Policy Evaluation with Nonlinear Trended Outcomes: COVID-19 Vaccination Rates in the US*

Lynn Bergeland Morgan[a], Peter C. B. Phillips[b,c,d] and Donggyu Sul[f]

[a] *University of Texas at Dallas,* [b] *Yale University,* [c] *University of Auckland,*
[d] *Singapore Management University,*
[f] *University of Texas at Dallas*

January 8, 2024

Abstract

This paper discusses pitfalls in two-way fixed effects (TWFE) regressions when the outcome variables contain nonlinear and possibly stochastic trend components. If a policy change shifts trend paths of outcome variables TWFE estimation can distort results and invalidate inference. A robust solution is proposed by allowing for dynamic club membership empirically using a relative convergence test procedure. The determinants of respective club memberships are assessed by panel ordered logit regressions. The approach allows for policy evolution and shifts in outcomes according to a convergence cluster framework with transitions over time and the possibility of eventual convergence to a single cluster as policy impacts mature. The long-run impact of a policy can thus be examined via its impact on convergence club membership. An application to new weekly US Covid-19 vaccination policy data reveals that Federal-level vaccine mandates produced a merger of state vaccination rates into a single convergence cluster by mid-September 2021.

Keywords: Two-way Fixed Effects Regression, Robust Clustering Algorithm, Relative Convergence, Automatic Clustering Mechanism, Panel logit Regression

JEL Classification: I18, C33, H51.

*Phillips acknowledges research support from the NSF (Grant No. SES 18-50860) at Yale University and a Kelly Fellowship at the University of Auckland.

Email: Lynn.Morgan@utdallas.edu; peter.phillips@yale.edu; d.sul@utdallas.edu.

1 Introduction

Two-way fixed effects (TWFE) regressions are commonly used in empirical work, particularly to evaluate the effectiveness of exogenous policy change. Several outcome variables are often examined in such regressions and policy effectiveness is measured through its potential impact on dependent variables that may themselves involve stochastic trends such as random walks or similar random processes with nonlinear trend components. Examples of such variables are personal wages, state per capita income, and life expectancy data, all of which are popular in practical work because of their immediate relevance in policy determination or more general interest in social science studies. For a recent estimate of such work we counted the number of published papers that employed nonstationary outcome variables in the January 2023 issue of the *AEJ: Applied Economics*. Of the 16 articles published in this issue all but one used TWFE regressions. Among these, 12 papers used nonstationary outcome variables including log earnings, log wages, median income and crop output; and some 86% of the papers using TWFE regressions employed nonstationary dependent or outcome variables.[1]

To fix ideas suppose y_{it} is a relevant outcome variable comprising a sequence of panel observations of interest across individuals ($i = 1, \cdots, n$) and over time ($t = 1, \cdots, T$) and let x_{it} be a relevant policy variable. A typical TWFE regression specification for analyzing such data takes the form

$$y_{it} = a_i^o + \theta_t^o + \beta x_{it} + \gamma' z_{it} + v_{it}, \tag{1}$$

where a_i^o and θ_t^o are individual and time specific effects, and z_{it} is a vector of control variables. When the policy x_{it} has the form of a panel dummy variable, some pitfalls in TWFE estimation have already been studied in recent work De Chaisemartin and d'Haultfoeuille (2020); Goodman-Bacon (2021). The present paper extends that research by considering more general and empirically relevant cases where the policy variable x_{it} may itself evolve over time and where the outcome variable y_{it} may have nonstationary characteristics sourced beyond the policy input. In complex applied work where the nonstationarity in outcomes is not fully captured by policy inputs or time specific effects, the regression residual v_{it} becomes nonstationary and the adequacy of TWFE regressions in accurately measuring policy effectiveness may be called into question.

Previous studies often bypass such complications by assuming a small fixed T, frequently just $T = 2$ observations, and allowing a large cross section sample size n to deliver asymptotics. In cases like the recent Covid-19 pandemic experience successive policy changes over time need to be accommodated in the data analysis and longer time series samples are available for assessing the

[1] Amongst individual household earnings or wages only a small fraction may have nonstationary characteristics but as these time series are aggregated nonstationary characteristics tend to stand out and become more dominant as the aggregation level successively rises from suburbs to cities and counties. See Sul (2019) for discussion.

evidence of successive policy effects in the observed outcomes. Our empirical study has $T = 40$ and $n = 51$ observations, so that cross section and time series sample sizes are comparable. This paper studies cases where trends are homogeneous and heterogeneous over individuals under the null of no treatment effect or no policy effectiveness. Under the alternative, policy variables influence outcome variables either in levels (C1) or in trends (C2). If a stationary policy variable affects a nonstationary outcome variable in levels (C1), the policy effectiveness does not last long, holding only temporarily. In this case, TWFE regression is valid only under a homogeneous trend with consistent estimation and valid standard errors as $n \to \infty$. But the properties of estimation and inference are less favorable when the outcome variables are nonstationary with heterogeneous trend behavior. A partial solution is straightforward under the assumption that policy variables affect outcomes *only in levels*. For level effects can be consistently estimated either under homogeneous or heterogeneous trends using the differenced regression

$$\Delta y_{it} = a_i + \Delta\theta_t^o + \beta\Delta x_{it} + \gamma'\Delta z_{it} + \Delta v_{it}, \tag{2}$$

However, it is unrealistic to expect that policy variables affect outcomes only in levels. When policy affects trend behavior of the outcome variables, the first difference TWFE regression (2) is misspecified, leading to inconsistent estimation of policy effects.

