

Pre-event trends in the panel event-study design

Simon Freyaldenhoven

Christian Hansen

Jesse M. Shapiro

Brown University

University of Chicago

Brown University and NBER*

This version: January 11, 2018

Abstract

We consider a linear panel event-study design in which unobserved confounds may be related both to the outcome and to the policy variable of interest. We provide sufficient conditions to identify the causal effect of the policy by exploiting covariates related to the policy only through the confounds. Our model implies a set of moment equations that are linear in parameters. The effect of the policy can be estimated by 2SLS, and causal inference is valid even when endogeneity leads to pre-event trends (“pre-trends”) in the outcome. Alternative approaches, such as estimation following a test for pre-trends, perform poorly.

JEL codes: C23, C26

KEYWORDS: pre-trends, event study, differences-in-differences

*Emails: simon.freyaldenhoven@brown.edu, chansen1@chicagobooth.edu, jesse_shapiro_1@brown.edu. We thank Isaiah Andrews, Jeff Clemens, Josh Gottlieb, Emily Oster, Bryce Steinberg, and seminar participants at Brown University for helpful comments. We thank Justine Hastings for sharing regression output. We acknowledge financial support from the National Science Foundation and the Brown University Population Studies and Training Center.

1 Introduction

We are interested in estimating the causal effect β of a policy variable z_{it} on an outcome y_{it} in a linear panel data model, where i indexes units and t indexes time. We are concerned that the strict exogeneity of z_{it} may fail due to the presence of a time-varying unobservable η_{it} that is correlated with both z_{it} and y_{it} . In the literature on the effects of the minimum wage, y_{it} is youth employment, i indexes states, t indexes calendar years, and z_{it} is an indicator for years after passage of a minimum-wage increase. The unobserved confound η_{it} is labor demand. The concern is that states tend to pass minimum-wage increases during good economic times (Neumark and Wascher [2006], Card and Krueger [1995]).

A common diagnostic approach in such settings is to look at whether the policy change appears to have an effect on the outcome before it actually occurs. The presence of such pre-event trends or “pre-trends” is taken as evidence against the strict exogeneity of the policy change.¹

This approach is incomplete. If pre-trends are not detected, it may be that there are no pre-trends, or that pre-trends are present but undetected due to limited statistical power. In the latter case, estimation under the assumption of strict exogeneity is typically inappropriate. If pre-trends are detected, it is understood that strict exogeneity is likely to fail, but it is not clear what to do.

In both cases, what is needed is a notion of magnitude: given some pre-trend in the outcome, how much of the apparent effect of the policy is due to confounds, and how much to the causal effect of the policy? Armed with such a notion, a researcher can conduct valid inference on β whether or not pre-trends are detected.

In this paper we propose to obtain such a notion from the behavior of a covariate x_{it} that is affected by the confound η_{it} but not by the policy. In the minimum wage context, adult employment responds to labor demand η_{it} but not to the minimum wage (Brown [1999]). Instead of using adult employment as a control variable, as is commonly done in the literature, we propose to look at its dynamics around minimum wage increases, and use these to infer the dynamics of η_{it} .

To fix ideas, suppose we observe the outcome y_{it} in periods $t = 1, \dots, T$ and the policy z_{it} in periods $t = 1 - L, \dots, T + L$ for some $L \geq 1$. Say that

$$y_{it} = \beta z_{it} + \gamma \eta_{it} + \varepsilon_{it} \tag{1}$$

$$\mathbb{E}(\varepsilon_{it} | \eta_{it}, \{z_{i,t}\}_{t=1-L}^{T+L}) = 0 \tag{2}$$

$$\mathbb{E}(x_{it} | \eta_{it}, \{z_{i,t}\}_{t=1-L}^{T+L}) = \lambda \eta_{it}, \tag{3}$$

¹Of the 16 papers in the 2016 *American Economic Review* that use a linear panel data model, 11 are concerned with the existence of pre-trends as a sign of endogeneity. Of these 11, 9 include a plot of pre-trends, of which 2 provide a formal test of whether pre-event effects are zero. In the minimum wage context, Allegretto et al. [2011] provide a plot of pre-trends.

where (1) defines the causal model, (2) is strict exogeneity of the policy with respect to the unobserved error ε_{it} , and (3) defines the relationship of the covariate x_{it} to the confound η_{it} up to the nonzero parameter λ . When the parameter γ is equal to zero, the confound does not affect the outcome, and identification of β is immediate.

Figure 1a plots coefficients from a regression of y_{it} on $\{\Delta z_{i,t+l}\}_{l=-L}^L$ in data simulated from an example of (1). Here and throughout Δ denotes the first difference operator. The plot resembles event-study plots in finance (Ball and Brown [1968], MacKinlay [1997]); hence, estimates like those underlying Figure 1a are sometimes called “event-study estimates” (e.g., Hoynes and Schanzenbach [2009], Duggan et al. [2016]).

Figure 1a shows a clear pre-trend in the outcome. What does this mean for the effect of the policy? Figure 1b shows an analogous plot for the covariate x_{it} . The covariate exhibits a pre-trend similar to that of the outcome, but a much smaller increase at the event time. Comparing the two plots, as in Figure 1c, allows us to decompose the change in the outcome at the event time into a component due to the causal effect of the policy and a component due to the dynamics of the confound η_{it} . The adjusted plot in Figure 1d removes the estimated effect of the pre-trend from Figure 1a, revealing the dynamics of the outcome net of the confound.

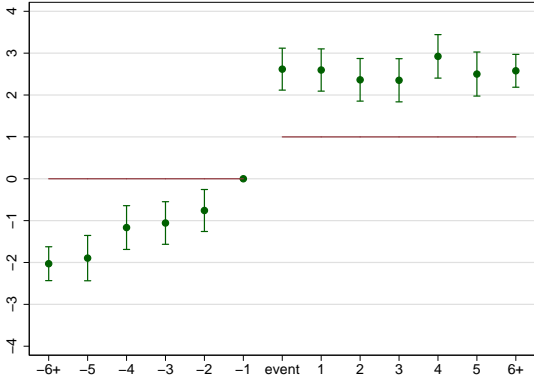
The geometry of these plots suggests an instrumental variables setup, in which Figure 1a plots the reduced form for the outcome and Figure 1b the first stage. Indeed, we show that β can be estimated by two-stage least squares (2SLS) regression of the outcome y_{it} on the policy z_{it} and covariate x_{it} , with leads (e.g., $z_{i,t+1}$) of the policy serving as excluded instruments.

Intuitively, from (3) we can think of the covariate x_{it} as a mismeasurement of the confound η_{it} . The use of an instrument removes the bias that would arise if we controlled directly for x_{it} . An essential assumption is that the dynamic relationship of x_{it} to z_{it} mirrors the dynamic relationship of η_{it} to z_{it} . This means, in particular, that x_{it} is affected by η_{it} but not by z_{it} .

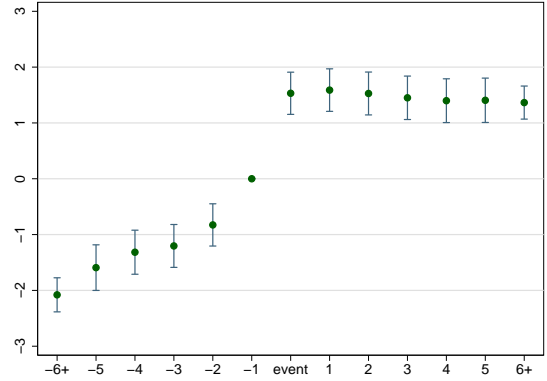
We also require that there be a pre-trend in η_{it} . As we argue in more detail in the paper, such a pre-trend arises naturally in the many economic settings in which the policy z_{it} changes when some unobserved state variable η_{it} crosses a threshold. Then the presence of a policy change in the future is informative about the current level of η_{it} . Our assumptions imply that a pre-trend in η_{it} manifests as a pre-trend in the covariate x_{it} , and may or may not manifest as a pre-trend in the outcome y_{it} .

Section 2 generalizes the setup in (1)-(3) to allow for multiple confounds, additive unit-specific fixed effects, and exogenous controls. We show that the model admits a GMM representation, from which standard results on estimation and inference (with large N and fixed T) are available.

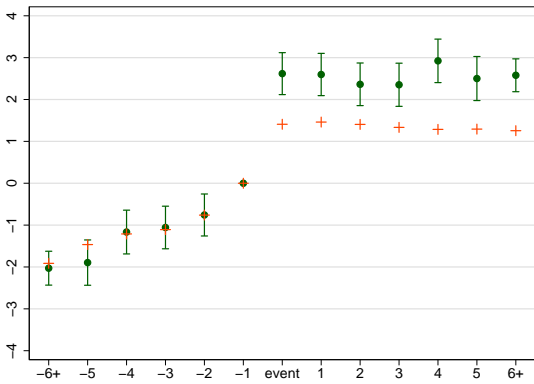
Section 3 presents Monte Carlo evidence on the finite-sample performance of our proposed estimator under a range of alternative data-generating processes, varying both the quality of the proxy x_{it} and the



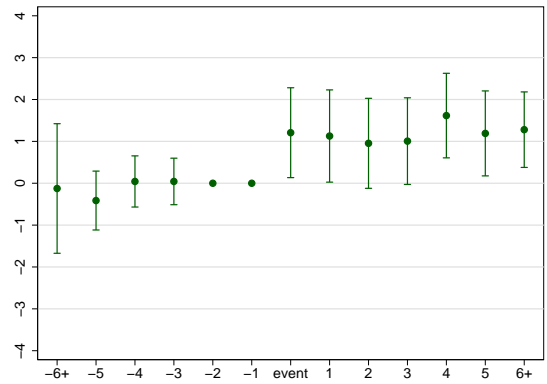
(a) Outcome of interest y_{it} around event time. Solid line depicts true causal effect.



(b) Unaffected covariate x_{it} around event time.



(c) Overlaying outcome of interest (with confidence intervals) y_{it} and unaffected covariate x_{it} (+, rescaled) around event time.



(d) Outcome of interest y_{it} around event time, using the behavior of the covariate to net out the effect of the confound.

Figure 1: Hypothetical event plots. An unobserved factor potentially causes endogeneity, manifested in a pre-trend in the outcome y_{it} . A covariate x_{it} affected by the confound, but not by the policy, permits us to learn the dynamics of the confound and adjust for it.

strength of identification. In these simulations, our estimator outperforms the approach of controlling directly for x_{it} except when x_{it} is a nearly perfect proxy for η_{it} . Our estimator also outperforms the approach of conducting a test for pre-trends before proceeding with estimation. We show that this “pre-test” approach is unreliable because it is vulnerable to undetectable (but important) trends in η_{it} .

Section 4 presents applications of our proposed approach to the effect of SNAP on household spending (Hastings and Shapiro [2017]), the effect of newspaper entry on voter turnout (Gentzkow et al. [2011]), and the effect of the minimum wage on youth employment (Neumark et al. [2014], Allegretto et al. [2017]). These applications illustrate a range of possibilities, including cases with clear pre-trends in the outcome, a case without meaningful pre-trends, and a case in which it is hard to tell. In some cases our proposed adjustment makes a small difference to point estimates, in some cases a larger difference, and in some cases it simply

implies greater statistical uncertainty.

