive. We show that implementing the generalized event study design yields statistically highly significant dynamic effects, which are well in line with the difference-in-difference estimates. Hence, implementing our preferred specification of an event study design strengthens the credibility of the novel measure of search effort suggested by Baker and Fradkin (2017) and provides even stronger support for their key empirical finding that PBD has a negative effect on search effort.

Our paper gives practical advice to researchers implementing event study designs. We pay attention to the role of data structure and the type of identifying variation (single vs. multiple treatments of different signs and varying intensities) that force researchers to make decisions that are often not discussed in theoretical work but are important for applied work. In this sense, our paper is similar to the study by Freyaldenhoven et al. (2022) who provide useful guidance for applied researchers when plotting the estimates obtained from event study designs.

We analyze event study designs under ideal conditions in two important dimensions: we assume (i) homogeneous treatment effects across cohorts, and (ii) parallel trends and no anticipation effects. The first condition (i) implies that we do not seek to contribute to an important recent econometric literature that discusses potential bias in two-way fixed effects models with heterogeneous treatment effects that differ across treatment cohorts. Such bias occurs in both event study specifications (Chaisemartin and D’Haultfœuille, 2022a, Sun and Abraham, 2021, Callaway and Sant’Anna, 2021, Borusyak et al., 2022) and difference-in-difference models (see Gibbons et al., 2019, Chaisemartin and D’Haultfœuille, 2020, Goodman-Bacon, 2021).³ The proposed alternative estimators address the problem of heterogeneous treatment effects by relying on a different weighting scheme than the one implied by OLS. There is also some recent work (Callaway et al., 2021, Chaisemartin and D’Haultfœuille, 2022a,c) that accounts for either continuous or multiple treatments which are common in public and labor economics.

We nevertheless believe that TWFE will remain the norm or the starting point for some time. First, the new “heterogeneity-robust” estimators are not available in all settings, notably when there are multiple treatments and/or treatments of different signs and varying intensities. Second, the insights from this paper remain relevant in cases, in which there is no staggered adoption and hence no bias from heterogeneity.⁴ Third, the intuitive appeal of the TWFE is one of the reasons for the success of the model in the credibility revolution. Intuitive and transparent models are likely to be preferred, but they have to yield consistent estimates. Hence, in settings where a new “heterogeneity-robust” estimator is available, these estimators should be applied and compared to the TWFE estimates as an informal test for heterogeneous treatment effects (Chaisemartin and D’Haultfœuille, 2020, Cengiz et al., 2019, Callaway and Sant’Anna, 2021, Sun and Abraham, 2021, Borusyak et al., 2022, Gardner, 2022). In settings with continuous and multiple treatments where no new “heterogeneity-robust” estimator is available, one option is to dichotomize the continuous

³ See Chaisemartin and D’Haultfœuille (2022b), Baker et al. (2022), and Roth et al. (2022) for recent surveys of the importance of heterogeneous treatment effects and the implications of binning in such contexts.

⁴ Covariates can cause problems in the case of treatment heterogeneity even if there is no staggered adoption (Powell, 2021).

treatment variable and apply an estimator that can account for multiple treatments.⁵

The second condition (ii) implies that we do not seek to contribute to the important literature discussing fixes to violations of the common trend assumption. Freyaldenhoven et al. (2019) show how to recover causal effects in a standard event study design where unobserved confounders generate a pre-treatment trend in the outcomes. Roth (2022) shows that treatment effects can be biased conditional on passing the flat pre-trend test. Rambachan and Roth (2022) propose inference methods for event study designs if the parallel trend assumption does not hold exactly. Malani and Reif (2015) point to the fact that non-flat pre-trends might also be due to anticipation rather than unobserved confounders.

The remainder of this paper is structured as follows. Section 2 introduces a standard version of an event study design in the simplest institutional environment and discusses how limiting the effect window and binning of endpoints imposes important parameter restrictions. In Section 3, we show that the event study model is equivalent to a standard distributed-lag model. In Section 4, we generalize the institutional environment and allow for multiple treatments of different signs and varying intensities. We show that event study designs can also be used in such settings and discuss the additional adjustments and assumptions that need to be made in these cases. In Section 5, we demonstrate the relevance of our methodological contributions by replicating and extending the study by Baker and Fradkin (2017). Section 6 concludes.

2 Standard Event Study Design

In this section, we set up an event study model in the simplest institutional environment. We refer to this model as the standard event study set-up throughout the paper. We use the standard case to highlight the importance of introducing parameter restrictions to identify the model and make a case to bin the endpoints of effect windows.

We assume a panel, in which each unit $i = 1, \dots, N$ is observed over time periods $t = \underline{t}, \dots, \bar{t}$.⁶

⁵ See Lichter et al. (2022) and Sieglöck et al. (2022) for examples of such a pragmatic approach and Section 4.3 for potential caveats when dichotomizing continuous treatment variables.

⁶ We assume a balanced panel for convenience but our remarks generalize to panels with unit-specific time windows.

There is a policy variable $T_{i,t}$, which we call *treatment status*. In this section, treatment status is a dummy variable indicating, for instance, if employment protection laws have yet been passed in a certain state or not (Autor, 2003). In the classic event study design, treatment is an absorbing state, i.e. there is a single treatment that is adopted at some unit-specific time E_i and lasts forever. Hence, treatment status changes at most once at some unit-specific time E_i from 0 to 1. *Treatment adoption* is indicated by $\Delta T_{i,t} = T_{i,t} - T_{i,t-1} = \mathbb{1}[E_i = t]$, a dummy variable. Some units i may never be treated. For these never-treated units $\Delta T_{i,t} = 0$ for all t . All treated units may receive treatment at the same time or treatment may be staggered over time with different units receiving treatment at different points in time. We seek to estimate the dynamic effects of this treatment on our dependent variable $Y_{i,t}$. Whenever we refer to *dynamic treatment effects* we are talking about the cumulative effect of treatment at a specific time relative to the event. In contrast, we refer to incremental effects as changes in the cumulative treatment effect.

In this set-up, the standard event study specification is given for all $t = \underline{t}, \dots, \bar{t}$ by

$$Y_{i,t} = \sum_{\ell=-\infty}^{\infty} \beta_{\ell} \Delta T_{i,t-\ell} + \mu_i + \theta_t + \varepsilon_{i,t} \quad (1)$$

where unit fixed effects are denoted by μ_i and time period fixed effects by θ_t . The parameter β_{ℓ} is the dynamic treatment effect ℓ periods after ($\ell \geq 0$). For the periods before an event ($\ell < 0$), the parameters β_{ℓ} either capture anticipation effects or – in their absence – are used to test for parallel pre-treatment trends and hence to assess the credibility of the identification strategy. All results derived in this paper also hold when the models include additional exogenous covariates $X_{i,t}$.

For the sake of clarity, we assume that equation (1) is correctly specified and the data-generating process follows standard panel data assumptions. In particular, we assume effect homogeneity in the parameters β_{ℓ} , μ_i and θ_t , independence across units i , and strict exogeneity of the treatment $T_{i,t}$. Treatment effects are hence assumed to be homogeneous across cohorts, in contrast to, e.g., Sun and Abraham (2021).

2.1 Normalizing the Effects

In a model with individual fixed effect, μ_i , parameters β_ℓ are only identified up to a constant. The intuition behind this is straightforward: adding a constant to β_ℓ for all $\ell = -\infty, \dots, \infty$ and subtracting the same constant from the unit fixed effect μ_i for units with an event between $-\infty$ and ∞ does not alter equation (1). This non-identification of the absolute size of the dynamic effects can be solved by setting a reference period to which the effects are compared.

Remark 1 (Normalization).

Dynamic treatment effects β_ℓ are typically expressed relative to some reference period, for example, one period before treatment adoption. The corresponding coefficient is normalized to zero, e.g. $\beta_{-1} = 0$. In practice, normalization is implemented by dropping the treatment adoption indicator for the reference period.

Typically $\beta_{-1} = 0$ is assumed, but any other periods in the effect window, including (binned) endpoints, can serve as a reference period (cf. Remark 2). Choosing a different reference period shifts point estimates of the dynamic effects and affects their standard errors. Generally, the choice of the reference point should depend on the parameter(s) of interest. In the absence of anticipation, policy-relevant post-treatment effects are typically given by normalizing to the immediate pre-treatment period.

2.2 Restricting the Effect Window (Binning)

In practice, researchers have to impose restrictions on the effect window to implement the event study design, since β_ℓ can never be estimated from the infinite past to the infinite future. Hence, the effect window has to be restricted to a finite number of leads and lags. There are two broad ways to restrict the effect window in practice. First, it is possible to define the last lag/first lead that should be estimated and exclude all unit-period observations outside of this window from the estimation sample. By construction, (pre)-treatment effects with a longer time horizon are not estimated. In the absence of never-treated units, at least one restriction on dynamic treatment effects in addition to

the normalization is needed to identify dynamic treatment effects (cf. Section 2.4 for more details).

