

Difference-in-Differences Designs: A Practitioner’s Guide[†]

ANDREW BAKER, BRANTLY CALLAWAY, SCOTT CUNNINGHAM,
ANDREW GOODMAN-BACON, AND PEDRO H. C. SANT’ANNA^{*}

Difference-in-differences (DiD) is arguably the most popular quasi-experimental research design. Its canonical form, with two groups and two periods, is well understood. However, empirical practices can be ad hoc when researchers go beyond that simple case. This article provides an organizing framework for discussing different types of DiD designs and their associated DiD estimators. It discusses covariates, weights, handling multiple periods, and staggered treatments. The organizational framework, however, applies to other extensions of DiD methods as well. (JEL C23, H75, I12, I38)

1. Introduction

Dating to the 1840s, difference-in-differences (DiD) is now the most common research design for estimating causal effects in the social sciences.¹ A basic DiD design requires two time periods, one before and one after some treatment begins, and two groups, one that receives a treatment and one that does not. The DiD estimate equals the change in outcomes for the treated group minus the change in outcomes for the untreated group: the difference of two differences. If the average change in the outcomes would have been the same in the two groups had the treatment not occurred, which is referred to as a “parallel trends” assumption, this comparison estimates the average treatment effect among treated units.

In practice, however, researchers apply DiD methods to situations that are more complicated than the classic two-period and two-group (2×2) setup. Most datasets cover multiple periods and units may enter (or exit) treatment at different times. Treatment might also vary in its amount or intensity. Other variables are often used to make treated and untreated units

^{*}Baker: University of California, Berkeley. Callaway: University of Georgia. Cunningham: Baylor University. Goodman-Bacon: Opportunity and Inclusive Growth Institute, Federal Reserve Bank of Minneapolis. Sant’Anna: Emory University.

[†] Go to <https://doi.org/10.1257/jel.20251650> to visit the article page and view author disclosure statement(s).

¹Currie, Kleven, and Zwierns (2020) find that almost 25 percent of all NBER empirical working papers and 17 percent of empirical articles in five leading general-interest economics journals in 2018 mention DiD. The earliest DiD applications we are aware of are from Ignaz Semmelweis from the 1840s (Semmelweis 1983) and Snow (1855). For a brief overview of the long history of DiD in economics, see section 2 of Lechner (2011).

more comparable. Today's typical DiD study includes at least one of these deviations from the canonical 2×2 setup.

For many years, the common practice in applied research was to estimate complex DiD designs using linear regressions with unit and time fixed effects (two-way fixed effects, henceforth TWFE). Their identifying assumptions and interpretation were informally traced to the fact that in the 2×2 case a TWFE estimator gives the same estimate as a DiD estimator calculated directly from sample means, and thus inherits a clear causal interpretation under a specific parallel trends identification assumption. This appeared to justify the use of a single technique for any type of design or specification. Recent research, however, has shown that simple regressions can fail to estimate meaningful causal parameters when DiD designs are complex and treatment effects vary, producing estimates that are not only misleading in their magnitudes but potentially of the wrong *sign*. The significance of these findings is substantial; given the prevalence of DiD analysis in modern applied econometrics work, common empirical practices have almost certainly yielded misleading results in several concrete cases (Baker, Larcker, and Wang 2022).

So, what should applied researchers do instead? This paper proposes a unified framework for discussing and conducting DiD studies that is rooted in the principles of causal inference in the presence of treatment effect heterogeneity. The central conclusion of recent methodological research is that even complex DiD studies can be understood as aggregations of 2×2 comparisons between one set of units for which treatment changes and another set for which it does not. This fact links a wide variety of DiD designs used in practice and guides methodological choices about estimating them. Viewing DiD studies through the lens of 2×2 “building blocks” aids in interpretability by clarifying that they yield causal quantities that aggregate the treatment effects identified by each 2×2 component. It also means that identification comes from the simple parallel trends assumptions required for each 2×2 building block. Practically, this framework suggests first estimating each 2×2 building block and then aggregating them. As long as the effective sample size is large, this approach allows for asymptotically valid inference using standard techniques.

This framework is a “forward-engineering” approach to DiD that embraces treatment effect heterogeneity and constructs estimators that recover well-motivated causal parameters under explicitly stated assumptions. By fixing the goals of the study (the target parameters) and deriving analytical techniques, forward engineering provides clear benefits over “reverse-engineering” approaches that begin with a familiar regression specification and derive the assumptions under which it has *some* causal interpretation. The methods we describe in this paper combine familiar techniques with some newer ones, but expressly avoid the difficulties of interpretation inherent in common regression estimators (Goodman-Bacon 2021; de Chaisemartin and D’Haultfœuille 2020; Sun and Abraham 2021; Borusyak, Jaravel, and Spiess 2024). Moreover, the interpretation of common regression estimators changes across specifications, which makes it hard to understand the difference between non-robustness and a shifting target parameter. In contrast, our proposed framework naturally leads to estimation procedures that target the same parameter under different transparent identification assumptions. Thus, two estimates can be distinguished easily by their identifying assumptions. Finally, the principles of the forward-engineering approach provide guidance to good econometric practices, even in settings without well-established methodological findings.

This paper is not designed to be a comprehensive literature review; its goal is to provide guidelines for practitioners who want to better understand DiD and its various forms.

Because of the tremendous variations in design, data, and specification that practitioners encounter, we opt to focus on three of the most common aspects of modern DiD studies: the use of weights, covariates, and staggered treatment timing. Table A1 includes a list of the acronyms that we, and the econometrics literature on DiD, use to distinguish different methods. We apply techniques to address these issues to a specific example: the causal effect of recent public health insurance expansions in the United States on county-level mortality. Our replication materials include data as well as R and Stata code that can serve as a template for any DiD study using these methods. In a Supplemental Appendix we briefly discuss related DiD designs with different treatment variables (ones that turn on and off or take many values), additional comparisons (i.e., triple-difference designs), distributional target parameters, or different data structures (repeated cross-sections or unbalanced panels). Several recent reviews follow the logic laid out here and cover additional DiD-related topics and technical details: Roth et al. (2023), de Chaisemartin and D’Haultfœuille (2023), and Callaway (2023).

The rest of the paper is structured as follows. Section 2 introduces the Medicaid example. Section 3 discusses the canonical 2×2 DiD setups with and without weights, and Section 4 discusses threats to the identification assumptions, how to assess them, and how to incorporate covariates. Section 5 extends the 2×2 setup to multiple periods with potentially staggered treatment adoption. Section 6 concludes and briefly discusses some extensions that involve more complex DiD designs.

2. *Medicaid and Mortality: The Running Example*

To make our methodological discussion concrete, we revisit a timely and important causal question: How did the expansion of public health insurance (Medicaid) under the Affordable Care Act (ACA) affect mortality?

Medicaid expansion is a great example of a staggered treatment adoption. The ACA originally mandated that in 2014 all states expand Medicaid eligibility to adults with incomes up to 138 percent of the federal poverty threshold. In upholding the law’s constitutionality in a 2012 decision, however, the Supreme Court made Medicaid expansion optional. As a result, many states expanded Medicaid after 2014, but several have not done so as of 2024. Columns 1 and 2 of Table 1 illustrate the variation in Medicaid expansion dates provided by the Kaiser Family Foundation (2026).

States expanded Medicaid largely because of economic and political considerations (Sommers and Epstein 2013), which created observable differences between expansion and non-expansion states. For instance, just four out of the 22 states that expanded Medicaid in 2014 are in the southern census region; conversely, seven out of 10 non-expansion states are in the South. This suggests a potential role for covariates when analyzing Medicaid expansion.

Finally, mortality is measured in jurisdictions like states and counties, which are of very different sizes. Choices about (population) weights determine not only how different estimation approaches average the units within a given expansion group, but also how a given estimation technique averages estimated effects across those groups. California, for example, represented 4.5 percent of the states that expanded Medicaid in 2014, 5.4 percent of the counties, but 27.7 percent of the adults ages 20–64; its contribution to “the” average outcome for the 2014 expansion group is very different with weights than without. The final three columns of Table 1 show that, in our data the entire 2014 expansion group contains 44 percent of the

TABLE 1
 MEDICAID EXPANSION UNDER THE AFFORDABLE CARE ACT

Expansion year	States	Share of states	Share of counties	Share of adults (2013)
Pre-2014	DE, MA, NY, VT	0.08	0.03	0.09
2014	AR, AZ, CA, CO, CT, HI, IA, IL, KY, MD, MI, MN, ND, NH, NJ, NM, NV, OH, OR, RI, WA, WV	0.44	0.36	0.45
2015	AK, IN, PA	0.06	0.06	0.06
2016	LA, MT	0.04	0.04	0.02
2019	ME, VA	0.04	0.05	0.03
2020	ID, NE, UT	0.06	0.04	0.02
2021	MO, OK	0.04	0.06	0.03
2023	NC, SD	0.04	0.05	0.03
Non-expansion	AL, FL, GA, KS, MS, SC, TN, TX, WI, WY	0.20	0.31	0.26

Notes: The table shows which states adopted the ACA's Medicaid expansion in each year as well as the share of all states, counties, and adults in each expansion year.

states, 36 percent of the counties, but 45 percent of all adults. Weighting will therefore change how important the estimated treatment effects are for the 2014 group.

Several recent papers study the effect of ACA Medicaid expansion on mortality rates for lower-income adults, who are most likely to gain insurance through Medicaid. Miller, Johnson, and Wherry (2021) and Wyse and Meyer (2025) use simple DiD methods to provide evidence that Medicaid reduced adult mortality rates for targeted subpopulations. Unfortunately, their analyses require restricted links between income and mortality data, which are important for overcoming the low statistical power in studies using aggregate mortality data (Black et al. 2022). Our goal is to pursue a replicable and shareable example based on a related analysis by Borgschulte and Vogler (2020). They use a sophisticated strategy to select and use covariates in a weighted TWFE regression using restricted access data and find that Medicaid expansion reduced aggregate county-level mortality rates. We use publicly available data, which we include in a fully reproducible replication package, and consider only a handful of intuitive demographic and economic covariates sufficient to illustrate several practical challenges that can arise with DiD.² This empirical exercise is meant solely to illustrate how to tackle several common features of DiD designs. The results are pedagogical in spirit and do not represent the best possible estimates of Medicaid's effect on adult mortality.

Our outcome variable is the crude adult mortality rate, $Y_{i,t}$, for people ages 20–64 (measured per 100,000) by county (i) from 2009 to 2019 released by the Centers for Disease Control and Prevention (2009–19).³ We denote county i 's adult population in 2013 by W_i and its socioeconomic covariates in year t (discussed below) by $X_{i,t}$. The information in

²These covariates include the percentages of a county's population that are female, white, or Hispanic from the Centers for Disease Control and Prevention (2009–19), the unemployment rate from the Bureau of Labor Statistics (2009–19), and county-level median income (in thousands of dollars) and the poverty rate from US Census Bureau (2009–19).

³It is common to adjust mortality rates by the county age distribution. Unfortunately, the CDC measurements of age-specific deaths are restricted for counties with fewer than 10 annual deaths. We aim to use publicly available and shareable data for pedagogical purposes; we follow Borgschulte and Vogler (2020) and use the crude mortality rate.

Table 1 defines the treatment group variable G_i that equals the year in which county i 's state expanded Medicaid; $G_i = \infty$ for the non-expansion states. Our final sample contains 2,604 counties in states with complete data on mortality rates from 2009 to 2019 and covariates for 2013 and 2014.

Faced with a setup such as this, researchers need to make a range of tightly related choices. Which treatment groups in Table 1 should be compared with each other and over what time horizons? What must be true for those comparisons to identify causal effects and how should one empirically evaluate their plausibility? How can other information, such as covariates or pre-period outcomes, be used to improve the credibility of the design? How do these methodological choices affect the causal interpretation of a given analysis? The aim of this review is to demonstrate to practitioners using DiD in realistic scenarios why and how to use state-of-the-art econometric tools to answer these questions.

3. 2×2 DiD Designs

We begin our discussion by focusing on the canonical 2×2 DiD setup, which has two time periods—one before and one after treatment—and two groups—one that remains untreated in both periods and one that becomes treated in the second period. In our Medicaid example, we focus on comparisons in 2014 versus 2013 between the 2014 expansion group (978 counties) and the group that had not expanded by 2019 (1,222 counties). When we consider more complex designs, this kind of comparison will still play a role: It will be one 2×2 “building block” among many.

Using these basic ingredients, we can now define a 2×2 DiD *design* composed of a causal target parameter, a treatment variable, an assumption under which it is identified, and an estimation approach, which will be the classic difference of two differences. This may be familiar territory in the simple case, but it is a crucial framework for building up appropriate techniques in more complex cases.

We first define a treatment group dummy, D_i , that equals one for the treated units (expansion states, $G_i = 2014$) and zero for the untreated units (states that had not expanded by 2019, $G_i > 2019$). The treatment *status* dummy, $D_{i,t} = D_i \times 1\{t \geq 2014\}$, then equals one for counties in 2014 expansion states during post-expansion years. To highlight how weights enter different kinds of DiD analyses, we use the following notation for expected values. For generic random variables A and C , for a given set of nonnegative weights, ω , define $\mathbb{E}_\omega[A|C] = \mathbb{E}[\omega A|C]/\mathbb{E}[\omega|C]$ as the ω -weighted population expectation of A given C . When $\omega = 1$ for all units in the population, we simply write $\mathbb{E}_\omega[A|C] = \mathbb{E}[A|C]$. Henceforth, unless otherwise noted, we assume that we have a balanced panel data random sample of $(Y_{t=1}, \dots, Y_{t=T}, G, X)$.

3.1 *Causal Effects and Target Parameters: The Average Treatment Effect on the Treated (ATT)*

The first step of any causal analysis is to define the causal quantity of interest, also called the target parameter. We use the potential outcomes framework of Rubin (1974) and Robins (1986) to do so. Let $Y_{i,t}(0,0)$ denote unit i 's potential outcome at time t if it remained untreated in both periods. Analogously, let $Y_{i,t}(0,1)$ denote unit i 's potential outcome at time t if untreated in the first period but exposed to treatment by the second period. In our example, $Y_{i,t}(0,0)$ is county i 's mortality rate in period t in a world in which Medicaid did not

expand in its state, and $Y_{i,t}(0,1)$ is its mortality rate in a world in which Medicaid did expand in 2014.⁴ To simplify notation, we will write $Y_{i,t}(0) = Y_{i,t}(0,0)$ and $Y_{i,t}(1) = Y_{i,t}(0,1)$, as the potential outcomes are defined by treatment exposure in period two (Medicaid expansion status by 2014). Nonetheless, it will be useful for later discussions that these potential outcomes correspond to treatment sequences.

In practice, we never observe $Y_{i,t}(1)$ and $Y_{i,t}(0)$ for the same unit. Instead, the data we observe, $Y_{i,t}$, are treated outcomes, $Y_{i,t}(1)$, for treated units and untreated outcomes, $Y_{i,t}(0)$, for untreated units, as in the following equation:

$$(1) \quad Y_{i,t} = (1 - D_i)Y_{i,t}(0) + D_iY_{i,t}(1).$$

We additionally assume that county mortality rates were not affected by the Medicaid expansion *before* Medicaid expanded, which is crucial to the validity of the DiD estimator (see, e.g., Abbring and van den Berg 2003; Malani and Reif 2015; and Roth et al. 2023). This standard “no anticipation” assumption ensures that we observe untreated potential outcomes before Medicaid expansion takes effect: $Y_{i,2013} = Y_{i,2013}(0)$. It also helps us define effective treatment dates. For instance, if the announcement of Medicaid expansion affected mortality before its actual expansion, “treatment” would begin when the policy was announced rather than implemented. We formally state this assumption for completeness and maintain it throughout the paper.

ASSUMPTION NA (No-Anticipation): *For all treated units i and all pretreatment periods t , $Y_{i,t}(1) = Y_{i,t}(0)$.*

The potential outcomes define a causal effect for every unit in every time period, $Y_{i,t}(1) - Y_{i,t}(0)$. These describe what Medicaid expansion did to mortality rates in a specific treated county or what it would have done in a specific untreated county. This framework allows for arbitrary heterogeneity in the effects across units and time; that is, it allows the effect of Medicaid expansion to be different in every county and year. But it is hard to learn about this degree of rich heterogeneity without additional strong conditions holding. Instead, DiD analyses typically seek to estimate (weighted) averages of heterogeneous treatment effects. In particular, most DiD designs target the average treatment effect on the treated at time t , or $ATT(t)$:

$$(2) \quad \begin{aligned} ATT(t) &= \mathbb{E}_\omega[Y_{i,t}(1) - Y_{i,t}(0) | D_i = 1] \\ &= \mathbb{E}_\omega[Y_{i,t} | D_i = 1] - \mathbb{E}_\omega[Y_{i,t}(0) | D_i = 1]. \end{aligned}$$

Equation (2) shows that $ATT(t)$ compares (weighted) average *observed* post-expansion mortality rates among treated counties ($\mathbb{E}_\omega[Y_{i,t} | D_i = 1]$) with the (weighted) average untreated mortality rates for the same treated counties ($\mathbb{E}_\omega[Y_{i,t}(0) | D_i = 1]$). The second quantity is

⁴We have implicitly introduced the stable unit treatment value assumption, which holds if the only treatment determining county i 's mortality rate is its own. In other words, Medicaid expansion in neighboring counties must not affect deaths in county i . If this fails, then there are effectively many different treatment variables and counterfactuals. The potential outcomes notation and ensuing analysis would then need to account for this.

counterfactual because untreated outcomes are never observed for treated counties. By the no-anticipation assumption, $ATT(t) = 0$ for all pretreatment periods; that is, $ATT(2013) = 0$ in our two-period Medicaid example. This ensures that cross-group outcome comparisons before treatment begins reflect *untreated* potential outcome gaps, which is central to the logic of DiD. Note that we abuse notation and omit the weight index when defining ATT 's; we do that to unclutter notation throughout the paper.

Equation (2) shows that weighting enters the analysis early on, as part of the definition of the causal parameter. In the Medicaid context, the unweighted $ATT(2014)$ answers the question, What was the average causal effect of Medicaid expansion on 2014 mortality rates among the 2014 expansion state counties? The weighted $ATT(2014)$ answers the question, What was the average causal effect of Medicaid expansion on 2014 mortality rates among adults in counties in states that expanded Medicaid in 2014? This point interacts with other justifications for weighting, such as improving precision. With heterogeneous treatment effects, adopting a weighting scheme designed to improve precision in the presence of heteroskedasticity in a constant-coefficient regression model will also change the target parameter, potentially by a lot when treatment effects are correlated with the weights (Solon, Haider, and Wooldridge 2015). Comparing weighted and unweighted estimates, therefore, does not show whether weighting matters for estimation or inference; these reflect different target parameters. In our example, the population-weighted parameter is probably more policy-relevant, but we conduct some of our empirical exercises both ways to show how weighting can affect a given DiD result. Which parameter is “of interest” is an argument about theoretical importance, policy relevance, and the use to which it will be put.

Other target parameters are also possible. Designs other than DiD identify different kinds of average treatment effects, and some DiD methods use quantile regression (Athey and Imbens 2006; Callaway and Li 2019) or distribution regression (Fernández-Val et al. 2024) approaches to target features of the marginal distributions of $Y_{i,t}(1)$ and $Y_{i,t}(0)$ among treated units. We focus on identification and estimation strategies that target ATT parameters but emphasize that the 2×2 building-block framework applies to DiD methods more broadly; see our Supplemental Appendix for more discussions about distributional parameters.

3.2 *Identifying Assumptions: Parallel Trends*

A research design is a strategy—a set of assumptions—to identify and estimate specific target parameters. Many different assumptions can identify the missing counterfactual for $ATT(2014)$ in the Medicaid example. For example, mean independence between $Y_{i,2014}(0)$ and D_i implies that the counterfactual equals average 2014 mortality rates in non-expansion counties ($\mathbb{E}_\omega[Y_{i,2014}(0) | D_i = 0]$). Under this assumption, which essentially entails assuming that Medicaid expansion is as good as random, the cross-sectional mortality gap in 2014 between expansion and non-expansion counties is the $ATT(2014)$. Similarly, time invariance of $Y_{i,t}(0)$ among expansion counties (plus the fact that we ruled out anticipatory behavior) implies that the counterfactual equals 2013 mortality rates in expansion counties ($\mathbb{E}_\omega[Y_{i,2013}(0) | D_i = 1]$). Under this assumption, which essentially rules out nontreatment-related changes in the outcome variable, the “time trend” in average mortality in expansion counties is the $ATT(2014)$.

DiD comes from an alternative assumption that identifies the relevant counterfactual even when the average untreated potential outcome differs across treatment groups (which violates mean independence) and changes over time (which violates time invariance). The parallel

trends assumption states that in the absence of treatment, the average outcome evolution is the same among treated and comparison groups. For general assumptions and results, we denote time periods by $t = 1, 2$, but continue to be explicit about which years are being used when we reference the Medicaid example.

ASSUMPTION PT (2×2 Parallel Trends): *The (weighted) average change of $Y_{i,t=2}(0)$ from $Y_{i,t=1}(0)$ is the same between treated and comparison groups; that is,*

$$(3) \quad \mathbb{E}_\omega[Y_{i,t=2}(0) | D_i = 1] - \mathbb{E}_\omega[Y_{i,t=1}(0) | D_i = 1] = \mathbb{E}_\omega[Y_{i,t=2}(0) | D_i = 0] - \mathbb{E}_\omega[Y_{i,t=1}(0) | D_i = 0].$$

If the parallel trends assumption holds, then it is easy to construct $\mathbb{E}_\omega[Y_{i,t=2}(0) | D_i = 1]$ from observable quantities—that is, to identify it:

$$(4) \quad \mathbb{E}_\omega[Y_{i,t=2}(0) | D_i = 1] = \mathbb{E}_\omega[Y_{i,t=1} | D_i = 1] + (\mathbb{E}_\omega[Y_{i,t=2} | D_i = 0] - \mathbb{E}_\omega[Y_{i,t=1} | D_i = 0]).$$

In the Medicaid example, assumption PT says that to calculate expansion counties' average 2014 mortality rate in a counterfactual world without Medicaid expansion, start with their average 2013 mortality rate and add the observed change in average mortality rates in non-expansion counties. Substituting (4) into the definition of $ATT(2014)$ and replacing potential outcomes with observed outcomes using (1) gives the 2×2 DiD estimand, an expression for the target parameter in terms of four estimable *population* averages:

$$(5) \quad ATT(2014) = \overbrace{\mathbb{E}_\omega[Y_{i,2014}(1) | D_i=1]}^{=\mathbb{E}_\omega[Y_{i,2014}(1) | D_i=1]} - \overbrace{(\mathbb{E}_\omega[Y_{i,2013} | D_i = 1] + (\mathbb{E}_\omega[Y_{i,2014} | D_i = 0] - \mathbb{E}_\omega[Y_{i,2013} | D_i = 0]))}^{=\mathbb{E}_\omega[Y_{i,2014}(0) | D_i=1]}$$

$$= \underbrace{(\mathbb{E}_\omega[Y_{i,2014} | D_i = 1] - \mathbb{E}_\omega[Y_{i,2013} | D_i = 1])}_{\text{(weighted) average change for } D_i=1} - \underbrace{(\mathbb{E}_\omega[Y_{i,2014} | D_i = 0] - \mathbb{E}_\omega[Y_{i,2013} | D_i = 0])}_{\text{(weighted) average change for } D_i=0}$$

Equation (5) highlights what makes DiD so attractive. It is intuitive, it has very mild data requirements (just four means), it answers ex post questions like what did the treatment do?, and its identifying assumption can be stated precisely.

The parallel trends assumption, sometimes just referred to as “parallel trends,” makes DiD distinct from causal designs that are based on statistical independence between treatment and potential outcomes. In designs like randomized trials or instrumental variables, the conditions—mean equalities across groups, for instance—that identify counterfactuals are often a statistical *consequence* of the randomness induced externally (Heckman 2000). In contrast, parallel trends is just a restriction on untreated potential outcome trends. It does not necessarily come from exogenous variation “outside the model.” In fact, because treatment adoption is often chosen by economic actors or policymakers “inside the model,” parallel trends need not hold. For this reason, DiD analyses (correctly) devote significant attention to evaluating parallel trends.