We are not aware of an existing formal proposal to use an unaffected covariate to adjust causal inference for pre-trends in a panel data model. In their appendix, Gentzkow et al. [2011] implement an estimator that is similar in spirit to the one that we propose, but that is not formally justified by our setup.² Borusyak and Jaravel [2016] study the identification and estimation of pre-trends in a dynamic panel data model but do not consider the use of a side covariate to address endogeneity.

Our framework is closely related to classical work on models with measurement error and on panel data models with strict exogeneity. Replacement of η_{it} with x_{it} produces a factor model or a measurement error model (see, e.g. Aigner et al. [1984]). A large literature, partially reviewed in Heckman and Vytlacil [2007], Abbring and Heckman [2007] and Matzkin [2007], shows how to establish identification in such models, typically by imposing covariance restrictions across equations governing multiple imperfect measurements of the latent factor. Instead, we impose the strict exogeneity of the policy variable z_{it} with respect to the measurement error in x_{it} to achieve identification with only a single covariate.³

There are other ways to address policy endogeneity in linear panel data models like (1). One is to find an instrument for policy changes (Besley and Case [2000]). This is an appealing approach when feasible, but in many settings such instruments are not readily available. Our approach replaces the requirement of an instrument that impacts the policy but not the outcome with the requirement of a covariate that is related to the confound but unaffected by the policy. Another approach is to impose dynamic restrictions on the relationship between x_{it} and η_{it} . In a panel data setting with mismeasured regressors, Griliches and Hausman [1986] propose to use lags of x_{it} to construct valid instruments for x_{it} . This approach requires either that the measurement errors are serially uncorrelated or that the correlation structure of the measurement error is known (see also Wansbeek [2001] and Xiao et al. [2010]). Our approach allows for arbitrary correlation in the measurement errors, but requires that the policy be strictly exogenous with respect to these errors.

²Specification (6) of Table B1 in Gentzkow et al. [2011] uses a dynamic first stage analogous to Figure 1b and a static second stage analogous to (1). Gentzkow et al. [2011] provide a heuristic justification of their estimator in their footnote 5 but do not justify it formally.

³In the applications we have in mind, the number of covariates is small. If there are instead many covariates that contain independent information about the unobserved confounds, one may alternatively adapt methods from the literature on factor models in high dimensions as in, e.g., Stock and Watson [2002], Bai [2003], Bai and Ng [2010], and Hansen and Liao [2016].

2 Setup and Proposed Estimator

2.1 Model and Identifying Assumptions

We consider a static linear panel data model:

$$y_{it} = \beta z_{it} + q'_{it}\theta + \eta'_{it}\gamma + \alpha_i + \varepsilon_{it} \quad (4)$$

$$x_{it} = \psi q_{it} + \Lambda \eta_{it} + \nu_i + u_{it}, \quad (5)$$

where y_{it} and z_{it} are observed scalars, q_{it} is an observed $Q \times 1$ vector, x_{it} is an observed $K \times 1$ vector, the $R \times 1$ vector η_{it} , the $K \times 1$ vector u_{it} and scalar ε_{it} are time-varying unobservables, α_i is a time-invariant unobserved scalar, ν_i is a time-invariant unobserved $K \times 1$ vector, and the remaining objects are conformably defined parameters. The target of interest is β . We require that $K \geq R$ and suppose for simplicity that $K = R$. We observe data $\{y_{it}, q_{it}, x_{it}\}_{i=1, t=1}^{N, T}$ and $\{z_{it}\}_{i=1, t=1-m}^{N, T+\ell}$ for $m \geq 0$ and $\ell \geq R$. We do not require that z_{it} is binary.

Vector q_{it} collects all observed exogenous variables (e.g., time period indicators) in the sense that we impose $\mathbb{E}[\eta_{it} | \{q_{it}\}_{t=1}^T] = \mathbb{E}[\varepsilon_{it} | \{q_{it}\}_{t=1}^T] = 0$ for all i and t . The vector is low-dimensional in the sense that $Q \ll N$. We do not impose any restrictions on the α_i and thus treat them as fixed effects.

We take two steps to simplify the presentation of the results. First, we set $\theta = \psi = 0$. Statements carry over to the more general case by interpreting all data matrices as residuals from the projection of the remaining variables onto the exogenous variables. Second, we remove the fixed effects. Let $\tilde{w}_{it} = w_{it} - \frac{1}{T} \sum_{s=1}^T w_{is}$ denote the within transformation for any variable w_{it} .⁴ Then we can simplify (4)-(5) to obtain

$$\tilde{y}_{it} = \beta \tilde{z}_{it} + \tilde{\eta}'_{it}\gamma + \tilde{\varepsilon}_{it} \quad (6)$$

$$\tilde{x}_{it} = \Lambda \tilde{\eta}_{it} + \tilde{u}_{it}. \quad (7)$$

We now state two assumptions that suffice to identify β .

Assumption 1 (Orthogonality conditions). *There exists a set of non-negative integers $l = 0, 1, \dots, L$ such that*

$$(a) \mathbb{E}[\tilde{z}_{it+l}\tilde{\varepsilon}_{it}] = 0 \quad \forall l$$

$$(b) \mathbb{E}[\tilde{z}_{it+l}\tilde{u}_{it}] = 0 \quad \forall l.$$

⁴We also use this convention for leads and lags of a variable, so, for example, $\tilde{w}_{it-m} = w_{it-m} - \frac{1}{T} \sum_{s=1}^T w_{is-m}$.

Assumption 2 (Rank conditions). Define a matrix $H = \mathbb{E}(w_{it}[\tilde{z}_{it}, \tilde{x}'_{it}])$ with $w_{it} = (\tilde{z}_{it}, \tilde{z}_{it+1}, \dots, \tilde{z}_{it+L})'$. Then:

(a) $\text{rank}(\Lambda) = R$.

(b) $\text{rank}(H) = (R + 1)$.

Assumptions 1 and 2 are analogous, respectively, to the exclusion and relevance conditions in a linear instrumental variables setup. Assumption 1(a) imposes strict exogeneity as is commonly done in panel event studies. Assumption 1(b) imposes strict exogeneity in the first stage relationship (7), allowing z_{it} and its leads to be correlated with x_{it} only through η_{it} .

Remark 1. We do not require the orthogonality of $\tilde{\varepsilon}_{it}$ and \tilde{u}_{it} . The covariates x_{it} may be correlated with the outcome y_{it} through channels other than the confound η_{it} .

Remark 2. We may think of (7) as structural, or as a projection of \tilde{x}_{it} on $\tilde{\eta}_{it}$. In principle, the latter interpretation permits the structural relationship between x_{it} and η_{it} to be nonlinear, provided the projection residuals respect Assumption 1(b).

Assumption 2(a) imposes that the covariates x_{it} contain information about all of the latent factors η_{it} . Assumption 2(b) is the equivalent of the usual instrumental variables relevance assumption and can in principle be checked in the data. It requires a nonzero correlation between the noisy proxy x_{it} and leads of z_{it} .

Remark 3. Although Assumption 2(a) rules out that $\Lambda = 0$, we allow that $\gamma = 0$. This means that Assumption 2(b) requires a pre-trend in x_{it} , but does not require one in y_{it} .

Remark 4. Suppose that $z_{it} = \mathbf{1}(\exists t^* \leq t : \eta_{it^*} > \eta^*)$ for some threshold η^* . Then Assumption 2(b) will hold for a wide range of processes. Intuitively, if η_{it} is autocorrelated, a threshold crossing at time $t + 1$ provides a signal that the latent η_{it} was already large (close to the threshold) in the previous period. Economic settings covered by this case include:

- **Means-tested program.** We are interested in the effect of a household's participation z_{it} in a means-tested program on some outcome y_{it} as in Hastings and Shapiro [2017]. Each household i becomes eligible for the program when the gap η_{it} between the household's income and a poverty line exceeds a threshold η^* .
- **Firm entry.** We are interested in the effect of firm entry into a market on some outcome y_{it} as in Gentzkow et al. [2011]. At any given time t , a single potential entrant can pay a one-time cost to enter market i and earn a stream of cash flows whose expected present discounted value is η_{it} . Under

appropriate assumptions on η_{it} (for example, that it evolves as a random walk with i.i.d. innovations) the firm enters the first time that η_{it} exceeds a threshold η^* (see, e.g., McDonald and Siegel [1986] or Dixit and Pindyck [1994]). The policy z_{it} is an indicator for the presence of a firm in the market.

- **State law change.** We are interested in the effect of passage of a law on some outcome y_{it} . A given state i passes the law when the underlying state η_{it} of its economy exceed some threshold η^* . The policy z_{it} is an indicator for periods after passage of the law.

Remark 5. It is straightforward to allow the dimension of z_{it} to be greater than one, and to allow for dynamic treatment effects such that $y_{it} = \sum_{j=0}^m \beta_j z_{it-j} + q'_{it}\theta + \eta'_{it}\gamma + \alpha_i + \varepsilon_{it}$. We do not pursue these in order to keep notation and statements simple.

Remark 6. Our model rules out anticipatory effects of the policy. Specifically, (6) and Assumption 1(a) exclude any direct effect of leads of z_{it} on y_{it} .

2.2 GMM Representation and 2SLS Estimator

Towards a GMM representation, use Assumption 2(a) to define the $R \times 1$ matrix $\tilde{\Gamma} = \Lambda(\Lambda'\Lambda)^{-1}\gamma$. Now define

$$\tilde{v}_{it} \equiv \tilde{\varepsilon}_{it} - \tilde{u}'_{it}\tilde{\Gamma} \quad (8)$$

$$= \tilde{y}_{it} - \beta\tilde{z}_{it} - \tilde{x}'_{it}\tilde{\Gamma} \quad (9)$$

where (9) follows from (6) and (7) given the definition of $\tilde{\Gamma}$. Now from Assumption 1:

$$\mathbb{E}[w_{it}\tilde{v}_{it}] = 0. \quad (10)$$

Assumption 2(b) guarantees that the moment conditions in (10) are sufficient to identify β (and, incidentally, $\tilde{\Gamma}$).

Estimation may proceed by GMM using the sample analogues of (10) as moment conditions. For the case where T is fixed and N grows large, estimation and inference results are available under standard regularity conditions (e.g., Newey and McFadden [1994]).

One convenient estimator justified by (10) is a 2SLS regression of \tilde{y}_{it} on \tilde{z}_{it} and \tilde{x}_{it} , treating the covariates \tilde{x}_{it} as mismeasured regressors and the leads of \tilde{z}_{it} as the excluded instruments. We will use the 2SLS estimator in our simulations and applications.

Remark 7. In principle any functions of the leads of \tilde{z}_{it} are valid instruments. In practice, we expect that T

will often be moderately sized, and that the closest leads will be most informative. As a default we therefore suggest choosing the R closest leads of \tilde{z}_{it} as instruments. Among formal methods for instrument selection, BIC will often be natural as the number of potential instruments will often be small and the instruments are ordered, with closer leads more likely to be informative.

Remark 8. Extending our framework to the case where $K > R$ provides many additional moment restrictions in principle, though one suspects the usual complications from using many moment conditions (e.g. Han and Phillips [2006], Newey and Windmeijer [2009]) would arise in this setting.

Remark 9. Suppose that we observe \tilde{y}_{it} and \tilde{z}_{it} in one sample and \tilde{x}_{it} and \tilde{z}_{it} in another. Then we may proceed with two-sample instrumental variables estimation (e.g., Angrist and Krueger [1992], Inoue and Solon [2010]) using the leads of \tilde{z}_{it} as instruments for \tilde{x}_{it} .