Dropping observations outside of the effect window is not straightforward when effect windows of multiple treatments overlap. (cf. Section 4). Second, it is possible to define the endpoints of the effect windows as open intervals, which also take into account past (future) treatments beyond the last lag (lead). We argue that this is the preferable approach in many applications. However, it goes along with an important but often unstated assumption about the dynamic treatment effects, which is summarized in the following Remark 2.

Remark 2 (Restricted effect window).

Restricting the effect window to a finite number of leads, $\underline{\ell}$, and/or lags, $\bar{\ell}$, requires assumptions about the nature of the effect outside of the window. It is often economically plausible to assume that dynamic treatment effects stay constant before $\underline{\ell}$ and/or after $\bar{\ell}$, i.e. $\beta_\ell = \beta_{\bar{\ell}}$ for all $\ell > \bar{\ell}$ and $\beta_\ell = \beta_{\underline{\ell}}$ for all $\ell < \underline{\ell}$. These assumptions should be explicitly stated and defended.

Note that restricting the effect window according to Remark 2 can be readily applied in settings with multiple treatments (cf. Section 4). Using Remark 2, we rewrite equation (1) as

$$Y_{i,t} = \beta_{\underline{\ell}} \sum_{\ell=-\infty}^{\underline{\ell}} \Delta T_{i,t-\ell} + \sum_{\ell=\underline{\ell}+1}^{\bar{\ell}-1} \beta_\ell \Delta T_{i,t-\ell} + \beta_{\bar{\ell}} \sum_{\ell=\bar{\ell}}^{\infty} \Delta T_{i,t-\ell} + \mu_i + \theta_t + \varepsilon_{i,t}$$

which simplifies to our preferred standard event study specification:

$$Y_{i,t} = \sum_{\ell=\underline{\ell}}^{\bar{\ell}} \beta_\ell D_{i,t}^\ell + \mu_i + \theta_t + \varepsilon_{i,t} \quad (2)$$

with

$$D_{i,t}^\ell = \begin{cases} \sum_{s=-\infty}^{\underline{\ell}} \Delta T_{i,t-s} & \text{if } \ell = \underline{\ell} \\ \Delta T_{i,t-\ell} & \text{if } \underline{\ell} < \ell < \bar{\ell} \\ \sum_{s=\bar{\ell}}^{\infty} \Delta T_{i,t-s} & \text{if } \ell = \bar{\ell}. \end{cases} \quad (3)$$

We refer to the variable $D_{i,t}^\ell$ as *binned treatment adoption indicator* which takes into account all observable past (future) treatments beyond the endpoints, i.e. the maximum lag (lead). Note that in

practice, the sums at the endpoints given in equation (3) can never be calculated until infinity using observed data but only until the first/last observation of $\Delta T_{i,t}$. We discuss the data requirements that follow in Section 2.2.3. The definition of endpoints in equation (3) is for example used in Smith et al. (2019) and Fuest et al. (2018). Endpoints can be equivalently defined as $D_{i,t}^{\underline{\ell}} = \sum_{s=t-\underline{\ell}}^{\infty} \Delta T_{is}$ and $D_{i,t}^{\bar{\ell}} = \sum_{s=-\infty}^{t-\bar{\ell}} \Delta T_{is}$. Another equivalent definition for the example used in McCrary (2007) is given by

$$D_{i,t}^{\ell} = \begin{cases} \mathbb{1}[t \leq E_i + \ell] & \text{if } \ell = \underline{\ell} \\ \mathbb{1}[t = E_i + \ell] & \text{if } \underline{\ell} < \ell < \bar{\ell} \\ \mathbb{1}[t \geq E_i + \ell] & \text{if } \ell = \bar{\ell}. \end{cases} \quad (4)$$

2.2.1 Length of the Effect Window

From the previous section it follows that it is a key question for applied researchers how to set the length of the effect window. First and foremost, this choice depends on the research question that the model tries to address. When do we expect the effects to have fully materialized? Are we interested in longer-run effects? What is a reasonable pre-treatment period given dependent and treatment variables and endogeneity concerns? Answers to all these questions can be based on theoretical considerations, existing empirical evidence, and economic intuition.

Second, the choice of the effect window is restricted by data availability. Increasing leads or lags reduces the estimation sample if treatment is not observed for these additional periods (see also Section 2.2.3). At some point, the sample size will be too small to estimate precise effects and/or have meaningful variation.

Last, researchers can use the empirical estimates post-estimation to assess whether the choice of size of the effect window was reasonable. By Remark 2, we assume $\beta_{\ell} = \beta_{\bar{\ell}}$ for all $\ell > \bar{\ell}$ at the end of the effect window. Hence, $\hat{\beta}_{\bar{\ell}}$ measures the long-run dynamic treatment effect. If treatment effects had fully materialized by $\bar{\ell}$ after the reform, we would expect estimates leading up to $\hat{\beta}_{\bar{\ell}}$ to level off and converge to $\hat{\beta}_{\bar{\ell}}$. A pronounced drop between $\hat{\beta}_{\bar{\ell}}$ and $\hat{\beta}_{\bar{\ell}-1}$ is instead an indication that dynamic

effects are still unfolding.⁷ We summarize these points in the following remark.

Remark 3 (Length of the effect window).

Researchers should experiment with different effect windows length (bearing in mind that the estimation sample might change), and assess empirically whether the underlying assumption of restricting the effect window at $\bar{\ell}$ given in Remark 2 is justified. Practically, researchers can check whether estimates leading up to the endpoint converge towards $\hat{\beta}_{\bar{\ell}}$.

2.2.2 Special cases of effect window restrictions.

As stated in Remark 2, it is possible to only restrict the effect window pre or post-treatment. In an extreme case, one could restrict all effects before the event and allow the treatment effect to continue indefinitely, i.e. $\underline{\ell} = -1$ (or equivalently $\beta_{\ell} = 0$ for all $\ell \leq -1$). But clearly, such a model does not allow us to estimate and visualize pre-trends.

The most extreme form of restricting the effect window is to bin both at the first pre-treatment and the first post-treatment period, i.e. $\underline{\ell} = -1$, and $\bar{\ell} = 0$. In this case, the binned treatment indicator $D_{i,t}^{\underline{\ell}}$ simplifies to $T_{i,t}$. The empirical model is given by $Y_{i,t} = \beta T_{i,t} + \mu_i + \theta_t + \varepsilon_{i,t}$. Note that this is a standard difference-in-differences model with staggered treatment and two-way fixed effects.

Another type of restriction, which is sometimes seen in the literature (cf. Baker and Fradkin, 2017 discussed in Section 5 and Online Appendix B), is to restrict the effect window but without binning of endpoints. If observations outside of the effect window are not dropped, such a model implicitly assumes that dynamic treatment effects drop to zero outside of the effect window. While such an assumption is typically hard to defend, it is often only stated implicitly by not spelling out the event study specification completely, i.e. not specifying what happens at the endpoints of the effect window.

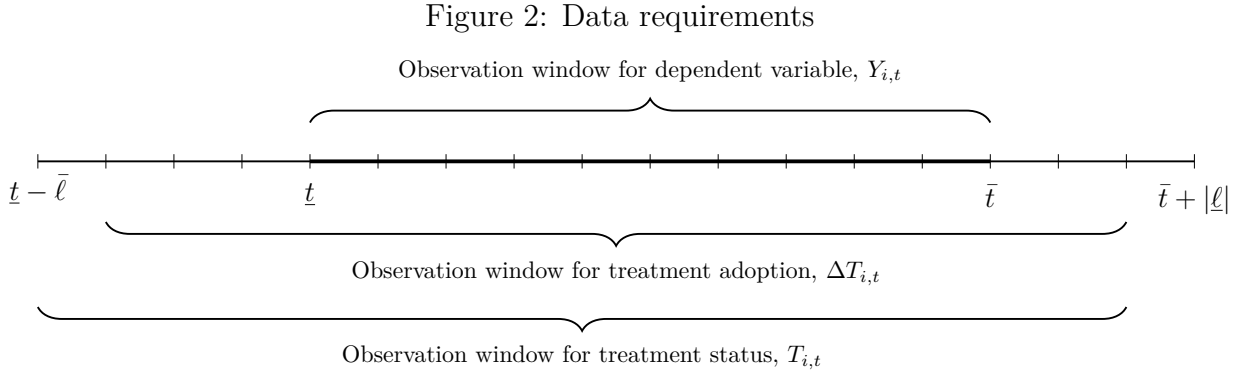
⁷ Freyaldenhoven et al. (2022) provide a similar discussion and propose a formal test for what they call the leveling-off at the end of the effect window. However, the model selection procedure affects the sampling properties of the post-model-selection estimator, which invalidates the usual inference (Leeb and Pötscher, 2005).

2.2.3 Data Requirements

Due to the leads and lags of the effect window, we need to observe the treatment variable for a longer observation window than the dependent variable. Restrictions on the effect window determine the requirements on data availability for dependent variable $Y_{i,t}$, treatment status $T_{i,t}$, and event indicator $\Delta T_{i,t}$. In the following remark, we summarize the data requirements for a given observation window of the dependent variable.

Remark 4 (Data requirements).