We discuss how to generate empirical evidence about the plausibility of parallel trends below, but new theoretical findings about treatment choice behaviors or selection mechanisms also inform the plausibility of parallel trends.⁵ These results explicitly connect DiD to behaviors, they provide grounding for stories about why a given DiD analysis is internally valid (or not), and they discipline empirical tests of parallel trends. A full review of this literature is outside the scope of this paper, but some helpful broad themes emerge.⁶ For instance, a trade-off exists between the time-series properties of untreated potential outcomes and the information that the agents who choose treatment have as well as how they use it. At one extreme, consider someone who knows $Y_{i,t}(0)$ before and after treatment and can opt into or out of treatment based on it. Parallel trends can only hold in this case if, other than common shifts for every unit, $Y_{i,t}(0)$ is constant (Ghanem, Sant’Anna, and Wüthrich 2026). In our Medicaid example, neither of these conditions—that state legislatures in 2013 knew their 2014 untreated mortality rates or that untreated mortality rates would all have shifted in parallel—is plausible. The opposite extreme is when treatment timing is random, in which case parallel trends holds without any time-series restrictions. Yet, in this case, DiD is not necessary, and more efficient estimators exist (Roth and Sant’Anna 2023a). Against the political and economic backdrop of the early 2010s, it is clear that Medicaid expansion decisions were not random.

Therefore, in realistic scenarios, parallel trends only holds if *some* restrictions on the way untreated outcomes enter the treatment selection mechanism hold as well. As one example, imagine that treatment selection depends on the permanent component of $Y_{i,t}(0)$ (fixed effects) but not on shorter-term fluctuations (“shocks”). For instance, if state legislatures knew and considered their long-run mortality levels only when making their expansion decisions they would be following this kind of selection mechanism. Expansion and non-expansion states might then have large differences in the permanent part of untreated outcomes, which cancel in equation (3), and so parallel trends would hold if shocks to $Y_{i,t}(0)$ had a time invariant mean conditional on the fixed effects.⁷ State legislatures, however, may have also known whether their 2013 mortality rates were especially high or low when considering expanding Medicaid. If the expansion choice is related to these 2013 mortality shocks as well, parallel trends would hold only if stronger time-series restrictions on $Y_{i,t}(0)$ hold. Ghanem, Sant’Anna and Wüthrich (2026) provide a fuller discussion of the selection/time-series trade-off and theory-driven templates to assess parallel trends, while Marx, Tamer, and Tang (2024) discuss economic models that are and are not compatible with parallel trends.

Another implication of the fact that DiD does not rely on statistical independence between $Y_{i,t}(0)$ and treatment status is that there is no guarantee that parallel trends holds across different transformations of $Y_{i,t}(0)$. As stated, it is simply an assumption about averages for a particular $Y_{i,t}(0)$. Roth and Sant’Anna (2023b) show that parallel trends is insensitive to functional form if and only if it holds between groups and across the distribution of $Y_{i,t}(0)$. This can only be true if Medicaid adoption is random, the mortality distribution is constant between 2013 and 2014,

⁵In fact, some of the earliest economic research on DiD methods examined exactly these questions (Ashenfelter and Card 1985; Heckman and Robb 1985).

⁶For details on selection models and outcome models that are consistent with parallel trends, see Ghanem, Sant’Anna, and Wüthrich (2026) and Marx, Tamer, and Tang (2024). Chabé-Ferret (2015) applies these arguments to earnings dynamics.

⁷Some researchers may find easier to understand these as “parallel changes” rather than “parallel trends.” However, the use of “parallel trends” is now firmly established in the literature, and other influential work has used “changes-in-changes” to refer to an alternative estimator to classical DiD estimation (Athey and Imbens 2006). To avoid confusion, we use “parallel trends” throughout this paper.

or a combination of the two cases. As these conditions are arguably ex ante implausible, our DiD analysis may depend on our choices to measure $Y_{i,t}$ in rates (deaths per 100,000) as opposed to logs, for example. One approach to this measurement choice is to propose and evaluate a theory that delivers it, though we recognize that this is not always possible. To assess whether cases where parallel trends holding for one functional form come at the cost of ruling other transformations, we recommend that researchers use the Roth and Sant'Anna's (2023b) falsification tests for the null that parallel trends are insensitive to functional form. In our application, we do not reject the null that parallel trends are insensitive to functional form, with p -values above 0.80.

The interplay between treatment selection and the properties of the outcome variable characterize the structural basis for a DiD analysis (see DiNardo and Lee 2011) and engaging with them is essential to any DiD application. While every study will have its own institutions, choices, and outcomes to consider, a rigorous DiD analysis must provide a transparent discussion about the reliability of the underlying identification assumptions. If parallel trends require implausible behavioral restrictions, one may be better off using an alternative research design.

3.3 Estimation and Inference: Four Means or One Regression?

Mapping the DiD estimand in equation (5) to the canonical 2×2 DiD estimator follows immediately from replacing population expectations with their sample analogs:

$$(6) \quad \widehat{ATT}(2014) = (\bar{Y}_{\omega,D=1,t=2014} - \bar{Y}_{\omega,D=1,t=2013}) - (\bar{Y}_{\omega,D=0,t=2014} - \bar{Y}_{\omega,D=0,t=2013}),$$

where $\bar{Y}_{\omega,D=g,t=t'} = \frac{\sum_{i=1}^n \mathbf{1}\{D_i = g\} \omega_i Y_{i,t'}}{\sum_{i=1}^n \omega_i \mathbf{1}\{D_i = g\}}$ is the ω -weighted sample mean of Y for treatment group g in period t' . Equation (6) is the classic difference of two differences written in terms of sample means. It is a direct recipe for actually estimating $ATT(t)$ and can be read directly from the following table of average mortality rates in 2013 and 2014 by expansion group.

The two across-time changes in equation (6) are in the third row of Table 2. Without weighting, average county-level mortality rates in expansion states rose by 9.3 deaths per 100,000 and 9.1 deaths in non-expansion states, so after rounding the DiD estimate of $ATT(2014)$ is 0.1 deaths per 100,000. This result implies that the average treatment effect of Medicaid expansion on mortality in 2014 among *counties* that are part of an expansion state was an increase of 0.1 deaths per 100,000. In contrast, the DiD result using population weights suggests that Medicaid expansion caused a reduction of 2.6 deaths per 100,000 for the average *adult* in expansion states.⁸

The same result can be obtained as the (weighted) least squares estimate of $\beta^{2 \times 2}$ in the following linear regression specification (which only uses data from $t = 2013$ and $t = 2014$):

$$(7) \quad Y_{i,t} = \beta_0 + \beta_1 \mathbf{1}\{D_i = 1\} + \beta_2 \mathbf{1}\{t = 2014\} + \beta^{2 \times 2} (\mathbf{1}\{D_i = 1\} \times \mathbf{1}\{t = 2014\}) + \varepsilon_{i,t},$$

⁸Columns 3 and 6 in Table 2 show cross-group gaps in average mortality in each year. These can also be used to construct the DiD estimate by rearranging equation (6): $(\bar{Y}_{D=1,t=2014} - \bar{Y}_{D=0,t=2014}) - (\bar{Y}_{D=1,t=2013} - \bar{Y}_{D=0,t=2013})$.

TABLE 2
SIMPLE 2×2 DiD

	Unweighted averages			Weighted averages		
	Expansion	No expansion	Gap/DiD	Expansion	No expansion	Gap/DiD
2013	419.2	474.0	-54.8	322.7	376.4	-53.7
2014	428.5	483.1	-54.7	326.5	382.7	-56.2
Trend/DiD	9.3	9.1	0.1	3.7	6.3	-2.6

Notes: This table reports average county-level mortality rates (deaths among adults aged 20–64 per 100,000 adults) in 2013 (row 1) and 2014 (row 2) in states that expanded adult Medicaid eligibility in 2014 (columns 1 and 4) and states that have not expanded by 2019 (columns 2 and 5). The first three columns present unweighted averages and the second three columns present population-weighted averages. Columns 1, 2, 4, and 5 in the third row show time trends in mortality between 2013 and 2014 for each group of states. The first two rows of columns 3 and 6 show the cross-sectional gap in mortality between expansion and non-expansion states in 2013 and 2014. The entries in bold red text in row 3 show the simple 2×2 difference-in-differences estimates without weights (column 3) and with them (column 6).

where β s are unknown coefficients and $\varepsilon_{i,t}$ is an idiosyncratic term uncorrelated with D_i . To see why, let's focus on the unweighted case ($\omega_i = 1$) and write each of the four means in $\widehat{ATT}(2)$ in terms of the estimated coefficients from (7):

- Sample average of $Y_{i,t}$ in post-period for treatment group is $\bar{Y}_{D=1,t=2014} = \hat{\beta}_0 + \hat{\beta}_1 + \hat{\beta}_2 + \hat{\beta}^{2 \times 2}$.
- Sample average of $Y_{i,t}$ in pre-period for treatment group is $\bar{Y}_{D=1,t=2013} = \hat{\beta}_0 + \hat{\beta}_1$.
- Sample average of $Y_{i,t}$ in post-period for comparison group is $\bar{Y}_{D=0,t=2014} = \hat{\beta}_0 + \hat{\beta}_2$.
- Sample average of $Y_{i,t}$ in pre-period for comparison group is $\bar{Y}_{D=0,t=2013} = \hat{\beta}_0$.

Substituting these expressions into the definition of $\widehat{ATT}(2)$ yields:

$$\widehat{ATT}(2014) = [(\hat{\beta}_0 + \hat{\beta}_1 + \hat{\beta}_2 + \hat{\beta}^{2 \times 2}) - (\hat{\beta}_0 + \hat{\beta}_1)] - [(\hat{\beta}_0 + \hat{\beta}_2) - \hat{\beta}_0] = \hat{\beta}^{2 \times 2}.$$

Table 3 demonstrates this equivalence for both unweighted (column 1) and weighted (column 4) regressions. In fact, with balanced panel data, the estimate of $\beta^{2 \times 2}$ is numerically the same if the regression instead contains fixed effects for each unit (columns 2 and 5) or if one regresses outcome changes on a constant and the treatment group dummy D_i (columns 3 and 6).⁹ Provided that one adopts the same notion of uncertainty, standard errors also coincide; such as when one clusters the standard errors at the county level.

The equivalence between calculating a 2×2 DiD by hand or with a regression has appealing features. Regressions are simple to run, and they do the averaging and differencing behind the scenes. They also allow the use of statistical inference tools from ordinary least squares

⁹When population weights vary over time, the equivalence between a by-hand DiD estimate and one that comes from a regression with unit fixed effects no longer holds.

TABLE 3
REGRESSION 2×2 DiD

	Unweighted			Weighted		
	Crude mortality rate		Δ	Crude mortality rate		Δ
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	474.0 (4.3)		9.1 (2.6)	376.4 (7.6)		6.3 (1.1)
Medicaid expansion	-54.8 (6.3)			-53.7 (11.5)		
Post	9.1 (2.6)			6.3 (1.1)		
Medicaid expansion \times post	0.1 (3.7)	0.1 (3.7)	0.1 (3.7)	-2.6 (1.5)	-2.6 (1.5)	-2.6 (1.5)
County fixed effects	No	Yes	No	No	Yes	No
Year fixed effects	No	Yes	No	No	Yes	No

Notes: This table reports the regression 2×2 DiD estimate comparing counties that expand Medicaid in 2014 to counties that do not expand Medicaid by 2019, using only data for the years 2013 and 2014. Columns 1–3 report unweighted regression results, while columns 4–6 weight by county population aged 20–64 in 2013. Columns 1 and 4 report results from regressing the crude mortality rate for adults ages 20–64 on indicators for expansion states (Treat) and post-expansion year (Post), with the DiD estimate being the coefficient on the interaction term. Columns 2 and 5 report the corresponding results for the interaction term using county and year fixed effects. Finally, columns 3 and 6 report the results of the long difference in county mortality rates on a treatment indicator. Robust standard errors (in parentheses) are clustered at the county level.

(OLS), which are themselves the subject of a large econometrics literature that is particularly important when it comes to standard error estimation (Wooldridge 2003; Bertrand, Duflo, and Mullainathan 2004; Donald and Lang 2007; Cameron, Gelbach, and Miller 2008; Conley and Taber 2011; Abadie et al. 2020, 2023).

Many inference procedures exist for DiD-type analyses, arising from a combination of choices about the target parameter, details of the data structure and sampling process, and maintained assumptions about the structure of outcomes. In practice, one needs to determine and discuss the forms of uncertainty the standard errors are designed to capture—that is, what is (conceptually) being resampled and what may or may not vary across those resamples. As discussed in Abadie et al. (2020), these details come from the nature of the parameter of interest—whether the focus is on sample-specific average treatment effects or population-level average treatment effects—and the stochastic elements of the model that make the estimator random. Heuristically, this involves a thought experiment (or stochastic process) hypothesized to generate the random components of the model (or the data-generating process). Different inferential frameworks highlight different sources of uncertainty by resampling distinct model components and treating other components as fixed (nonrandom).

Inferential frameworks on two extremes help cement these concepts. Design-based frameworks treat potential outcomes and covariates as nonrandom, focus on finite-population parameters (e.g., sample average treatment effects), and consider the allocation of treatment

as the only source of the randomness in the model (Imbens and Rubin 2015).¹⁰ The only thing that is random and thus varies across the hypothetical resamples from this point of view is the treatment allocation. On the other hand, a traditional sampling-based approach to inference presumes that we independently sample units from a super-population. In this case, it is customary to focus on population parameters (like $ATT(2)$), treat all variables in the model as random variables, and compute standard errors by clustering at the level in which the (hypothesized) sampling was conducted. In this framework, every variable in the analysis—outcomes, covariates, and treatment—is randomly redrawn across the hypothetical resamples. A drawback of the sampling approach is that sometimes it is unnatural to think of the data as a random sample from a well-defined population.

A third popular approach to inference—the model-based approach—is more structural and involves taking a stand on the structure of the error component of the model (e.g., imposing a putative model for how shocks affect outcomes and their relationship with treatment and other variables in the model). The uncertainty reflected in this model-based setting entails a thought experiment in which different values of these shocks and the other random variables in the model are drawn from their joint distribution (Abadie et al. 2023). This model-based approach is common in econometrics, and it almost always takes the linear regression specification (or model, in this case) as the starting point of the analysis. Although this is often convenient, it is important to note that imposing model restrictions on the error component of the model necessarily imposes restrictions on treatment effect heterogeneity and on the relationship between potential outcomes; see section 5 and appendix A of Roth et al. (2023) for a discussion. Another challenge with the model-based approach is that it is hard to use this framework when adopting estimation strategies other than linear regressions—for example, inverse probability weighting or doubly robust procedures—which we will discuss in Section 4.4.

Ultimately, as the discussion above highlights, each approach has pros and cons, and discussions about the best way to compute standard errors are complex and intrinsically context-specific. A detailed treatment of the topic is outside the scope of this paper. It requires information about the sampling process, the research design and target parameters, what is treated as fixed and random, and the structure of the error components of their models (e.g., the presence of spatial or serial correlation), among other factors. We refer interested readers to Abadie et al. (2020, 2023) and section 5 of Roth et al. (2023) for discussions on these topics, though we also emphasize that further methodological research in this area is warranted. For the remainder of this article, we adopt a sampling perspective for uncertainty and cluster our standard errors at the county level. In our context, this is compatible with treating all variables as random, including treatment groups and potential outcomes. It also allows us to avoid (a) making time-series dependence restrictions on potential (and realized) outcomes—as we are in a short-panel framework with a large number of units and a fixed number of time periods—and (b) taking an explicit stand on the structure of error components of the model, which is particularly appealing as the starting point of our analysis is potential outcomes rather than regression models. It is also worth mentioning that as our treatment in the empirical example is assigned at the state level, clustering at the county level would also be compatible with treating state-specific shocks as fixed (or conditioned on) and assessing if they lead to violations of

¹⁰Traditionally, design-based inference procedures are justified when treatment assignment is fully random, which is a much stronger requirement than parallel trends. See Rambachan and Roth (2025) for a discussion on design-based inference for quasi-experimental designs, including a discussion of the Medicaid expansion.

parallel trends (Roth et al. 2023, Section 5.1). Clustering at the state level would be justified if we were using a design-based perspective (Rambachan and Roth 2025), though that would require us to treat potential outcomes as fixed (which we do not do in this paper). Our choice of inference procedure is not without controversies, and other inferential approaches may also be rationalized.

We conclude this section by stressing that the appeal of using regressions like (7) to estimate ATT in DiD designs comes from their numerical equivalence to the “by-hand” DiD estimator (6), which was explicitly derived from the $ATT(t)$ and the parallel trends assumption. This ensures that the regression specification respects the underlying identifying assumptions and estimates the desired target parameter. Unfortunately, the tight connection between (TWFE) regressions and DiD designs breaks under more complex setups that are ubiquitous in practice. We now turn to some of these issues and how approaching them from the point of view of 2×2 building blocks can guide good econometric practices.

4. Incorporating Covariates into 2×2 DiD

So far, we have focused on 2×2 DiD designs that do not leverage any information about covariates, but researchers frequently incorporate them into DiD analyses in one of three ways: checking for balance in variables thought to influence $Y_{i,t}(0)$, controlling for those variables in the main estimates, and estimating treatment effect heterogeneity. For example, in the absence of Medicaid expansion, mortality rates likely would have evolved differently in poorer and richer counties; they certainly did before 2014 (Currie and Schwandt 2016). Therefore, parallel trends may fail if poverty rates differ between expansion and non-expansion counties. If they do, then one may want to “control for” poverty rates when estimating ATT parameters. Finally, because Medicaid expansion reached more people in higher-poverty counties, its average effects on overall mortality may be larger there. This heterogeneity may be of interest in its own right or may be used to assess the plausibility of the overall DiD design.

This section discusses how to use auxiliary information on covariates to evaluate parallel trends, identify ATT parameters under potentially weaker assumptions, and study heterogeneity. These approaches stem from viewing a DiD with covariates in terms of 2×2 building blocks that themselves condition on those variables, which creates a clear link to the assumptions, estimators, and interpretation of unconditional 2×2 designs. We also discuss how regression estimators that control for covariates impose extra assumptions and can fail to identify ATT parameters.

4.1 Covariate Balance: Are Unconditional Parallel Trends Plausible?

Assumption PT is fundamentally untestable, as it contains an unobserved counterfactual component. Therefore, “tests” of these assumptions are necessarily indirect and rely on other *observed* variables thought to be related to untreated potential outcome trends. For example, during the 2010s, demographics and economic conditions were strongly correlated with mortality levels and trends. If these relationships would have held in the absence of Medicaid expansion, and if expansion and non-expansion counties differ in those demographic and economic characteristics, then the parallel trends assumption (3) may fail to hold. Checking balance in observable determinants of changes in $Y_{i,t}(0)$ is thus a common and sensible way to evaluate parallel trends.

TABLE 4
COVARIATE BALANCE STATISTICS

Variable	Unweighted			Weighted		
	Non-adopt	Adopt	Norm. diff.	Non-adopt	Adopt	Norm. diff.
<i>2013 covariate levels</i>						
% female	49.43	49.33	-0.03	50.48	50.07	-0.24
% White	81.64	90.48	0.59	77.91	79.54	0.11
% Hispanic	9.64	8.23	-0.10	17.01	18.86	0.11
Unemployment rate	7.61	8.01	0.16	7.00	8.01	0.50
Poverty rate	19.28	16.53	-0.42	17.24	15.29	-0.37
Median income	43.04	47.97	0.43	49.31	57.86	0.68
<i>2014–2013 covariate differences</i>						
% female	-0.02	-0.02	0.00	0.02	0.01	-0.09
% White	-0.21	-0.21	0.01	-0.32	-0.33	-0.04
% Hispanic	0.20	0.21	0.04	0.25	0.33	0.29
Unemployment rate	-1.16	-1.30	-0.21	-1.08	-1.36	-0.55
Poverty rate	-0.55	-0.28	0.14	-0.41	-0.35	0.05
Median income	0.98	1.11	0.06	1.10	1.74	0.32

Notes: This table reports the covariate balance between adopting and non-adopting states. In the top panel, we report the averages and standardized differences of each variable, measured in 2013, by adoption status. All variables are measured in percentage values, except for median household income, which is measured in thousands of US dollars. In the bottom panel we report the average and standardized differences of the county-level long differences between 2014 and 2013 of each variable. We report both weighted and unweighted measures of the averages to correspond to the different estimation methods of including covariates in a 2×2 setting.

Most DiD analyses check for balance across groups in baseline covariate levels ($\mathbb{E}_\omega[X_{i,t=1}|D_i = 1] - \mathbb{E}_\omega[X_{i,t=1}|D_i = 0]$) or covariate trends before and after treatment ($\mathbb{E}_\omega[\Delta X_{i,t=2}|D_i = 1] - \mathbb{E}_\omega[\Delta X_{i,t=2}|D_i = 0]$), where $\Delta X_{i,t=2} = X_{i,t=2} - X_{i,t=1}$). We consider the following variables, $X_{i,t}$: the percentages of a county's population that are female, White, or Hispanic; the unemployment rate; the poverty rate; and county-level median income (in thousands of dollars).¹¹ The top panel of Table 4 reports averages of these variables by group in 2013 with and without population weights. We also report a measure of imbalance that is comparable across variables: the normalized difference in means between the treatment and comparison groups (Imbens and Rubin 2015, Chapter 14),

$$\text{Norm. Diff}_\omega = \frac{\bar{X}_{\omega,T} - \bar{X}_{\omega,C}}{\sqrt{(S_{\omega,T}^2 + S_{\omega,C}^2)/2}},$$

where $\bar{X}_{\omega,T}$ and $\bar{X}_{\omega,C}$ are the sample weighted or unweighted averages for the treatment and comparison groups, respectively, and $S_{\omega,T}^2$ and $S_{\omega,C}^2$ are the sample weighted or unweighted variances

¹¹We focus on these variables for convenience. Borgschulte and Vogler (2020) use a least absolute shrinkage and selection operator (LASSO) procedure that selects more and different covariates to include in their analysis. We replicate their findings when we follow their methodology but diverge here for the sake of brevity.

of the covariates for the treatment and comparison group. As a general rule of thumb, values of the normalized difference in excess of 0.25 in absolute value indicate a potentially problematic imbalance between the two groups (Imbens and Rubin 2015, p. 277).¹²

We find meaningful imbalance in several baseline measures. Expansion counties in 2013 were Whiter and had higher unemployment rates despite lower poverty and higher median income than non-expansion counties. Because DiD uses *changes* in outcomes, researchers sometimes argue that the effect of pretreatment variables is differenced out. This logic does not hold, though, if baseline covariates are related to untreated potential outcome trends themselves. The imbalance in the top panel of Table 4 will lead to violations of parallel trends to the extent that counties with different racial composition or income distributions would have had different mortality changes even without Medicaid expansion.

Nevertheless, checks of balance in covariate changes can be informative about parallel trends as well. The bottom panel of Table 4 reports average changes by group between 2013 and 2014 as well as normalized differences. Many of the imbalances evident in baseline levels change or even flip signs when measured in changes. Unemployment, for example, was higher in expansion states in 2013 but fell faster. To the extent that these changes are important determinants of $\Delta Y_{i,t}(0)$, then these results could suggest that Assumption PT is violated.

Why do we say “could”? A major challenge in interpreting cross-group gaps in $\Delta X_{i,t}$ involves deciding which variables are truly covariates and which are mechanisms/outcomes. If an element of $X_{i,t}$ cannot be affected by the treatment, it is a (strictly exogenous) covariate, and differential changes in exogenous covariates may indicate a PT violation. Since the treatment cannot have caused $X_{i,t}$ to change (by assumption), something else that differs across groups and over time must have. Since little research suggests an effect of Medicaid expansion on unemployment, this may be a good assumption. On the other hand, if Medicaid expansion can change the demographic and economic composition of its counties, then differential changes in these variables may actually be a consequence of the expansion itself.¹³ If so, then differential posttreatment changes in them would not necessarily indicate a parallel trends violation; they could partially reflect a causal effect. As with the plausibility of Assumption PT itself, whether something is a covariate or a mechanism is not a data question per se. It requires context-specific knowledge about how the treatment works.