3 Simulations

This section presents results from a number of Monte Carlo simulations. These allow us to compare the performance of alternative estimators and to assess the adequacy of standard asymptotic approximations of the finite-sample distributions of our estimator.

3.1 Data-generating Processes and Estimators

Definition 1 (Data-generating processes). *Throughout we consider the following data-generating processes (DGPs):*

- $\eta_{it} = \rho\eta_{it-1} + \zeta_{it}$ where $\zeta_{it} \sim N(0, \sigma_\zeta^2)$ are *i.i.d* across i and t
- $z_{it} = \mathbf{1}(\{\exists t^* \leq t : \eta_{it^*} > \eta^*\})$, where η^* is chosen so that the average number of events is approximately constant across different values of the simulation parameters.⁵
- $K = 1$ and

$$x_{it} = \lambda\eta_{it} + u_{it} \tag{11}$$

where $u_{it} \sim N(0, \sigma_u^2)$ are *i.i.d.* across i and t .

- The outcome is generated by:

$$y_{it} = \beta z_{it} + .25\eta_{it} + .2t + \alpha_i + \varepsilon_{it} \tag{12}$$

where $\beta = 1$, $\alpha_i \sim N(0, 1)$ are *i.i.d.* across i , $\varepsilon_{it} \sim N(0, 1)$ are *i.i.d.* across i and t , and α_i and ε_{it} are independent for all i and t .

⁵Specifically $\eta^* = \mathbf{1}(\rho < .9)(1.96 + .2\rho) + \mathbf{1}(\rho = .9)1.85 + \mathbf{1}(\rho = 1)4$

All of the simulations are based on the DGPs specified in Definition 1. Section 3.2 presents benchmark results for a design with $\lambda = \rho = 1$, $\sigma_\zeta^2 = 1$ and $\sigma_u^2 = 3$. We initialize η_{it} with $\eta_{i1} = 0$ and generate 20 time series observations for each i . We then use the ten time periods $t \in \{6, 15\}$ as the sample for estimation.

Section 3.3 presents more extensive results for a variety of designs with $\rho \in [0, 1)$. For these we choose σ_ζ^2 and σ_u^2 such that $\text{Var}(\tilde{\eta}_{it}) = 1$ and $\text{Var}(\tilde{x}_{it}) = 2$. To simulate these designs, we generate 20 time-series observations for each of 1000 cross-sectional units i . We initialize $\eta_{i,-19}$ as i.i.d. draws from a standard normal random variable and use the initial 20 observations $t = -19, -18, \dots, 0$ as burn-in. We then keep an estimation sample of 10 time-series observations consisting of the periods $t = 6, 7, \dots, 15$, retaining the full history of z_{it} so that we can construct leads and lags. Applying this procedure leaves us with $T = 10$ time-series observations on $N = 1000$ units, of which approximately 200 experience an event.⁶

To vary the strength of identification we will consider varying values of ρ in $[0, .9]$. As ρ increases, our instruments, the leads of z_{it} , will become stronger, resulting in better identification. On the other hand, as the autocorrelation in η_{it} approaches zero, we lose identification. Within this design, stronger persistence in η_{it} will tend to exacerbate bias that arises from failing to account for η_{it} .

To vary the quality of x_{it} as a proxy for η_{it} , we vary λ to control the population R^2 from the infeasible regression of x_{it} on η_{it} in (11). When this R^2 is one, x_{it} is a perfect proxy, and the best possible control for η_{it} is x_{it} . As this R^2 approaches 0, the proxy x_{it} provides no signal about the latent variable η_{it} , and identification fails.

We consider four different feasible estimators for the policy effect β and its dynamic counterparts. The first estimator we consider ignores η_{it} entirely and simply regresses the outcome on the event indicator plus individual and time fixed effects (“No control”). The second estimator uses x_{it} as a proxy for η_{it} and corresponds to the regression of the outcome onto the event indicator, x_{it} , and a full set of individual and time effects (“Controlling for x ”). The third estimator is from our proposed 2SLS regression of the outcome on the event indicator, x_{it} , and a full set of individual and time dummies, using z_{it+1} as an excluded instrument for x_{it} (“2SLS - one lead”). Online Appendix Figure 5 presents corresponding results using the BIC to choose the number of first-stage leads.

The last estimator that we consider formalizes the idea of testing for pre-trends that is common in applied work (“Pre-testing for trend”). To compute this estimator, we first compute the typical event study estimates, normalized so that the coefficient on z_{it+1} is equal to zero. We then perform a conventional test that the coefficient on z_{it+2} is equal to 0 at the 5% level. If we fail to reject the hypothesis, we conclude that

⁶Online Appendix Figure 3 shows the mean number of cross-sectional observations in which an event occurs, across the design space considered in the stationary case. Within each set of simulation parameters the number of units with an event generally lies between 160 and 240. We include in our analysis all cross-sectional units, including those in which an event does not occur (Borusyak and Jaravel [2016]).

there is no pre-trend and proceed with the analysis as in “No control.” If we reject the null, we conclude that there is a pre-trend and “give up.” We formalize the notion of “giving up” by returning a confidence interval of $(-\infty, \infty)$ and no point estimate. When evaluating point estimates for this procedure, we consider only those cases where we do not give up. Online Appendix Figure 4 summarizes the rejection frequency of the pre-test for trend.

3.2 Results for a Benchmark Data-generating Process

Figure 2 presents event-study estimates for a single realization from the DGP with $\rho = 1$. Specifically, each panel of Figure 2 depicts estimates of the coefficients δ_k from a different method of estimating the parameters of the following model:

$$y_{it} = \delta_{-6+}(1 - z_{it+5}) + \delta_{6+z_{it-6}} + \sum_{k=-5}^5 \delta_{-k} \Delta z_{it+k} + \omega_t + \alpha_i + \eta_{it} + \varepsilon_{it}, \quad (13)$$

where $(1 - z_{it+5})$ indicates that the event is more than 5 time periods in the future, and z_{it-6} indicates the event took place more than 5 periods in the past. We use the normalization that $\delta_{-1} = 0$.

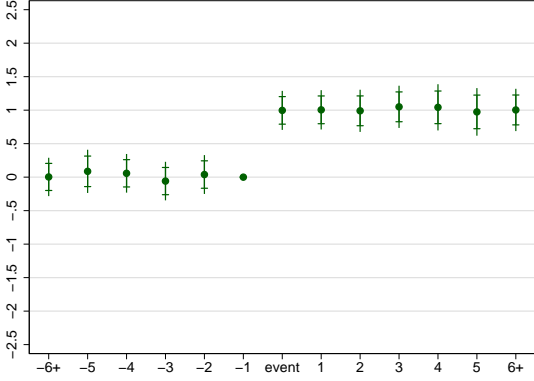
Figure 2 shows both pointwise 95% confidence intervals and uniform 95% sup-t confidence bands (Olea and Plagborg-Møller [2017]). Applied papers commonly include pointwise confidence intervals in event plots. These permit testing only of preselected pointwise hypotheses. Uniform bands such as those we show here are designed to contain the true path of the coefficients 95 percent of the time, and are therefore arguably more useful for the goal of giving readers a sense of what kinds of pre-trends are consistent with the data.

Figure 2a reports results from estimating (13) including η_{it} as an additional regressor. Because η_{it} is unobserved, this approach is infeasible, but it provides a useful benchmark of best-case performance. Point estimates of pre-event trends are reasonably small and well-estimated. Estimates of the policy effects (δ_k for $k > 0$) are reasonably close to one, the true value.

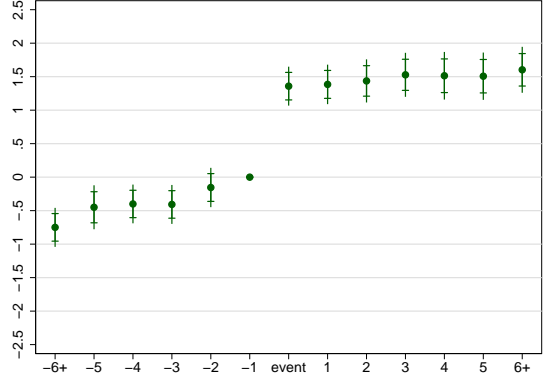
Figure 2b reports estimates without any control for η_{it} and shows both strong pre-trends and substantial bias in the estimated effects of the policy. Figure 2c reports estimates based on including the observable x_{it} in place of the latent variable η_{it} . As x_{it} is a noisy measure of η_{it} , controlling for x_{it} only partially mitigates the pre-trends, and the bias in the estimated policy effects, relative to Figure 2b.

Figure 2d shows the event plot using our proposed 2SLS estimator to account for the unobserved factor η_{it} . Specifically, we proxy for η_{it} with x_{it} and instrument for x_{it} with $z_{i,t+1}$.⁷ As expected, the proposed estimator delivers sensible estimates of pre-trends and policy effects, though there is a loss of precision relative to the infeasible benchmark in Figure 2a.

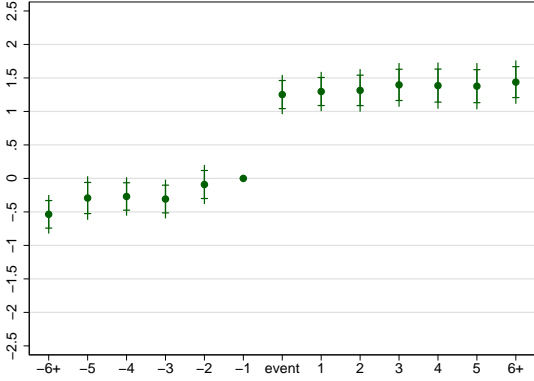
⁷Using $z_{i,t+1}$ as an instrument means that we need to normalize δ_k for an additional k . In Figure 2d we set $\delta_{-2} = 0$.



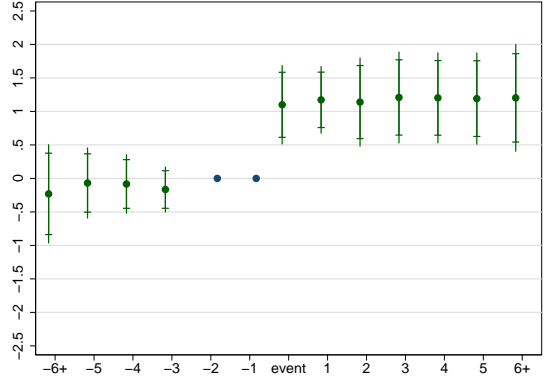
(a) Controlling directly for η_{it} (infeasible)



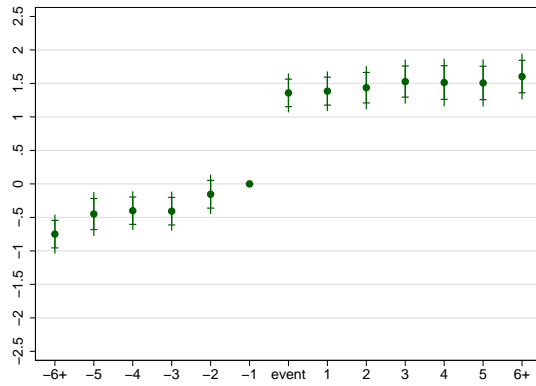
(b) Failing to control for η_{it}



(c) Using x_{it} as proxy for η_{it}



(d) Proposed 2SLS estimator, with closest lead of z_{it} as excluded instrument



(e) Pre-testing for trend

Figure 2: Exemplary event plots in the presence of a confounding factor using simulated data. Each plot shows estimates of the coefficients δ_k from (13) using either the infeasible estimator or one of the four feasible estimators defined in Section 3.1. Inner confidence sets as indicated by the dashes correspond to 95% pointwise confidence intervals, while outer confidence sets are the uniform 95% sup-t bands (with critical values obtained via simulation). Standard errors are clustered at the individual level. Data are a single draw from the benchmark DGP defined in Section 3.1.