For a given balanced panel of the dependent variable from $[\underline{t}, \bar{t}]$ and a limited effect window $[\underline{\ell}, \bar{\ell}]$,⁸ we need to observe treatment adoption from $\underline{t} - \bar{\ell} + 1$ to $\bar{t} + |\underline{\ell}| - 1$. If treatment adoption is derived from changes in treatment status we need to observe treatment status from $\underline{t} - \bar{\ell}$ to $\bar{t} + |\underline{\ell}| - 1$. The following figure visualizes the required width of the observation window for a given, limited effect window.



To understand the intuition behind Remark 4, it is first important to note that an event that happens before \underline{t} , i.e. the first data year of the dependent variable, can affect the outcome like any other event happening between \underline{t} and \bar{t} and needs to be taken into account. Likewise, events happening after \bar{t} might affect pre-treatment effects. This rationale prompts the second question of how long before \underline{t} and after \bar{t} we need to observe events. By Remark 2, dynamic treatment effects are constant beyond the endpoints of the effect window $\bar{\ell}$ and $\underline{\ell}$. Hence, we need to observe events

⁸ We assume an identical observation window for the dependent variable for all units i for notational simplicity. Remark 4 also holds for unit-specific observation windows $[\underline{t}_i, \bar{t}_i]$ which leads to unit-specific observation windows for treatment status, $[\underline{t}_i - \bar{\ell}, \bar{t}_i + |\underline{\ell}| - 1]$, and treatment adoption, $[\underline{t}_i - \bar{\ell} + 1, \bar{t}_i + |\underline{\ell}| - 1]$.

at least from $\underline{t} - \bar{\ell}$ to $\bar{t} + |\underline{\ell}|$. Remark 4 states that it suffices to observe events in one fewer year at each end of the data window, hence from $\underline{t} - \bar{\ell} + 1$ to $\bar{t} + |\underline{\ell}| - 1$. To see this, consider a case where the event takes place at $\underline{t} - \bar{\ell}$. The binned treatment adoption indicator $D_{i,t}^{\bar{\ell}}$ will be equal to one for this unit for all t . Conditional on unit fixed effects, this unit cannot be used to identify treatment effects. An analogous argument applies at the other end of the data window. The binned treatment adoption indicator in equation (3) is therefore in practice implemented as

$$D_{i,t}^{\underline{\ell}} = \begin{cases} \sum_{s=\underline{t}-\bar{\ell}+1}^{\underline{\ell}} \Delta T_{i,t-s} & \text{if } \ell = \underline{\ell} \\ \Delta T_{i,t-\ell} & \text{if } \underline{\ell} < j < \bar{\ell} \\ \sum_{s=\bar{\ell}}^{\bar{t}+|\underline{\ell}|-1} \Delta T_{i,t-s} & \text{if } \ell = \bar{\ell}. \end{cases} \quad (5)$$

The top left panel of the numerical example shown in Online Appendix A.1 shows the construction of the binned treatment indicator in practice. The matrix immediately demonstrates the need for normalization (cf. Remark 1) as all row sums in the left matrix are equal to one. Binning also implies that if the reform had happened on or before $t = \underline{t} - \bar{\ell} = 2000 - 4 = 1996$ (rather than in $t = 2005$ as assumed in the example), $D_{i,t}^{\bar{\ell}} = D_{i,t}^4 = 1$ for all t from 2000 to 2010. In this case of an always-treated unit, $D_{i,t}^{\bar{\ell}} = D_{i,t}^4$ is a constant, and its effect is absorbed in the unit fixed effect μ_i . By the analogous argument, events on or after $t = \bar{t} + |\bar{\ell}| = 2010 + 3 = 2013$ imply $D_{i,t}^{\underline{\ell}} = D_{i,t}^{-3} = 1$ for all t from 2000 to 2010 whose effect is absorbed in μ_i . It, therefore, suffices to know all events from time period $\underline{t} - \bar{\ell} + 1 = 1997$ to $t = \bar{t} + |\bar{\ell}| - 1 = 2012$ to estimate the model, see Remark 4.

In the case of an absorbing single treatment, missing data can be filled. If unit i is observed to adopt treatment in period E_i , then we know that there was no treatment adoption in any time period $t < E_i$ or $t > E_i$ and hence $\Delta T_{i,t} = 0$ for all $t \neq E_i$. This assumption is implicitly made when the binned treatment adoption indicators are constructed by equation (4) rather than (3). For units that are never observed to adopt treatment, we may assume that treatment is not adopted outside the observed window including the yet unknown future, either (see the discussion on assuming never-treated units in Section 2.4.1). This assumption is made when equation (4) is used and E_i is set to ∞ for units that are not observed to have been treated.

The data requirements in Remark 4 can also be derived from the equivalent distributed-lag specification described in Section 3. However, they are immediately spotted in the estimation equations (10) and (11) without the involved discussion of the endpoints here. The data requirements described for the standard event study design also apply to the generalized design in Section 4.

2.3 Estimation

The dynamic treatment effects β_ℓ in the standard event study model of equation (2) can be estimated with ordinary least squares including two-way fixed effects for units and time periods. This is a regression of the dependent variable $Y_{i,t}$ on the binned treatment adoption indicators with leads and lags

$$Y_{i,t} = \sum_{\ell=\underline{\ell}}^{-2} \beta_\ell D_{i,t}^\ell + \sum_{\ell=0}^{\bar{\ell}} \beta_\ell D_{i,t}^\ell + \mu_i + \theta_t + \varepsilon_{i,t}, \quad (6)$$

where we drop $D_{i,t}^1$ as the reference period. As the treatment adoption indicator $D_{i,t}^\ell = \Delta T_{i,t}$ is the first difference of the treatment status $T_{i,t}$ for $\underline{\ell} < \ell < \bar{\ell}$, this is a regression of levels on (binned) changes.

Alternatively, dynamic cumulative treatment effects β_ℓ can be estimated as a regression in first differences without unit fixed effects

$$\Delta Y_{i,t} = \sum_{\ell=\underline{\ell}}^{-2} \beta_\ell \Delta D_{i,t}^\ell + \sum_{\ell=0}^{\bar{\ell}} \beta_\ell \Delta D_{i,t}^\ell + \phi_t + \Delta \varepsilon_{i,t}, \quad (7)$$

where $\phi_t = \theta_t - \theta_{t-1}$ and where we again drop $D_{i,t}^1$ as the reference period. As $\Delta D_{i,t}^\ell = \Delta \Delta T_{i,t}$ for $\underline{\ell} < \ell < \bar{\ell}$, this is a regression of changes on changes of (binned) changes. Estimation of the event study model in first differences may therefore look somewhat disturbing and is rarely estimated in practice. However, the equivalent distributed-lag specification shown later in equation (11) in Section 3 is numerically equivalent to equation 7 and commonly applied.

Under standard assumptions of the fixed effects regression model, i.e. effect homogeneity, independence across units i , and strict exogeneity of the treatment adoption indicator $\Delta T_{i,t}$, the

estimated regression coefficients $\widehat{\beta}_\ell$ are consistent and asymptotically normally distributed as the number of units $N \rightarrow \infty$ while the number of time periods is fixed (see e.g. Wooldridge, 2010 or Cameron and Trivedi, 2005).

For a cohort that experiences treatment at the same point in time and in the presence of never-treated units, two-way fixed effects estimates for this cohort can be interpreted as average treatment effects (ATT) when individual treatment effects are heterogeneous. Sun and Abraham (2021) call these effects cohort-specific average treatment effects (CATT) and show that they are consistently estimated by two-way fixed effects assuming parallel trends in baseline outcomes and no anticipatory behavior before treatment. The interpretation holds for effects at a specific lag/lead but also for binned endpoints. It therefore also holds for the standard difference-in-differences (DiD) model which is the most extreme form of binning. The estimated ATTs for the endpoints are averages over individual heterogeneity *as well as over all lags/leads* beyond the binned endpoint.

If, however, treatment adoption is staggered, the recent literature on heterogeneous treatment effects in event studies has shown that the two-way fixed effects estimates of β_ℓ cannot be interpreted as average treatment effects (Chaisemartin and D’Haultfœuille, 2022a, 2020, Sun and Abraham, 2021, Callaway and Sant’Anna, 2021, Borusyak et al., 2022, Gibbons et al., 2019, Goodman-Bacon, 2021). Even if individual effects were homogeneous (as we assume throughout this paper), the two-way fixed effect estimates of binned endpoints cannot be interpreted as ATTs when the true dynamic treatment effects are not constant beyond the endpoints. Imagine now that the dynamic treatment effects have not fully materialized and are increasing beyond the endpoint. In this situation, the estimate of the binned endpoint will correctly pick up the level of the effects at and beyond the endpoint. But the increasing trend around the level cannot be picked up by a single parameter and is instead picked up by the secular time period effects. These biased secular time period effects, in turn, pin down the counterfactual for other treated units whose lags/leads of outcomes are still within the effect window. Consequently, all estimated dynamic treatment effects are biased because two-way fixed effects estimation fails to correctly separate dynamic causal effects from time trends when treatment is staggered. This problem is more severe in the absence of never-treated units,

where the secular time period effects are only identified by the restrictions imposed by binning (see Section 2.4.2 and Remark 5). We discuss in Section 3.2 that estimation in first differences does not solve this problem.