4.2 *DiD with Covariates: Identification under Conditional Parallel Trends*

Having detected covariate imbalance that casts doubt on Assumption PT, how should we proceed to estimate $ATT(2)$? Because the imbalance documented in Table 4 suggested that unconditional parallel trends may not hold, our goal is to develop a DiD identification strategy based on an assumption that accounts for this imbalance. Working from a conditional parallel

¹²As discussed in Austin (2009, section 3.2), the normalized difference was initially proposed in the psychological literature and is sometimes referred to as Cohen's effect size index. In this context, normalized differences of 0.2, 0.5, and 0.8 are sometimes used to represent small, medium, and large imbalances; however, normalized differences as small as 0.1 are sometimes considered worrisome, depending on how important the covariate in question is. Ultimately, there is no universally accepted threshold for what value indicates important imbalances. See Ho et al. (2007) and Austin (2009) for additional discussions.

¹³In fact, comparing mean covariate changes in expansion and non-expansion is the same as using $X_{i,t}$ as the outcome in a 2×2 DiD estimator.

trends assumption shows how to construct $ATT(2)$ from 2×2 comparisons that are each conditioned on specific covariate values, thus addressing the imbalance problem.

Let X_i be a vector of observed determinants of changes in $Y_{i,t}(0)$. Here, we purposefully omit the time subscript on X_i because the covariates in this section can be time invariant, such as fixed variations or baseline values ($X_{i,t=1}$), or time varying in the sense of including values from the second period, $X_{i,t=2}$. The empirical content of a “new” identification assumption that incorporates X_i , henceforth the conditional parallel trends (CPT) assumption, is formalized as follows.

ASSUMPTION CPT (2×2 Conditional Parallel Trends): *The (weighted) average change of $Y_{i,t=2}(0)$ from $Y_{i,t=1}(0)$ is the same between treated and comparison units that share the same covariate values,*

$$(8) \quad \mathbb{E}_\omega[Y_{i,t=2}(0) - Y_{i,t=1}(0) | X_i, D_i = 1] = \mathbb{E}_\omega[Y_{i,t=2}(0) - Y_{i,t=1}(0) | X_i, D_i = 0].$$

Assumption CPT has the same structure as Assumption PT but states that PT holds within each covariate-specific stratum rather than across the whole population. With respect to the baseline covariates in Table 4, this amounts to assuming parallel trends in $Y_{i,t}(0)$ between expansion and non-expansion counties that in 2013 had the same female, White, and Hispanic shares, as well as the same unemployment and poverty rates and median income. This does not restrict *how* $Y_{i,t}(0)$ changes in different covariate strata, nor is not necessary to estimate these trends. Assumption CPT only requires that covariate-specific trends are common between expansion and non-expansion counties.

For both expectations in Assumption CPT to be well-defined for all values of X_i , there must be both untreated and treated units in the population at each covariate value. If, for some covariate values, there are only treated units, for example, then the right-hand side of (8) is undefined. A formal statement of this assumption, called “common support” or “strong overlap,” is as follows.¹⁴

ASSUMPTION SO (Strong Overlap): *The conditional (weighted) probability of belonging to the treatment group, given observed covariates, X_i , which are determinants of untreated potential outcome growth, is uniformly bounded away from zero and one. That is, for some $\epsilon > 0$, $\epsilon < P_\omega[D_i = 1 | X_i] < 1 - \epsilon$.*

Under assumptions CPT and SO the $ATT(2)$ is identified:

$$(9) \quad \begin{aligned} ATT(2) &= \mathbb{E}_\omega[Y_{i,t=2}(1) | D_i = 1] - \mathbb{E}_\omega[Y_{i,t=2}(0) | D_i = 1] \\ &= \mathbb{E}_\omega[Y_{i,t=2} | D_i = 1] - \mathbb{E}_\omega[\mathbb{E}_\omega[Y_{i,t=2}(0) | X_i, D_i = 1] | D_i = 1] \\ &= \mathbb{E}_\omega[Y_{i,t=2} | D_i = 1] - \mathbb{E}_\omega[\mathbb{E}_\omega[Y_{i,t=1}(0) | X_i, D_i = 1] \\ &\quad + \mathbb{E}_\omega[\Delta Y_{i,t=2}(0) | X_i, D_i = 0] | D_i = 1] \\ &= \mathbb{E}_\omega[\Delta Y_{i,t=2} | D_i = 1] - \mathbb{E}_\omega[\mathbb{E}_\omega[\Delta Y_{i,t=2} | X_i, D_i = 0] | D_i = 1]. \end{aligned}$$

¹⁴ $ATT(2)$ is still identified under conditional parallel trends if some values of X_i have only untreated observations. We require $P_\omega[D_i = 1 | X_i]$ to be bounded away from zero and one to avoid irregular inference procedures; see Khan and Tamer (2010) for additional details.

The first line restates the definition of $ATT(2)$, and the second line uses the law of iterated expectations to write the counterfactual mean for the whole treatment group as an average of counterfactual means conditional on X_i . These quantities are exactly the ones that appear in the conditional parallel trends assumption (and the overlap condition). Therefore, as in Section 3, the third line uses Assumption CPT to rewrite the conditional counterfactuals in terms of observable population quantities under no anticipation (Assumption NA). Equation (9) then uses the law of iterated expectations again to group terms and express $ATT(2)$ in terms of observed variables $(Y_{i,t=2}, Y_{i,t=1}, G_i, X_i)$; that is, it establishes that $ATT(2)$ is nonparametrically identified under our assumptions. This expression has a clear intuition: The $ATT(2)$ is equal to the path of outcomes experienced by the treated group (the term on the left) minus the average path of outcomes in the comparison group for each value of the covariates, averaged over the treated group's distribution of covariates (the term on the right).

4.3 DiD Estimation with Covariates: TWFE

Moving from the population identification result in equation (9) to sample analogs is a challenge unless the covariates are discrete and the conditional expectations themselves are easily calculable. This difficulty does not arise in an unconditional DiD. With continuous covariates, or many discrete ones, it may not be feasible to construct $\mathbb{E}_\omega[\Delta Y_{i,t=2} | X_i, D_i = 0]$. Conditional DiD estimation, therefore, uses additional econometric techniques to bridge this gap. We begin, however, by discussing how regression DiD estimators that include covariates relate to the assumptions used for identification in (9).

Because the TWFE specification in (7) recovers the $ATT(2)$ in 2×2 DiD setups without covariates, it is natural to extend this logic to regressions with covariates. Indeed, this is by far the most popular approach adopted by practitioners, arguably because it is both easy and familiar. A typical regression specification is

$$(10) \quad Y_{i,t} = \theta_t + \eta_i + \beta_{treat} D_{i,t} + X'_{i,t} \beta_{covs} + e_{i,t}$$

where the unit and time fixed effects, treatment status, and covariates have already been defined, $e_{i,t}$ is an error term, and β_{treat} is interpreted as the parameter of interest. A related specification explicitly controls for baseline covariates by replacing $X_{i,t}$ with interactions of the pretreatment covariates $X_{i,t=1}$ and a posttreatment dummy,

$$(11) \quad Y_{i,t} = \theta_t + \eta_i + \beta_{treat,2} D_{i,t} + (\mathbf{1}\{t = 2\} X_{i,t=1}) \beta_{covs,2} + e_{i,t}$$

In Table 5, we report the OLS and weighted least squares estimates of the unconditional 2×2 DiD estimate, β_{treat} from (10), $\beta_{treat,2}$ from (11), and their cluster-robust standard errors, using the covariates from Table 4.

Although only one covariate-adjusted estimate in Table 5 is (marginally) statistically significant, the point estimates differ noticeably. In the unweighted case, adjusting for the 2013 levels of the covariates decreases the estimated effect of Medicaid expansion on short-run mortality rates from a point estimate of roughly 0.12 to -2.35 . However, if we include their time-varying values instead, we estimate an effect of -0.49 , a large difference. We find a similar result when using weighted regressions; while the coefficient remains constant (-2.56) when using 2013 values of the covariates, it attenuates to -1.37 if we use (10).

TABLE 5
REGRESSION 2×2 DiD WITH COVARIATES

	Unweighted			Weighted		
	No covs	$X_{i,t=2013}$	$X_{i,t}$	No covs	$X_{i,t=2013}$	$X_{i,t}$
	(1)	(2)	(3)	(4)	(5)	(6)
Medicaid expansion	0.12 (3.75)	-2.35 (4.29)	-0.49 (3.83)	-2.56 (1.49)	-2.56 (1.78)	-1.37 (1.62)

Notes: This table reports the regression 2×2 DiD estimate comparing counties that expand Medicaid in 2014 to counties that do not expand Medicaid, adjusting for the inclusion of covariates (percent female, percent White, percent Hispanic, the unemployment rate, the poverty rate, and median household income). Columns 1–3 report unweighted regression results, while columns 4–6 weight by county population aged 20–64 in 2013. Columns 1 and 4 report results for expansion states without covariates, columns 2 and 5 adjust for the baseline levels of the covariates in 2013, and columns 3 and 6 control for the time-varying covariate values in 2014 and 2013. Robust standard errors (in parentheses) are clustered at the county level.

The jump from the conditional DiD identification result in (9) to the TWFE *estimators* in (10) and (11) skips a crucial question about β_{treat} or $\beta_{treat,2}$: do they equal the target parameter $ATT(2)$ under the conditional parallel trends assumption? It turns out that the close relationship among regression DiD, $ATT(2)$, and parallel trends in a design without covariates does not hold with covariates. The issues come from exactly what kinds of covariates are effectively being “controlled for” in these specifications and how the regression estimator combines outcome trends for covariate subgroups.

Note that in our two-period setup, (10) and (11) are respectively equivalent to (with some abuse of notation),

$$\Delta Y_{i,t=2} = \alpha + \beta_{treat} D_i + \Delta X'_{i,t=2} \beta_{covs} + \Delta e_{i,t=2},$$

$$\Delta Y_{i,t=2} = \alpha + \beta_{treat,2} D_i + X'_{i,t=1} \beta_{covs,2} + \Delta e_{i,t=2}.$$

The first thing that is clear from these representations is that because time-invariant variables drop out of equation (10), a TWFE specification can account for differential trends related to baseline covariate levels only if they enter as interactions with the posttreatment dummy, as in equation (11). The exact regression specification, therefore, determines the implied conditional parallel trends assumption. Controlling for annual poverty rates really means controlling for poverty changes, and areas that are poor are not the same as areas that are becoming poor.

Another limitation evident in (10) relates to “bad controls.” Whenever $X_{i,t=2}$ is affected by the treatment, then conditioning on it (in any way) can bias estimates of the $ATT(2)$. If Medicaid expansion lowered poverty rates, for example, then including 2014 poverty rates or the 2013–14 change in poverty rates as a covariate is problematic. This echoes our discussion about testing balance in $\Delta X_{i,t}$ in the sense that time-varying covariates must be unaffected by treatment in order to interpret imbalance in their trends as a source of bias and to be able to control for them to address that bias. See Caetano et al. (2024) for a discussion.

Suppose we have decided on which variables to include in a conditional parallel trends assumption and whether to measure them in levels or changes. If Assumptions CPT and SO hold with respect to this set of covariates, does β_{treat} recover the $ATT(2)$? In the DiD context, Caetano and Callaway (2024) tackle exactly this question. They show that β_{treat} equals a weighted average of conditional average treatment effects, defined as $ATT_{X_k}(2) \equiv \mathbb{E}_\omega[Y_{i,t=2}(1) - Y_{i,t=2}(0) | D_i = 1, X_i = x_k]$, with weights that may not be convex, plus three bias terms reflecting misspecification either in the set of control variables or the fact that they are included linearly. These conclusions relate to recent findings about the properties of regression estimators in the presence of heterogeneity in other contexts, including instrumental variables (Mogstad and Torgovitsky 2024), cross-sectional designs (Angrist 1998; Aronow and Samii 2016; Słoczyński 2022; Goldsmith-Pinkham, Hull, and Kolesár 2024), and panel data (Goodman-Bacon 2021; de Chaisemartin and D'Haultfœuille 2020; Sun and Abraham 2021; Poirier and Słoczyński 2025).

The weighting results suggest that even when the covariates are correctly selected, measured, and added with the correct functional form, β_{treat} could be negative even when $ATT_{X_i}(2)$ is positive for all values of covariates. Short of this extreme sign-reversal case, the weighting scheme of β_{treat} generally does not yield the $ATT(2)$ target parameter and instead puts too much weight on $ATT_{X_i}(2)$ for X_i s that are relatively uncommon among the treated group relative to the untreated group and puts too little weight on $ATT_{X_i}(2)$ for X_i s that are relatively common among the treated group relative to the untreated group (Słoczyński 2022; Caetano and Callaway 2024).

Taken together, these results imply that β_{treat} identifies $ATT(2)$ under the *additional* assumption that treatment effects across covariate strata are constant. To see why, write the conditional ATT in period two given $\Delta X_{i,t=2}$ as

$$ATT_{\Delta X_{i,t=2}}(2) = \mathbb{E}_\omega[Y_{i,t=2}(1) - Y_{i,t=2}(0) | D_i = 1, \Delta X_{i,t=2}],$$

and note that, under Assumptions CPT and SO, it is identified by

$$ATT_{\Delta X_{i,t=2}}(2) = \mathbb{E}_\omega[\Delta Y_{i,t=2} | D_i = 1, \Delta X_{i,t=2}] - \mathbb{E}_\omega[\Delta Y_{i,t=2} | D_i = 0, \Delta X_{i,t=2}].$$

If we take (10) to be a correctly specified regression, then

$$ATT_{\Delta X_{i,t=2}}(2) = (\beta_{treat} + \Delta X'_{i,t=2} \beta_{covs}) - (\Delta X'_{i,t=2} \beta_{covs}) = \beta_{treat}.$$

In other words, (10) implicitly rules out that treatment effects can vary across covariate strata, which makes the weighting issues identified by Caetano and Callaway (2024) irrelevant to the interpretation of β_{treat} . Research on the Medicaid expansion using data on mortality rates by income, however, shows clear evidence of heterogeneous effects (Miller, Johnson, and Wherry 2021; Wyse and Meyer 2025).

One way to avoid these limitations would be to make (11) (or (10)) more flexible by including interactions of the covariates with treatment group, time, and treatment group by time. An alternative possibility is to adopt a “forward-engineering” perspective (Mogstad and Torgovitsky 2024) and derive an estimator for $ATT(2)$ that directly leverages Assumptions NA, CPT and SO. In some situations, the forward-engineering approach will also use regressions. Still, the target parameter and the identifying assumptions will guide the specification we should use (and not the other way around).

4.4 DiD Estimators with Covariates That Target the $ATT(2)$

Fortunately, TWFE is not the only way to bring covariates into DiD estimation. We now discuss alternative ways to use covariates in a DiD analysis that start from the identification result in (9):

$$ATT(2) = \mathbb{E}_\omega[\Delta Y_{i,t=2} | D_i = 1] - \mathbb{E}_\omega[\mathbb{E}_\omega[\Delta Y_{i,t=2} | X_i, D_i = 0] | D_i = 1].$$

This equation provides an intuitive recipe for estimating $ATT(2)$. The first term in $ATT(2)$ is the same as in an unconditional 2×2 design and can be replaced by its sample analog as in equation (6): $(\bar{Y}_{\omega,D=1,t=2} - \bar{Y}_{\omega,D=1,t=1})$.

One way to obtain the second term is to first estimate the inner conditional expectation, $\mathbb{E}_\omega[\Delta Y_{i,t=2} | X_i, D_i = 0]$. This object is just an (unknown) function that relates average outcome trends for untreated units to their covariates. The most common way to proceed, especially in cases where X_i contains many variables or continuous ones, is to specify a working model, $\mu_{\omega,\Delta,D=0}(X_i)$, with parameters that are simple to estimate. A natural and empirically friendly choice is a linear model, $\mu_{\omega,\Delta,D=0}(X_i) = X_i' \beta_{D=0}$, whose parameters come from a (weighted) regression of $\Delta Y_{i,t=2}$ on X_i in the sample of untreated units. The results of this regression describe untreated outcome trends as a function of X_i . The fitted model then generates predicted values, $\hat{\mu}_{\omega,\Delta,D=0}(X_i) = X_i' \hat{\beta}_{D=0}$, for all units in the sample, including treated units. With these fitted values we can estimate $\mathbb{E}_\omega[\mathbb{E}_\omega[\Delta Y_{i,t=2} | X_i, D_i = 0] | D_i = 1]$ using the plug-in principle; that is, by replacing $\mathbb{E}_\omega[\Delta Y_{i,t=2} | X_i, D_i = 0]$ with its fitted value $\hat{\mu}_{\omega,\Delta,D=0}(X_i)$, and replacing population (weighted) expectations with their sample analogs: $\frac{\sum_{i=1}^n D_i \omega_i \hat{\mu}_{\omega,\Delta,D=0}(X_i)}{\sum_{i=1}^n D_i \omega_i}$.

Putting the pieces together gives the following estimator for $ATT(2)$:

$$(12) \quad \widehat{ATT}_{ra}(2) = \frac{\sum_{i=1}^n D_i \omega_i (\Delta Y_{i,t=2} - \hat{\mu}_{\omega,\Delta,D=0}(X_i))}{\sum_{i=1}^n D_i \omega_i}.$$

This strategy is often referred to as the regression-adjustment (RA) or outcome regression approach to DiD; see, for example, Heckman, Ichimura, and Todd (1997).

We apply this strategy in our application using only the baseline covariates from Table 4, and report the parameters $\hat{\beta}_{D=0}$ of our working model in columns 1 and 3 of Table 6. In practice, mortality changes in non-expansion states are only weakly related to our baseline covariates. Multiplying the weighted coefficients in Table 6 times the weighted treatment group 2013 means in Table 4 gives a predicted change in untreated mortality rates for the average treated county of 7.2 deaths per 100,000, the second term in equation (12), compared to the observed (weighted) change among untreated counties of 6.3 deaths. The observed weighted trend in mortality for expansion counties from Table 2 is 3.7 deaths. Together, these results imply that this approach, based only on assumptions CPT, SO, NA, the choice of baseline covariates, and a linear model for $\mathbb{E}_\omega[\Delta Y_{i,t=2} | X_i, D_i = 0]$, yields an estimated $ATT(2014)$ of -3.5 . This matches the formal RA DiD estimate we report in column 4 of Table 7 (labeled as “Regression”). Column 1, however, gives an unweighted estimate from the same procedure of -1.62 .

To compute standard errors, we need to account for the fact that (12) is a two-step estimation procedure and take into account the uncertainty associated with estimating the working

TABLE 6
OUTCOME REGRESSION AND PROPENSITY SCORE MODELS

	Unweighted		Weighted	
	Regression	Propensity score	Regression	Propensity score
	OLS (1)	logit (2)	OLS (3)	logit (4)
Constant	-20.91 (74.56)	-10.00 (1.34)	-4.62 (41.88)	-8.17 (4.39)
% female	0.04 (0.86)	-0.04 (0.02)	-0.09 (0.65)	-0.19 (0.06)
% White	0.15 (0.25)	0.06 (0.01)	0.20 (0.10)	0.04 (0.01)
% Hispanic	-0.08 (0.19)	-0.02 (0.00)	-0.08 (0.06)	-0.02 (0.01)
Unemployment rate	1.14 (1.99)	0.32 (0.03)	0.88 (0.95)	0.68 (0.08)
Poverty rate	0.21 (1.13)	0.03 (0.02)	-0.13 (0.46)	0.11 (0.05)
Median income	0.09 (0.46)	0.08 (0.01)	-0.05 (0.19)	0.15 (0.02)

Notes: This table reports the outcome regression propensity score models that enter into the estimator from Sant'Anna and Zhao (2020) and Callaway and Sant'Anna (2021). The first two columns report the results for unweighted regressions and the second two report results from weighted regression models. The regression model predicts changes in the outcome variable (mortality rates) on the 2013 covariate values for just the counties that do not expand Medicaid in 2014. The propensity score model uses data for 2013 and estimates a logit model of an expansion indicator variable on the 2013 covariate levels. Robust standard errors (in parentheses) are clustered at the county level for the outcome regression models.

TABLE 7
DiD ESTIMATES WITH COVARIATES

	Unweighted			Weighted		
	Regression	IPW	Doubly robust	Regression	IPW	Doubly robust
Medicaid expansion	-1.62 (4.68)	-0.86 (4.61)	-1.23 (4.94)	-3.46 (2.39)	-3.84 (3.39)	-3.76 (3.24)

Notes: This table reports the 2×2 DiD estimate comparing counties that expand Medicaid in 2014 to counties that do not expand Medicaid, adjusting for the inclusion of 2013 covariate values using the methodologies discussed in Sant'Anna and Zhao (2020) and Callaway and Sant'Anna (2021). The first column reports results using regression adjustment, the second column uses inverse probability weighting based on a propensity score model using the included covariates, and the third column uses the doubly robust combination of the two approaches. Standard errors (in parentheses) are clustered at the county level. IPW is inverse probability weighted.

model $\mu_{\omega, \Delta, D=0}(X_i)$. This is standard, though, and most statistical software automates this process; see, for example, Sant'Anna and Zhao (2020) and Callaway and Sant'Anna (2021).

A linear working model for $\mathbb{E}_\omega[\Delta Y_{i,t=2}|X_i, D_i = 0]$ is a familiar choice, but not the only one. More flexible or even fully nonparametric working models are possible, and the procedure is the same: Estimate the model on untreated units, get its fitted values for the covariate values of the treated units, and then estimate the $ATT(2)$ using (12). An important consideration when choosing the working models is sample size. Large samples permit more flexible estimators that do not sacrifice too much precision. In smaller samples, a parametric linear model may be more appealing. Ultimately, the reliability of DiD RA estimators for $ATT(2)$ depends on how well $\hat{\mu}_{\omega, \Delta, D=0}(X_i)$ approximates $\mathbb{E}_\omega[\Delta Y_{i,t=2}|X_i, D_i = 0]$. If the working model is misspecified—for example by omitting relevant nonlinear terms—the resulting DiD RA estimator will be biased.

Table 6 showed that our covariates did not strongly predict untreated mortality trends and thus do little to change our potential biased unconditional DiD estimates. However, $ATT(2)$ can be estimated conditional on covariates in a different way, without needing to specify the outcome model explicitly. Instead, one can improve the comparability of the comparison group directly by selecting a model for the conditional probability of being treated and applying an IPW DiD procedure (Abadie 2005). The logic of IPW builds on the balance checks we conducted in Table 4: If imbalance in covariates is the source of parallel trends violations, then adjusting the comparison group to be balanced on covariates can address that bias. The adjustment takes the form of re-weighting the observed changes in adult mortality rates for non-expansion counties to ensure that the expansion and non-expansion counties are similar on covariates, thus addressing the compositional source of bias.

To implement the IPW DiD procedure and construct these “balancing weights,” we need to model $p_\omega(X_i) = P_\omega(D_i = 1|X_i)$, the (weighted) conditional probability of belonging to the treatment group, known as the “propensity score.” IPW weights are a function of $p_\omega(X_i)$, and for estimating $ATT(2)$ they take a form that forces the underlying weighted distribution of covariates for comparison units to match the distribution for treatment units (Rosenbaum and Rubin 1983). Intuitively, these weights are formed so that if we find some units that were likely to be observed in the treatment groups (based on their covariate values) but ended up in the comparison group, we give these untreated observations “extra” weight.