Figure 2e reports estimates after pre-testing for a trend. As no trend is detected in this particular realization, this plot is identical to Figure 2b.

Figure 3 shows the median and uniform confidence band for the estimates in Figure 2 across repeated simulations from the same benchmark DGP. Figure 3 reinforces the conclusion from Figure 2 that, among the feasible estimators, only the 2SLS approach is centered at the true value.⁸

3.3 Results for a Set of Data-generating Processes

We turn next to an exploration of the full space covered by the stationary variant of the DGPs. We consider estimates $\hat{\beta}$ from

$$y_{it} = \beta z_{it} + \omega_t + \alpha_i + \eta_{it}\gamma + \varepsilon_{it}, \quad (14)$$

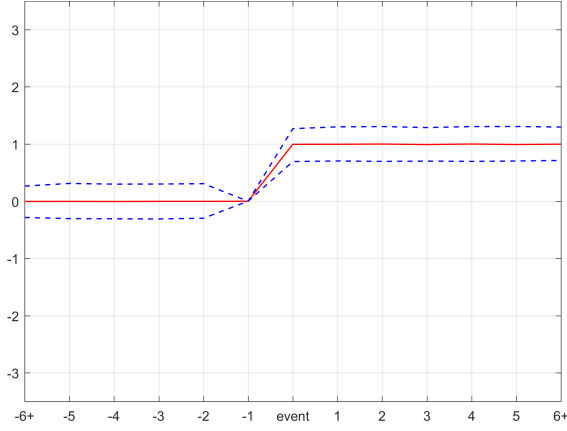
where ω_t are time effects. We consider the four feasible estimators defined in Section 3.1.

Figure 4 depicts the absolute median bias of each estimator. As expected, the presence of the unobserved confound severely biases the estimator that completely fails to control for η_{it} (top left panel). Using x_{it} directly to control for η_{it} also results in severe bias except when the R^2 from the infeasible regression of x_{it} on η_{it} is very large, in which case x_{it} is a nearly perfect proxy for η_{it} (top right panel). Also in line with our expectations, the median of our proposed 2SLS estimator is close to the true value across most of the parameter space (bottom left panel). The exceptions occur when the relationship between x_{it} and η_{it} becomes weak or there is little autocorrelation in η_{it} , both of which result in weak identification. Finally, the test for pre-trends leads to little improvement relative to no controls at all (bottom right panel). The reason is that the test often fails to detect trends, even though they are present.

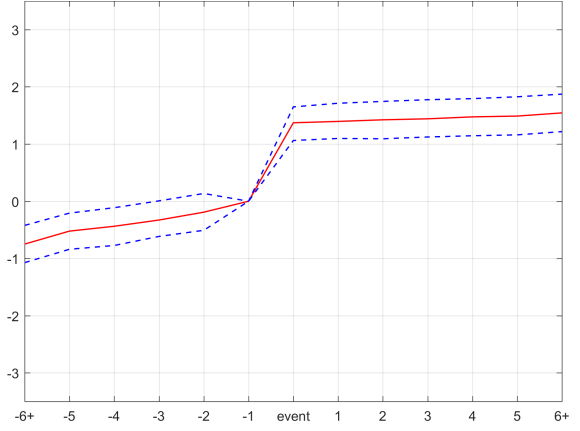
Figure 5 depicts the median absolute deviation of each estimator from the true parameter value. The sampling distribution of estimators other than our proposed 2SLS estimator are dominated by bias. Therefore, for these estimators, the plots in Figure 5 closely resemble those in Figure 4. In contrast, our proposed estimator performs well except in regions of the parameter space in which identification is weak.

Figure 6 depicts the coverage of the 95% confidence interval for each estimator constructed from the usual asymptotic approximation assuming the underlying sampling distribution is approximately normal and correctly centered. Failing to do anything to account for η_{it} results in severe size distortions across the entire parameter space (top left panel). Coverage is likewise poor when x_{it} is used directly as a proxy for η_{it} , except when x_{it} proxies η_{it} very well (top right panel). In contrast, empirical coverage for the 2SLS

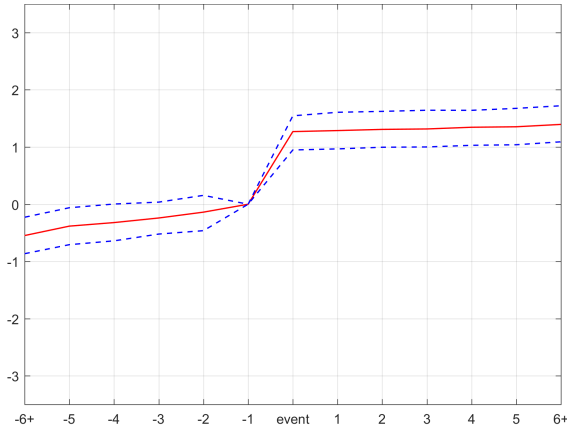
⁸Online Appendix Figure 2 shows how the behavior of plots based on our proposed 2SLS estimator depends on the choice of normalization.



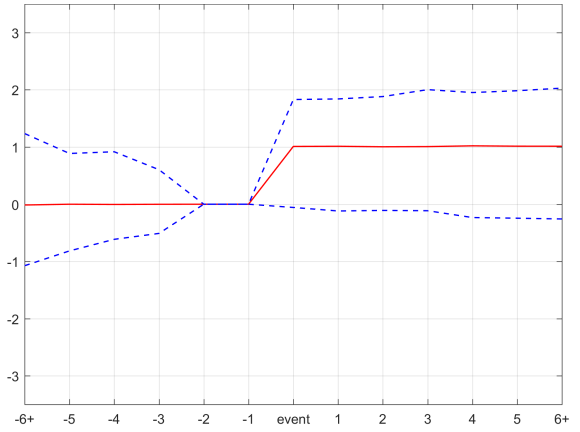
(a) Controlling directly for η_{it} (infeasible)



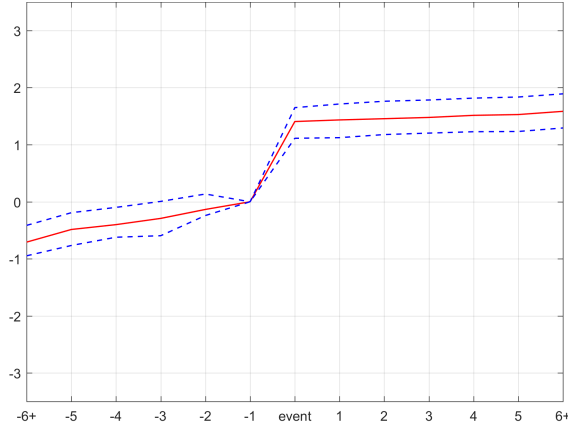
(b) Failing to control for η_{it}



(c) Using x_{it} as proxy for η_{it}



(d) Proposed 2SLS estimator, with closest lead of z_{it} as excluded instrument



(e) Pre-testing for pre-trend. Ignoring η_{it} if no pre-trend accepted. Estimates based on the 584 realisations with no detected pre-trend.

Figure 3: Distribution of event plots under the presence of a confounding factor using simulated data from the benchmark DGP defined in Section 3.1. Each plot shows estimates of the coefficients δ_k from (13) using either the infeasible estimator or one of the four feasible estimators defined in Section 3.1. The red solid line in the center represents the median estimate across 1000 realizations, while the blue dotted lines depict the uniform 95% confidence band: 95% of the estimated sets of coefficients lie within this band.

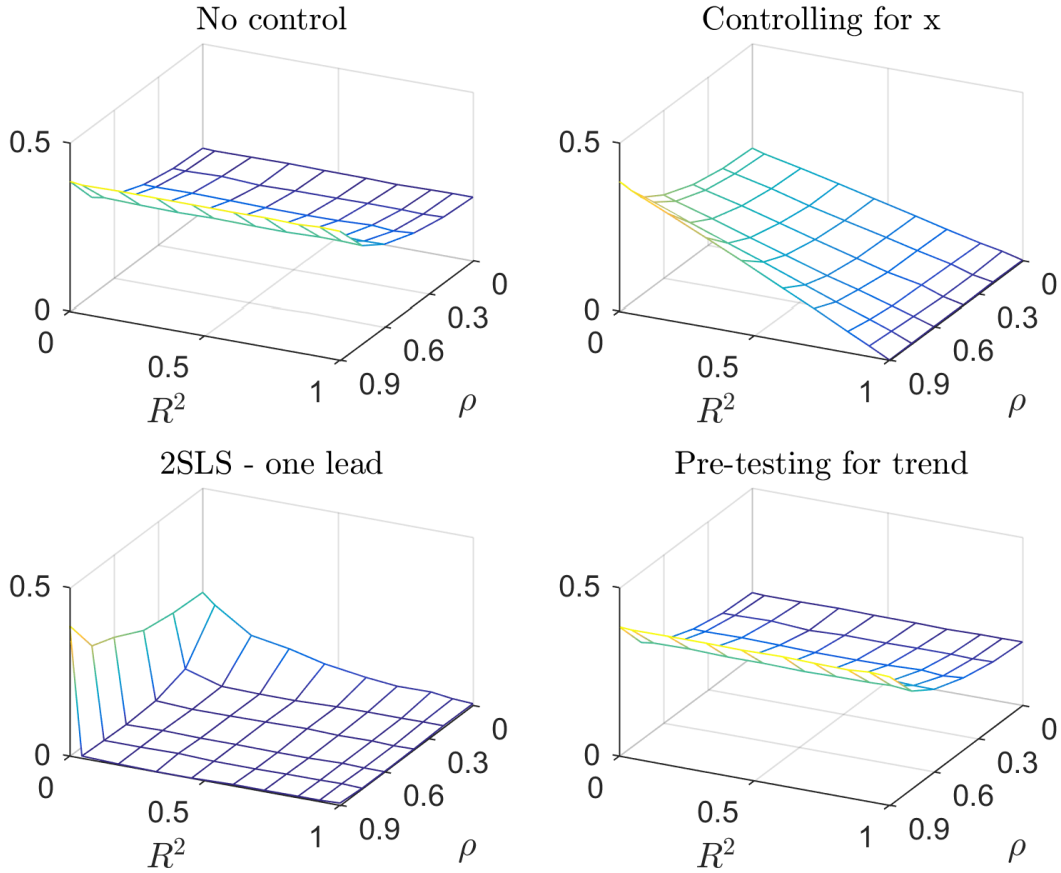


Figure 4: Median bias for each estimator defined in Section 3.1. Each point represents the median bias across 2000 simulation replications from the DGPs in Definition 1 with $\rho \in [0, 0.9]$. The horizontal axes in each panel correspond to the different values of ρ and of the population R^2 from the infeasible regression of x_{it} onto η_{it} in (11).

estimator is close to 95% throughout the parameter space, except where identification is weak (bottom left panel).⁹

The pre-test estimator exhibits uniformly poor coverage in this simulation design (bottom right panel). The observed coverage is a consequence of two offsetting patterns. When we reject the null of no pre-trend, coverage is necessarily 1 as we conclude we cannot use the data to learn about β . When we fail to detect a pre-trend and proceed as if no confound is present, coverage is close to 0 as the estimator is severely biased.