2.4 Identification

It is important to assure that the model is econometrically identified such that the dynamic treatment effects β_ℓ are disentangled from the secular time fixed effects θ_t . Throughout the paper, we understand by identification the identifiability of the regression coefficients β_ℓ in a given sample, i.e. the design matrix is of full column rank.⁹ In particular, we do not mean the identification of average treatment effects in the population. Recall that we assume effect homogeneity of the parameters β_ℓ (see also the discussion in the introduction).

2.4.1 Identification with a Never-Treated Group

Assume that there are units that are never treated. These never-treated units will serve as a control group, which uniquely identifies the secular time trends θ_t if there is at least one control group observation for each time period t . To identify the dynamic treatment effects β_ℓ , we need to observe at least one treated unit for each lag and lead j of the effect window. The number of leads and lags can be purely based on data availability; there is no need to restrict the effect window to overcome the underidentification problem discussed in Subsection 2.4.2.

The existence of a never-treated control group is, however, less trivial than it may appear at first sight. *Observing* that a unit is not treated does not imply that this unit is never-treated. A unit could have been treated outside of the observation window – even in the yet unrealized future. When treatment effects are allowed to have effects into the infinite future (or affect the infinite past), no unit is known to be never-treated with certainty. There are situations in which some units cannot possibly be treated. For example, in a study on the effects of giving birth to a child, men

⁹ The identification problem discussed in this section could also be avoided by dropping either time effects θ_t or unit fixed effects μ_i . However, dropping either dimension of fixed effects typically leads to omitted variable bias unless treatment is randomized.

will constitute a natural never-treated group. But if a unit cannot possibly be treated for its special characteristics, the very same special characteristics are potential confounders. Or in the logic of the Rubin causal model, all units can be manipulated, at least conceptually, so that treatment happens. Hence, relying on a never-treated control group is not a fully convincing strategy in many applications as it raises the question of whether units that are never treated are comparable to units that receive treatments.

2.4.2 Identification without a Never-Treated Group

Identification of dynamic treatment effects is more challenging in the absence of a never-treated group. Borusyak et al. (2022) show that with an infinite effect window, $[\underline{\ell}, \bar{\ell}] = [-\infty, \infty]$, dynamic effects are only identified up to a linear trend. Formally, we can cast this underidentification problem in our notation as follows: $Y_{i,t} = \sum_{\ell=-\infty}^{\infty} \beta_{\ell} \Delta T_{i,t-\ell} + \mu_i + \theta_t + \varepsilon_{i,t} = \sum_{\ell=-\infty}^{\infty} \tilde{\beta}_{\ell} \Delta T_{i,t-\ell} + \tilde{\theta}_t + \tilde{\mu}_i + \varepsilon_{i,t}$ where $\tilde{\beta}_{\ell} = \beta_{\ell} + \lambda \cdot \ell$, $\tilde{\theta}_t = \theta_t - \lambda \cdot t$ and $\tilde{\mu}_i = \mu_i + \lambda \cdot E_i$. Hence, one can add a linear trend $\lambda \cdot \ell$ to the dynamic treatment effects and adjust secular time fixed effects θ_t and unit fixed μ_i to maintain the same predicted values of the model.¹⁰ This violates the assumption of full rank of the design matrix. In practice, this underidentification can easily be overlooked as many statistical packages automatically drop regressors in the case of multicollinearity. A non-identified linear trend leads to dropping either one event dummy or one time dummy.

In the following, we show that restricting the effect window as proposed in Remark 2 introduces restrictions that allow separately identifying dynamic effects, β_{ℓ} , and secular time trends, θ_t . Restricting the effect window as in Remark 2 leads to our equation (2), $Y_{i,t} = \sum_{\ell=\underline{\ell}}^{\bar{\ell}} \beta_{\ell} D_{i,t}^{\ell} + \mu_i + \theta_t + \varepsilon_{i,t}$. In this model, adding a linear trend $\lambda \cdot \ell$ to the restricted number of dynamic treatment effects does not offset adding a linear trend to the secular time trend for the observations outside of the effect window, i.e. if $t < E_i - |\underline{\ell}| + 1$ or $t > E_i + \bar{\ell}$.¹¹

¹⁰ The underidentification problem arises because all units are treated at some point in the absence of never-treated units and hence $\sum_{\ell=-\infty}^{\infty} \ell \Delta T_{i,t-\ell} = \sum_{\ell=-\infty}^{\infty} \ell \mathbb{1}[t = E_i + \ell] = t - E_i$ for all units i and all time periods t . This underidentification problem also holds in the presence of never-treated units when unit- or group-specific linear trends are included in the model. Wolfers (2006) shows that it is challenging to identify the treatment effect in the standard difference-in-differences specification, which is an extreme form of binning (see Section 2.2.2).

¹¹ Adding a linear trend to both dynamic effects and to secular time trends leads to

In the absence of never-treated units, observations of treated units outside the effect window are used to estimate the secular time trends. These observations, therefore, serve as the control group for observations of units that are within the effect window at the same time period. Restricting the effect window may even produce units i for which we only observe the outcome $|\underline{\ell}|$ or more periods before the event. The observed outcomes of such units will all be affected by the same constant effect $\beta_{\underline{\ell}}$. Hence all changes over time can be attributed to secular time fixed effects and orthogonal noise. The analogous argument holds if the outcome is only observed $\bar{\ell}$ or more periods after the event. Hence, binning of one or both endpoints allows for identification with or without the presence of never-treated units. We summarize this in the following remark.

Remark 5 (Identification and restricted effect window).

In the absence of both never-treated units and unit-period observations outside of the effect window, dynamic treatment effects and secular time fixed effects are not separately identified. Binning of one or both endpoints generates unit-period observations outside of the effect window, which serve as control group observations. The length of the effect window hence directly affects identification and helps to separately identify dynamic treatment and secular time fixed effects.

As an alternative to binning, identification is also achieved by dropping observations outside of the effect window and setting another dynamic treatment effect $\beta_{\underline{\ell}}$ to zero, e.g. $\beta_{-2} = 0$ in addition to $\beta_{-1} = 0$ (as proposed in earlier versions of Borusyak et al., 2022). This assumption affects the shape of the estimated pre-trends as well as the estimated secular time trends. This identifying assumption must be defended by a convincing a priori argument.

$$\begin{aligned}
Y_{i,t} &= \left(\sum_{\ell=\underline{\ell}}^{\bar{\ell}} (\beta_{\ell} + \lambda \cdot \ell) D_{i,t}^{\ell} \right) + (\theta_t - \lambda \cdot t) + \mu_i + \lambda \cdot E_i + \varepsilon_{i,t} \\
&= \begin{cases} \left[\sum_{\ell=\underline{\ell}}^{\bar{\ell}} \beta_{\ell} D_{i,t}^{\ell} + \theta_t + \mu_i + \varepsilon_{i,t} \right] + \lambda [\underline{\ell} - t + E_i] & \text{if } t < E_i - |\underline{\ell}| \\ \sum_{\ell=\underline{\ell}}^{\bar{\ell}} \beta_{\ell} D_{i,t}^{\ell} + \theta_t + \mu_i + \varepsilon_{i,t} & \text{if } E_i - |\underline{\ell}| \leq t \leq E_i + \bar{\ell} \\ \left[\sum_{\ell=\underline{\ell}}^{\bar{\ell}} \beta_{\ell} D_{i,t}^{\ell} + \theta_t + \mu_i + \varepsilon_{i,t} \right] + \lambda [\bar{\ell} - t + E_i] & \text{if } t > E_i + \bar{\ell} \end{cases}
\end{aligned}$$

3 Event Studies and Distributed-lag Models

In this section, we show that event study (ES) models with binned endpoints and distributed-lag (DL) models yield identical parameter estimates. The derivation of this equivalence results mainly serves two purposes. First, there is still some confusion about the two types of parametrizations and their similarities. Second, the isomorphism helps to build intuition about data requirements and identification of the underlying model. In Subsection 3.1, we formally demonstrate under which restrictions ES and DL models yield equivalent dynamic treatment effects. In Subsection 3.4, we discuss the practical implications of this isomorphism and argue that DL models are easier to implement and less error-prone. In Online Appendix A.1, we illustrate all formal claims using a simple numerical example.