Formally, we can show that, under Assumptions CPT and SO, when panel data are available,

$$(13) \quad ATT(2) = \mathbb{E}\left[(w_{\omega, D=1}(D_i) - w_{\omega, D=0}(D_i, X_i))\Delta Y_{i,t=2}\right],$$

where

$$(14) \quad w_{\omega, D=1}(D) = D\omega/\mathbb{E}[\omega D],$$

and

$$w_{\omega, D=0}(D, X) = \frac{\omega(1-D)p_\omega(X)}{1-p_\omega(X)} \bigg/ \mathbb{E}\left[\frac{\omega(1-D)p_\omega(X)}{1-p_\omega(X)}\right];$$

see, for example, Abadie (2005) and Sant’Anna and Zhao (2020). The structure of the weights in equation (14) and the way they enter equation (13) highlight several intuitive features of how the IPW estimator works. First, $w_{\omega, D=1}(D)$ is only nonzero for treated units, and $w_{\omega, D=0}(D, X)$ is only nonzero for untreated units, which means the estimator subtracts a particular mean of

outcome trends for untreated units from a particular mean of outcome trends for treated units. Second, the $w_{\omega,D=1}(D)$ weights do not involve X_i and simply lead to a (ω -weighted) mean for the treatment group. Third, the $w_{\omega,D=0}(D, X)$ weights are functions of X_i that, as (14) shows, give increasingly more weight to untreated units with high propensity scores. This specific IPW weighting function builds a comparison group whose covariate distribution matches the treatment group. Fourth, equation (14) clarifies that two types of weights both enter an IPW analysis. We already discussed how the ω weights shape the parameter of interest, parallel trends assumptions, and estimation. These simply multiply the IPW weights, which act to address imbalance (and their product is rescaled to integrate to one within group). Finally, the appeal of (13) is that if we have a better sense of how units sort into treatment than of the factors that shape outcome trends, we may be more comfortable modeling $p_{\omega}(X_i)$ than $\mathbb{E}_{\omega}[\Delta Y_{i,t=2} | X_i, D_i = 0]$.

To leverage the characterization of the $ATT(2)$ in (13) for estimation and inference purposes, we need a working model for the true propensity score, $p_{\omega}(X_i)$. When X s are all discrete and low dimensional, this is simple and does not involve functional form restrictions: Create covariate-specific strata and then, within each strata, compute the proportion of treated units, and call these estimates $\hat{\pi}(x_k)$, x_k being a strata indicator. When X s have continuous components, or when there are too many strata relative to the available sample size, one can adopt a flexible working model, $\pi_{\omega}(X_i)$, for the propensity score. A common choice for $\pi_{\omega}(X_i)$ is a (weighted) logistic model whose parameters can be estimated using maximum likelihood—or an alternative estimation procedure such as inverse probability tilting (Graham, Pinto, and Egel 2012). We follow this strategy in our Medicaid application and report in columns 2 and 4 of Table 6 the unweighted and weighted maximum likelihood logit coefficients from our propensity score model. Our covariates appear to explain expansion decisions better than untreated outcome trends, suggesting that an estimation approach based on propensity scores may change our $ATT(2014)$ estimate more than the RA approach did. We then use these logit coefficients to get the fitted values for all observations, $\hat{\pi}_{\omega}(X_i)$, that will serve as our estimates of $p_{\omega}(X_i)$.

From these fitted values, we can estimate $ATT(2)$ by

$$(15) \quad \widehat{ATT}_{ipw}(2) = \frac{1}{n} \sum_{i=1}^n (\hat{w}_{\omega,D=1}(D_i) - \hat{w}_{\omega,D=0}(D_i, X_i)) \Delta Y_{i,t=2},$$

where

$$(16) \quad \hat{w}_{\omega,D=1}(D) = D\omega \left/ \frac{1}{n} \sum_{i=1}^n \omega_i D_i \right.,$$

$$\hat{w}_{\omega,D=0}(D, X) = \frac{\omega(1-D)\hat{\pi}_{\omega}(X)}{1 - \hat{\pi}_{\omega}(X)} \left/ \frac{1}{n} \sum_{i=1}^n \omega_i \frac{(1-D_i)\hat{\pi}_{\omega}(X_i)}{1 - \hat{\pi}_{\omega}(X_i)} \right..$$

We report the $ATT(2014)$ estimates and their standard errors using this IPW DiD procedure in Table 7 (labeled as “IPW”) using both unweighted and population-weighted procedures, where we use a logistic regression that is linear in covariates as our propensity score working model. We use the delta method to account for the estimation uncertainty inherited in this two-step estimation procedure when computing standard errors. The IPW DiD estimates are

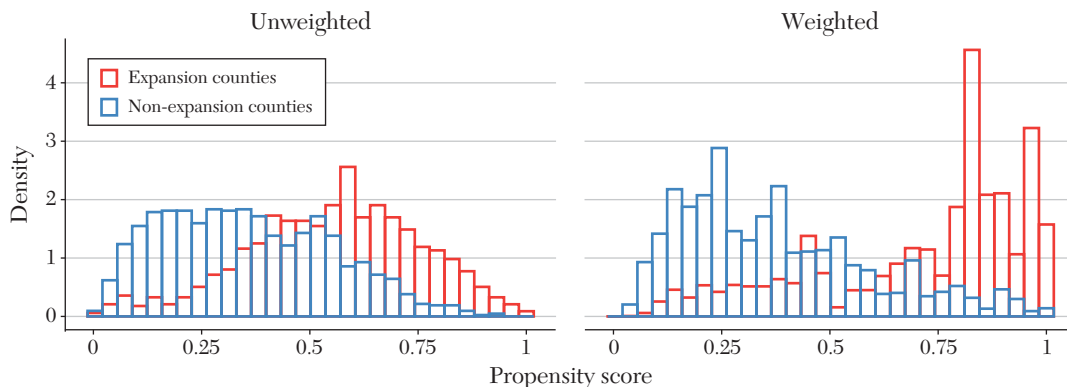


Figure 1. Distribution of Propensity Scores

Notes: This figure shows the distribution of propensity scores using both the weighted and unweighted logit propensity score estimates in our 2×2 Medicaid example.

generally similar to the RA DiD estimates in that we obtain negative estimates that are larger in magnitude when using population weights. However, despite neither being statistically significant, the unweighted IPW estimate is less than half the size of the RA estimate. This is consistent with the broad conclusion from Table 6 that our covariates explain Medicaid expansion better than mortality trends.

IPW estimators for $ATT(2)$ tend to be noisy when fitted propensity score estimates are too close to 1 among untreated units, a condition related to the Assumption SO. The reason for this is simple: When $\hat{\pi}_\omega(X_i)$ is close to one among units with $D_i = 0$, the (estimated) inverse probability weights, $\hat{w}_{\omega, D=0}(D, X)$, become more volatile, as one is essentially dividing by zero. A good practice when using IPW estimators is to check the plausibility of strong overlap using estimated propensity scores. Figure 1 plots the propensity scores that come from the logit models in Table 6 for the two groups of counties. Few untreated units have very high estimated propensity scores, so extreme weighting is not a significant concern.¹⁵ In addition, propensity scores of non-expansion counties seem to lie within the support of the expansion counties' propensity scores, supporting strong overlap. Trimming high- or low-propensity score observations from the sample may be warranted when overlap is weak; for a discussion, see, for example, Crump et al. (2009), Sasaki and Ura (2022), and Ma et al. (2023).

The reliability of these two approaches to estimating conditional DiD designs centers on selecting good models for two different functions: the conditional expectation of untreated outcome changes and the true, unknown propensity score $p_\omega(X_i)$. So which approach should one pick? In practice, there are major advantages to combining both in a way that leads to estimators for the $ATT(2)$ that are more robust against model misspecification (Sant'Anna and Zhao 2020). This is the so-called doubly robust (DR) approach, sometimes referred to as the augmented inverse probability weighting approach; see, for instance, Sant'Anna and Zhao (2020) and Chang (2020).

¹⁵ By default, the did and DRDID R packages from Callaway and Sant'Anna (2021) and Sant'Anna and Zhao (2020) trim untreated units with propensity scores in excess of 0.995.

The key idea of the DR DiD approach is to express the $ATT(2)$ in terms of both $p(X_i)$ and $\mathbb{E}_\omega[\Delta Y_{i,t=2}|X_i, D_i = 0]$ in a way that it provides some “protection” in case the working models for these functions, $\hat{\pi}(X_i)$ and $\hat{\mu}_{\Delta, C=0}(X_i)$, are wrong. The resulting DR estimator for the $ATT(2)$ is consistent when *either* of these nuisance working models is correctly specified. If one is exactly right, it does not matter if the other is wrong. Furthermore, if both working models are only slightly wrong, their errors will multiply and DR will perform (asymptotically) better than either one alone.¹⁶ Following the steps in Sant’Anna and Zhao (2020), we can express the $ATT(2)$ as

$$(17) \quad ATT(2) = \mathbb{E}\left[\left(w_{\omega, D=1}(D_i) - w_{\omega, D=0}(D_i, X_i)\right)\left(\Delta Y_{i,t=2} - \mathbb{E}_\omega[\Delta Y_{i,t=2}|X_i, D_i = 0]\right)\right],$$

with the weights as defined in (14). Notice that when we omit the $\mathbb{E}_\omega[\Delta Y_{i,t=2}|X_i, D_i = 0]$ term from (17) we are back to the IPW estimand (13). When we omit the $w_{\omega, D=0}(D_i, X_i)$ term from (17) we are back to the regression adjustment estimand (9).

Constructing a DR DiD estimator for the $ATT(2)$ based on (13) is straightforward. Choose a flexible working model for the propensity score and compute its fitted values, $\hat{\pi}_\omega(X_i)$. Then, choose a flexible working model for the outcome evolution of untreated units and compute its fitted values for all treated and untreated units, $\hat{\mu}_{\omega, \Delta, D=0}(X_i)$. Then use the plug-in principle to estimate the $ATT(2)$ by

$$(18) \quad \widehat{ATT}_{dr}(2) = \frac{1}{n} \sum_{i=1}^n \left(\hat{w}_{\omega, D=1}(D_i) - \hat{w}_{\omega, D=0}(D_i, X_i)\right) \left(\Delta Y_{i,t=2} - \hat{\mu}_{\omega, \Delta, D=0}(X_i)\right),$$

with the estimated weights as defined in (16).

We report the $ATT(2014)$ estimates and standard errors using the DR DiD procedure with a linear in covariates working model for $\mu_{\omega, \Delta, D=0}(X_i)$ and (weighted) logistic regression working model that is also linear in covariates for $p_\omega(X_i)$ in Table 7 (labeled as “Doubly robust”). Population-weighted results are fairly similar across our methods, but the unweighted DR DiD estimates are about halfway between the RA and IPW results. However, an advantage of the DR DiD procedures is that model mis-specifications are arguably less of a concern.

Overall, in this section we discussed different “forward-engineered” estimators for the $ATT(2)$ that respect conditional parallel trends, correctly incorporate covariates, and do not restrict treatment effect heterogeneity. Between RA, IPW, and DR DiD estimators, we recommend practitioners favor the *doubly robust* one compared with the RA and IPW approaches on their own, as this procedure adds additional “protection.” When strong overlap fails, though, RA DiD estimators can still work because they extrapolate the outcome model to obtain predicted counterfactual outcome changes even for treated units with covariate values not observed in the comparison group. The credibility of this extrapolation, however, rests on the accuracy of the working model outside the support of X_i in the untreated group.

¹⁶We usually arrive at such estimands by deriving the efficient influence function; see, for example, Sant’Anna and Zhao (2020) for a discussion on DiD setups. See also Seaman and Vansteelandt (2018) for an overview. When we adopt non-parametric or ML-based estimators for the nuisance functions, we get a different type of double robustness, “rate double robustness,” where we can trade off precision between the different working models. For more details, see, Kennedy et al. (2017) and Smucler, Rotnitzky, and Robins (2019).

If this extrapolation is not reliable, one can make the DR approach more robust against weak overlap by trimming “extreme” propensity scores and performing a bias correction to ensure that the target parameter (the $ATT(2)$ in our case) remains the same. See Ma et al. (2023) for details.

4.5 Heterogeneity Analysis

A different motivation for using covariates is to estimate heterogeneous treatment effects by the values of X_{it} . The identification result in equation (9) can be altered slightly to show that $ATT(2)$ is simply an aggregation of covariate-specific 2×2 DiD estimands:

$$ATT(2) = \mathbb{E}_\omega[\mathbb{E}_\omega[\Delta Y_{i,t=2} | X_i, D_i = 1] - \mathbb{E}_\omega[\Delta Y_{i,t=2} | X_i, D_i = 0] | D_i = 1].$$

Thus, Assumptions CPT and SO also imply that we can identify conditional ATT parameters:

$$\begin{aligned} ATT_{X_i}(2) &\equiv \mathbb{E}_\omega[Y_{i,t=2}(1) - Y_{i,t=2}(0) | D_i = 1, X_i] \\ &= \mathbb{E}_\omega[\Delta Y_{i,t=2} | X_i, D_i = 1] - \mathbb{E}_\omega[\Delta Y_{i,t=2} | X_i, D_i = 0]. \end{aligned}$$

This not only demonstrates the building-block structure of conditional designs, but also connects the strategies we discussed for estimating $ATT(2)$ itself to the underlying treatment effect heterogeneity by covariate values. When this heterogeneity is of interest in its own right it can also be targeted, identified, and estimated (at least under some additional conditions).

When all covariates are discrete, estimating $ATT_{X_i}(2)$ parameters is fairly straightforward. One can form saturated partitions, $x_k, k = 1, \dots, K$, and then subset the data to contain only information from the units from the specified partition k . At this stage, the analysis is analogous to the unconditional DiD setup (except that you need to repeat this step from all partitions of interest), as each $ATT_{x_k}(2)$ is identified by a (conditional on a discrete-variable) comparison of means, that is,

$$ATT_{x_k}(2) = \mathbb{E}_\omega[\Delta Y_{i,t=2} | D_i = 1, X_i = x_k] - \mathbb{E}_\omega[\Delta Y_{i,t=2} | D_i = 0, X_i = x_k].$$

One can estimate these $ATT_{x_k}(2)$ s using their sample analogs or using two-way fixed effects regression specifications analogous to (7). Inference is also standard, provided that each partition is sufficiently large. This type of exercise is commonly used to conduct heterogeneity analysis. For example, suppose one wants to see if the effect of Medicaid expansion on adult mortality rate varies across US census regions. Then, one would partition the data into US regions and run one (unconditional) DiD analysis for each region, provided we have treated and untreated units in each region.

When some covariates are continuous or there are so many partitions that each one contains few observations, one can still *identify* the $ATT_{X_i}(2)$ s, though they are hard to estimate unless auxiliary assumptions hold. In such cases, it is customary to identify and estimate a more aggregated conditional ATT parameter than $ATT_{X_i}(2)$. For instance, one may want to assess if the effect of Medicaid expansion on adult mortality is higher (or lower) in counties with above- versus below-median unemployment rates. Similar partitions could be made for any

other covariates. To formalize this notion of “partition specific” ATT, let $\text{PART}(X_i)$ be some user-specified partition of the covariate space such that $\text{PART}(X_i) \in \{1, 2, \dots, K\}$, and define

$$\text{ATT}_k(2) = \mathbb{E}_\omega[Y_{i,t=2}(1) - Y_{i,t=2}(0) | D_i = 1, \text{PART}(X_i) = k].$$

Under Assumptions CPT and SO, it follows that

$$\begin{aligned} \text{ATT}_k(2) &= \mathbb{E}_\omega[\Delta Y_{i,t=2} | D_i = 1, \text{PART}(X_i) = k] \\ &\quad - \mathbb{E}_\omega[\mathbb{E}_\omega[\Delta Y_{i,t=2} | D_i = 0, X_i, \text{PART}(X_i) = k] | D_i = 1, \text{PART}(X_i) = k], \end{aligned}$$

which implies that

$$\text{ATT}(2) = \sum_{k=1}^K P_\omega(\text{PART}(X_i) = k | D_i = 1) \text{ATT}_k(2).$$

Thus, for heterogeneity analysis with covariates, one can partition the data with a user-specified partition map, $\text{PART}(X_i)$, and then, within each of these partitions, use arguments like the ones we used to establish (9) to guarantee that each partition-specific ATT is identified (which is precisely what we did to get $\text{ATT}_k(2)$ above). Regarding estimation and inference, one can use regression adjustment, inverse probability weighting, or doubly robust estimators as in Section 4.4; the difference is that these are implemented “locally” on each partition.

Obtaining the overall $\text{ATT}(2)$ is just a matter of aggregating these partition-specific ATTs using weights equal to the relative partition size among treated units. All these parameters have clear causal interpretations, can be used to answer different policy questions, and are identified under the same identification assumptions already discussed. Other types of heterogeneity analysis are also possible and even attractive. For instance, Abadie (2005) discusses how one can highlight how the ATT varies across a subset of the covariates required for conditional parallel trends to hold. For example, this would entail checking if the average effect of Medicaid expansion among expansion counties varies with the 2013 poverty rate and/or median income. One can also get the best linear approximation of these conditional ATT curves, which involves estimating fewer parameters. These heterogeneity analyses are generally more granular than the partition-based ones discussed above and do not require discretizing the data. They complement each other well.

We close this section with a remark on the choice of partition and the type of heterogeneity analysis to conduct. Heterogeneity analysis has the potential to offer policymakers and researchers novel insights about the treatment of interest and its mechanisms, opening the door for more informed policy recommendations and targeted expansions. But how should one define the subgroups in which to estimate heterogeneous effects? If researchers knew the relevant partition that policymakers care about, they could aim to estimate them. But this is rarely the case. Taking no stand about heterogeneity implies reporting unit-level effects, which requires incredibly strong assumptions and could also lead to noisy estimates. On the other hand, taking a too-coarse partition may mask important types of treatment effect heterogeneity. So, it seems that a balance between these extremes is important for

a policy-relevant heterogeneity analysis. Estimating conditional ATTs across one or two covariate dimensions is useful, but it may mix some other interactive effects. Another potential avenue is to go beyond averages and adapt the sorted-effects procedure proposed by Chernozhukov, Fernández-Val, and Luo (2018) to DiD designs. It would also be interesting and practically relevant to extend the heterogeneity tools for experimental data discussed in Chernozhukov et al. (2023) to DiD setups. In our view, this is an area in which applied econometrics practice would benefit from more thorough methodological guidance.

5. *DiD Designs with Multiple Time Periods*

The previous sections focused on fundamental identification issues and estimation approaches for 2×2 building blocks, but generally did not build them into anything because they targeted the only feasible $ATT(t)$, $ATT(2)$. DiD designs with multiple periods are notably different. They can target ATTs in each period after treatment to trace out dynamics and they can produce cross-group outcome trends from *before* treatment starts to evaluate the plausibility of parallel trends. Multiple periods also admit the possibility of more complex treatment variation. We will focus on staggered-timing designs in which treatment “turns on” at different times for different units and stays on, as the Medicaid expansion did. It is possible, however, to use DiD methods to study treatments that “turn off” or occur more than once, though some modifications are warranted; see, for example, de Chaisemartin and D’Haultfoeuille (2020, 2024).

We need to expand the notation to define the relevant concepts with multiple periods. When treatment can happen only at one point in time (so we still have only two treatment groups with D_i equal to one or zero), the only changes we need to make are to acknowledge that time runs from $t = 1, 2, \dots, T$, and that the map between potential outcomes and observed outcomes is

$$Y_{i,t} = D_i Y_{i,t}(1) + (1 - D_i) Y_{i,t}(0).$$

By Assumption NA, we have that $Y_{i,t} = Y_{i,t}(0)$ for all pretreatment periods, and that, in posttreatment periods, we observe $Y_{i,t}(0)$ if the group remains untreated by $t = T$, and $Y_{i,t}(1)$ for groups treated at the unique treatment date, $t = g$.

5.1 *Simple Event Studies* ($2 \times T$)

The term “event study” refers to estimating and reporting effects across a range of time periods before and after treatment. A design with one treatment timing group and multiple time periods ($2 \times T$) is the simplest case in which to discuss event studies. We continue to focus on the same two county groups—those in 2014 expansion states or in non-expansion states—but include data from every year between 2009 and 2019. We report population-weighted results for brevity. Figure 2 plots the time series of the weighted mortality rates for the 2014 and non-expansion counties. It is the analog of Table 2 in the sense that it presents the raw data elements necessary to construct event study estimates. The treatment year 2014 naturally divides the x -axis into two windows: posttreatment (2014–19) and pretreatment (2009–13). An event study constructs DiD-type estimates in both windows, but they have different interpretations.

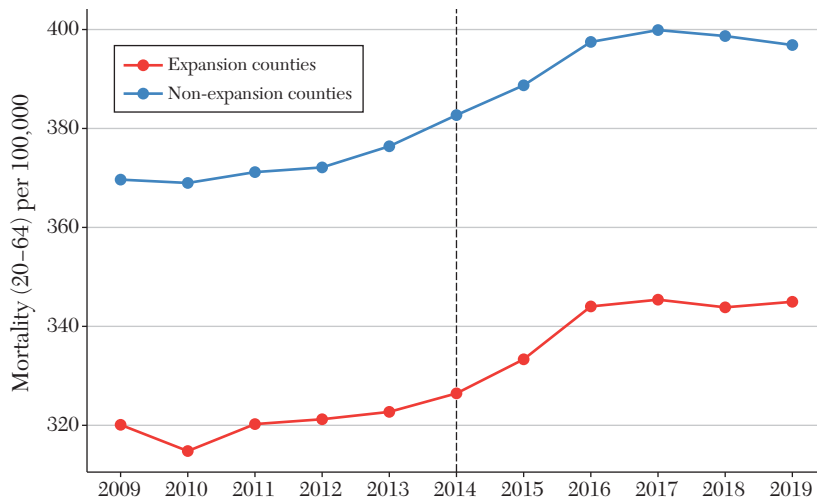


Figure 2. County Mortality Trends by Expansion Decision

Notes: This figure shows county population-weighted average mortality rates for adults ages 20–64 for counties that expanded Medicaid in 2014 and counties that did not expand Medicaid by 2019, from 2009 to 2019.

5.1.1 Event Study Estimates in the Posttreatment Periods

The target parameters in a simple event study are the $ATT(t)$ s with the same definition as in (5); there are just more of them to identify, estimate, and interpret than in a 2×2 design. The $ATT(t)$ s in the posttreatment period, $t \geq g$, reflect treatment effect dynamics. For example, economic models that view health as a stock imply that $ATT(t)$ may grow if Medicaid stimulates health investments. Furthermore, people and institutions may take time to adjust their behavior after Medicaid expands, suggesting a dynamic treatment effect. Event study parameters answer these kinds of subtle questions.

Identification of each $ATT(t)$ follows from the same arguments outlined in Section 3. Note, however, that an event study analysis requires parallel trends in *every* post-period, as in the following assumption.

ASSUMPTION PT-ES ($2 \times T$ Parallel Trends): *The average change of $Y_{i,t}(0)$ from $Y_{i,t=g-1}(0)$ is the same between treated and comparison units for all posttreatment periods $t \geq g$; that is,*

$$(19) \quad \mathbb{E}_\omega[Y_{i,t}(0) - Y_{i,t=g-1}(0) | D_i = 1] = \mathbb{E}_\omega[Y_{i,t}(0) - Y_{i,t=g-1}(0) | D_i = 0] \forall t \geq g.$$

Assumption PT-ES suggests that learning about long-run effects requires stronger assumptions than learning about short-run effects. That is, to identify the average effect of Medicaid expansion in 2019 among expansion counties, Assumption PT-ES requires parallel trends to hold in every year from 2014 to 2019. On the other hand, if we are interested in learning only




Figure 3. $2 \times T$ Event Study

Notes: This figure shows the population-weighted event study estimates in the $2 \times T$ case, comparing counties in states that expanded in 2014 to counties in states that had not yet expanded by 2019. It uses the unconditional estimator from Callaway and Sant'Anna (2021). The outcome variable is the crude mortality rate for adults ages 20–64 and the standard errors are clustered at the county level. The point estimate is reported by the circles, and both point-wise (black) and simultaneous (red) confidence intervals are reported with the vertical lines (see Section 5.1.3).

about short-run effects—say effects up until 2015—we would require Assumption PT-ES to hold from 2014 and 2015 only.

If Assumption PT-ES holds (as well as no anticipation), then each $ATT(t)$ is identified by a DiD comparison between period $g - 1$ and t , as in (5), and the $ATT(t)$ s can be estimated by the familiar comparison of four sample averages:

$$(20) \quad \widehat{ATT}(t) = (\bar{Y}_{\omega, D=1, t} - \bar{Y}_{\omega, D=1, t=g-1}) - (\bar{Y}_{\omega, D=0, t} - \bar{Y}_{\omega, D=0, t=g-1}).$$

Figure 3 plots (weighted) event study estimates for the 2014 Medicaid expansion group. $\widehat{ATT}(t)$ s lie to the right of the vertical dashed line. The estimate for event-time 0 (that is, $t = g = 2014$) equals the 2×2 results from Table 2 (-2.6). The other estimates have the same DiD form: comparisons of cross-group changes in different post-2014 years (t), but always relative to 2013 ($g - 1$). The point estimates do not suggest large mortality effects from Medicaid expansion among expansion counties.