Although we present quantitative results for one particular form of pre-testing, we expect similar issues would arise for any analogous attempt to weed out bad cases in advance of estimation. The reason is that such approaches are vulnerable to the presence of trends that are small enough that they cannot be reliably detected in finite samples, but still large enough to significantly bias results obtained assuming such

⁹Poor coverage in regions of weak identification could be corrected by applying appropriate weak-identification robust procedures (e.g., Stock et al. [2002], Andrews and Mikusheva [2016]).

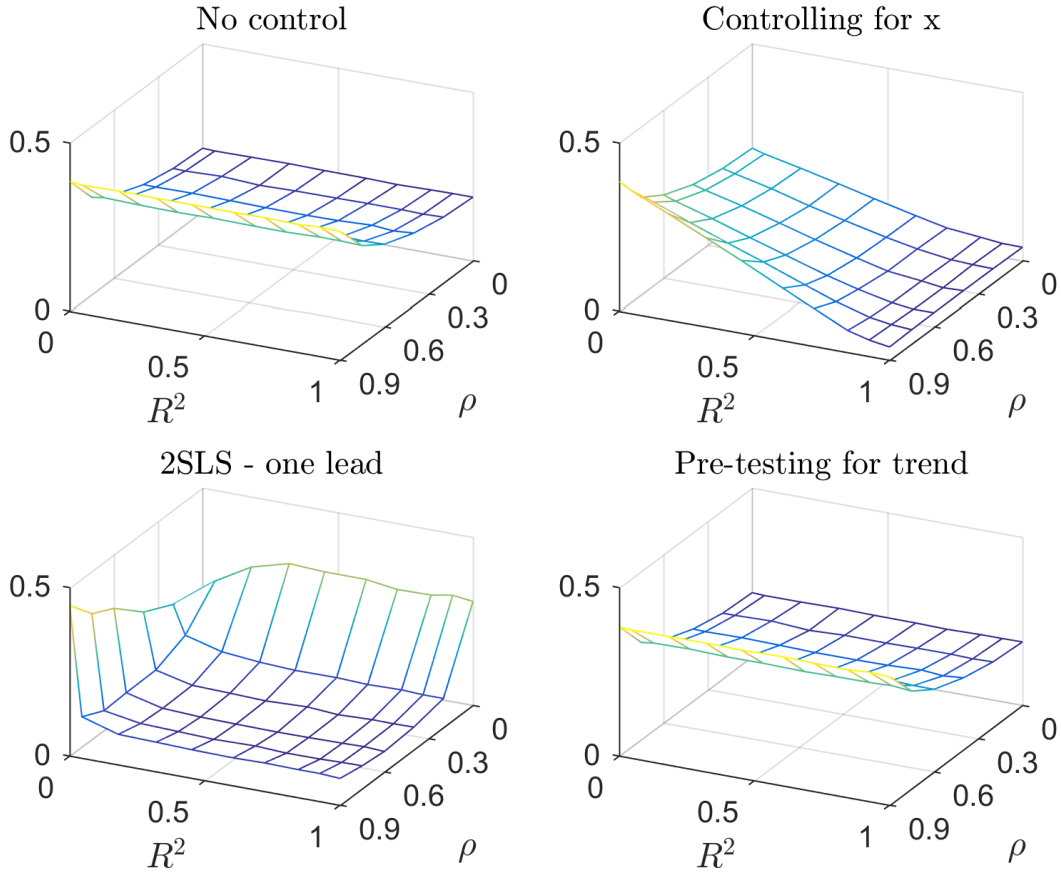


Figure 5: Median absolute deviation from the true parameter value for each estimator defined in Section 3.1. Each point represents the median absolute deviation across 2000 simulation replications from the DGPs in Definition 1 with $\rho \in [0, 0.9]$. The horizontal axes in each panel correspond to the different values of ρ and of the population R^2 from the infeasible regression of x_{it} onto η_{it} in (11).

trends are absent. By contrast, under the assumptions in Section 2, our proposed estimator delivers valid inferential statements when there are large trends (in scenarios where pre-trends are detected), when there are modest trends (where pre-trends are present but too small to be detected), and when there are no trends (where omitting η_{it} would also be fine). Our recommendation to use the 2SLS estimator is in line with the large statistics and econometrics literature regarding the use of pre-tests for model/specification choice. Guggenberger [2010], in particular, makes a very similar argument in the context of choosing between OLS or IV estimation.

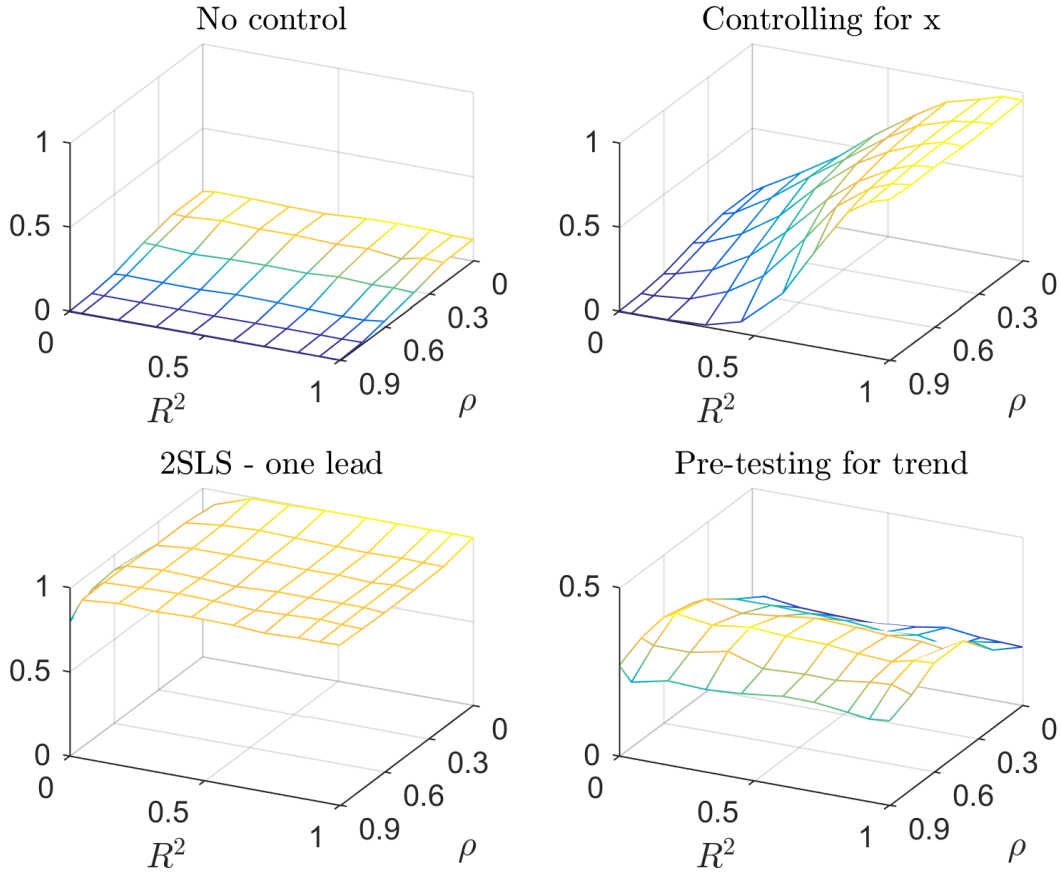


Figure 6: Coverage of the 95% confidence interval for each estimator defined in Section 3.1. Each point represents the coverage of the 95% confidence interval across 2000 simulations replications from the DGPs defined in Definition 1 with $\rho \in [0, .9]$. The confidence interval is constructed from the usual asymptotic approximation, with standard errors clustered at the individual level. The horizontal axes in each panel correspond to the different values of ρ and of the population R^2 from the infeasible regression of x_{it} onto η_{it} in (11).

4 Applications

In this section we apply our proposed estimator to empirical settings corresponding to the three examples discussed in Remark 4. Together these capture many of the scenarios a practitioner might encounter:

- A clear pre-trend in the outcome variable and a clear pre-trend in the covariate (Section 4.1).
- No pre-trend in the outcome variable and a clear pre-trend in the covariate (Section 4.1).
- An unclear pre-trend in the outcome variable and a clear pre-trend in the covariate (Section 4.2).
- An unclear pre-trend in the outcome variable and an unclear pre-trend in the covariate (Section 4.3).

4.1 The Effects of SNAP Participation on Household Spending Patterns

Hastings and Shapiro [2017] study the effect of participation in the Supplemental Nutrition Assistance Program (SNAP) on household spending in a panel event-study design. Here i indexes households and t indexes calendar quarters. The outcome y_{it} is either at-home food expenditures or the share of food spending going to store-brand items. The policy z_{it} is an indicator for time periods following entry into the program. SNAP is means-tested, so households become eligible when income η_{it} is sufficiently low. Past research shows that lower household income is associated with lower at-home food expenditures (e.g., Castner et al. [2010]) and greater store-brand share (e.g., Bronnenberg et al. [2015]), so income is a potential confound. Hastings and Shapiro [2017] have access to Rhode Island administrative data which includes SNAP participation z_{it} and a measure x_{it} of household income, and separate data from a grocery retailer which includes SNAP participation z_{it} and the outcomes y_{it} .¹⁰

Figure 7 reproduces from Hastings and Shapiro [2017] a plot of the time path of household income around the adoption of SNAP. Specifically, denoting average monthly household income during the quarter as x_{it} , we depict estimates $\hat{\delta}$ from

$$x_{it} = \delta_{-5+}(1 - z_{it+4}) + \delta_{5+}z_{it-5} + \sum_{k=-4}^4 \delta_{-k}\Delta z_{it+k} + \omega_t + \nu_i + u_{it}, \quad (15)$$

where ω_t denotes time effects.

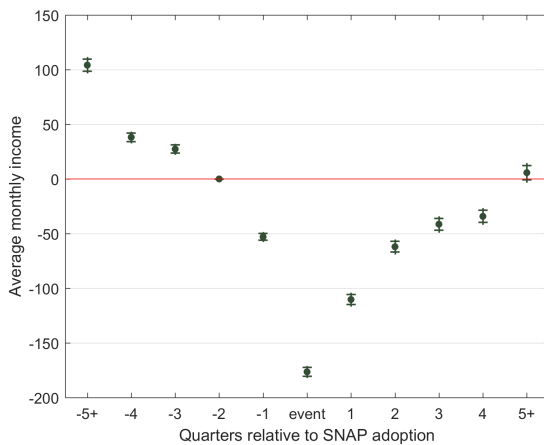


Figure 7: Estimated changes in household income at quarters around SNAP adoption. Each figure plots estimates of coefficients δ from (15), with the time period two quarters prior to SNAP adoption (“-2”) as the omitted category. Inner confidence sets as indicated by the dashes correspond to 95% pointwise confidence intervals, while outer confidence sets are the uniform 95% sup-t bands (with critical values obtained via simulation). Standard errors are clustered at the household level.

¹⁰The results in this section are based on regression output obtained from the authors at <http://www.brown.edu/Research/Shapiro/data/government.zip> and <http://www.brown.edu/Research/Shapiro/data/retailer.zip> on January 11, 2018.