3.1 Equivalence

We start by showing the equivalence of event study and distributed-lag models in the general case with an infinite effect window. Taking first differences of the standard event study specification given in equation (1), we can rewrite the event study specification into a distributed-lag model:

$$\begin{aligned}
 \Delta Y_{i,t} &= Y_{i,t} - Y_{i,t-1} \\
 &= \sum_{\ell=-\infty}^{\infty} \beta_{\ell} \Delta T_{i,t-\ell} - \sum_{\ell=-\infty}^{\infty} \beta_{\ell} \Delta T_{i,t-1-j} + \phi_t + \Delta \varepsilon_{i,t} \\
 &= \sum_{\ell=-\infty}^{\infty} \beta_{\ell} \Delta T_{i,t-\ell} - \sum_{\ell=-\infty}^{\infty} \beta_{\ell-1} \Delta T_{i,t-\ell} + \phi_t + \Delta \varepsilon_{i,t} \\
 &= \sum_{\ell=-\infty}^{\infty} \gamma_{\ell} \Delta T_{i,t-\ell} + \phi_t + \Delta \varepsilon_{i,t}, \tag{8}
 \end{aligned}$$

where $\gamma_{\ell} = \beta_{\ell} - \beta_{\ell-1}$ and $\phi_t = \theta_t - \theta_{t-1}$ are time fixed effects. Parameters γ_{ℓ} are the incremental changes of the dynamic (cumulative) treatment effects β_{ℓ} , measuring the slope of treatment effects from one time period to the next. The distributed-lag specification in equation (8) is the first

difference of the following distributed-lag specification in levels

$$Y_{i,t} = \sum_{\ell=-\infty}^{\infty} \gamma_{\ell} T_{i,t-\ell} + \mu_i + \theta_t + \varepsilon_{i,t} \quad , \quad (9)$$

where μ_i denotes unit fixed effects. Note that the distributed-lag specification is either a regression of levels on levels (equation 9) or of changes on changes (equation 8) while the event-study specification is a regression of levels on (binned) changes (equation 6) or a regression of changes on changes of (binned) changes (equation 7).

At first sight, it might not seem intuitive to estimate an event study specification in first differences since the explanatory variable is the second difference of the treatment status, i.e. $\Delta^2 T_{i,t} = \Delta T_{i,t} - \Delta T_{i,t-1} = T_{i,t} - 2\Delta T_{i,t-1} + T_{i,t-2}$. At the same time, it is completely natural to estimate the numerically equivalent distributed-lag specification in first differences where both the dependent and the explanatory variables are first differences. Hence, it is important to keep the equivalence in mind and freely choose the estimator that is most appropriate for the question under study (see also Remark 6 and Section 3.4.)

We have proven the equivalence between event study and distributed-lag models in the general case without restricting the effect window. Next, we show that the equivalence between ES and DL models for restricted effect windows holds if and only if endpoints of the effect window are binned as in Remark 2. The distributed-lag parameters γ_{ℓ} are related to the event study parameters β_{ℓ} by $\gamma_{\ell} = \beta_{\ell} - \beta_{j-1}$. Binning the upper endpoint, $\beta_{\ell} = \beta_{\bar{\ell}}$ for all $j > \bar{\ell}$, is therefore equivalent to assuming that $\gamma_{\ell} = 0$ for all $j > \bar{\ell}$; for the lower endpoint $\beta_{\ell} = \beta_{\underline{\ell}}$ for all $j < \underline{\ell}$ is equivalent to $\gamma_{\ell} = 0$ for all $j \leq \underline{\ell}$. The event study model with a restricted effect window between $\underline{\ell}$ and $\bar{\ell}$ and binned endpoints in equation (2)

$$Y_{i,t} = \sum_{\ell=\underline{\ell}}^{\bar{\ell}} \beta_{\ell} D_{i,t}^{\ell} + \mu_i + \theta_t + \varepsilon_{i,t}$$

is therefore equivalent to a distributed-lag specification with $\bar{\ell}$ lags and $|\underline{\ell}| - 1$ leads

$$Y_{i,t} = \sum_{\ell=\underline{\ell}+1}^{\bar{\ell}} \gamma_{\ell} T_{i,t-\ell} + \mu_i + \theta_t + \varepsilon_{i,t}. \quad (10)$$

Without binning, ES and DL specifications are based on different parameter restrictions and yield different parameter estimates. We summarize this result in the following remark:

Remark 6 (Equivalence of Event Study and Distributed-Lag Model).

The event study specification with binned endpoints at $\ell = \bar{\ell}$ and $\ell = \underline{\ell}$ as specified in equations (2) and (3) is equivalent to a distributed lag models with $\bar{\ell}$ lags and $|\underline{\ell}| - 1$ leads as given by equation (10).

An implication of Remark 6 is that if one regards the distributed-lag model with a limited number of leads and lags as a sensible econometric model, binning of the endpoints in event study designs must be equally sensible.

3.2 Estimation of the distributed-lag model

The dynamic incremental treatment effects γ_{ℓ} of the distributed-lag model in equation (10) can be estimated with ordinary least squares including two-way fixed effects for units and time periods. This is a regression of the dependent variable $Y_{i,t}$ in levels on the treatment status in levels with leads and lags. Alternatively, the dynamic incremental treatment effects γ_{ℓ} can be estimated as a regression in first differences without unit fixed effects

$$\Delta Y_{i,t} = \sum_{\ell=\underline{\ell}+1}^{\bar{\ell}} \gamma_{\ell} \Delta T_{i,t-\ell} + \phi_t + \Delta \varepsilon_{i,t}, \quad (11)$$

where $\phi_t = \theta_t - \theta_{t-1}$. This is a regression of changes on changes.

Estimation in both levels and first differences is consistent (with $N \rightarrow \infty$) as discussed in Subsection 2.3 when the model is correctly specified. However, the incremental effect at the endpoint of the effect window $\gamma_{\bar{\ell}}$ will pick up all effects after ℓ periods when estimated in levels while it only

picks up the incremental effect after exactly ℓ periods when estimated in first differences. When the effect window is too short and dynamic effects have hence not fully materialized after ℓ periods, the dynamic effect $\beta_{\bar{\ell}}$ at the endpoint often shows a pronounced spike when estimated in levels but not when estimated in first differences. This is exemplified in the application in Section 5 and Online Appendix B, Figure B.4. This feature also leads to the often observed difference between fixed effects and first difference estimates in standard difference-in-difference models if the treatment effect builds up over time.

As already discussed in Section 2.3, dynamic treatment effects are not consistently estimated by two-way fixed effects regression if treatment adoption is staggered and the effects are not constant beyond the endpoints. This finding also holds when the binned model is estimated in first differences. This is because changes in the outcome variable beyond the endpoint are fully attributed to the secular time period effects. As above, these biased secular time fixed effects will lead to biases when pinning down the counterfactual for all incremental effects. As in levels, the problem is more severe in the absence of never-treated units, where the secular time period effects are only identified by the restrictions imposed by binning (see Section 2.4.2 and Remark 5).

3.3 Recovery of cumulative treatment effects from the distributed-lag model

The distributed-lag coefficients γ_ℓ measure incremental changes in the dynamic (cumulative) treatment effects. Hence, we can construct these dynamic cumulative treatment effects from the estimated incremental effects. Note that we need to estimate one fewer incremental effect γ_ℓ than event study coefficient β_ℓ since we always normalize one dynamic treatment effect to zero (see Remark 1 and Remark 7 below). Thus, we include DL leads and lags γ_ℓ running from $\underline{\ell} + 1$ (not $\underline{\ell}$ as in the event study design) to $\bar{\ell}$. The event study parameters β_ℓ can be calculated from the distributed-lag parameters γ_ℓ by using the difference $\gamma_\ell = \beta_\ell - \beta_{\ell-1}$. The starting point for this difference equation is given by the normalization in Remark 1. Normalizing to one period before treatment, i.e. $\beta_{-1} = 0$, treatment effects β_ℓ can be uniquely recovered recursively as $\beta_\ell = \beta_{\ell-1} + \gamma_\ell$ for $\ell = 1, 2, \dots, \bar{\ell}$ and as

$\beta_\ell = \beta_{\ell+1} - \gamma_{\ell+1}$ for $\ell = -1, -2, \dots, \underline{\ell}$:

$$\beta_\ell = \begin{cases} -\sum_{k=\ell+1}^{-1} \gamma_k & \text{if } \ell \leq -2 \\ 0 & \text{if } \ell = -1 \\ \sum_{k=0}^{\ell} \gamma_k & \text{if } \ell \geq 0. \end{cases} \quad (12)$$

We summarize this result in the following remark:

Remark 7 (Recovery of dynamic cumulative treatment effects from a distributed-lag model).

Dynamic treatment effects β_ℓ are recovered from distributed-lag parameters γ_ℓ as cumulative effects starting from a reference period, typically the period before treatment adoption, according to equation (12).

As in the event study model, we need a normalization in the distributed-lag model since parameters β_ℓ are only identified up to a constant due to the individual fixed effect μ_i (cf. Remark 1). Equation (12) shows how to recover the dynamic treatment effects β_ℓ as the sums of distributed-lag parameters γ_ℓ . Concretely, for post-treatment effects $\ell > -1$, we intuitively cumulate upwards: $\beta_\ell = \beta_{\ell-1} + \gamma_\ell$ with $\beta_{-1} = 0$. Importantly, for pre-treatment effect $j \leq -1$, we cumulate *downwards* with a *negative* sign: $\beta_{\ell-1} = \beta_\ell - \gamma_\ell$ with $\beta_{-1} = 0$. For instance, $\beta_{-2} = -\gamma_{-1}$; we *must not* assume $\gamma_{-1} = 0$.¹²

3.4 Practical Implications

While the derivation of equivalence results of Remark 6 is mathematically straightforward, it has useful practical implications. We briefly discuss the most important ones in the following subsection.