5.1.2 Event Study Estimates in the Pre-periods

Multiple time periods also allow for falsification/placebo tests based on DiD-type comparisons between pretreatment periods. The no-anticipation assumption implies that all $ATT(t)$ s before time g are equal to zero, which means that a 2×2 estimand between periods $t = g - k$

and $t = g - 1, k > 1$, equals a difference in weighted average *untreated* potential outcome trends (as they are all from pretreatment periods):

$$\begin{aligned}\tau_{-k} &= \mathbb{E}_\omega[Y_{i,t=g-k}(0) - Y_{i,t=g-1}(0) | D_i = 1] - \mathbb{E}_\omega[Y_{i,t=g-k}(0) - Y_{i,t=g-1}(0) | D_i = 0] \\ &= \mathbb{E}_\omega[Y_{i,t=g-k} - Y_{i,t=g-1} | D_i = 1] - \mathbb{E}_\omega[Y_{i,t=g-k} - Y_{i,t=g-1} | D_i = 0].\end{aligned}$$

The τ_{-k} terms are usually called “differential trends” or “pre-trends,” and they appear to the left of the vertical dashed line in Figure 3. The credibility of event study analyses rests in large part on finding small estimates of τ_{-k} , but Figure 3 illustrates how challenging it can be to draw conclusions from informally looking at these pre-trends. No individual τ_{-k} is statistically significant, so we fail to reject the null hypothesis that the pre-trend estimates equal zero (individually or jointly). This kind of result is often interpreted to mean that parallel trends hold. But the τ_{-k} 's also tend to be positive with a mean of about 2.3, which is larger in magnitude than all post-period point estimates except one. Sometimes, this kind of result is used to argue that parallel trends fail. Do these results support or refute parallel trends?

Recent methodological work suggests three practical lessons about using pre-trend estimates to assess parallel trends; examples include Bilinski and Hatfield (2026), Manski and Pepper (2018), Kahn-Lang and Lang (2020), Roth (2022), Rambachan and Roth (2023), Dette and Schumann (2024), and Freyaldenhoven et al. (2024). The most fundamental, which will tend to shape language more than practice, is that Assumption PT-ES is not testable, as it only makes restrictions on untreated potential outcomes in posttreatment periods, $t \geq g$. Under no anticipation, pre-trends do measure differences in untreated outcome trends between treated and untreated units, but they necessarily measure them in the “wrong” periods, $t < g$. This does not mean parallel pre-trends are not informative; it just means they are not the *same* as the parallel trends assumption in Assumption PT-ES.

When specifically are observed pre-trends uninformative about the parallel trends assumptions required for identification?⁹ One case is when pre-trends are measured too far before treatment starts. The conditions or shocks that jointly shape treatment decisions and untreated outcomes may differ many periods before treatment, but this need not imply that they generate bias in the periods after treatment. A second case, discussed in Section 3.2, is when units select into treatment based on time-varying pretreatment unobservables. As treatment selection depends on pretreatment unobservables, nonparallel pre-trends may appear as a consequence of the selection mechanism, though parallel trends may still be plausible. Ghanem, Sant'Anna, and Wüthrich (2026) show that in this case, a *necessary* condition for parallel trends is that the untreated potential outcomes satisfy a martingale property. When this martingale property does not hold, parallel trends cannot hold either, *regardless* of the shape of pre-trends. In such situations, it is worthwhile to assess the potential bias components of the DiD estimator using the “selection-aware” benchmarking tools presented in Ghanem, Sant'Anna, and Wüthrich (2026), as well as to resort to their theory-based templates for justifying parallel trends based on selection mechanisms. Ultimately, how informative pre-trends are for the plausibility of Assumption PT-ES is case-specific, though we would of course always prefer to have parallel pre-trends.

When Assumption PT-ES holds in all periods (pre- and posttreatment), it is worth noting that one can construct estimators for the $ATT(t)$ that are more precise than those in (20); see, for example, Borusyak, Jaravel, and Spiess (2024), Gardner (2022), Harmon (2024), Marcus

and Sant’Anna (2021), Wooldridge (2025), and Chen, Sant’Anna, and Xie (2025). In such cases, though, one would require parallel trends to hold in all pretreatment periods, parallel pre-trends become directly testable (as the model is overidentified), and τ_{-k} could be used to assess its plausibility directly. Other types of overidentification tests could also be used in this case to assess the validity of the DiD model; see Marcus and Sant’Anna (2021) and Chen, Sant’Anna, and Xie (2025) for a discussion.

The second lesson is that statistical precision shapes the usefulness of pre-trend estimates. The hypothesis tests for parallel pre-trends in Figure 3 are low powered to detect practically important violations (Roth 2022 and Freyaldenhoven et al. 2024). The τ_{-k} estimates fail to rule out flat pre-trends, but as we shall see below, they also fail to rule out large pre-trends that would indicate serious bias in the $\widehat{ATT}(t)$ s. They simply do not say very much.¹⁷ Roth (2022) discusses how conditioning the analysis on these kinds of low-powered tests can exacerbate biases and should be interpreted with care.

The third lesson is that researchers can make better use of pre-trend estimates by taking a stand on the size of plausible and/or problematic parallel trends violations. For example, Bilinski and Hatfield (2026) propose selecting a value for differential trends that *would* fully explain the estimated treatment effects and then using it instead of zero as the null hypothesis when conducting statistical tests of the estimated pre-trends; see also Dette and Schumann (2024). Rambachan and Roth (2023) develop inference methods for an approach that bounds the $ATT(t)$ s under assumptions about the maximum size of parallel trends violations based on pretreatment periods (see also Manski and Pepper 2018). One can either select a magnitude using contextual knowledge or set it equal to a multiple of the largest period-to-period pre-trend estimate. One then constructs an identified set that contains $ATT(t)$ not under parallel trends but under the weaker assumption that parallel trends are not violated by more than this predetermined maximum in each post-period. Importantly, Rambachan and Roth (2023) also show how to use the estimated covariance matrix of the event study estimates to construct confidence intervals around the set.¹⁸ Given the existence of these methods, which are theoretically grounded in how to consider violations of pre-trends, we caution producers (and consumers) of DiD work against the common practice of using only a simple “eye-test” for whether pre-trends differ substantially from zero.

In our application, Rambachan and Roth’s (2023) method underscores how little information the pre-trend estimates convey. The largest one-period pre-trend in Figure 3 is between event-time -5 and -4 , when outcomes fall by roughly four deaths more in the expansion group versus the non-expansion group. If we assume that parallel trend violations are no bigger than this, the identified set for $ATT(2014)$ is $-2.6 \pm 4 = [-6.6, 1.4]$, and given the size of the pre-period standard errors, we obtain a robust confidence interval of $[-11.1, 5.1]$, which spans implausibly large effects in both directions. Assuming smaller parallel trend violations would shrink this interval, but applying the method to subsequent event times widens it. In general, Rambachan and Roth (2023) provide a flexible method to use information about potential

¹⁷Another possibility is that very precisely estimated pre-trends are distinguishable from zero, yet also rule out even small violations. The magnitude of the violations also matters: A precisely estimated but small violation of parallel trends in pretreatment periods is “better” than an imprecisely potentially large estimated pre-trend that does not rule out zero.

¹⁸In cases where a measured variable accounts for the pre-trends, Freyaldenhoven, Hansen, and Shapiro (2019) develop two-stage-least-squares estimators that recover $ATT(t)$ parameters by extrapolating the pre-period relationship between outcomes and the covariate into the post-period.

parallel trends violations drawn either from external knowledge or the τ_{-k} estimates, as well as the precision of the pre-trend estimates, to gauge (statistically) the robustness of the $\widehat{ATT}(t)$ s.

When pre-trends suggest that Assumption PT-ES fails, a way forward is to assume that it holds only after conditioning on covariates and proceeding similarly to Section 4.4; we discuss this path in detail in Section 5.1.4. Alternatively, one can attempt to parametrically model the violations of parallel trends. Usually this is done by including unit-specific linear trends; see, for instance, Mora and Reggio (2019), Wooldridge (2025, section 7), Lee and Wooldridge (2023), and Freyaldenhoven et al. (2024). We note, however, that the practice of using unit-specific linear trends deviates from the standard DiD procedures: It relies on alternative identification assumptions involving an explicit parametric model for unit-specific trends. We also note that sensitivity analysis procedures that do not rely on such models are available—Rambachan and Roth (2023) is one example—and we encourage practitioners to consider them.

5.1.3 Estimation and Aggregating across Time in Event Studies

The link between a $2 \times T$ event study and a series of 2×2 DiD building blocks makes estimation simple. The point estimates in Figure 3 are the $ATT(t)$ estimates based on (20).¹⁹ An equivalent way to obtain all the $\widehat{ATT}(t)$ s in one step is to run a TWFE regression with time fixed effects, θ_t , unit fixed effects, η_i , and a set of interactions between the treatment group dummy and the time dummies. Omitting the treatment interaction for $t = g - 1$ avoids multicollinearity and fixes $g - 1$ as the baseline period for all β_t estimates, which matches Assumption PT-ES. This generalizes the TWFE regression equation for a single $ATT(t)$ in (7) to

$$(21) \quad Y_{i,t} = \theta_t + \eta_i + \sum_{k=1}^{g-2} \beta_k (\mathbf{1}\{G_i = g\} \cdot \mathbf{1}\{t = k\}) \\ + \sum_{k=g}^T \beta_k (\mathbf{1}\{G_i = g\} \cdot \mathbf{1}\{t = k\}) + \varepsilon_{i,t}.$$

This regression produces identical estimates to those obtained “by hand” via (20): $\hat{\beta}_t = \widehat{ATT}(t)$. It also generates (point-wise) confidence intervals based on clustered standard errors, as discussed in Section 3.3. An additional issue in event study inference, however, involves the fact that we are now estimating many treatment effect parameters. Thus, when we compare across event study estimates we are conducting many hypothesis tests, and the usual normal critical values used to construct confidence intervals do not account for these. Asymptotically correct inferences about the entire event study curve require “inflating” critical values to perform a multiple hypothesis test adjustment. In Figure 3, the thick black bars represent the standard point-wise confidence intervals from clustered standard errors at the county level, while the red line shows uniform confidence bands that cover the 95 percent confidence interval for the entire treatment path of the event study coefficients after accounting for multiple testing. These are produced by default in the Callaway and Sant’Anna (2021) statistical packages, using a multiplier bootstrap procedure to compute critical values of the sup- t test

¹⁹We estimate these with the did R package from Callaway and Sant’Anna (2021).

statistic. Alternatively, one can construct these using the estimated variance-covariate matrix of all $\hat{\beta}_i$'s paired with the Montiel Olea and Plagborg-Møller (2019) simulation procedure. Alternative bootstrap procedures, such as the nonparametric bootstrap, multiplier bootstrap, and weighted/Bayesian bootstrap, can also be used to compute the sup- t critical values that account for multiple testing.²⁰

A final estimation issue arises when targeting aggregations of the $ATT(t)$ s. For example, the average treatment effect in the post-period, $ATT_{avg} = \frac{1}{T - (g - 1)} \sum_{t=g}^T ATT(t)$, is a convenient scalar measure that improves statistical precision, especially when $ATT(t)$ is relatively constant. The easiest way to get an estimate of \widehat{ATT}_{avg} is just to construct it from the 2×2 $\widehat{ATT}(t)$ building-block estimates. Standard post-estimation commands achieve this if the event study estimates come from a regression, and newer DiD packages report this parameter automatically. A common shortcut, however, is to run a second regression that replaces the event study dummies with the treatment status dummy. $D_{i,t} = \mathbf{1}\{D_i = 1\} \times \mathbf{1}\{t \geq g\}$:

$$(22) \quad Y_{i,t} = \theta_t + \eta_i + \beta^{OLS} D_{i,t} + \varepsilon_{i,t}.$$

Unfortunately, the (weighted) least squares estimator $\hat{\beta}^{OLS}$ does not generally equal the \widehat{ATT}_{avg} . The reason is that $\hat{\beta}^{OLS}$ is equivalent to first collapsing the multiple-periods data to averages in the post- and pre-periods and then estimating a 2×2 DiD on the resulting means. This in turn is the same as subtracting the average pre-period τ_{-k} estimate (including the zero at time $g - 1$) from the average post-period $\widehat{ATT}(t)$ estimate. This implies that interpreting both β_i s from (21) and β^{OLS} from (22) requires Assumption PT-ES to hold in every time period, not just in the posttreatment periods.²¹ To the extent that the gap in mean outcomes over the whole pre-period differs from the gap in outcomes in period $t = g - 1$, the two summary parameters will not be equal. In Figure 3, the two summary parameters are quite different: \widehat{ATT}_{avg} equals -0.70 while $\hat{\beta}^{OLS}$ equals -2.53 .²²

5.1.4 Covariates in Event Studies

Another advantage of seeing event studies as collections of 2×2 building blocks is that all the tools for incorporating covariates from Section 4.4 immediately apply to each event study estimate. In fact, the only difference is that instead of using “short differences,” $\Delta Y_{i,t=2}$, one would now use “long differences,” $Y_{i,t} - Y_{i,t=g-1}$. This would imply that using the regression-adjustment procedure would require estimating a working model for $\mathbb{E}_\omega[Y_{i,t} - Y_{i,t=g-1} | D_i = 0, X_i]$ for each time t . The propensity score working model used to construct the IPW DiD estimate (13), on the other hand, is exactly the same as in a 2×2 analysis of the same groups. Since the DR

²⁰The sup- t critical value governs the width of the uniform confidence band that yields simultaneous coverage probabilities for a given confidence level (Montiel Olea and Plagborg-Møller 2019). In our context, its main idea is constructing asymptotically valid critical values for the entire event-study trajectory based on the maximum of all t -statics (one for each event-time considered). This procedure avoids the conservativeness of multiple-testing corrections such as Bonferroni's. See Montiel Olea and Plagborg-Møller (2019) for a general discussion.

²¹The decision to estimate each $\widehat{ATT}(t)$ relative to period $t = g - 1$ comes directly from the choice to define PT that way. If parallel trends hold in every period, one can typically form more efficient estimators than those discussed above. See Marcus and Sant'Anna (2021) for a discussion.

²²Interestingly, this quantity is almost identical to our estimate of $ATT(2014)$, but rather than representing anything reassuring, it comes from the offsetting effects of positive pre-period estimates and small post-period ones.




Figure 4. $2 \times T$ Event Study with Covariates

Notes: This figure shows event study estimates that include covariates. The outcome variable is the crude mortality rate for adults ages 20–64 and the covariates include 2013 values of the percentage of the county population that is female, the percentage of the county population that is White, the percentage of the county population that is Hispanic, the unemployment rate, the poverty rate, and county-level median income. The sample includes counties in states that expanded Medicaid by 2014 and in states that did not expand Medicaid by 2019. The point estimate is reported by the circles and both point-wise (black) and simultaneous (red) confidence intervals are reported with the vertical lines. All procedures use population weights.

DiD estimation procedure builds on both RA and IPW procedures, it would involve estimating different outcome regression working models for each time t . We also note that the potential pitfalls of controlling for covariates in a TWFE specification still apply with multiple periods and actually become more complex (Caetano and Callaway 2024).

For completeness and ease of access, we list the RA, IPW, and DR estimands for $ATT(t)$:

$$\begin{aligned}
 ATT_{ra}(t) &= \mathbb{E}_{\omega} [Y_{i,t} - Y_{i,t=g-1} | D_i = 1] - \mathbb{E}_{\omega} [\mathbb{E}_{\omega} [Y_{i,t} - Y_{i,t=g-1} | X_i, D_i = 0] | D_i = 1], \\
 ATT_{ipw}(t) &= \mathbb{E} \left[(w_{\omega, D=1}(D_i) - w_{\omega, D=0}(D_i, X_i)) (Y_{i,t} - Y_{i,t=g-1}) \right], \\
 ATT_{dr}(t) &= \mathbb{E} \left[(w_{\omega, D=1}(D_i) - w_{\omega, D=0}(D_i, X_i)) (Y_{i,t} - Y_{i,t=g-1} \right. \\
 &\quad \left. - \mathbb{E}_{\omega} [Y_{i,t} - Y_{i,t=g-1} | X_i, D_i = 0]) \right],
 \end{aligned}$$

where $w_{\omega, D=1}(D_i)$ and $w_{\omega, D=0}(D_i, X_i)$ are as defined in (14).

Figure 4 shows weighted event study estimates using our three preferred covariate strategies: regression adjustment, inverse propensity weighting, or doubly-robust estimation. In our case, covariates do little to change the unadjusted estimates. Note, however, that Borgschulte and Vogler (2020) use an IPW estimator with different covariates selected by a LASSO procedure and obtain notably stronger evidence of mortality reductions. Evidently, the exact set of covariates one conditions on matters a great deal in this analysis and, as we have discussed previously, one should attempt to include in X_i all the determinants of the change in untreated potential outcome or of the treatment assignment.

5.2 Staggered Treatment Adoption ($G_{\#} \times T$)

Viewing $2 \times T$ event studies as a collection of 2×2 DiD building blocks makes the jump to staggered timing designs straightforward. The key distinction is that with staggered timing, each treatment date defines a distinct treatment group, and each of these has its own set of simple event study parameters. New choices arise about the comparison units used to identify and estimate these group-specific event studies as well as about how to aggregate the estimates across timing groups. The 2×2 structure, as well as all the tools we have developed to evaluate parallel trends (covariate balance and pre-trends) and to estimate (with or without weights and covariates) carry over.

When treatment start dates can vary across units, we need to allow the potential outcomes, and thus the target parameters and identifying assumptions, to reflect this richer notion of treatment. We therefore index potential outcomes by the time treatment begins, g : $Y_{i,t}(g)$, and use $Y_{i,t}(\infty)$ to denote never-treated potential outcomes.²³ We use G_i to denote each unit's treatment date, and with some abuse of terminology, we call units not exposed to treatment by period T the “never-treated” group.²⁴ Finally, we use \mathcal{G} to represent the set of all treatment times (rows of Table 1 in our example). With these modifications we can map potential outcomes to observed outcomes using a generalization of (1):

$$Y_{i,t} = \sum_{g \in \mathcal{G}} Y_{i,t}(g) \mathbf{1}\{G_i = g\}.$$

With multiple treatment groups, we also need to extend our notion of no anticipation (though its empirical content is exactly the same).

ASSUMPTION NA-S (No Anticipation with Staggered Treatment Timing): *For all units i that are eventually treated and all pretreatment periods t , $Y_{i,t}(g) = Y_{i,t}(\infty)$.*

Like Assumption NA, Assumption NA-S imposes that treatment effects are zero in all pretreatment periods as a consequence of units not acting on the potential knowledge of future treatment dates before they are actually exposed to treatment. We maintain this assumption throughout this section.

Finally, we assume that a “never-treated” group always exists in our staggered DiD setup. If all units are eventually treated, we drop all the data from when the last cohort is treated, so the last-treated cohort becomes the “never-treated” cohort, and T here denotes the number

²³Let $\mathbf{0}_s$ and $\mathbf{1}_s$ be s -dimensional vectors of zeros and ones, respectively, and denote the potential outcome for unit i at time t if first exposed to the treatment at time g by $Y_{i,t}(\mathbf{0}_{g-1}, \mathbf{1}_{T-g+1})$, and denote by $Y_{i,t}(\mathbf{0}_T)$ the outcome if untreated by time $t = T$. We discussed the two-period treatment this way in Section 3.1, when we defined potential outcomes as a function of the period one and period two treatment. These treatment paths define the potential outcomes we work with: $Y_{i,t}(g) = Y_{i,t}(\mathbf{0}_{g-1}, \mathbf{1}_{T-g+1})$ and $Y_{i,t}(\infty) = Y_{i,t}(\mathbf{0}_T)$. Writing potential outcomes as functions of treatment paths helps with transparency regarding causal parameters of interest and the DiD design we are in. When treatment can turn on and off, writing potential outcomes in terms of the entire path becomes crucial to avoid “hidden” assumptions that rule out treatment effect heterogeneity and dynamics. Because of space constraints, we do not cover these cases in this article.

²⁴In practice, “never-treated” really means “not observed to be treated by $t = T$.” Given more data, units untreated at T could, in many cases, take up the treatment. In fact, this is the case with the Medicaid expansion. We use data through 2019 but include states that expanded Medicaid in 2020, 2021, and 2023 as “never treated,” alongside states that have not expanded as of 2024.




Figure 5. County Mortality Trends by Expansion Decision with Staggered Timing

Notes: This figure shows county population-weighted average mortality rates for adults ages 20–64 from 2009 to 2019. There are 978 counties in states in the 2014 expansion group, 171 counties in the 2015 expansion group, 93 counties in the 2016 expansion group, 140 counties in the 2019 expansion group, and 1,222 counties that did not expand Medicaid by 2019.

of available periods in the subset of the data that we will use in our analysis. This is essentially without loss of generality, because under standard DiD assumptions we cannot identify any *ATT* for periods where all units are treated. We also dropped data from units treated in the first available period, $G_i = 1$, as such a treatment group does not have any pretreatment data, preventing us from conducting a DiD analysis.

Figure 5 plots average weighted mortality data by the year of Medicaid expansion (G_i) and time. As in Table 2 and Figure 2, these are all the means necessary to calculate a staggered DiD estimate.

5.2.1 Building-Block Parameters with Staggered Adoption

Staggered treatment timing affects the structure of a DiD analysis because it changes the definition of treatment. Until now, we have used counties in the 2014 expansion states as the only treatment group represented with a single treatment dummy, D_i . But Table 1 shows that as of 2019, there were four different groups of expansion states defined by whether they expanded Medicaid in 2014, 2015, 2016, or 2019. Therefore, D_i is not rich enough to capture the relevant definition of treatment groups in staggered setups because there are many treatment groups, not just two. Fortunately, we can use the treatment timing notation, G_i , to define *ATT* parameters, parallel trends assumptions, and estimators, just as we have done so far.

A simple 2×2 DiD design had one target parameter ($ATT(2)$), and a $2 \times T$ DiD design had $T - 1$ of them: $T - (g - 1)$ posttreatment parameters, $ATT(t), t \geq g$, and $g - 2$ pre-trend

parameters. In staggered DiD designs, each treatment group (sometimes referred to as a cohort), defined by its treatment date g , has its own set of $T - 1$ event study parameters. We call these group-time average treatment effects:

$$ATT(g, t) = \mathbb{E}_\omega[Y_{i,t}(g) - Y_{i,t}(\infty) | G_i = g].$$

Each $ATT(g, t)$ is the average treatment effect of starting treatment at period g relative to never starting it, at time period t , among units that actually started treatment in period g . It is simply a set of event study parameters for treatment timing group g . The unobserved counterfactuals are now untreated potential outcome means for each treatment group in each period, $\mathbb{E}_\omega[Y_{i,t}(\infty) | G_i = g]$.

5.2.2 Identification with Staggered Designs

Identifying the $ATT(g, t)$ s works exactly as in the previous sections because they are just group-specific event studies. Under no anticipation, a set of parallel trends assumptions for $t \geq g$ identifies the causal posttreatment parameters. DiD comparisons for $t < g$ represent differential pre-trends in untreated potential outcomes.