The patterns in Figure 7 are consistent with a model in which household income is a determinant of SNAP eligibility as in Remark 4. We see a clear downward trend in income in the time periods leading up to a household’s adoption of SNAP. Following the adoption, we observe an increase in household income.

Figure 8 depicts estimates $\hat{\delta}$ from two specifications of

$$y_{it} = \delta_{-5+}(1 - z_{it+4}) + \delta_{5+}z_{it-5} + \sum_{k=-4}^4 \delta_{-k}\Delta z_{it-k} + \gamma\eta_{it} + \alpha_i + \omega_t + \varepsilon_{it}, \quad (16)$$

with all variables on the right side defined as in (15) and outcome y_{it} representing either monthly at-home food expenditure (Figures 8a and 8b) or store-brand share of food expenditures (Figures 8c and 8d).

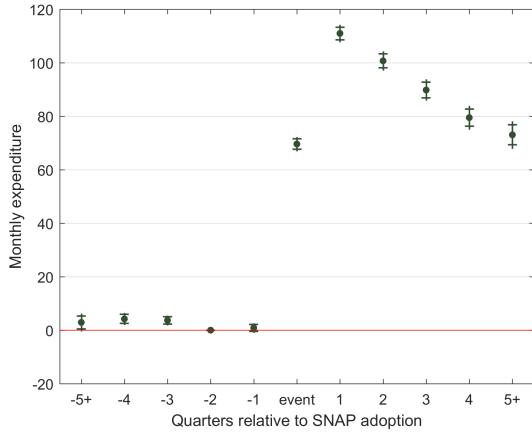
In Figures 8a and 8c, the term in (16) involving η_{it} is ignored and so no attempt is made to control for confounds. Figure 8a shows that there is no economically meaningful trend in monthly at-home food expenditure prior to the adoption of SNAP. This is consistent with the argument in Hastings and Shapiro [2017] that the effect of cash income on food spending is small. By contrast, 8c shows a clear pre-trend in store-brand share that is small in absolute terms but large relative to the change on adoption. We note that, since SNAP adoption can occur at any time in the quarter, period 0 is “partially treated”.

Figures 8b and 8d use our proposed estimator, with the closest lead of z_{it} serving as an excluded instrument for x_{it} . Because y_{it} and x_{it} are not observed jointly, we use a two-sample instrumental variables estimator (Angrist and Krueger [1992], Inoue and Solon [2010]). In the case of at-home food expenditures, Figure 8b shows that, as expected, taking the income confound into account does not alter the conclusions from the uncorrected plot in Figure 8a.

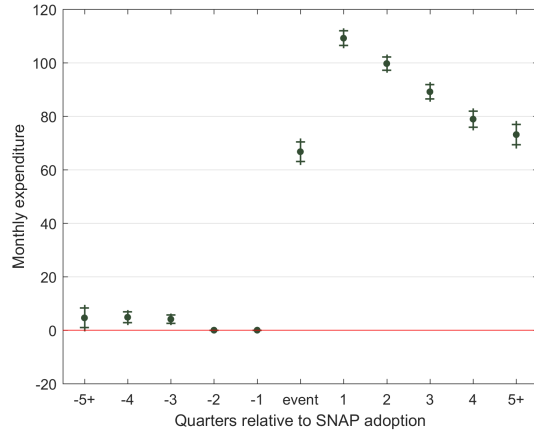
By contrast, Figure 8d differs markedly from Figure 8c. The relatively large pre-trend in store-brand share implies a significant response to changes in income. The instrumental variable approach we take accounts for this trend through the presence of the confound η_{it} , and eliminates the pre-trend from the plot. The dynamics of store-brand share that we observe following adoption likely reflects households’ gradual exit from the program following adoption.

Figure 9 provides a geometric intuition for our proposed procedure. It combines a rescaled version of Figure 7 with Figure 8c. Our proposed estimator uses the dynamics in both the household income and monthly expenditure in the two quarters prior to the event to infer the size of the effect of the confound. Geometrically this amounts to aligning the two plots in the two-period window prior to the event. We interpret the remaining difference, depicted in Figure 8d, as an approximation of the causal effect of SNAP adoption on the store-brand share.

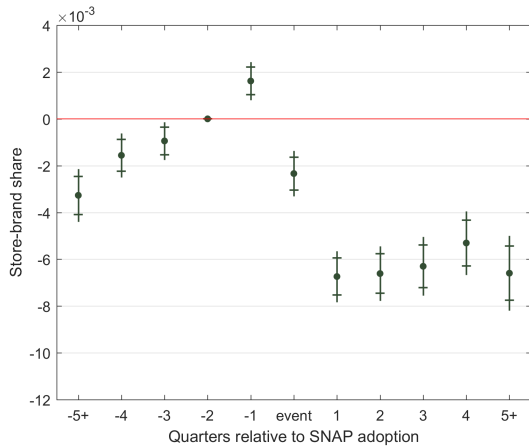
Figure 8 illustrates two possible scenarios for applying our approach in the presence of a clear potential confound. In the first scenario, there is no meaningful pre-trend in the outcome, and our proposed method



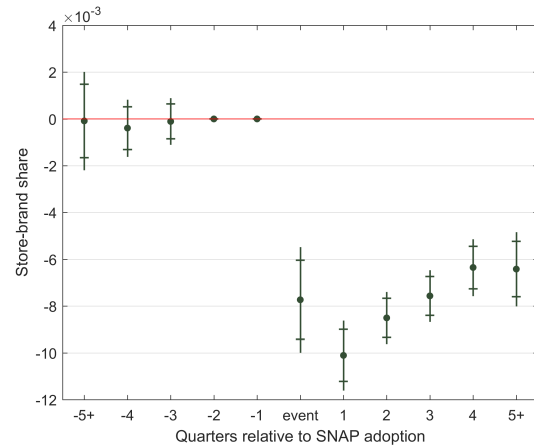
(a) At-home food expenditure around SNAP adoption, not controlling for household income.



(b) At-home food expenditure around SNAP adoption. Proposed 2SLS estimator, with z_{it+1} as excluded instrument.



(c) Store-brand share of food expenditures around SNAP adoption, not controlling for household income.



(d) Store-brand share of food expenditures around SNAP adoption. Proposed 2SLS estimator, with z_{it+1} as excluded instrument.

Figure 8: Estimated changes in outcomes at quarters around SNAP adoption. Each figure plots estimates of coefficients δ from (16), with the time period two quarters prior to SNAP adoption (“-2”) as the omitted category. Inner confidence sets as indicated by the dashes correspond to 95% pointwise confidence intervals, while outer confidence sets are the uniform 95% sup-t bands (with critical values obtained via simulation). Standard errors are clustered at the household level.

formalizes the intuitive notion that the confound does not cause significant bias in the estimation of the policy effect. In the second scenario, there is a clear pre-trend in the outcome, and our method adjusts causal inference for the presence of the confound.

Table 1 presents estimates $\hat{\beta}$ from the static analogue of (16). Although a static model does not capture the post-treatment dynamics of the outcomes, it is a common way to summarize the effect size. The first row shows that, with no control for household income, the estimated effect of adopting SNAP on monthly expenditure is 86 dollars, while SNAP adoption leads to a decrease in the store-brand share of 0.4%. The

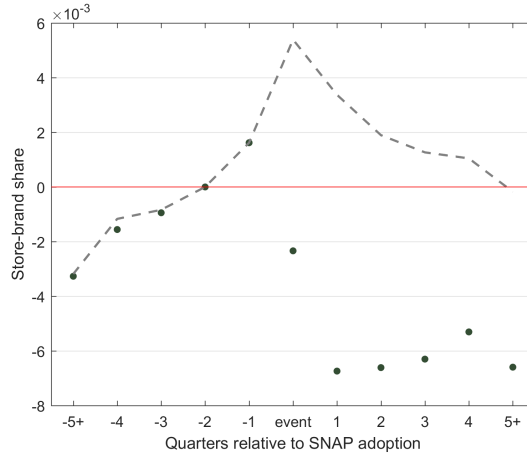


Figure 9: Geometric illustration of proposed estimator. Dashed line depicts household income (Figure 7, rescaled). Round markers depict monthly expenditure on SNAP-eligible items (Figure 8c) around SNAP adoption respectively. Difference depicted in Figure 8d.

second row notes that controlling for household income directly is infeasible, as household income and the outcomes of interest are not observed in the same data. The third row shows that, using our proposed 2SLS estimator, the estimated effect of SNAP adoption on monthly food expenditure is 84 dollars, similar to the first row, and the estimated effect on the store-brand share is a decrease of 0.7%, an increase in magnitude of almost 60% compared to the first row. As expected from Figure 7, the first stage is highly significant. The last row shows the estimate from pre-testing for a pre-trend. This estimate coincides with that in the first row in the case of monthly food expenditure, where there is no detectable pre-trend, and is undefined in the case of the store-brand share, where there is a clear pre-trend.

4.2 The Effect of Newspaper Entry and Exit on Electoral Politics

Gentzkow et al. [2011] study the effect of newspapers on voter turnout, exploiting variation generated by the entries and exits of daily newspapers to local markets in the US. Here, i indexes local markets (counties) and t indexes presidential election years. The outcome y_{it} is voter turnout. The policy z_{it} is the number of English-language daily newspapers in the market. Following Remark 4, it is reasonable to expect the entry of a newspaper to coincide with an improvement in market profitability η_{it} . Because the state of the local economy could also affect voter turnout, there is a potential confound. Gentzkow et al. [2011] have proxies for profitability, including a measure x_{it} of the log of the voting-eligible population.¹¹

¹¹We use the authors' original data, available at <http://www.icpsr.umich.edu/icpsrweb/ICPSR/studies/30261/version/5>, in our analysis.

Estimator	Effect of SNAP adoption on		Coefficient on lead in first-stage
	monthly expenditure	store-brand share	
No control	85.97 (1.23)	-0.0044 (0.0004)	
Controlling for x_{it}	infeasible	infeasible	
Proposed 2SLS estimator (one lead)	84.35 (1.11)	-0.0070 (0.0004)	-151.81 (2.55)
Pre-testing for pre-trend	85.97 (1.23)	— (—)	

Table 1: Estimates of the effect of SNAP adoption. In the first two columns, each row corresponds to a different estimate $\hat{\beta}$ from $y_{it} = \beta z_{it} + \omega_t + \gamma \eta_{it} + \alpha_i + \varepsilon_{it}$. The first row uses no control for household income. The second row reports that controlling directly for household income is infeasible. The third row uses our proposed 2SLS estimator, treating the closest lead of SNAP adoption as an excluded instrument for household income. The last row depicts our pre-test estimator. For monthly expenditure, we cannot reject $\delta_{-2} = 0$ at the 5% level in (16) (see Figure 8a), and therefore proceed with estimation as in the first row. In contrast, for store-brand share, we reject $\delta_{-2} = 0$ at the 5% level in (16) (see Figure 8c), and therefore do not proceed with estimation. The last column shows the coefficient on the excluded instrument in the first stage of the 2SLS estimator. Standard errors in parentheses are clustered at the household level.

Figure 10 depicts estimates of the coefficients δ_k from

$$\Delta x_{it} = \sum_{k=-5}^5 \delta_{-k} \Delta z_{it+k} + \Delta \omega_{st} + \Delta u_{it}, \quad (17)$$

where ω_{st} is a state-year fixed effect. The patterns in the figure are consistent with a model in which the voting-eligible population approximates newspaper profitability: We see a clear upward trend in population growth in the time periods leading up to a market entry, and then population growth flattens out again after an entry has occurred.