Model choice. As dynamic treatment effects estimated from an event study design and dynamic treatment effects recovered from a distributed-lag model are numerically equivalent, the choice of the empirical model is purely a question of convenience.

¹² In our example 1 with effect window from $\underline{\ell} = -3$ to $\bar{\ell} = 4$, the coefficients are $\beta_{-3} = -(\gamma_{-1} + \gamma_{-2})$, $\beta_{-2} = -\gamma_{-1}$, $\beta_{-1} = 0$, $\beta_0 = \gamma_0$, $\beta_1 = \gamma_0 + \gamma_1$, $\beta_2 = \gamma_0 + \gamma_1 + \gamma_2$, $\beta_3 = \gamma_0 + \gamma_1 + \gamma_2 + \gamma_3$, $\beta_4 = \gamma_0 + \gamma_1 + \gamma_2 + \gamma_3 + \gamma_4$.

Binning vs. cumulating. In the event study design, treatment variables have to be binned at the endpoints of the effect window according to equation (3). Consequently, the event study model delivers direct estimates of the dynamic treatment effects and therefore readily interpretable parameters. In contrast, the coefficients from the distributed-lag model $\gamma = [\gamma_{\underline{\ell}+1}, \dots, \gamma_{\bar{\ell}}]'$ have to be cumulated following equation (12) to obtain the event study parameters $\beta = [\beta_{\underline{\ell}}, \dots, \beta_{-2}, \beta_0, \dots, \beta_{\bar{\ell}}]'$. This linear transformation transfers the statistical properties (consistency and asymptotic normality) of $\hat{\gamma}$ to the calculated $\hat{\beta}$. Standard errors of $\hat{\beta}_{\ell}$ can be calculated from the variances and covariances of the vector $\hat{\gamma}$ by the usual formula for linear combinations and are identical to the direct event study estimates.

Data Requirements. The isomorphism between ES and DL models is also insightful when thinking about the data requirements to estimate the model because obviously the same data requirements have to apply to both the ES and DL model. The DL specification with $\bar{\ell}$ lags and $|\underline{\ell}| - 1$ leads given in equation (10) immediately reveals that we need to observe treatment status $T_{i,t}$ from $(\underline{t} - \bar{\ell})$ to $(\bar{t} + |\underline{\ell}| - 1)$ if all observations of the dependent variable $Y_{i,t}$ are to be included in the estimation sample. In case we observe treatment adoption $\Delta T_{i,t} = T_{i,t} - T_{i,t-1}$, we only need to observe it from $(\underline{t} - \bar{\ell} + 1)$ to $(\bar{t} + |\underline{\ell}| - 1)$.

Fixed effect vs. first difference estimator. We have shown that the ES and the DL model can both be estimated in levels or in first differences. Both estimators are consistent and asymptotically normal as $N \rightarrow \infty$ under standard assumptions for panel data models. In finite samples, the estimates obtained with the fixed effect estimator differ from the ones obtained with the first difference estimator. If the dynamic nature of the effect is modeled correctly, i.e. if the effect is truly constant $\bar{\ell}$ periods after the event, the two estimators will be numerically similar. If, however, the true dynamic treatment effect continues to unfold beyond $\bar{\ell}$, fixed effects and first difference estimates can differ strongly. As an example, assume that the true treatment effect is negligible in the short run and only materializes after several post-treatment periods. Further, assume that the researcher includes too few post-treatment parameters to capture the full dynamic treatment effects.

In this case, the first difference estimator will be close to zero while the fixed effect estimator will pick up some average of the delayed response. Both estimators will underestimate the true long-term response. In contrast, if the number of lags is specified such that the dynamic treatment effects have fully materialized within the effect window (cf. Remarks 2 and 3), both the fixed effects and the first difference estimator will correctly pick up the dynamic effects and estimate the long-run effect and the point estimates should be similar.¹³

4 Generalized Event Study Design

In many applications, treatments may occur repeatedly and be of different signs and varying intensities across units and/or time. In this section, we show that the standard event study design can be generalized to accommodate institutional set-ups where multiple treatments with known but varying treatment intensities take place. We formally derive the generalized event study in Subsection 4.1. Subsection 4.2 discusses four typical institutional environments in which the generalized event study can be applied. Online Appendices A.2 to A.4 provide numerical examples for these cases.

4.1 Set-up and Equivalence in the Generalized Design

In the following, we set up a generalized event study design that can be used in the case of multiple treatments of identical intensities, single treatments with different signs and varying treatment intensities, and multiple treatments of different signs and varying intensities. The set-up also nests the standard event study design set up in Section 2 as a special case.

In the generalized design, the treatment status $T_{i,t}$ does no longer need to be a dummy variable but can be a continuous measure. The change of the treatment status $\Delta T_{i,t} = T_{i,t} - T_{i,t-1}$ can therefore also be a continuous measure of the change in the intensity of the treatment. We call $\Delta T_{i,t}$ *scaled treatment adoption* in the generalized design. Scaled treatment adoption $\Delta T_{i,t}$ measures the

¹³ There is also an efficiency trade-off: the fixed effect estimator is more efficient when the errors ε_{it} are not serially correlated, whereas the first difference estimator is more efficient when the errors follow a random walk (see e.g. Wooldridge, 2010).

exact size of a change in a certain policy variable (e.g. a tax rate) from unit $t - 1$ to unit t .

The event study specification in equation (1) does also apply to the generalized interpretation of $T_{i,t}$ and $\Delta T_{i,t}$. The equivalence between event study and distributed-lag models, therefore, holds in the generalized design exactly as in the standard events study specification. As in the standard event study design, the equivalence also holds when the effect window is restricted through binning of the endpoints. Importantly, the generalization does only work with the definition of the binned treatment adoption measure $D_{i,t}^{\ell}$ in equation (3) and not with the more common definition in equation (4). We, therefore, call this variable *binned treatment adoption measure* rather than binned treatment adoption indicator in the generalized design.

Remark 1 on normalization, Remarks 2 and 5 on restricting the event window, and the practical implications discussed in Section 3.4 on estimating the event study vs. the distributed-lag models also hold in the general case.

Estimating dynamic treatment effects using the generalized event study only produces consistent estimates under a linearity and additivity assumption, which is summarized in the following remark.

Remark 8 (Applicability of the Generalized Event Study Design).

Assuming that the treatment effect is proportional to the observed treatment intensity, the generalized event study described by equations (2) and (3) delivers consistent estimates of the dynamic treatment effects. Numerically identical treatment effects can also be obtained from a distributed-lag model as specified in equation (9).

Note that it may be the case that the proportionality assumption of Remark 8 is violated. For example, treatments with opposite signs may have asymmetric effects on the outcome (see, e.g. Fuest et al., 2018, Benzarti et al., 2020). It is therefore always advisable to test for symmetric effects. Even if treatments of opposite signs do not have symmetric effects, it is still valuable to exploit treatment variation within the set of positive and negative treatments (see the application in Section 5 and Online Appendix B).

As the generalized event study specification incorporates the intensity of treatment, estimated effects can be interpreted as the effect of a one-unit increase akin to the interpretation in a generalized

differences-in-differences model. This way, event study estimates can be used to infer long-term effects on an intuitive scale.

4.2 Typical Cases and Applications

In this subsection, we discuss typical cases of the generalized event study design and provide selected examples from recently published applications.

Case 1: Single Treatments of Identical Intensities. This is the standard case discussed in Section 2; a numerical example is given in Online Appendix A.1.

Case 2: Multiple Treatments of Identical Intensities. Consider the case in which events of identical intensity take place repeatedly for a unit. This implies that treatment adoption $\Delta T_{i,t} = T_{i,t} - T_{i,t-1}$ is a dummy variable that takes value 1 in *any* time period where an event took place and 0 in other periods. Treatment status $T_{i,t}$ increases accordingly by one unit with every event (see Online Appendix A.2 for a numerical example).

Few analyses have applied event studies in such an institutional context (see Dube et al., 2011, for a notable exception). However, many institutional set-ups, such as hospital admissions or firm switches, fit the case. Sometimes, only the first of potentially many events is considered in a standard event study framework as developed in Section 2. This approach leads to inconsistent estimates unless the second and subsequent events are known to have no additional effect.

Case 3: Single Treatments of Different Signs and Varying Intensities. Next, consider the case where each unit receives one treatment in period E_i , but treatment intensity s_i differs across units. Hence, scaled treatment adoption is $\Delta T_{i,t} = \mathbb{1}[t = E_i] \cdot s_i$. A numerical example is given in Online Appendix A.3.

This case is quite frequently applied as it fits an institutional setting where a shock at some aggregate level hits units at a disaggregate level with varying intensities (see, e.g., Alsan and Wanamaker, 2018, Charles et al., 2018, Clemens et al., 2018, Goodman-Bacon, 2018). Many applications of this type formally refer to the standard event study model but discuss generalization

and treatment of endpoints only verbally if at all.

Case 4: Multiple Treatments of Different Signs and Varying Intensities. Last, we consider the most general case, in which events may occur multiple times per unit and their treatment intensity varies. A numerical example is given in Online Appendix A.4.