The most important way that staggered DiD changes this approach is that having access to multiple treatment groups with different treatment starting dates allows one to use alternative sets of comparison groups. For example, our Medicaid analysis so far has used counties in states that did not expand Medicaid by 2019 as the comparison group. For estimating, say, the ATT for the 2014 expansion group in 2015, counties in states that did not expand until 2016, 2017, or 2018 could also serve as comparison units. Choosing which comparison groups to use to identify an $ATT(g, t)$ is directly tied to which form(s) of parallel trends hold. Relative to simple event studies and especially 2×2 setups, staggered timing creates many potential parallel trends assumptions. Here, we discuss three types of staggered parallel trends. The first two use either the never-treated units or any not-yet-treated groups as the comparisons for all eventually treated groups (Callaway and Sant'Anna 2021), and the third option assumes that parallel trends hold in all periods and in all groups, which is something that several DiD methods require; see de Chaisemartin and D'Haultfœuille (2020), Sun and Abraham (2021), Wooldridge (2025), Borusyak, Jaravel, and Spiess (2024), and Harmon (2024).²⁵

ASSUMPTION PT-GT-NEV (Parallel Trends Based on Never-Treated Groups): *For every eventually treated group g and posttreatment time period $t \geq g$,*

$$\mathbb{E}_\omega[Y_{i,t}(\infty) - Y_{i,t-1}(\infty) | G_i = g] = \mathbb{E}_\omega[Y_{i,t}(\infty) - Y_{i,t-1}(\infty) | G_i = \infty].$$

²⁵ While these three types of parallel trends are fairly intuitive choices, others are possible. For instance, Cengiz et al. (2019) use a comparison group of units treated at least $\delta + 1$ periods after time g . Thus, their parallel trends are tailored to this particular choice. As a result, all of their $ATT(g, t)$ estimates from time g to time $g + \delta$ use the same comparison group. Marcus and Sant'Anna (2021) also discuss other alternative parallel trends assumptions.

ASSUMPTION PT-GT-NYT (Parallel Trends Based on Not-Yet-Treated Groups): For every eventually treated group g , not-yet-treated group g' and time periods t such that $t \geq g$ and $g' > t$,

$$\mathbb{E}_\omega[Y_{i,t}(\infty) - Y_{i,t-1}(\infty) | G_i = g] = \mathbb{E}_\omega[Y_{i,t}(\infty) - Y_{i,t-1}(\infty) | G_i = g'].$$

ASSUMPTION PT-GT-ALL (Parallel Trends for Every Period and Group): For every treatment group g and g' and time periods t ,

$$\mathbb{E}_\omega[Y_{i,t}(\infty) - Y_{i,t-1}(\infty) | G_i = g] = \mathbb{E}_\omega[Y_{i,t}(\infty) - Y_{i,t-1}(\infty) | G_i = g'].$$

Assumption PT-GT-NEV is the analog of the PT assumptions we used in the $2 \times T$ design. It uses the never-treated units as the relevant comparison group for all eventually treated units, and it imposes parallel trends in posttreatment periods only. In our Medicaid application, this would entail using the non-expansion counties as the comparison group for the 2014, 2015, 2016, and 2019 expansion groups. In addition, since Assumption PT-GT-NEV imposes parallel trends only for the future, the farthest we can go into pretreatment periods is $t = g - 1$, which will serve as the only (justifiable) baseline period. More formally, under Assumption PT-GT-NEV, it is straightforward to show that for posttreatment periods,

$$(23) \quad ATT(g, t) = \mathbb{E}_\omega[Y_{i,t} - Y_{i,t=g-1} | G_i = g] - \mathbb{E}_\omega[Y_{i,t} - Y_{i,t=g-1} | G_i = \infty],$$

which shows that $ATT(g, t)$ is identified (Callaway and Sant'Anna 2021; Sun and Abraham 2021). Note that (23) highlights that we are essentially back to a 2×2 design when it comes to learning about $ATT(g, t)$: It leverages data from only two periods, t (post) and $g - 1$ (pre), and two treatment groups, $G_i = g$ (treated) and $G_i = \infty$ (comparison).

Under assumption PT-GT-NYT, one can use not only the never-treated units but any group of units that are not-yet-treated by time t . In our Medicaid example, we could now use non-expansion counties and 2016 and 2019 expansion counties as comparison groups when estimating $ATT(2014, 2015)$. Using Assumption PT-GT-NYT, Callaway and Sant'Anna (2021) have shown that we can identify the $ATT(g, t)$ for posttreatment periods $t \geq g$ by²⁶

$$(24) \quad ATT(g, t) = \mathbb{E}_\omega[Y_{i,t} - Y_{i,t=g-1} | G_i = g] - \mathbb{E}_\omega[Y_{i,t} - Y_{i,t=g-1} | G_i > \max\{g, t\}].$$

Like equation (23), this follows a 2×2 setup: It leverages data from only two periods, t (post) and $g - 1$ (pre), and two treatment groups, $G_i = g$ (treated), and $G_i > t$ (comparison).²⁷

²⁶de Chaisemartin and D'Haultfoeuille (2020) have derived the same results but restrict attention to instantaneous treatment effects—that is, $ATT(g, g)$ s. They do allow for treatment turning on and off, but also impose that being exposed to a treatment today does not affect outcomes tomorrow (a no-carryover assumption). See de Chaisemartin and D'Haultfoeuille (2024) for some extensions.

²⁷We note that more general results are also possible. In fact, for any not-yet-treated group, $g' > t$, it is easy to show that under Assumption PT-GT-NYT, for $t \geq g$, $ATT(g, t) = \mathbb{E}_\omega[Y_{i,t} - Y_{i,t=g-1} | G_i = g] - \mathbb{E}_\omega[Y_{i,t} - Y_{i,t=g-1} | G_i = g']$. One can then flexibly combine the different comparison groups by leveraging user-specified or efficiency-oriented weights. Chen, Sant'Anna, and Xie (2025) discuss how to efficiently explore all the information implied by the identification assumptions to form semiparametrically efficient DiD estimators—that is, estimators that asymptotically enjoy the shortest possible (theoretically justified) confidence intervals without making strong functional form or model-based assumptions related to error terms, such as homoskedasticity and restrictions on serial dependence.

Finally, under assumption PT-GT-ALL, one can make use of any not-yet-treated units as a comparison group, as well as any pretreatment period as a baseline period. For instance, to identify the ATT for 2015 expansion counties, we can now use the never-treated group and the 2016 and 2019 expansion groups and use any or all of the years from 2009 to 2014 as a baseline. Using Assumption PT-GT-ALL, it is easy to show that, for any pretreatment period, $t_{pre} < g$, and any not-yet-treated group, $g' > t$, we can identify the $ATT(g, t)$ for group g 's posttreatment periods $t \geq g$ by

$$(25) \quad ATT(g, t) = \mathbb{E}_\omega[Y_{i,t} - Y_{i,t_{pre}} | G_i = g] - \mathbb{E}_\omega[Y_{i,t} - Y_{i,t_{pre}} | G_i = g'],$$

which again maps back to the 2×2 DiD setup, as it leverages two periods, t (post) and t_{pre} (pre), and two groups, $G_i = g$ (treated) and $G_i = g'$ (comparison). This representation makes it clear that, in practice one can use various pretreatment periods and not-yet-treated comparison groups to characterize the $ATT(g, t)$ under Assumption PT-GT-ALL. We can also combine several of these to form an $ATT(g, t)$ estimand that uses more data (Wooldridge 2025; Gardner 2022; Liu, Wang, and Xu 2024; Borusyak, Jaravel, and Spiess 2024; Chen, Sant'Anna, and Xie 2025). An intuitive estimand that naturally extends (24) and allows us to use more pretreatment data is given by

$$(26) \quad ATT(g, t) = \mathbb{E}_\omega[Y_{i,t} - \bar{Y}_{i,t \leq g-1} | G_i = g] - \mathbb{E}_\omega[Y_{i,t} - \bar{Y}_{i,t \leq g-1} | G_i > \max\{g, t\}],$$

where $\bar{Y}_{i,t \leq g-1} = \sum_{s=1}^{g-1} Y_{i,s} / (g-1)$ is the time average of group g pretreatment periods for each unit i ; see, for example, Callaway (2023, section 3.2), Lee and Wooldridge (2023), and the discussion in de Chaisemartin and D'Haultfœuille (2023, section 3.2.4). Although the estimand does look different from the previous one, it too resembles a 2×2 design: It effectively leverages two periods, t (post) and the average of all period $t < g$ (pre), and two groups, $G_i = g$ (treated) and $G_i > t$ (comparison). In practice, however, there is no general econometrics guarantee that estimators based on (26) will be more precise than estimators based on (24); see, for instance, Harmon (2024). This arises because it is not always optimal to weigh all pretreatment periods equally when forming estimators for $ATT(g, t)$. In fact as discussed in Chen, Sant'Anna, and Xie (2025), the optimal way (or, more formally, the semiparametric efficient way) to aggregate information across pretreatment periods and comparison groups under Assumption PT-GT-ALL depends on the correlation structure of how the outcome changes over time across different comparison groups. And an effective way to leverage this information consists of constructing DiD estimators for $ATT(g, t)$ s that efficiently weigh several 2×2 DiD estimators for the $ATT(g, t)$ that use different comparison groups and different baseline periods. These efficiency weights do not depend on additional hard-to-motivate assumptions (e.g., homoskedasticity or restriction on the serial correlation), arise as a consequence of the information content of the identification assumptions, and can be transparently visualized using the tools provided by Chen, Sant'Anna, and Xie (2025). What is perhaps most important here is that even these more complex DiD estimators closely resemble the approach we took in the 2×2 design.

In the end, a natural question arises: Which parallel trends assumption should one use? This context-specific question is hard to answer, as each assumption has pros and cons. For instance, Assumption PT-GT-ALL leads to more precise $ATT(g, t)$ estimators because it uses

data from multiple pretreatment periods and multiple comparison groups. Given that power is important when conducting causal inference, this is appealing. On the other hand, it imposes parallel pre-trends, an assumption that is not *required* for identification of $ATT(g, t)$ and that we have not imposed in $2 \times T$ DiD designs (see our discussion of equation (21)). If pre-trends are not parallel, then estimates of $ATT(g, t)$ based on Assumption PT-GT-ALL can be biased.

The other extreme is to make Assumption PT-GT-NEV and use only comparison groups made up of never-treated units. This avoids compositional changes in the comparison group over time, does not restrict pre-trends, and identifies all the $ATT(g, t)$ s.²⁸ It also avoids using as a comparison group units that may have chosen to begin treatment in a given period because of pretreatment outcomes, which potentially violates parallel trends. For instance, states that expanded Medicaid in 2016 may have done so on the basis of county mortality rates from previous years, likely violating the parallel trends assumption for this group. On the other hand, never-treated units may have remained untreated for reasons related to trends in $Y_{i,t}(0)$. Non-expansion counties may be too different from expansion counties for them to reflect the relevant counterfactual. This could be, in part, justified after examining the differences in covariate levels and trends between treatment and control groups, as we saw in the discussion of Table 4. Also, depending on how widespread treatment is, there may be too few never-treated units to obtain precise estimates.

We view Assumption PT-GT-NYT as a middle step that uses all not-yet-treated units as a comparison group without restricting all pretreatment trends to be parallel. It uses more information than Assumption PT-GT-NEV, which can lead to gains in precision and helps to incorporate covariates. While it uses less information than Assumption PT-GT-ALL, it is also less susceptible to bias from violations of parallel pre-trends. For our Medicaid application, we favor Assumption PT-GT-NYT, as we prefer not to impose parallel pre-trends from 2009 to 2014 or to rely exclusively on comparisons to the set of states that have not expanded Medicaid as of 2024. On the other hand, if parallel trends are not plausible for a particular group of eventually treated units, perhaps owing to selection based on time-varying unobservables (Ghanem, Sant'Anna, and Wüthrich 2026), it is important to remove these units from the DiD analysis to retain interpretability. Ultimately, the plausibility of each parallel trends assumption may vary across different contexts. At the very least, we strongly recommend that researchers clearly state the specific parallel trends assumption they are actually imposing in their analysis to allow readers to discuss its plausibility in a scientifically grounded manner.

5.2.3 *Estimators for Staggered Designs without Covariates*

The identification results for $ATT(g, t)$ discussed in Section 5.2.2 suggest very simple and intuitive estimators for the $ATT(g, t)$. Given the estimand that comes from the chosen parallel trends assumption, the estimators replace the population (weighted) expectations with their sample analogs. The principle is the same as in the 2×2 setup of Section 3.3.

²⁸Without never-treated units, we cannot estimate $ATT(g, t)$ for the last observed treatment date, which shapes the feasible target parameters. For example, in our Medicaid expansion example, without the presence of never-treated (by 2019) counties we would not be able to estimate the treatment effects in 2019 (i.e., $ATT(g = 2015, t = 2019)$).

For example, under Assumption PT-GT-NYT we can leverage (24) and form plug-in estimators for $ATT(g, t)$ using

$$(27) \quad \widehat{ATT}_{nyt}(g, t) = \frac{\sum_{i=1}^n \mathbf{1}\{G_i = g\} \omega_i (Y_{i,t} - Y_{i,t=g-1})}{\sum_{i=1}^n \mathbf{1}\{G_i = g\} \omega_i} - \frac{\sum_{i=1}^n \mathbf{1}\{G_i > t\} \omega_i (Y_{i,t} - Y_{i,t=g-1})}{\sum_{i=1}^n \mathbf{1}\{G_i > t\} \omega_i}.$$

This simple estimator is what Callaway and Sant'Anna (2021) propose when one uses the not-yet-treated group as the comparison group.²⁹

When Assumption PT-GT-NEV holds, it is straightforward to build on (23) and estimate $ATT(g, t)$ by

$$(28) \quad \widehat{ATT}_{never}(g, t) = \frac{\sum_{i=1}^n \mathbf{1}\{G_i = g\} \omega_i (Y_{i,t} - Y_{i,t=g-1})}{\sum_{i=1}^n \mathbf{1}\{G_i = g\} \omega_i} - \frac{\sum_{i=1}^n \mathbf{1}\{G_i = \infty\} \omega_i (Y_{i,t} - Y_{i,t=g-1})}{\sum_{i=1}^n \mathbf{1}\{G_i = \infty\} \omega_i}.$$

This estimator was proposed by Callaway and Sant'Anna (2021) and Sun and Abraham (2021) when using never-treated units as a comparison group, though Sun and Abraham (2021) arrive at this using a fully saturated regression specification and estimating the regression coefficients $\beta_{g,e}^{SA}$ with (weighted) least squares,

$$(29) \quad Y_{i,t} = \theta_t + \eta_i + \sum_{g \neq \infty} \sum_{e \neq -1} \beta_{g,e}^{SA} \mathbf{1}\{G_i = g\} \mathbf{1}\{G_i + e = t\} + \epsilon_{i,t}.$$

It is straightforward to show that $\beta_{g,e}^{SA} = \widehat{ATT}_{never}(g, g + e)$, emphasizing that (29) is just a way to contrast sample means across groups and periods that respect Assumption PT-GT-NEV.

When Assumption PT-GT-ALL holds instead, one can construct plug-in estimators for (26):

$$\widehat{ATT}_{avg}(g, t) = \frac{\sum_{i=1}^n \mathbf{1}\{G_i = g\} \omega_i (Y_{i,t} - \bar{Y}_{i,t \leq g-1})}{\sum_{i=1}^n \mathbf{1}\{G_i = g\} \omega_i} - \frac{\sum_{i=1}^n \mathbf{1}\{G_i > t\} \omega_i (Y_{i,t} - \bar{Y}_{i,t \leq g-1})}{\sum_{i=1}^n \mathbf{1}\{G_i > t\} \omega_i}.$$

²⁹ Recently, Dube et al. (2025) show that one can also get $ATT(g, t)$ estimates that are equivalent to $\widehat{ATT}_{nyt}(g, t)$ by using local projections. One can also get similar estimators using a “stacked DiD” procedure akin to what Fadlon and Nielsen (2021), Deshpande and Li (2019), and Cengiz et al. (2019) have implemented.

Alternatively, Wooldridge (2025) proposed constructing estimators for $ATT(g, t)$ based on Assumption PT-GT-ALL, using the following “extended” TWFE specification:

$$(30) \quad Y_{i,t} = \theta_t + \eta_i + \sum_{g \neq \infty} \sum_{s=g}^T \beta_{g,s}^W \mathbf{1}\{G_i = g\} \mathbf{1}\{s = t\} + \epsilon_{i,t},$$

where the $\beta_{g,t}^W$ s are estimated using (weighted) least squares. Wooldridge (2025) shows that $\hat{\beta}_{g,t}^W$ consistently estimates $ATT(g, t)$ under Assumption PT-GT-ALL, though we do not know the exact way that $\hat{\beta}_{g,t}^W$ combines pretreatment periods and not-yet-treated units, which is to say that we do not know the statistical estimand associated with $\beta_{g,t}^W$. Wooldridge (2025) also shows that $\hat{\beta}_{g,t}^W$ is numerically the same as the “imputation” estimators proposed by Gardner (2022), Liu, Wang, and Xu (2024), and Borusyak, Jaravel, and Spiess (2024) with balanced panel data (and no covariates).³⁰

Overall, these different $ATT(g, t)$ estimators highlight that we can leverage our DiD expertise built in the 2×2 setup to estimate heterogeneous $ATT(g, t)$ parameters. The exact estimator one uses should be shaped by which parallel trends assumption holds. In our Medicaid application, we report in Figure 6 our $ATT(g, t)$ estimates based on Assumption PT-GT-NYT and (27). As we have four sets of counties defined by their Medicaid expansion timing, we report four sets of event studies, one for each expansion group. For the 2014, 2016, and 2019 expansion groups, we find that Medicaid did not lead to significant changes in adult mortality rates. For the 2015 expansion group, adult mortality rates rose after expansion. Regarding pre-trends, Figure 6 suggests that there may be some non-negligible pre-trends for the 2016 expansion group, though these are not statistically different from zero.

In Section 4, we highlighted that unconditional assumptions may fail in our Medicaid context. One reason is that important determinants of changes in untreated adult mortality rates are imbalanced across treatment and comparison groups. In such cases, one should interpret the results in Figure 6 with great care. In addition, it is also important to highlight that, as indicated in Table 1, the 2015, 2016, and 2019 expansion groups are relatively small and represent only 6 percent, 2 percent, and 3 percent of the US population, respectively. Thus, analyzing these groups separately may be “too noisy” and not representative of the overall effect for the US population. This does not mean that these $ATT(g, t)$ s are not useful. They are actually an essential part of our DiD analysis, but we may want to aggregate the cohorts to get more informative target parameters for the overall treated population. We now turn to how to aggregate the $ATT(g, t)$ s and then discuss how to incorporate covariates.

³⁰Imputation procedures work in two steps. The first step uses all untreated observations to run the regressions $Y_{it} = \theta_t + \eta_i + \epsilon_{it}$ using only data from the (i, t) pairs that satisfy this criterion and get the fitted values $\hat{Y}_{i,t}(0) = \hat{\theta}_t + \hat{\eta}_i$ for all eventually treated observations. The second step estimates $ATT(g, t)$ by $\widehat{ATT}_{imp}(g, t) = \frac{\sum_{i=1}^n \mathbf{1}\{G_i = g\} (Y_{i,t} - \hat{Y}_{i,t}(0))}{\sum_{i=1}^n \mathbf{1}\{G_i = g\}}$. See Gardner (2022), Liu, Wang and Xu (2024), and Borusyak, Jaravel, and Spiess (2024) for details. Note that the specification in (30) is similar to Sun and Abraham's (2021) specification (29), but it omits the pretreatment event-time dummies. This is justified because (30) imposes parallel pre-trends (Assumption PT-GT-ALL) while (29) does not (it effectively relies on Assumption PT-GT-NEV). Thus, in general, one should not expect $\hat{\beta}_{g,t}^W$ to be equal to $\hat{\beta}_{g,t}^{SA}$.

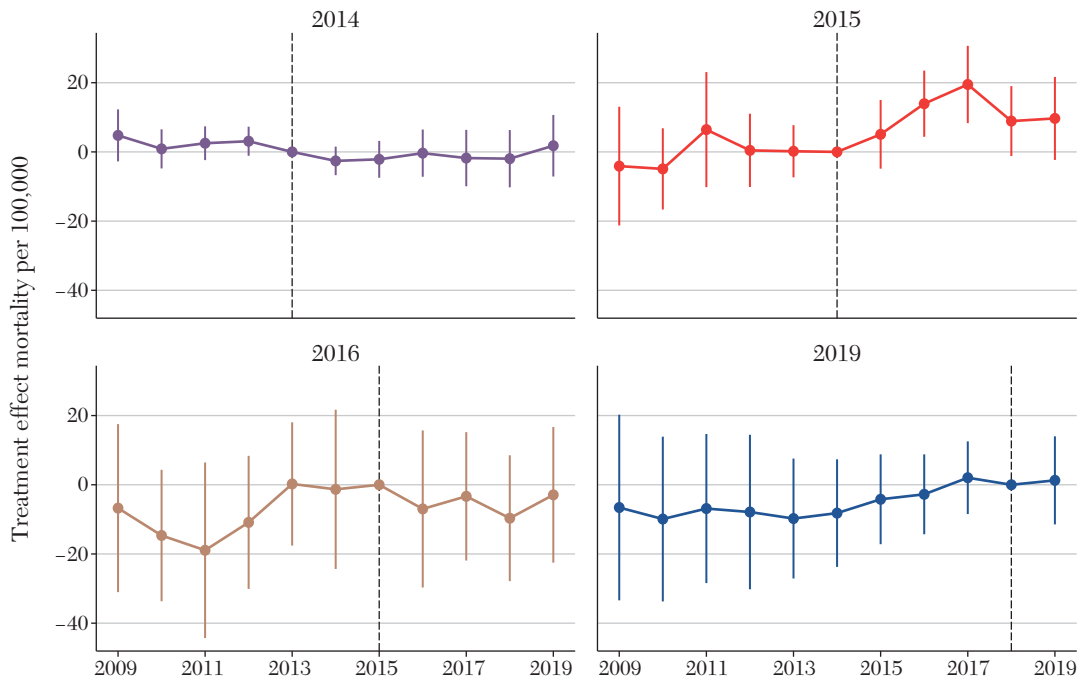


Figure 6. $ATT(g,t)$ s for Each Expansion Group

Notes: This figure shows the group-time ATT estimates ($ATT(g,t)$) in calendar time for the four groups of counties that expanded Medicaid by 2019. Each panel uses 1,222 not-yet-treated counties as the comparison group and shows their uniform confidence intervals at the 95 percent significance level. The outcome variable is the crude mortality rate for adults ages 20–64, and standard errors are clustered at the county level. The vertical line represents the year before Medicaid expansion (i.e., $g - 1$) for the timing group. All results use population weights.

5.2.4 Aggregating Group-Time Average Treatment Effects

The previous section highlighted that, in many applications, estimating all $ATT(g,t)$ s precisely and attaching policy-relevant interpretations to them may be challenging. Aggregating them into a summary treatment effect measure therefore has clear benefits: It improves precision, reduces the number of results, and yields a parameter that averages over all treated units like the $ATT(2)$ identified in 2×2 designs.

Aggregation in staggered designs involves a notion of time (either calendar time t or event-time $e = t - g$), a length of time (how many periods to aggregate across), and group weights (so larger treatment groups can “matter more” than smaller ones). Given some set of weights, it is simple to average the $ATT(g,t)$ building blocks into many kinds of summary parameters:

$$ATT_{aggte} = \sum_{g,t} w_{\omega,g,t} ATT(g,t),$$




Figure 7. $ATT(g,t)$ in Event Time

Notes: This figure shows the group-time ATT estimates ($ATT(g,t)$) in relative event time for the four treatment timing groups of counties that expanded Medicaid by 2019, using not-yet-treated units as the comparison group. The outcome variable is the crude mortality rate for adults ages 20–64. All estimates use population weights.

where $w_{\omega,g,t}$ is a “generic” (ω -weighted) group and time-specific (nonnegative) weight that sums up to one. Specific choices for the $w_{\omega,g,t}$ weights map to different ways to aggregate and present interpretable causal effects in this kind of complex setting.

As highlighted in Section 5.1, an appealing feature of having access to data from multiple periods is that we can assess how average treatment effects evolve with the time since treatment, or event time $e = t - g$. Figure 6 displays event studies for each expansion group, and the aggregation question is how to combine them into a single summary event study.

A basic observation about weighting is that we can give positive weights only to the $ATT(g,t)$ s that we actually identified and estimated. Earlier treated groups have $ATT(g,t)$ estimates for later event times by definition (and later treated groups have estimates for earlier pre-trends), so it will not be possible to include every group in an aggregated event-study parameter at every event time. To see which $ATT(g,t)$ s will contribute to our event study, Figure 7 recenters our $ATT(g,t)$ estimates in event time instead of calendar time. That is, we plot $ATT(g, g + e)$ against e for each expansion group.