Figure 11 depicts estimates of the coefficients δ_k from three specifications of the equation:

$$\Delta y_{it} = \sum_{k=-5}^5 \delta_{-k} \Delta z_{it+k} + \Delta \omega_{st} + \gamma \Delta \eta_{it} + \Delta \varepsilon_{it}. \quad (18)$$

In Figure 11a the term involving η_{it} is omitted from (18). This specification therefore does not control for newspaper profitability. Figure 11b controls for market profitability by directly substituting the observed x_{it} for η_{it} in (18). Figure 11c uses our proposed 2SLS estimator, with the closest lead of Δz_{it} serving as an excluded instrument for Δx_{it} . Based on the uniform bands, we reject the hypothesis of no pre-trends when no control is used, but fail to reject in the other two specifications.

Table 2 presents estimates $\hat{\beta}$ from the static analogue of (18), which represents the causal effect of an additional newspaper on voter turnout. The first row shows that with no controls the estimated effect is

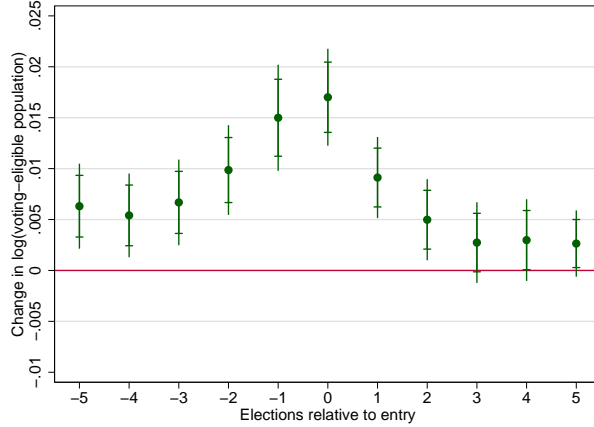


Figure 10: Estimated changes in population at election years around newspaper entries/exits. The plot shows estimates of coefficients δ_k from (17). Inner confidence sets as indicated by the dashes correspond to 95% pointwise confidence intervals, while outer confidence sets are the uniform 95% sup-t bands (with critical values obtained via simulation). Standard errors are clustered at the county level.

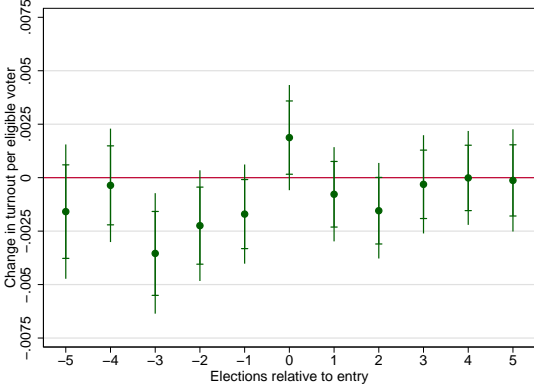
0.26 percentage points per newspaper. The second row shows that controlling for the log of voting-eligible population leads the estimate to increase to 0.37 percentage points per newspaper. The third row shows that our proposed 2SLS estimator gives an estimate of 0.34 percentage points per newspaper, which is statistically and economically similar to the estimate in the second row.¹²

The fourth row shows the results from the pre-test procedure defined in Section 3.1. A test of the hypothesis that the coefficient on Δz_{it+1} in Figure 11a is significantly different from zero yields a t -statistic

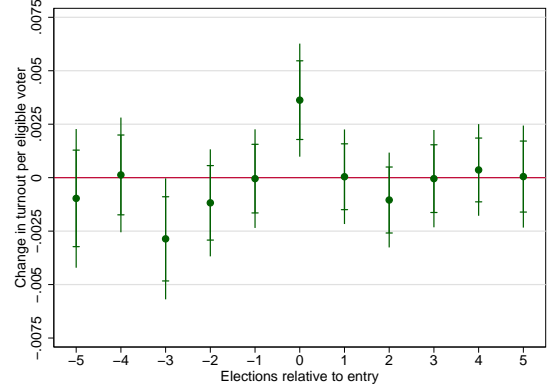
¹²The p-values for equality of estimates relative to controlling for x_{it} directly are 0.0026 for the estimator with no control and 0.752 for our proposed 2SLS estimator. These p-values are based on 100 cluster-bootstrap replications.

Estimator	Effect of newspaper entry	Coefficient on lead in first-stage
No control	0.0026 (0.0009)	
Controlling for x_{it}	0.0037 (0.0010)	
Proposed 2SLS estimator (one lead)	0.0034 (0.0013)	0.0128 (0.0017)
Pre-testing for pre-trend	— (—)	

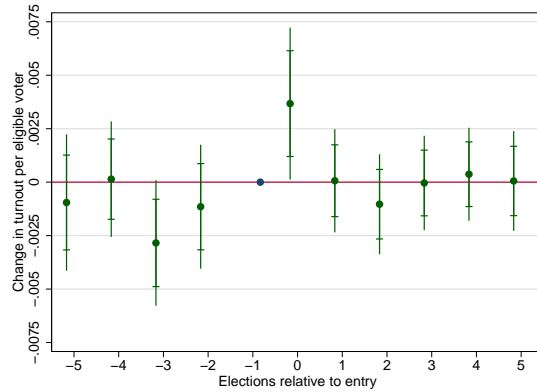
Table 2: Estimates of the effect of newspapers on voter turnout. In the first column, each row corresponds to a different estimate $\hat{\beta}$ from $\Delta y_{it} = \beta \Delta z_{it} + \Delta \omega_{st} + \gamma \Delta \eta_{it} + \Delta \varepsilon_{it}$. The first row uses no control for market profitability. The second row uses the log of the voting-eligible population as a proxy. The third row uses our proposed 2SLS estimator, treating the closest lead of the number of newspapers as an excluded instrument for the log of the voting-eligible population. The last row depicts our pre-test estimator. Because we reject $\delta_{-1} = 0$ at the 5% level in (18) (see Figure 11a), we do not proceed with estimation. The second column shows the coefficient on the excluded instrument in the first stage of the 2SLS estimator. Standard errors in parentheses are clustered at the county level.



(a) Not controlling for market profitability



(b) Using the log of voting-eligible population as a proxy for market profitability



(c) Proposed 2SLS estimator, with Δz_{ct+1} as excluded instrument

Figure 11: Estimated effects on presidential turnout at election years around newspaper entries/exits. The plot shows estimates of coefficients δ_k from (18). Inner confidence sets as indicated by the dashes correspond to 95% pointwise confidence intervals, while outer confidence sets are the uniform 95% sup-t bands (with critical values obtained via simulation). Standard errors are clustered at the county level.

of -2.07 with a p-value of 0.039. Absent controls, the pre-test procedure therefore suggests not to proceed with estimation.

4.3 The Effect of the Minimum Wage on Youth Employment

There is an ongoing debate about the effect of the minimum wage on youth employment (e.g. Neumark et al. [2014], Allegretto et al. [2017]). Here i indexes states and t indexes quarters. The outcome y_{it} is the log of the teen (16-19) employment-to-population ratio. The policy z_{it} is the log of the state minimum wage. The control q_{it} is the share of teenagers in the population. We may be concerned that states implement minimum-wage increases when demand η_{it} for labor is strong (Neumark and Wascher [2006], Card and Krueger [1995]). We proxy for labor market conditions using a measure x_{it} of the log of the prime-age (25-

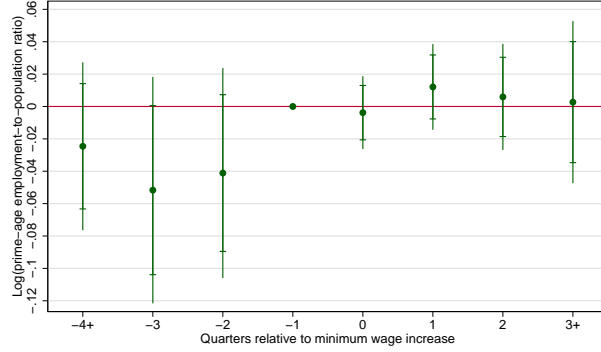


Figure 12: Prime-age employment at quarters around minimum wage increases. The plot shows estimates of coefficients δ from (19). Inner confidence sets as indicated by the dashes correspond to 95% pointwise confidence intervals, while outer confidence sets are the uniform 95% sup-t bands (with critical values obtained via simulation). Standard errors are clustered at the state level.

55) employment-to-population ratio. For prime-age workers the effect of minimum wages is minor compared to other sources of variation (Brown [1999]). However, directly controlling for x_{it} , as is commonly done, fails to allow for mismeasurement of the true demand for youth labor.

We construct data on y_{it} , x_{it} , and q_{it} from the CPS Outgoing Rotation Groups for the years 1985–2014.¹³ We obtain data on z_{it} from David Neumark’s Minimum Wage Dataset¹⁴ All regressions in this section are weighted by teen population.

Figure 12 depicts the time path of our proxy, the log of prime-age employment, around minimum wage increases. Specifically, the figure depicts estimates of the coefficients δ_k from

$$x_{it} = \delta_{-4+}(1 - z_{it+3}) + \delta_{3+}z_{it-3} + \sum_{k=-2}^3 \delta_{-k}\Delta z_{it+k} + q'_{it}\theta + \omega_t + \nu_i + u_{it}. \quad (19)$$

Consistent with our expectation, the point estimates indicate that increases in the minimum wage tend to occur following an increase in prime-age employment. However, the estimates are imprecise, and based on the uniform confidence intervals we cannot reject the hypothesis of no pre-trends.

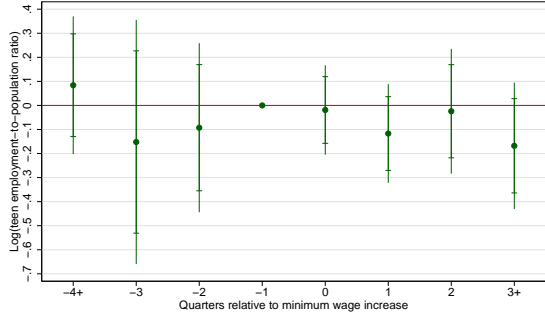
Figure 13 depicts estimates $\hat{\delta}$ from three specifications of the equation:

$$y_{it} = \delta_{-4+}(1 - z_{it+3}) + \delta_{3+}z_{it-3} + \sum_{k=-2}^3 \delta_k\Delta z_{it+k} + \gamma\eta_{it} + q'_{it}\theta + \omega_t + \alpha_i + \varepsilon_{it}. \quad (20)$$

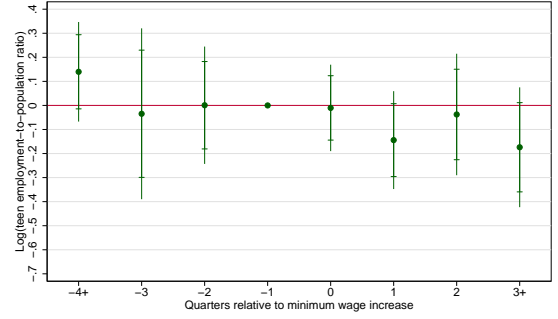
In Figure 13a the term involving η_{it} is omitted from (20). This specification therefore does not control for

¹³The Current Population Survey data is available at <http://www.nber.org/data/morg.html>. We construct the employment-to-population ratios as the proportion of individuals in the corresponding age category who self-report as either “Working” or “With a job, not at work.” We weight individual observations using the final weight variable to obtain state-level aggregates.