Many settings fit this model, such as multiple tax changes or minimum wage hikes, and correspondingly many applications. Traditionally, the respective models were framed as distributed-lag models rather than event study designs (Suárez Serrato and Zidar, 2016, Drechsler et al., 2017, Fuest et al., 2018, Brülhart et al., 2022). Note that the notion of events becomes less pertinent the more such events happen. In the extreme case, treatment may change each period and $T_{i,t}$ becomes a constantly changing continuous variable.

4.3 Dichotomizing treatment variables.

A common alternative empirical specification used when treatments have varying intensities is to dichotomize treatment variables and use a dummy variable that is only switched on for relatively large events (see, e.g., Simon, 2016, Fuest et al., 2018). However, the parameter estimates of such a dichotomization may be difficult to interpret both in magnitude and direction. To see this, consider the following case: each unit is treated once, and there are two types of treatment: a small reform $\Delta T_{i,t}^S = 1$ and a large reform $\Delta T_{i,t}^L = 2$; treatments are distributed randomly in time and treatment effects are linear in the intensity of the reform. Ignoring small events and applying the standard event dummy set-up yields $\Delta T_{i,t}^S = 0$, $\Delta T_{i,t}^L = 1$. In this case, units with small reforms become part of the control group although they respond to the reform. This induces a bias in the time fixed effects and thereby also in the treatment coefficients. In such a specification, estimated dynamic treatment effects can be larger, smaller, or identical to a specification that considers all reforms with their intensity. A possible fix for this ambiguity is to exclude units with small events from the sample, in which case, the model is, however, estimated on a different and possibly selected sample. Moreover, the dichotomization of the treatment variable eliminates valuable information which could otherwise be used to identify the magnitude of the effect.

5 Application

In this section, we demonstrate the relevance of the results derived in Sections 2 to 4 by replicating and extending the study by Baker and Fradkin (2017) (BF2017).¹⁴ We will particularly focus on the importance of restricting the effect window and on the power of the generalized event study design. In the following, we discuss the high-level results and refer to Online Appendix B for a more detailed, self-contained analysis.


BF2017 study the effect of the changes in the potential duration of unemployment benefits (PBD) across US states between 2006 and 2015 on job search effort using Google Search data. They apply a two-way fixed effects model with state and year-month fixed effects and estimate a standard difference-in-difference model. They find a highly significant effect implying that a ten-week increase in PBD leads to 2.07% drop in aggregate job search. However, their dynamic event studies only weakly support this finding.

When setting up their event study model, BF2017 make several modeling choices that directly speak to many of the Remarks developed above. First, there is no normalization to a reference period (Remark 1). Second, the authors assume a rather small effect window around the reform of 3 to 4 months (Remark 3) but do not bin endpoints – implicitly assuming that treatment effects are zero beyond the effect window (Remarks 2 and 5). Finally, they dichotomize the changes in PBD into increase and decrease dummy variables and do not exploit that state-level PBD reforms had varying intensities (Remark 8).

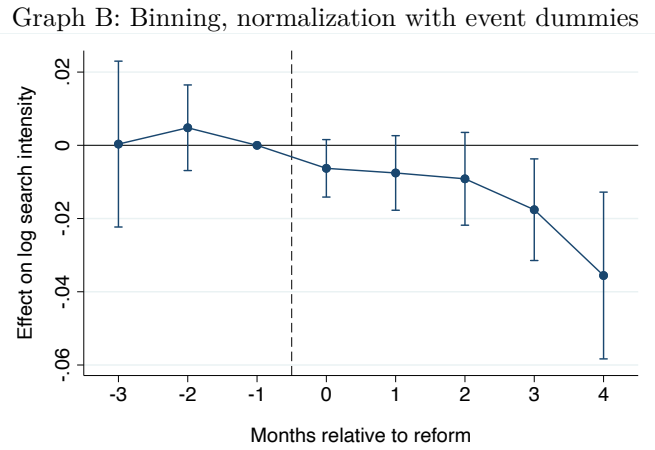
In Online Appendix B, we show step-by-step that making these adjustments leads to convincing and highly significant dynamic estimates. Figure 3 summarizes the main results. Graph A replicates BF2017’s main event study results; in graph B we adjust BF’s model by normalizing to a reference period and binning endpoints. While Graph B still uses a dummy variable specification that only considers large reforms as events, we implement the generalized event study design in Graph C, exploiting all PBD reforms with their varying intensities. The estimates of the generalized event

¹⁴ Replication code in Stata and R is available in the Harvard Dataverse at <https://doi.org/10.7910/DVN/LXMYV6> and in the JAE data archive at <https://dx.doi.org/10.15456/jae.2023031.1844400605>.

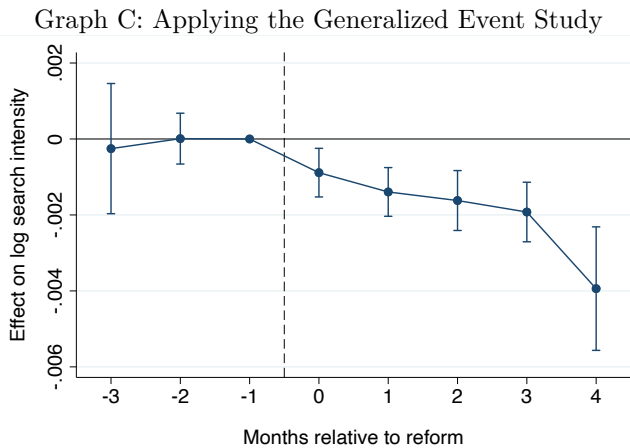
Figure 3: Main Replication Results



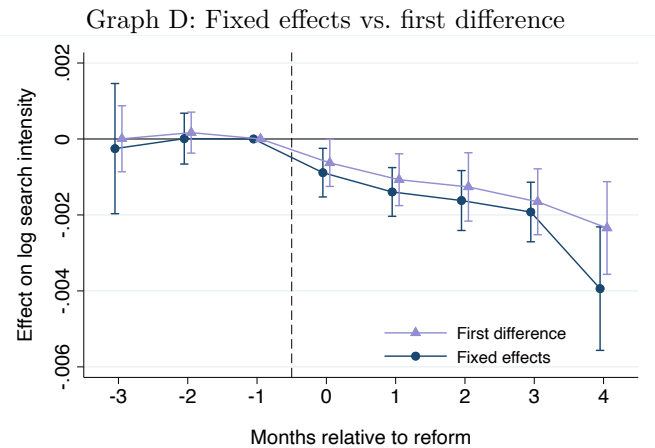
Observations: 3672, states: 51, periods: 72 (1/2006 - 12/2011).



Observations: 3468, states: 51, periods: 68 (5/2006 - 12/2011).



Observations: 3468, states: 51, periods: 68 (5/2006 - 12/2011).



Observations: 3468, states: 51, periods: 68 (5/2006 - 12/2011).

Notes: The figure replicates and extends the main event study estimates reported in (Baker and Fradkin, 2017), BF2017. Graphs show point estimates and 95%-confidence intervals based on standard errors clustered by states. Graphs A replicates the estimates reported in specification (3) of BF2017-Table 8 and plotted in the left panels of BF2017-Figure 4. It shows the dynamic effect of a large increase (at least 13 weeks) in potential benefit duration (PBD) on log search intensity as measured with the newly proposed Google Job Search Index (GJSI). States that experience no changes in a certain month or smaller changes, including negative ones are in the control group. In Graph B, we extend the original specification by binning endpoints according to Remark 2 and by normalizing the effect at the pre-event period to zero according to Remark 1. In Graph C, we apply the generalized event study according to Remark 8, accounting for treatment intensity. Models underlying Graphs A to C are estimated in levels with state and time fixed effects. In Graph D, we contrast estimates from the generalized event study in a fixed effects model to a specification in first differences.

study imply that a ten-week increase in potential benefit duration leads to a decrease in log job search activity of -0.027 (s.e. = 0.007), i.e. 2.7%, after 4 months, which corresponds to the difference-in-difference estimates reported above. Hence, the event study results reinforce BF2017's important main result that longer potential benefit duration in unemployment insurance reduces search effort.

In Graph D we contrast the baseline fixed effects model with a first difference specification. At the endpoint of the effect window, the first difference model only accounts for the change happening from $\bar{\ell} - 1$ to $\bar{\ell}$, while the fixed effects model takes into account a weighted average of the remaining changes. As a result, coefficients from the fixed effects and the first difference specification will deviate if the effect has not fully materialized within the given effect window. The figure clearly shows a deviation between the two specifications, which is an indication that the effect window might be too short.¹⁵ In Figure B.3 in Online Appendix B, we show that dynamic treatment effects have fully materialized approximately after 16 months, with a long-run effect of PBD on search intensity of -0.005 (0.001).

6 Conclusion

This paper makes three interrelated methodological contributions, which are important to bear in mind when setting up event study designs in economics. The points are valid in general and might be particularly helpful when applying the event study technique to settings in public and labor economics with multiple policy shocks of varying intensities.