We will take “vertical” weighted averages of each available $ATT(g, g + e)$ for each event time e based on Figure 7. For instance, to estimate an aggregate event study in event time 0 (a measure of instantaneous treatment effects), we would average estimates of $ATT(2014, 2014)$, $ATT(2015, 2015)$, $ATT(2016, 2016)$ and $ATT(2019, 2019)$. When we are interested in event time 1, we would now average $ATT(2014, 2015)$, $ATT(2015, 2016)$, $ATT(2016, 2017)$. The same logic applies to other event times.

When constructing timing-group weights at a given event time, it is also important to account for group sizes so that the resulting parameters equal sensible averages of treated units. Table 1 contains all the information necessary for this. The 2014 expansion group

accounts for 80 percent of treated adults in the groups we consider, while the 2016 expansion group accounts for 3.5 percent. If we would like our aggregate event study to be a representative summary of the dynamic effects among treated counties, we should choose weights that are proportional to the treatment group size.

Putting these pieces together, we can formally state the exact summarized causal parameter that highlights treatment effect dynamics in terms of event time:

$$(31) \quad \begin{aligned} ATT_{es}(e) &= \mathbb{E}_{\omega} [ATT(G, G + e) | G + e \in [1, T], G \leq T] \\ &= \sum_{g < \infty} w_{\omega, g, e}^{es} ATT(g, g + e), \end{aligned}$$

where each weight $w_{g, e}^{es}$ gives the share of a group $G = g$ among treated units that have been exposed to treatment for exactly e periods (the groups that we have data for event time e in Figure 7), and is formally defined as

$$w_{\omega, g, e}^{es} = \mathbf{1}\{g + e \in [1, T]\} P_{\omega}(G = g | G + e \in [1, T], G \leq T).$$

Note that $ATT_{es}(e)$ gives the average treatment effect among the units that have been exposed to treatment for exactly e periods, conditional on being observed having participated in the treatment for that number of periods (the condition that $G + e \in [1, T]$) and ever-participating in the treatment by period T ($G \leq T$). One can also take a simple average of all available posttreatment event times, $ATT_{es}(e), e \geq 0$, and report an overall ATT measure. See Callaway and Sant’Anna (2021) for a discussion of alternative aggregations based on calendar time and groups.

Estimating $ATT_{es}(e)$ is straightforward and, once again, relies on the plug-in principle: We need to replace $ATT(g, g + e)$ with its sample analogs (which we already computed in Figure 6), and use the relative adult population share of expansion group g among eventually treated units as estimates of the event study weights. Figure 8 reports our population-weighted estimates of the event-study aggregation from event times -5 to 5 , which respectively correspond to five years before the Medicaid expansion and five years after the Medicaid expansion. We also report point-wise and simultaneous confidence intervals at the 95 percent significance level. Overall, the results suggest that Medicaid expansion has no effect on adult mortality rates among counties that eventually experience a Medicaid expansion. The pre-trends are also fairly close to zero, suggesting that our parallel trends assumption may be reasonable.

We conclude this section by stressing that the way we have constructed the event study parameters in Figure 8 uses all available information from Figure 6. A potential drawback of this strategy is that we do not always use the same set of groups across all event times. Practitioners usually refer to this as “imbalance in event time.” For instance, the 2019 expansion group contributes only to event time $e = 0$, not $e = 1$ or later event times. When compositional changes are a concern, one can impose balance in event time and estimate a balanced event study aggregation:

$$(32) \quad \begin{aligned} ATT_{es, bal, [\underline{e}, \bar{e}]}(e) &= \mathbb{E}_{\omega} [ATT(G, G + e) | G + \bar{e} \in [1, T], G + \underline{e} \in [1, T], G \leq T] \\ &= \sum_{g \in \mathcal{G}_{treat}} w_{\omega, g, [\underline{e}, \bar{e}]}^{es, bal} ATT(g, g + e), \end{aligned}$$

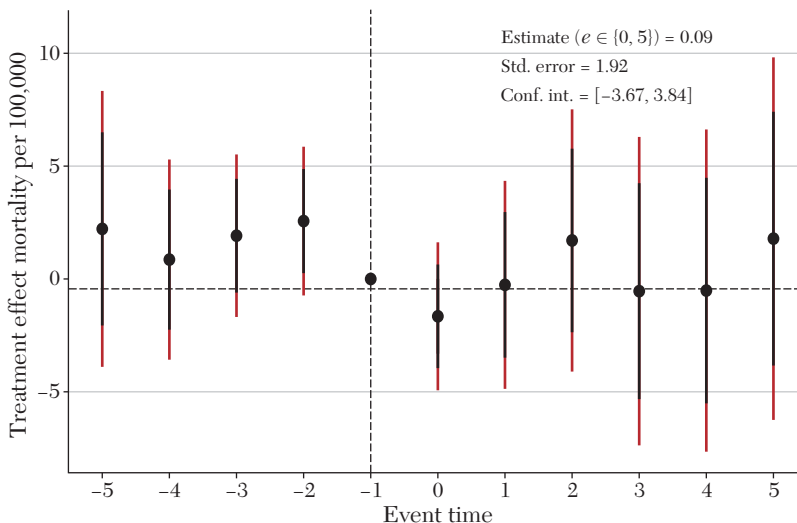


Figure 8. $G \times T$ Event Study without Covariates

Notes: This figure shows the event-study estimates with staggered treatment timing using the doubly-robust estimation method from Callaway and Sant’Anna (2021), using 1,222 not-yet-treated units as the comparison group. The outcome variable is the crude mortality rate for adults ages 20–64. The point estimate is reported by the circles and both 95 percent point-wise (black) and simultaneous (red) confidence intervals are reported with the vertical lines. We also report the simple average of all nonnegative event times as a summary of the overall ATT (together with their standard errors and 95 percent confidence interval). All results use population weights.

where balanced event-time weights $w_{g, [\underline{e}, \bar{e}]}^{es, bal}$ are given by

$$w_{g, [\underline{e}, \bar{e}]}^{es, bal} = \mathbf{1}\{g + \bar{e} \leq T\} \mathbf{1}\{g + \underline{e} \geq 1\} P_{\omega}(G = g | G + \bar{e} \in [1, T], G + \underline{e} \in [1, T], G \leq T).$$

Although intimidating, $w_{g, [\underline{e}, \bar{e}]}^{es, bal}$ just measures the relative size of a particular treatment group that was kept in the balanced data. We can interpret $ATT_{es, bal, [\underline{e}, \bar{e}]}(e)$ as the average group-time average treatment effect among units whose event time is equal to e and is observed to participate in the treatment for at least \bar{e} periods, and have at least \underline{e} available pretreatment periods (if \underline{e} is negative).³¹

³¹This discussion assumes that there are no “holes” in event times for each treatment group. It is straightforward to adjust the interpretation to those more complicated cases, as the same logic can be applied.

5.2.5 Estimators for Staggered Designs with Covariates

As the discussions in Section 5.2.2 made it clear, we can view the staggered DiD setups as a collection of simpler 2×2 DiD building blocks. A benefit of this interpretation is that, if parallel trends holds only after conditioning on the covariates that determine untreated potential outcome changes, we can easily leverage all the results discussed in Section 4 to identify, estimate, and make inference about the $ATT(g, t)$ s, using regression adjustment, inverse probability weighting, or doubly robust methods.

Of course, to proceed in this manner, we would need to adopt conditioned on covariates versions of Assumption PT-GT-NEV, PT-GT-NYT, or PT-GT-ALL as well as impose an overlap condition. Given that all these are fairly similar to each other, here we state only an extension of Assumption PT-GT-NYT and a strong overlap condition that can be used for all three cases.

ASSUMPTION CPT-GT-NYT (Conditional Parallel Trends Based on Not-Yet-Treated Groups): For every eventually treated group g , not-yet-treated group g' , time periods t such that $t \geq g$ and $g' > t$, and every covariate value X_i ,

$$\mathbb{E}_\omega[Y_{i,t}(\infty) - Y_{i,t-1}(\infty) | G_i = g, X_i] = \mathbb{E}_\omega[Y_{i,t}(\infty) - Y_{i,t-1}(\infty) | G_i = g', X_i].$$

ASSUMPTION SO-GT (Strong Overlap with Staggered Adoption): For every group $g \in \mathcal{G}$, the conditional (weighted) probability of belonging to a treatment group g , given observed covariates X_i that are determinants of untreated potential outcome growth, is uniformly bounded away from zero and one. That is, for some $\epsilon > 0$ and for every group $g \in \mathcal{G}$, $\epsilon < P_\omega[G_i = g | X_i] < 1 - \epsilon$.

Using these identifying assumptions and building on the results in Sections 4 and 5.2.2, we can follow the arguments in Callaway and Sant'Anna (2021) and establish that the posttreatment $ATT(g, t)$ s are identified by the regression-adjusted, inverse-probability-weighted, and doubly-robust estimands given by³²

$$\begin{aligned} ATT_{ra}(g, t) &= \mathbb{E}_\omega[Y_{i,t} - Y_{i,t=g-1} | G_i = g] - \mathbb{E}_\omega[\mathbb{E}_\omega[Y_{i,t} - Y_{i,t=g-1} | X_i, G_i > t] | G_i = g], \\ ATT_{ipw}(g, t) &= \mathbb{E}\left[\left(w_{\omega, G=g}(G_i) - w_{\omega, g, t}(G_i, X_i)\right)(Y_{i,t} - Y_{i,t=g-1})\right], \\ ATT_{dr}(g, t) &= \mathbb{E}\left[\left(w_{\omega, G=g}(G_i) - w_{\omega, g, t}(G_i, X_i)\right)(Y_{i,t} - Y_{i,t=g-1} \right. \\ &\quad \left. - \mathbb{E}_\omega[Y_{i,t} - Y_{i,t=g-1} | X_i, G_i > t])\right], \end{aligned}$$

³²Wooldridge (2025) proposes alternative estimators for the $ATT(g, t)$ that incorporate covariates and their interactions with group and time dummies into (30). Although Borusyak, Jaravel, and Spiess (2024) and de Chaisemartin and D'Haultfœuille (2020) also allow for covariates in their estimation procedure, they affect the outcome levels and not their changes over time. Thus, to some extent, these procedures do not build on conditional parallel trends like Assumption CPT-GT-NYT.




Figure 9. $G \times T$ Event Study with Covariates

Notes: This figure shows the event-study estimates with staggered treatment timing, using the doubly-robust estimation method from Callaway and Sant’Anna (2021). The outcome variable is the crude mortality rate for adults ages 20–64 and the covariates include the percentage of the county population that is female, the percentage of the county population that is White, the percentage of the county population that is Hispanic, the unemployment rate, the poverty rate, and the county-level median income. The point estimate is reported by the circles, and both 95 percent point-wise (black) and simultaneous (red) confidence intervals are reported with the vertical lines. We also report the simple average of all nonnegative event times as a summary of the overall ATT (together with their standard errors and 95 percent confidence interval). All results use population weights.

where $w_{\omega,G=g}(G_i)$ and $w_{\omega,g,G>t}(G_i, X_i)$ are the analogs of the weights in (14) and are defined as

$$w_{\omega,G=g}(G) = \frac{\omega \mathbf{1}\{G = g\}}{\mathbb{E}[\omega \mathbf{1}\{G = g\}]},$$

$$w_{\omega,g,t}(G, X) = \frac{\omega \mathbf{1}\{G > t\} \mathbf{1}\{G \neq g\} p_{\omega,g,t}(X)}{1 - p_{\omega,g,t}(X)} \bigg/ \frac{\mathbb{E}[\omega \mathbf{1}\{G > t\} \mathbf{1}\{G \neq g\} p_{\omega,g,t}(X)]}{1 - p_{\omega,g,t}(X)},$$

and $p_{\omega,g,t}(X) = \mathbb{E}_{\omega}[\mathbf{1}\{G_i = g\} | X, \mathbf{1}\{G_i = g\} + \mathbf{1}\{G_i > t\} = 1]$ denote the (weighted) probability of belonging to the group, g , given covariates, X , and that the unit belongs to either to group g —the treated group for the $ATT(g, t)$ of interest—or the not-yet-treated group $G_i > t$ —the comparison group.

Estimating the $ATT(g, t)$ s follows exactly as in Section 4, and event study aggregations follow from the arguments in Section 5.2.4. Figure 9 reports event study summary estimates incorporating covariates into the Medicaid analysis. Like Figure 8, it suggests that Medicaid expansion had no effect on adult mortality among counties that expanded Medicaid by 2019. Given the point estimates and uniform confidence interval, we can be reasonably confident

that the treatment effects are not greater than 6 or less than -11 deaths per 100,000 adults for the 6-year period following Medicaid expansion.

5.3 Limitations of TWFE Regressions

Our framework emphasizes building an estimator from 2×2 components, each of which targets a well-defined *ATT* parameter under a specific parallel trends assumption. The most common estimator for staggered designs, a TWFE regression, comes instead from extending convenient estimation tools that work well in the 2×2 case. A TWFE specification that estimates a summary treatment effect parameter is:

$$(33) \quad Y_{i,t} = \theta_t + \eta_i + \beta^{twfe} D_{i,t} + e_{i,t}.$$

In this section, we abstract from weights.

A major breakthrough in recent DiD research has been to demonstrate two potentially large problems with β^{twfe} (de Chaisemartin and D’Haultfœuille 2020; Goodman-Bacon 2021; Sun and Abraham 2021; Borusyak, Jaravel, and Spiess 2024). The primary issue comes from the fact that TWFE implicitly uses already-treated comparison groups. Even if PT holds for all groups and all periods, the resulting estimand can actually put negative weight on certain $ATT(g,t)$ parameters. The only way TWFE avoids the problem is if treatment effects do not change over time, a strong additional assumption.

To isolate this issue, consider a setting with two time periods and three groups: a group that enters treatment in the first period ($G_i = 1$), a group that becomes treated in the second time period ($G_i = 2$), and a never-treated group ($G_i = \infty$). This is a staggered design because group 1 and group 2 are treated at different times, but because there are only two time periods, we can rewrite the TWFE specification as

$$\Delta Y_{i,2} = \Delta \theta_2 + \beta^{twfe} \Delta D_{i,2} + \Delta e_{i,2}.$$

Because $\Delta D_{i,2}$ takes only two values—1 for units whose treatment status increases ($G_i = 2$) and 0 for units whose treatment status does not change ($G_i = 1$ and $G_i = \infty$)—the TWFE estimand is the following simple comparison of means:

$$(34) \quad \begin{aligned} \beta^{twfe} &= \mathbb{E}[\Delta Y_{i,2} | \Delta D_{i,2} = 1] - \mathbb{E}[\Delta Y_{i,2} | \Delta D_{i,2} = 0] \\ &= (\mathbb{E}[\Delta Y_{i,2} | G_i = 2] - \mathbb{E}[\Delta Y_{i,2} | G_i = \infty]) (1 - w_1) \\ &\quad + (\mathbb{E}[\Delta Y_{i,2} | G_i = 2] - \mathbb{E}[\Delta Y_{i,2} | G_i = 1]) w_1, \end{aligned}$$

where $w_1 = \frac{p_1}{p_1 + p_\infty}$ and $p_g = P(G = g)$ is group g ’s share of units. The TWFE coefficient, in this case, is a weighted average of two DiD terms that use the already-treated units or the never-treated units as comparisons. We have already discussed how, under PT between group $G_i = 2$ and never-treated units, the first term in the $\hat{\beta}^{twfe}$ decomposition, $\mathbb{E}[\Delta Y_{i,2} | G_i = 2] - \mathbb{E}[\Delta Y_{i,2} | G_i = \infty]$, equals $ATT(2,2)$. But what about the second term with the already-treated comparison group? It turns out that under parallel trends and no anticipation,

if we add and subtract different terms, this type of estimand generally equals a combination of treatment effects for both groups:

$$\begin{aligned}
\mathbb{E}[\Delta Y_{i,2} | G_i = 2] - \mathbb{E}[\Delta Y_{i,2} | G_i = 1] &= \mathbb{E}[Y_{i,2} | G_i = 2] - \mathbb{E}[Y_{i,1} | G_i = 2] \\
&\quad - (\mathbb{E}[Y_{i,2} | G_i = 1] - \mathbb{E}[Y_{i,1} | G_i = 1]) \\
&= \mathbb{E}[Y_{i,2}(2) - Y_{i,2}(\infty) | G_i = 2] \\
&\quad + \mathbb{E}[Y_{i,2}(\infty) - Y_{i,1}(\infty) | G_i = 2] \\
&\quad - (\mathbb{E}[Y_{i,2}(1) - Y_{i,2}(\infty) | G_i = 1] \\
&\quad\quad - \mathbb{E}[Y_{i,1}(1) - Y_{i,1}(\infty) | G_i = 1]) \\
&\quad - \mathbb{E}[Y_{i,2}(\infty) - Y_{i,1}(\infty) | G_i = 1] \\
&= ATT(2,2) - (ATT(1,2) - ATT(1,1)) \\
&\quad + \mathbb{E}[Y_{i,2}(\infty) - Y_{i,1}(\infty) | G_i = 2] \\
&\quad - \mathbb{E}[Y_{i,2}(\infty) - Y_{i,1}(\infty) | G_i = 1] \\
&= ATT(2,2) - (ATT(1,2) - ATT(1,1)),
\end{aligned}$$

where the first equality follows from the linearity of expectations. In the second, we add and subtract averages of untreated potential outcomes and explore no anticipation, and the observation rules that $Y_{i,t}(g)$ is observed for units with $G_i = g$. In the third equality we rearrange terms, and the last equality follows from a parallel trends for units in group $G = 1$ and $G = 2$.

The TWFE estimand in this case equals:

$$(35) \quad \beta^{twfe} = ATT(2,2) + ATT(1,1)w_1 - ATT(1,2)w_1.$$

The $ATT(1,2)$ building block receives *negative* weight in the overall TWFE estimand unless there are no treatment effect dynamics ($ATT(1,1) = ATT(1,2)$). In this example, the problem is easy to fix: drop the always-treated units and target $ATT(2,2)$. With multiple periods, however, the problem is more complex.

There are two primary results regarding TWFE for general staggered timing designs. One is a decomposition of the TWFE *estimator*. Goodman-Bacon (2021) expresses the TWFE estimator as a weighted average of all possible 2×2 DiD comparisons between pairs of groups and time periods during which one group enters treatment and the other does not. The terms in his decomposition are not two-period $ATT(g,t)$ -type estimators; they aggregate over the relevant pre- and post-periods just as β^{OLS} from (22) does. They include many comparisons to already-treated units like (34). This mechanical decomposition always has strictly positive weights that are larger for larger groups and for groups treated closer to the middle of the panel, which have a larger variance of $D_{i,t}$ conditional on the fixed effects. This result shows

how TWFE estimators actually function and creates a clear link to the estimation approaches we outlined above for summary parameters. Both are averages of 2×2 DiD terms, but they differ in which comparisons they use and how they aggregate.

The second result is a decomposition of the TWFE *estimand*, which clarifies why the TWFE estimator is not guaranteed to identify a desirable parameter in staggered designs, even if parallel trends hold. The reason is their use of already-treated comparisons like the ones in (35) (Goodman-Bacon 2021; de Chaisemartin and D’Haultfœuille 2020; Borusyak, Jaravel, and Spiess 2024; Imai and Kim 2021; Strezhnev 2018; Sun and Abraham 2021). This will tend to bias β^{twfe} away from the sign of ATT_{avg} , and β^{twfe} can even have the opposite sign of ATT_{avg} (Baker, Larcker, and Wang 2022). It may appear that a more flexible regression specification could solve this problem, but Sun and Abraham (2021) show that a TWFE event study specification suffers from a similar bias when the dynamics of the $ATT(g, t)$ s differ across cohorts. Moreover, the “variance-weighting” feature of OLS means that β^{twfe} has nonintuitive weights even when $ATT(g, t) = ATT(g)$.

While TWFE remains common, it has well-understood, potentially serious, and easily remedied problems, and we do not recommend using it. In many cases, especially those with many untreated units or minimal treatment effect dynamics, TWFE estimates may be similar to those derived from the theoretically grounded estimators discussed above. The only way to be sure, however, is to estimate both. In that case, it is unclear why a researcher would not report the estimates from a procedure motivated by a desirable target parameter and a credible PT assumption.

6. Conclusion

The starting point of this paper was a 2×2 DiD design that researchers have been using for almost 200 years. The end point was a design with five treatment timing groups, 11 years of data, six covariates, three types of parallel trends assumptions, and four estimation techniques. Our fundamental message is that without understanding how complex designs are built up from simpler ones, it is exceedingly difficult to navigate all the empirical tools now available for DiD designs. This lesson applies not only to the design details we considered here—weighting, covariates, and staggered designs—but to any DiD design.

The forward-engineering philosophy we followed in this paper suggests a set of steps that researchers can follow in any DiD study:

- Step 1: *Define target parameters.* Adopt a potential outcomes notation that fits the study’s specific setting and use it to define causal target parameters that answer the study’s motivating question. Building-block parameters usually aggregate across units using (conditional) weighted averages. Summary target parameters then aggregate across the building blocks. This step fixes the study’s goals in terms of causal quantities and facilitates comparisons with related studies.
- Step 2: *State (formally) the identification assumptions.* DiD studies leverage parallel trends assumptions, but they also rely on no anticipation and, in some cases, overlap conditions, or more. Be explicit about which form of these assumptions is required for identification in the study. Engage with the theoretical arguments necessary for them to hold and generate appropriate empirical evidence, such as pretreatment differential trends, that can falsify or (indirectly) support their plausibility.

- Step 3: *Determine the appropriate estimation method.* In some DiD designs, estimation is as simple as replacing population expectations with sample means. In others, such as conditional DiD designs, estimation involves choosing econometric techniques (e.g., a regression adjustment, inverse probability weighting, or doubly robust procedure) to map theoretical quantities to estimable sample quantities. Each of these strategies relies on additional modeling restrictions that should be stated clearly.
- Step 4: *Discuss sources of uncertainty.* Statistical inference procedures for DiD designs stem from basic assumptions about where randomness comes from in a given design. Some researchers may adopt a sampling approach to inference, whereas others may be more comfortable with a design-based perspective. It is important to discuss what variables of the model are being treated as fixed and what variables are considered random, as well as to use inference techniques that are compatible with the model structure and assumptions.
- Step 5: *Estimate.* Steps 1–4 provide a specific structure for using data to estimate the causal parameters of interest.
- Step 6: *Conduct sensitivity analysis.* A clear statement of the identification and estimation assumptions also facilitates a clear statement of what violations of those assumptions might mean. No study is robust to all the ways its assumptions may fail, but a good study should be robust against likely violations of plausible magnitudes. Combine context-specific knowledge about how the assumptions from step 2 might be violated, and by how much, with the structure of the estimator from step 3 to evaluate how much the DiD estimates vary if the key identification assumptions are not exactly true.
- Step 7: *Conduct heterogeneity analysis.* Sometimes aggregate parameters defined in step 1 mask important heterogeneity, in which case the forward-engineering approach simply suggests targeting subgroup parameters as well. This can include variation in parameters over time, between groups of units with different characteristics, or across different sources of treatment variation. Be clear about which types of heterogeneous effects are relevant and how they are identified and estimated.
- Step 8: *Keep learning.* DiD is not the only or the best research design in all settings; it is just one of many causal inference techniques. If the assumptions required for a DiD analysis appear implausible *ex ante* or are refuted by evidence or non-robustness in practice, then explore different designs. If existing DiD methods do not provide enough guidance, then use a forward-engineering approach to deduce what advances would help.

Some researchers may still prefer to use standard regression tools to conduct DiD studies. The properties and pitfalls of some popular regression specifications are now well understood and one can easily explain how this choice fits with (and perhaps satisfies) the steps above. But using simple regressions in any DiD-type setting is an implicit choice to reverse engineer a research design from the statistical method, rather than forward engineer a reliable estimator

from a substantive question and transparent assumptions. Ultimately, important questions and credible identification strategies should guide DiD analyses (regression based or otherwise), not the other way around.