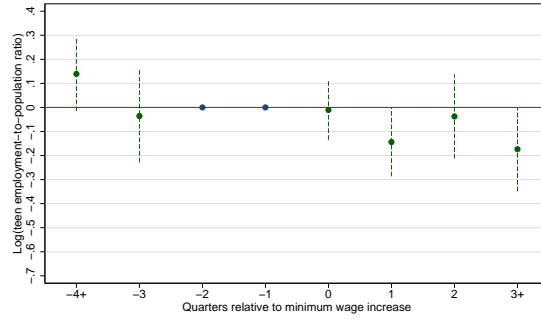
¹⁴The minimum wage data is available at <http://www.socsci.uci.edu/~dneumark/datasets.html>. We use the higher of the federal or state minimum wage as the prevailing minimum wage.



(a) Not controlling for state of economy



(b) Using prime-age employment as a proxy for the state of the economy



(c) Proposed 2SLS estimator, with z_{it+1} as excluded instrument

Figure 13: Teen employment at quarters around minimum wage increases. The plot shows estimates of coefficients δ_k from (20). In the top row, inner confidence sets as indicated by the dashes correspond to 95% pointwise confidence intervals, while outer confidence sets are the uniform 95% sup-t bands (with critical values obtained via simulation). In the bottom figure, dashed confidence intervals correspond to 95% pointwise confidence intervals, ignoring weak-identification. Standard errors are clustered at the state level.

the state of the labor market. Figure 13b uses prime-age employment x_{it} directly as a control. Figure 13c depicts the results from our proposed estimator, in which we use the closest lead of the policy, z_{it+1} , as an excluded instrument for x_{it} . Note that the first stage for this model is weak, and the confidence set for γ based on inversion of the Anderson-Rubin (AR) test consists of the entire real line. A projection argument therefore implies that valid confidence sets in Figure 13c also include the entire real line.

Table 3 presents estimates $\hat{\beta}$ from the following static model, represented in first differences:

$$\Delta y_{it} = \beta \Delta z_{it} + \Delta \mathbf{q}'_{it} \theta + \Delta \omega_t + \gamma \Delta \eta_{it} + \Delta \varepsilon_{it}. \quad (21)$$

The first row of 3 shows the estimate with no control for η_{it} . The second row shows the estimate when controlling directly for x_{it} as a proxy for η_{it} . The third row shows our proposed 2SLS estimator, using Δz_{it+1} as an excluded instrument for Δx_{it} . We present both conventional standard errors and a confidence interval constructed by projection based on inversion of the AR test for γ . The last row shows the estimate

	Effect of log(MW)	Coefficient on lead in first stage
No control	-0.0114 (0.0743)	
Controlling for prime-age employment	-0.0094 (0.0708)	
Proposed 2SLS estimator (one lead)	0.0003 (0.0668)	0.0314 (0.0136)
	$[-\infty, \infty]$	
Pre-testing for pre-trend	-0.0114 (0.0743)	

Table 3: Estimates of the effect minimum wage on teen (16-19) employment. Dependent variable: $\log(\text{employment}/\text{population})$. Each row corresponds to a different estimate $\hat{\beta}$ from the model in first differences given by (21). The first row uses no control for the state of the economy. The second row uses the prime-age employment as a proxy. The third row uses our proposed 2SLS estimator, treating the change in the first lead of the log of the minimum wage as an excluded instrument for the change in the log of prime-age employment. We present both conventional standard errors and a confidence interval (in square brackets) constructed by projection based on an inversion of the AR test for γ . The last row depicts our pre-test estimator. Because we cannot reject $\delta_{-2} = 0$ at the 5% level in (20) (see Figure 13a), we proceed with estimation as in “No control”. All regressions are weighted by teen population. Standard errors are clustered at the state level.

from the pre-test procedure. Because we cannot reject the hypothesis of no pre-trends in y_{it} , this estimate coincides with the one in the first row.

This last application demonstrates the limitations of our proposed estimator. Instrument relevance requires a strong pre-trend in the covariate x_{it} . Absent such a pre-trend, our proposed estimator is not strongly identified and our approach implies that the econometrician cannot learn about the parameter of interest. Arguably, however, that is a valid conclusion if we are concerned about a confound η_{it} and do not believe that x_{it} is a perfect proxy for that confound.

5 Conclusion

We consider a linear panel data model with possible endogeneity. The common approach of examining plots of event-study estimates, or formally testing for the presence of pre-event effects, is inadequate. We show how to exploit a covariate related to the confound but unaffected by the policy of interest to perform causal inference in this setting. We validate our proposal in simulations from a range of data-generating processes, and apply it to three economic settings of interest.

References

- Jaap H Abbring and James J Heckman. Econometric evaluation of social programs, part iii: Distributional treatment effects, dynamic treatment effects, dynamic discrete choice, and general equilibrium policy evaluation. *Handbook of econometrics*, 6:5145–5303, 2007.
- Dennis J Aigner, Cheng Hsiao, Arie Kapteyn, and Tom Wansbeek. Latent variable models in econometrics. *Handbook of econometrics*, 2:1321–1393, 1984.
- Sylvia Allegretto, Arindrajit Dube, Michael Reich, and Ben Zipperer. Credible research designs for minimum wage studies: A response to neumark, salas, and wascher. *ILR Review*, 70(3):559–592, 2017.
- Sylvia A Allegretto, Arindrajit Dube, and Michael Reich. Do minimum wages really reduce teen employment? accounting for heterogeneity and selectivity in state panel data. *Industrial Relations: A Journal of Economy and Society*, 50(2):205–240, 2011.
- Isaiah Andrews and Anna Mikusheva. A geometric approach to nonlinear econometric models. *Econometrica*, 84(3):1249–1264, 2016.
- Joshua D Angrist and Alan B Krueger. The effect of age at school entry on educational attainment: an application of instrumental variables with moments from two samples. *Journal of the American statistical Association*, 87(418):328–336, 1992.
- Jushan Bai. Inferential theory for factor models of large dimensions. *Econometrica*, 71(1):135–171, 2003.
- Jushan Bai and Serena Ng. Instrumental variable estimation in a data rich environment. *Econometric Theory*, 26(06):1577–1606, 2010.
- Ray Ball and Philip Brown. An empirical evaluation of accounting income numbers. *Journal of accounting research*, pages 159–178, 1968.
- Timothy Besley and Anne Case. Unnatural experiments? estimating the incidence of endogenous policies. *The Economic Journal*, 110(467):672–694, 2000.
- Kirill Borusyak and Xavier Jaravel. Revisiting event study designs. Working paper, 2016.
- Bart J. Bronnenberg, Jean-Pierre Dub, Matthew Gentzkow, and Jesse M. Shapiro. Do pharmacists buy bayer? informed shoppers and the brand premium *. *The Quarterly Journal of Economics*, 130(4): 1669–1726, 2015. doi: 10.1093/qje/qjv024.

- Charles Brown. Minimum wages, employment, and the distribution of income. *Handbook of labor economics*, 3:2101–2163, 1999.
- David Edward Card and Alan B Krueger. *Myth and Measurement: The New Economics of the Minimum Wage*. Princeton University Press, 1995.
- Laura Castner, James Mabli, et al. Low-income household spending patterns and measures of poverty. *Washington, DC: Mathematica Policy Research*, 2010.
- Avinash K Dixit and Robert S Pindyck. *Investment under uncertainty*. Princeton university press, 1994.
- Mark Duggan, Craig Garthwaite, and Aparajita Goyal. The market impacts of pharmaceutical product patents in developing countries: Evidence from india. *The American Economic Review*, 106(1):99–135, 2016.
- Keith Finlay, Leandro Magnusson, and Mark E Schaffer. Weakiv: Stata module to perform weak-instrument-robust tests and confidence intervals for instrumental-variable (iv) estimation of linear, probit and tobit models. *Statistical Software Components*, 2016.
- Matthew Gentzkow, Jesse M Shapiro, and Michael Sinkinson. The effect of newspaper entry and exit on electoral politics. *The American Economic Review*, 101(7):2980–3018, 2011.
- Zvi Griliches and Jerry A Hausman. Errors in variables in panel data. *Journal of econometrics*, 31(1):93–118, 1986.
- Patrik Guggenberger. The impact of a hausman pretest on the asymptotic size of a hypothesis test. *Econometric Theory*, 26(02):369–382, 2010.
- Chirok Han and Peter C. B. Phillips. Gmm with many moment conditions. *Econometrica*, 74(1):147–192, 2006.
- Christian Hansen and Yuan Liao. The factor-lasso and k-step bootstrap approach for inference in high-dimensional economic applications. *arXiv preprint arXiv:1611.09420*, 2016.
- Justine S Hastings and Jesse M Shapiro. How are snap benefits spent? evidence from a retail panel. Technical report, National Bureau of Economic Research, 2017.
- James J Heckman and Edward J Vytlacil. Econometric evaluation of social programs, part ii: Using the marginal treatment effect to organize alternative econometric estimators to evaluate social programs, and to forecast their effects in new environments. *Handbook of econometrics*, 6:4875–5143, 2007.

- Hilary W. Hoynes and Diane Whitmore Schanzenbach. Consumption responses to in-kind transfers: Evidence from the introduction of the food stamp program. *American Economic Journal: Applied Economics*, 1(4): 109–139, 2009.
- Atsushi Inoue and Gary Solon. Two-sample instrumental variables estimators. *The Review of Economics and Statistics*, 92(3):557–561, 2010.
- A Craig MacKinlay. Event studies in economics and finance. *Journal of economic literature*, 35(1):13–39, 1997.
- Rosa L Matzkin. Nonparametric identification. *Handbook of Econometrics*, 6:5307–5368, 2007.
- Robert McDonald and Daniel Siegel. The value of waiting to invest. *The Quarterly Journal of Economics*, pages 707–728, 1986.
- David Neumark and William Wascher. Minimum wages and employment: A review of evidence from the new minimum wage research. Technical report, National Bureau of Economic Research, 2006.
- David Neumark, JM Ian Salas, and William Wascher. More on recent evidence on the effects of minimum wages in the united states. *IZA Journal of Labor policy*, 3(1):24, 2014.
- Whitney K Newey and Daniel McFadden. Large sample estimation and hypothesis testing. *Handbook of econometrics*, 4:2111–2245, 1994.
- Whitney K. Newey and Frank Windmeijer. Generalized method of moments with many weak moment conditions. *Econometrica*, 77(3):687–719, 2009.
- José Luis Montiel Olea and Mikkel Plagborg-Møller. Simultaneous confidence bands: Theoretical comparisons and suggestions for practice. 2017. URL https://scholar.harvard.edu/files/plagborg/files/conf_band.pdf.
- James H Stock and Mark W Watson. Forecasting using principal components from a large number of predictors. *Journal of the American statistical association*, 97(460):1167–1179, 2002.
- James H. Stock, Jonathan H. Wright, and Motohiro Yogo. A survey of weak instruments and weak identification in generalized method of moments. *Journal of Business and Economic Statistics*, 20(4):518–529, 2002.
- Tom Wansbeek. Gmm estimation in panel data models with measurement error. *Journal of Econometrics*, 104(2):259–268, 2001.

Zhiguo Xiao, Jun Shao, and Mari Palta. Instrumental variable and gmm estimation for panel data with measurement error. *Statistica Sinica*, pages 1725–1747, 2010.