First, researchers need to define an effect window, i.e. the window within which the effect is studied. While this choice is a practical necessity due to limited data availability, it is far from being innocuous. Setting the number of leads and lags to a finite number practically requires defining the last lag (lead) as an open interval capturing all known events that (will) have happened in the past (future). We refer to this practice as binning. We show that binning determines which unit-period observations are assigned to the treatment or control group and thus directly affects the identifying assumption. At the same time, binning introduces important parameter restrictions, which help to identify the model econometrically.

Second, we demonstrate that event study designs and distributed-lag models are equivalent. To be precise, the distributed-lag model is a reparametrization of an event study design *with binned*

¹⁵ Note that there is no ex-ante prediction on whether the effect of the first-difference or fixed effects model should be smaller or larger. As the sample size grows both models should eventually yield identical estimates.

endpoints. Event study estimates can be recovered from distributed-lag models by cumulating the post-treatment and pre-treatment effects away from zero. We use this isomorphism to reinforce the necessity and importance of limiting the effect window properly and critically discuss the plausibility of alternative parameter restrictions used in the literature. The distributed-lag model is in our view less error-prone in the practical implementation.

Third, we generalize the simple event study with single treatments to account for multiple treatments and/or treatments of different signs and varying intensities. We show that the event study methodology is perfectly applicable to such environments and that the equivalence between event study and distributed-lag models also holds in the general case. We point to the necessary underlying assumptions and briefly discuss where generalized event study designs could be implemented in light of current empirical research.

In the final part of the paper, we demonstrate the practical relevance of our three methodological contributions by replicating and extending the event study design in Baker and Fradkin (2017).

References

- Alsan, M. and M. Wanamaker (2018). “Tuskegee and the Health of Black Men”. *Quarterly Journal of Economics* 133 (1), 407–455.
- Angrist, J. D. and J.-S. Pischke (2010). “The Credibility Revolution in Empirical Economics: How Better Research Design Is Taking the Con out of Econometrics”. *Journal of Economic Perspectives* 24 (2), 3–30.
- Autor, D. (2003). “Outsourcing at Will: The Contribution of Unjust Dismissal Doctrine to the Growth of Employment Outsourcing”. *Journal of Labor Economics* 21 (1), 1–42.
- Baker, A. C., D. F. Larcker, and C. C. Y. Wang (2022). “How Much Should We Trust Staggered Difference-In-Differences Estimates”. *Journal of Financial Economics* 2, 370–395.
- Baker, S. R. and A. Fradkin (2017). “The Impact of Unemployment Insurance on Job Search: Evidence from Google Search Data”. *The Review of Economics and Statistics* 99 (5), 756–768.

- Benzarti, Y., D. Carloni, J. Harju, and T. Kosonen (2020). “What Goes Up May Not Come Down: Asymmetric Pass through of Value Added Taxes”. *Journal of Political Economy* 128 (12), 4438–4474.
- Borusyak, K., X. Jaravel, and J. Spiess (2022). “Revisiting Event Study Designs: Robust and Efficient Estimation”. *arXiv:2108.12419*.
- Brülhart, M., J. Gruber, M. Krapf, and K. Schmidheiny (2022). “Behavioral Responses to Wealth Taxes: Evidence from Switzerland”. *American Economic Journal: Economic Policy* 4 (14).
- Callaway, B., A. Goodman-Bacon, and P. H. C. Sant’Anna (2021). “Difference-in-Differences with a Continuous Treatment”. *arXiv:2107.02637*.
- Callaway, B. and P. H. C. Sant’Anna (2021). “Difference-in-Differences with Multiple Time Periods”. *Journal of Econometrics* 225 (2), 200–230.
- Cameron, A. C. and P. K. Trivedi (2005). *Microeconometrics: Methods and Applications*. New York: Cambridge University Press.
- Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019). “The Effect of Minimum Wages on Low-Wage Jobs”. *The Quarterly Journal of Economics* 134 (4), 1405–1454.
- Chaisemartin, C. de and X. D’Haultfoeuille (2020). “Two-way fixed effects estimators with heterogeneous treatment effects”. *American Economic Review* 110 (9), 2964–2996.
- (2022a). “Difference-in-Differences Estimators of Intertemporal Treatment Effects”. *arXiv:2007.04267*.
- (2022b). “Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey”. *Econometrics Journal* (*forthcoming*).
- (2022c). “Two-way Fixed Effects Regressions with Several Treatments”. *arXiv:2012.10077*.
- Charles, K. K., E. Hurst, and M. Notowidigdo (2018). “Housing Booms and Busts, Labor Market Opportunities, and College Attendance”. *American Economic Review* 108 (10), 2947–2994.
- Clemens, M. A., E. G. Lewis, and H. M. Postel (2018). “Immigration Restrictions as Active Labor Market Policy: Evidence from the Mexican Bracero Exclusion”. *American Economic Review* 108 (6), 1468–1487.

- Dolley, J. C. (1933). “Characteristics and Procedure of Common Stock Split-Ups”. *Harvard Business Review* (11), 316–326.
- Drechsler, I., A. Savov, and P. Schnabl (2017). “The Deposits Channel of Monetary Policy”. *The Quarterly Journal of Economics* 132 (4), 1819–1876.
- Dube, A., E. Kaplan, and S. Naidu (2011). “Coups, Corporations, and Classified Information”. *Quarterly Journal of Economics* 126 (3), 1375–1409.
- Freyaldenhoven, S., C. Hansen, J. Pérez Pérez, and J. M. Shapiro (2022). “Visualization, identification, and estimation in the linear panel event-study design”. *Advances in Economics and Econometrics: Twelfth World Congress*.
- Freyaldenhoven, S., C. Hansen, and J. M. Shapiro (2019). “Pre-event Trends in the Panel Event-study Design”. *The American Economic Review* 109 (9).
- Fuest, C., A. Peichl, and S. Sieglöcher (2018). “Do Higher Corporate Taxes Reduce Wages? Micro Evidence from Germany”. *American Economic Review* 108 (2), 393–418.
- Gardner, J. (2022). “Two-stage differences in differences”. *arXiv:2207.05943*.
- Gibbons, C. E., J. C. Suárez Serrato, and M. B. Urbancic (2019). “Broken or Fixed Effects?” *Journal of Econometric Methods* 8 (1), 1–12.
- Goodman-Bacon, A. (2018). “Public Insurance and Mortality: Evidence from Medicaid Implementation”. *Journal of Political Economy* 126 (1), 216–262.
- (2021). “Difference-in-Differences with Variation in Treatment Timing”. *Journal of Econometrics* 225 (2), 254–277.
- Leeb, H. and B. M. Pötscher (2005). “Fairness and Redistribution”. *Econometric Theory* 21, 21–59.
- Lichter, A., M. Löffler, I. Isphording, T.-V. Nguyen, F. Pöge, and S. Sieglöcher (2022). “Profit Taxation, R&D Spending, and Innovation”. *ECONtribute Discussion Paper No. 202/2022*.
- MacKinlay, A. C. (1997). “Event Studies in Economics and Finance”. *Journal of Economic Literature* 35 (1), 13–39.
- Malani, A. and J. Reif (2015). “Interpreting pre-trends as anticipation: Impact on estimated treatment effects from tort reform”. *Journal of Public Economics* 124, 1–17.

- McCrary, J. (2007). “The Effect of Court-Ordered Hiring Quotas on the Composition and Quality of Police”. *The American Economic Review* 97 (1), 318–353.
- Powell, D. (2021). “The Labor Supply Consequences of the Opioid Crisis”. *RAND Discussion Paper*.
- Rambachan, A. and J. Roth (2022). “A More Credible Approach to Parallel Trends”. *Review of Economic Studies* (forthcoming).
- Roth, J. (2022). “Pre-test with Caution: Event-study Estimates After Testing for Parallel Trends”. *American Economic Review: Insights* 4 (3), 305–322.
- Roth, J., P. H. C. Sant’Anna, A. Bilinski, and J. Poe (2022). “What’s Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature”. *Journal of Econometrics* (forthcoming).
- Siegloch, S., N. Wehrhöfer, and T. Etzel (2022). “Spillover, Efficiency and Equity Effects of Regional Firm Subsidies”. *ECONtribute Discussion Paper No. 210/2022*.
- Simon, D. (2016). “Does Early Life Exposure to Cigarette Smoke Permanently Harm Childhood Welfare? Evidence from Cigarette Tax”. *American Economic Journal: Applied Economics* 8 (4), 128–159.
- Smith, M., D. Yagan, O. Zidar, and E. Zwick (2019). “Capitalists in the Twenty-First Century”. *The Quarterly Journal of Economics* 134 (4), 1675–1745.
- Suárez Serrato, J. C. and O. Zidar (2016). “Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms”. *The American Economic Review* 106 (9), 2582–2624.
- Sun, L. and S. Abraham (2021). “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects”. *Journal of Econometrics* 225 (2), 175–199.
- Wolfers, J. (2006). “Did Unilateral Divorce Laws Raise Divorce Rates? A Reconciliation and New Results”. *The American Economic Review* 96 (5), 1802–1820.
- Wooldridge, J. (2010). *Econometric analysis of cross section and panel data*. 2nd Edition. Cambridge: MIT Press.