Although this paper is by no means an exhaustive guide to DiD practice, the eight steps above are a rigorous framework for tackling all the DiD topics that we did not cover. The Supplemental Appendix briefly discusses DiD methods for (i) treatments that turn on and off over time, (ii) continuous and multi-valued treatments, (iii) triple differences, (iv) distributional parameters, and (v) repeated cross-section or unbalanced panel data. While each of these designs differs from what we covered in the main text, a forward-engineering approach that moves from defining parameters and assumptions, to settling on estimation and inference techniques, to probing robustness applies equally to all of them. While the specifics of any given DiD analysis may change across research questions, treatment variables, econometric techniques, and data structures, the principles by which one can conduct reliable and transparent causal inference stay the same.³³

³³There are other DiD topics of interest that we do not cover, including fuzzy DiD and instrumented DiD designs (de Chaisemartin and D'Haultfœuille 2018; Miyaji 2025), nonlinear DiD models (Wooldridge 2023; Tchetgen Tchetgen, Park, and Richardson 2024), issues related to few clusters (Roth et al. 2023, section 5), and situations with multiple treatments (de Chaisemartin and D'Haultfœuille 2024; Yanagi 2023). We also do not cover some methods that address violations of parallel trends (Freyaldenhoven, Hansen, and Shapiro 2019; Arkhangelsky et al. 2021; Callaway and Karami 2023; Imbens, Kallus, and Mao 2021), nor do we examine setups that impose as-good-as-random treatment timing (Athey and Imbens 2022; Roth and Sant'Anna 2023a; Arkhangelsky et al. 2024).

APPENDIX

TABLE A1
LIST OF ACRONYMS

Acronym	Definition
2×2	Two-group two-time-periods DiD
$2 \times T$	Two-group T -time-periods DiD
ACA	Affordable Care Act
ATT	Average treatment effect on the treated
CPT	Conditional parallel trends
CPT-GT-NYT	Conditional parallel trends based on not-yet-treated groups
DiD	Difference-in-differences
DR	Doubly robust
ETWFE	Extended two-way fixed effects
IPW	Inverse probability weighted
NA	No anticipation
NA-S	No anticipation with staggered treatment timing
OLS	Ordinary least squares
PT	Parallel trends
PT-ES	Parallel trends event study
PT-GT-ALL	Parallel trends for every period and group
PT-GT-NEV	Parallel trends based on never-treated groups
PT-GT-NYT	Parallel trends based on not-yet-treated groups
RA	Regression adjustment
SO	Strong overlap
SO-GT	Strong overlap with staggered adoption
TWFE	Two-way fixed effects

REFERENCES

- Abadie, Alberto. 2005. "Semiparametric Difference-in-Difference Estimators." *Review of Economic Studies* 72 (1): 1–19.
- Abadie, Alberto, Susan Athey, Guido W. Imbens, and Jeffrey M. Wooldridge. 2020. "Sampling-Based versus Design-Based Uncertainty in Regression Analysis." *Econometrica* 88 (1): 265–96.
- Abadie, Alberto, Susan Athey, Guido W. Imbens, and Jeffrey M. Wooldridge. 2023. "When Should You Adjust Standard Errors for Clustering?" *Quarterly Journal of Economics* 138 (1): 1–35.
- Abbring, Jaap H., and Gerard J. van den Berg. 2003. "The Nonparametric Identification of Treatment Effects in Duration Models." *Econometrica* 71 (5): 1491–517.
- Angrist, Joshua D. 1998. "Estimating the Labor Market Impact of Voluntary Military Service Using Social Security Data on Military Applicants." *Econometrica* 66 (2): 249–88.
- Arkhangelsky, Dmitry, Susan Athey, David A. Hirshberg, Guido W. Imbens, and Stefan Wager. 2021. "Synthetic Difference-in-Differences." *American Economic Review* 111 (12): 4088–118.
- Arkhangelsky, Dmitry, Guido W. Imbens, Lihua Lei, and Xiaoman Luo. 2024. "Design-Robust Two-Way-Fixed-Effects Regression for Panel Data." *Quantitative Economics* 15 (4): 999–1034.
- Aronow, Peter M., and Cyrus Samii. 2016. "Does Regression Produce Representative Estimates of Causal Effects?" *American Journal of Political Science* 60 (1): 250–67.
- Ashenfelter, Orley, and David Card. 1985. "Using the Longitudinal Structure of Earnings to Estimate the Effect of Training Programs." *Review of Economics and Statistics* 67 (4): 648–60.

- Athey, Susan, and Guido W. Imbens. 2006. "Identification and Inference in Nonlinear Difference-in-Differences Models." *Econometrica* 74 (2): 431–97.
- Athey, Susan, and Guido W. Imbens. 2022. "Design-Based Analysis in Difference-in-Differences Settings with Staggered Adoption." *Journal of Econometrics* 226 (1): 62–79.
- Austin, Peter C. 2009. "Balance Diagnostics for Comparing the Distribution of Baseline Covariates between Treatment Groups in Propensity-Score Matched Samples." *Statistics in Medicine* 28 (25): 3083–107.
- Baker, Andrew, Brantly Callaway, Scott Cunningham, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna. 2026. *Data and Code for "Difference-in-Differences Designs: A Practitioner's Guide."* American Economic Association; distributed by Inter-university Consortium for Political and Social Research. <https://doi.org/10.3886/E239028V1>.
- Baker, Andrew C., David F. Larcker, and Charles C. Y. Wang. 2022. "How Much Should We Trust Staggered Difference-in-Differences Estimates?" *Journal of Financial Economics* 144 (2): 370–95.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan. 2004. "How Much Should We Trust Differences-in-Differences Estimates?" *Quarterly Journal of Economics* 119 (1): 249–75.
- Bilinski, Alyssa, and Laura A. Hatfield. 2026. "Nothing to See Here? Non-inferiority Approaches to Parallel Trends and Other Model Assumptions." Preprint, arXiv. <https://doi.org/10.48550/arXiv.1805.03273>.
- Black, Bernard, Alex Hollingsworth, Leticia Nunes, and Kosali Simon. 2022. "Simulated Power Analyses for Observational Studies: An Application to the Affordable Care Act Medicaid Expansion." *Journal of Public Economics* 213: 104713.
- Borgschulte, Mark, and Jacob Vogler. 2020. "Did the ACA Medicaid Expansion Save Lives?" *Journal of Health Economics* 72: 102333.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2024. "Revisiting Event Study Designs: Robust and Efficient Estimation." *Review of Economic Studies* 91 (6): 3253–85.
- Bureau of Labor Statistics. 2009–2019. *Local Area Unemployment Statistics*. United States Department of Labor. <https://www.bls.gov/lau/tables.htm>. (accessed December 5, 2025).
- Caetano, Carolina and Brantly Callaway. 2024. "Difference-in-Differences When Parallel Trends Holds Conditional on Covariates." Preprint, arXiv. <https://doi.org/10.48550/arXiv.2406.15288>.
- Caetano, Carolina, Brantly Callaway, Stroud Payne, and Hugo Sant'Anna Rodrigues. 2024. "Difference in Differences with Time-Varying Covariates." Preprint, arXiv. <https://doi.org/10.48550/arXiv.2202.02903>.
- Callaway, Brantly. 2023. "Difference-in-Differences for Policy Evaluation." In *Handbook of Labor, Human Resources and Population Economics*, edited by Klaus F. Zimmermann. Springer.
- Callaway, Brantly, and Sonia Karami. 2023. "Treatment Effects in Interactive Fixed Effects Models with a Small Number of Time Periods." *Journal of Econometrics* 233 (1): 184–208.
- Callaway, Brantly, and Tong Li. 2019. "Quantile Treatment Effects in Difference in Differences Models with Panel Data." *Quantitative Economics* 10 (4): 1579–618.
- Callaway, Brantly, and Pedro H. C. Sant'Anna. 2021. "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics* 225 (2): 200–230.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller. 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *Review of Economics and Statistics* 90 (3): 414–27.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. "The Effect of Minimum Wages on Low-Wage Jobs." *Quarterly Journal of Economics* 134 (3): 1405–54.
- Centers for Disease Control and Prevention. 2009–2019. *National Vital Statistics System—Mortality Data: Underlying Cause of Death Database*. Department of Health and Human Services. <https://wonder.cdc.gov/ucd-icd10.html> (accessed December 5, 2025).
- Chabé-Ferret, Sylvain. 2015. "Analysis of the Bias of Matching and Difference-in-Difference under Alternative Earnings and Selection Processes." *Journal of Econometrics* 185 (1): 110–23.
- Chang, Neng-Chieh. 2020. "Double/Debiased Machine Learning for Difference-in-Differences." *Econometrics Journal* 23 (2): 177–91.
- Chen, Xiaohong, Pedro H. C. Sant'Anna, and Haitian Xie. 2025. "Efficient Difference-in-Differences and Event Study Estimators." Preprint, arXiv. <https://doi.org/10.48550/arXiv.2506.17729>.
- Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Iván Fernández-Val. 2023. "Fisher-Schultz Lecture: Generic Machine Learning Inference on Heterogenous Treatment Effects in Randomized Experiments, with an Application to Immunization in India." Preprint, arXiv. <https://doi.org/10.48550/arXiv.1712.04802>.
- Chernozhukov, Victor, Iván Fernández-Val, and Ye Luo. 2018. "The Sorted Effects Method: Discovering Heterogeneous Effects Beyond Their Averages." *Econometrica* 86 (6): 1911–38.
- Conley, Timothy G., and Christopher R. Taber. 2011. "Inference with 'Difference in Differences' with a Small Num-

- ber of Policy Changes." *Review of Economics and Statistics* 93 (1): 113–25.
- Crump, Richard K., V. Joseph Hotz, Guido W. Imbens, and Oscar A. Mitnik. 2009. "Dealing with Limited Overlap in Estimation of Average Treatment Effects." *Biometrika* 96 (1): 187–99.
- Currie, Janet, Henrik Kleven, and Esmée Zviers. 2020. "Technology and Big Data Are Changing Economics: Mining Text to Track Methods." *AEA Papers and Proceedings* 110: 42–48.
- Currie, Janet, and Hannes Schwandt. 2016. "Mortality Inequality: The Good News from a County-Level Approach." *Journal of Economic Perspectives* 30 (2): 29–52.
- de Chaisemartin, Clément, and Xavier D'Haultfœuille. 2018. "Fuzzy Differences-in-Differences." *Review of Economic Studies* 85 (2): 999–1028.
- de Chaisemartin, Clément, and Xavier D'Haultfœuille. 2020. "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects." *American Economic Review* 110 (9): 2964–96.
- de Chaisemartin, Clément, and Xavier D'Haultfœuille. 2023. "Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey." *Econometrics Journal* 26 (3): C1–30.
- de Chaisemartin, Clément, and Xavier D'Haultfœuille. 2024. "Difference-in-Differences Estimators of Intertemporal Treatment Effects." *Review of Economics and Statistics*. https://doi.org/10.1162/rest_a_01414.
- Deshpande, Manasi, and Yue Li. 2019. "Who Is Screened Out? Application Costs and the Targeting of Disability Programs." *American Economic Journal: Economic Policy* 11 (4): 213–48.
- Dette, Holger, and Martin Schumann. 2024. "Testing for Equivalence of Pre-trends in Difference-in-Differences Estimation." *Journal of Business and Economic Statistics* 42 (4): 1289–1301.
- DiNardo, John, and David S. Lee. 2011. "Program Evaluation and Research Designs." In *Handbook of Labor Economics*, Vol. 4A, edited by Orley Ashenfelter and David Card, 463–536. North-Holland.
- Donald, Stephen G., and Kevin Lang. 2007. "Inference with Difference-in-Differences and Other Panel Data." *Review of Economics and Statistics* 89 (2): 221–33.
- Dube, Arindrajit, Daniele Girardi, Óscar Jordà, and Alan M. Taylor. 2025. "A Local Projections Approach to Difference-in-Differences." *Journal of Applied Econometrics* 40 (7): 741–58.
- Fadlon, Itzik, and Torben Heien Nielsen. 2021. "Family Labor Supply Responses to Severe Health Shocks: Evidence from Danish Administrative Records." *American Economic Journal: Applied Economics* 13 (3): 1–30.
- Fernández-Val, Iván, Jonas Meier, Aico van Vuuren, and Francis Vella. 2024. "Distribution Regression Difference-in-Differences." Preprint, arXiv. <https://doi.org/10.48550/arXiv.2409.02311>.
- Freyaldenhoven, Simon, Christian Hansen, and Jesse M. Shapiro. 2019. "Pre-event Trends in the Panel Event-Study Design." *American Economic Review* 109 (9): 3307–38.
- Freyaldenhoven, Simon, Christian Hansen, Jorge Pérez Pérez, and Jesse M. Shapiro. 2024. "Visualization, Identification, and Estimation in the Linear Panel Event-Study Design." In *Advances in Economics and Econometrics: Twelfth World Congress*, edited by Victor Chernozhukov, Johannes Hörner, Eliana La Ferrara, and Iván Werning, 225–68. Cambridge University Press.
- Gardner, John. 2022. "Two-Stage Differences in Differences." Preprint, arXiv. <https://doi.org/10.48550/arXiv.2207.05943>.
- Ghanem, Dalia, Pedro H. C. Sant'Anna, and Kaspar Wüthrich. 2026. "Selection and Parallel Trends." Preprint, arXiv. <https://doi.org/10.48550/arXiv.2203.09001>.
- Goldsmith-Pinkham, Paul, Peter Hull, and Michal Kolesár. 2024. "Contamination Bias in Linear Regressions." *American Economic Review* 114 (12): 4015–51.
- Goodman-Bacon, Andrew. 2021. "Difference-in-Differences with Variation in Treatment Timing." *Journal of Econometrics* 225 (2): 254–77.
- Graham, Bryan S., Cristine Campos De Xavier Pinto, and Daniel Egel. 2012. "Inverse Probability Tilting for Moment Condition Models with Missing Data." *Review of Economic Studies* 79 (3): 1053–79.
- Harmon, Nikolaj A. 2024. "Difference-in-Differences and Efficient Estimation of Treatment Effects." Unpublished.
- Heckman, James J. 2000. "Causal Parameters and Policy Analysis in Economics: A Twentieth Century Retrospective." *Quarterly Journal of Economics* 115 (1): 45–97.
- Heckman, James J., Hidehiko Ichimura, and Petra E. Todd. 1997. "Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme." *Review of Economic Studies* 64 (4): 605–54.
- Heckman, James J., and Richard Robb Jr. 1985. "Alternative Methods for Evaluating the Impact of Interventions: An Overview." In *Longitudinal Analysis of Labor Market Data*, Vol. 30, edited by James J. Heckman and Burton Singer, 156–246. Cambridge University Press.
- Ho, Daniel E., Kosuke Imai, Gary King, and Elizabeth A. Stuart. 2007. "Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference." *Political Analysis* 15 (3): 199–236.

- Imai, Kosuke, and In Song Kim. 2021. "On the Use of Two-Way Fixed Effects Regression Models for Causal Inference with Panel Data." *Political Analysis* 29 (3): 405–15.
- Imbens, Guido, Nathan Kallus, and Xiaojie Mao. 2021. "Controlling for Unmeasured Confounding in Panel Data Using Minimal Bridge Functions: From Two-Way Fixed Effects to Factor Models." Preprint, arXiv. <https://doi.org/10.48550/arXiv.2108.03849>.
- Imbens, Guido W., and Donald B. Rubin. 2015. *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge University Press.
- Kahn-Lang, Ariella, and Kevin Lang. 2020. "The Promise and Pitfalls of Differences-in-Differences: Reflections on 16 and Pregnant and Other Applications." *Journal of Business and Economic Statistics* 38 (3): 613–20.
- Kaiser Family Foundation. 2026. "Status of State Medicaid Expansion Decisions." <https://www.kff.org/status-of-state-medicaid-expansion-decisions>. (accessed March 11, 2025).
- Kennedy, Edward H., Zongming Ma, Matthew D. McHugh, and Dylan S. Small. 2017. "Non-parametric Methods for Doubly Robust Estimation of Continuous Treatment Effects." *Journal of the Royal Statistical Society Series B: Statistical Methodology* 79 (4): 1229–45.
- Khan, Shakeeb, and Elie Tamer. 2010. "Irregular Identification, Support Conditions, and Inverse Weight Estimation." *Econometrica* 78 (6): 2021–42.
- Lechner, Michael. 2011. "The Estimation of Causal Effects by Difference-in-Difference Methods." *Foundations and Trends in Econometrics* 4 (3): 165–224.
- Lee, Soo Jeong, and Jeffrey M. Wooldridge. 2023. "A Simple Transformation Approach to Difference-in-Differences Estimation for Panel Data." Preprint, SSRN. <http://dx.doi.org/10.2139/ssrn.4516518>.
- Liu, Licheng, Ye Wang, and Yiqing Xu. 2024. "A Practical Guide to Counterfactual Estimators for Causal Inference with Time-Series Cross-Sectional Data." *American Journal of Political Science* 68 (1): 160–76.
- Ma, Yukun, Pedro H. C. Sant'Anna, Yuya Sasaki, and Takuya Ura. 2023. "Doubly Robust Estimators with Weak Overlap." Preprint, arXiv. <https://doi.org/10.48550/arXiv.2304.08974>.
- Malani, Anup, and Julian Reif. 2015. "Interpreting Pre-trends as Anticipation: Impact on Estimated Treatment Effects from Tort Reform." *Journal of Public Economics* 124: 1–17.
- Manski, Charles F., and John V. Pepper. 2018. "How Do Right-to-Carry Laws Affect Crime Rates? Coping with Ambiguity Using Bounded-Variation Assumptions." *Review of Economics and Statistics* 100 (2): 232–44.
- Marcus, Michelle, and Pedro H. C. Sant'Anna. 2021. "The Role of Parallel Trends in Event Study Settings: An Application to Environmental Economics." *Journal of the Association of Environmental and Resource Economists* 8 (2): 235–75.
- Marx, Philip, Elie Tamer, and Xun Tang. 2024. "Parallel Trends and Dynamic Choices." *Journal of Political Economy Microeconomics* 2 (1): 129–71.
- Miller, Sarah, Norman Johnson, and Laura R. Wherry. 2021. "Medicaid and Mortality: New Evidence from Linked Survey and Administrative Data." *Quarterly Journal of Economics* 136 (3): 1783–829.
- Miyaji, Sho. 2025. "Instrumented Difference-in-Differences with Heterogeneous Treatment Effects." Preprint, arXiv. <https://doi.org/10.48550/arXiv.2405.12083>.
- Mogstad, Magne, and Alexander Torgovitsky. 2024. "Instrumental Variables with Unobserved Heterogeneity in Treatment Effects." In *Handbook of Labor Economics*, Vol. 5, edited by Christian Dustmann and Thomas Lemieux, 1–114. Elsevier.
- Montiel Olea, José Luis, and Mikkel Plagborg-Møller. 2019. "Simultaneous Confidence Bands: Theory, Implementation, and an Application to SVARs." *Journal of Applied Econometrics* 34 (1): 1–17.
- Mora, Ricardo, and Iliana Reggio. 2019. "Alternative Diff-in-Diffs Estimators with Several Pretreatment Periods." *Econometric Reviews* 38 (5): 465–86.
- Poirier, Alexandre, and Tymon Słoczyński. 2025. "Quantifying the Internal Validity of Weighted Estimands." Preprint, arXiv. <https://doi.org/10.48550/arXiv.2404.14603>.
- Rambachan, Ashesh, and Jonathan Roth. 2023. "A More Credible Approach to Parallel Trends." *Review of Economic Studies* 90 (5): 2555–91.
- Rambachan, Ashesh, and Jonathan Roth. 2025. "Design-Based Uncertainty for Quasi-Experiments." *Journal of the American Statistical Association*. <https://doi.org/10.1080/01621459.2025.2526700>.
- Robins, James. 1986. "A New Approach to Causal Inference in Mortality Studies with a Sustained Exposure Period—Application to Control of the Healthy Worker Survivor Effect." *Mathematical Modelling* 7 (9–12): 1393–512.
- Rosenbaum, Paul R., and Donald B. Rubin. 1983. "The Central Role of the Propensity Score in Observational Studies for Causal Effects." *Biometrika* 70 (1): 41–55.
- Roth, Jonathan. 2022. "Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends." *American*

- Economic Review: Insights* 4 (3): 305–22.
- Roth, Jonathan, and Pedro H. C. Sant'Anna. 2023a. "Efficient Estimation for Staggered Rollout Designs." *Journal of Political Economy Microeconomics* 1 (4): 669–709.
- Roth, Jonathan, and Pedro H. C. Sant'Anna. 2023b. "When Is Parallel Trends Sensitive to Functional Form?" *Econometrica* 91 (2): 737–47.
- Roth, Jonathan, Pedro H. C. Sant'Anna, Alyssa Bilinski, and John Poe. 2023. "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature." *Journal of Econometrics* 235 (2): 2218–44.
- Rubin, Donald B. 1974. "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies." *Journal of Educational Psychology* 66 (5): 688–701.
- Sant'Anna, Pedro H. C., and Jun Zhao. 2020. "Doubly Robust Difference-in-Differences Estimators." *Journal of Econometrics* 219 (1): 101–22.
- Sasaki, Yuya, and Takuya Ura. 2022. "Estimation and Inference for Moments of Ratios with Robustness against Large Trimming Bias." *Econometric Theory* 38 (1): 66–112.
- Seaman, Shaun R., and Stijn Vansteelandt. 2018. "Introduction to Double Robust Methods for Incomplete Data." *Statistical Science* 33 (2): 184–97.
- Semmelweis, Ignaz. 1983. *The Etiology, Concept and Prophylaxis of Childbed Fever*. Translated by K. Codell Carter. University of Wisconsin Press.
- Słoczyński, Tymon. 2022. "Interpreting OLS Estimands When Treatment Effects Are Heterogeneous: Smaller Groups Get Larger Weights." *Review of Economics and Statistics* 104 (3): 501–09.
- Smucler, Ezequiel, Andrea Rotnitzky, and James M. Robins. 2019. "A Unifying Approach for Doubly-Robust ℓ_1 Regularized Estimation of Causal Contrasts." Preprint, arXiv. <https://doi.org/10.48550/arXiv.1904.03737>.
- Snow, John. 1855. *On the Mode of Communication of Cholera*. John Churchill.
- Solon, Gary, Steven J. Haider, and Jeffrey M. Wooldridge. 2015. "What Are We Weighting For?" *Journal of Human Resources* 50 (2): 301–16.
- Sommers, Benjamin D., and Arnold M. Epstein. 2013. "US Governors and the Medicaid Expansion No Quick Resolution in Sight." *New England Journal of Medicine* 368 (6): 496–99.
- Strezhnev, Anton. 2018. "Semiparametric Weighting Estimators for Multi-period Difference-in-Differences Designs." Unpublished.
- Sun, Liyang, and Sarah Abraham. 2021. "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects." *Journal of Econometrics* 225 (2): 175–99.
- Tchetgen Tchetgen, Eric J., Chan Park, and David B. Richardson. 2024. "Universal Difference-in-Differences for Causal Inference in Epidemiology." *Epidemiology* 35 (1): 16–22.
- US Census Bureau. 2009–2019. *Small Area Income and Poverty Estimates (SAIPE) Program*. US Department of Commerce. <https://www.census.gov/programs-surveys/saipe.html> (accessed December 5, 2025).
- Wooldridge, Jeffrey M. 2003. "Cluster-Sample Methods in Applied Econometrics." *American Economic Review* 93 (2): 133–38.
- Wooldridge, Jeffrey M. 2023. "Simple Approaches to Nonlinear Difference-in-Differences with Panel Data." *Econometrics Journal* 26 (3): C31–66.
- Wooldridge, Jeffrey M. 2025. "Two-Way Fixed Effects, the Two-Way Mundlak Regression, and Difference-in-Differences Estimators." *Empirical Economics* 69 (5): 2545–87.
- Wyse, Angela, and Bruce D. Meyer. 2025. "Saved by Medicaid: New Evidence on Health Insurance and Mortality from the Universe of Low-Income Adults." NBER Working Paper 33719.
- Yanagi, Takahide. 2023. "An Effective Treatment Approach to Difference-in-Differences with General Treatment Patterns." Preprint, arXiv. <https://doi.org/10.48550/arXiv.2212.13226>.