Our study considers several issues that bear directly on TWFE and difference-in-difference (DiD) treatment effect analysis. The primary concern examines outcome variables when stochastic trend components are omitted from the regression and when heterogeneous trend behavior over time or across individuals affects outcomes. In both cases, the TWFE system is misspecified, leading to growing uncertainty and inconsistent estimation. A typical DiD approach[2] to resolving this issue is to employ first difference outcomes or growth rates while maintaining the same policy variable, so that the TWFE regression now takes the form

$$\Delta y_{it} = a_i + \theta_t + \beta x_{it} + \gamma' z_{it} + u_{it}, \tag{3}$$

in which the influence of x_{it} on the growth rate of y_{it} is studied. This approach leads to a further pitfall which arises when the outcome y_{it} has nonlinear trend behavior. First difference TWFE regressions, either in (2) or (3), typically examine short run responses by removing presumed long run stochastic trend (unit root process) relationships, which may not be well suited to other long run behavior. For instance, suppose that a certain policy change has a positive effect only for a

[2] In the DiD or event-study regression literature, several devices have been proposed, which are well surveyed in the recent review by Roth et al. (2023) and an alternative solution allowing for violation of the parallel trends assumption is considered in Rambachan and Roth (2023). Our approach allows, in addition, for policy evolution and shifts in outcomes over time according to a convergence cluster framework with transitions over time and the possibility of eventual convergence to a single cluster as policy impacts mature.

short period but produces a negative impact on the outcome variable in the long run. In such a case, TWFE estimation as in (3) can mistakenly lead to an implied positive impact even for large T, as demonstrated later.

One of the main contributions of the paper is to provide a methodology that organizes non-stationary panel data into ordered panel multinomial variables by means of a dynamic clustering mechanism that allows for shifts in clusters over time. To fix ideas, define C_{it} to be the convergence club membership in period t of the i-th individual and suppose all individuals within a certain convergence club share the same stochastic trend. If an individual grows faster over time and joins another convergence club that attains higher outcomes, then C_{it} membership changes over time. The convergence clustering mechanism (CCM hereafter) proposed in earlier work by Phillips and Sul (2007a, hereafter, P-S) transforms statistics from panel observations to clustered cross-sections, but potential dynamic changes amongst club memberships are ignored. In practice, club membership can change over time, particularly as relevant policies are introduced or evolve over time, and such evolution can itself be a natural focus of interest concerning policy impacts. Our approach in the present paper is to develop the CCM of P-S into a dynamic version that accommodates such possibilities. Once dynamic group membership is estimated, panel logit (or multinomial logit) regression enables estimation and inference concerning driver variables and the determining mechanisms of the groups. The proposed method can be used with either event-study (DiD) or continuous policy interventions. The approach is also robust against functional forms with heterogeneous trends as long as C_{it} membership is well defined.

A second main contribution of the paper is to create a new weekly database that tracked state and District of Columbia announcements of the numerous vaccination policies[3] implemented over the period from March 2021 to February 2022. This database enables a detailed empirical study of the impact of federal and state vaccination policies on state vaccination rates. A final contribution is technical and relates to the use of logarithmic transformation of the data. In particular, we consider the use of logarithms of ratios in panel data and provide a simple procedure for dealing with the practical problem of logarithmic representations when a few data points are zero or negative.

The rest of the paper is organized as follows. Section 2 considers issues arising from nonstationary outcome variables and pitfalls in the use of first differences in policy evaluation. Dynamic mechanisms for club membership are developed in Section 3. Section 4 consists of two subsections. The first subsection discusses data preparation and some useful new adjustments in logarithmic transforms and then second subsection provides an application of this mechanism to state vaccination rates over time in the United States. Panel logit modeling is applied to assess the effects of U.S. federal level mandate announcements and various state level COVID-19 vaccine policies on

[3] These policies included lotteries, cash for vaccination incentives, community outreach programs, vaccine mandates for state employees and or healthcare workers, indoor vaccine mandates or mandates for gatherings over a certain number of people, mask mandates, bans on proof of vaccination, and bans on mask mandates.

state convergence club membership. Section 5 concludes. Federal and state level mandate policies are described in Appendices A and B, and additional logit regression specification results are given in Appendix C. Technical background, derivations, proofs and further simulations are provided in the Online Supplement to this paper.

2 Pitfalls of TWFE Regressions with Nonstationary Outcomes

Throughout the paper we assume that the dependent or outcome variables are nonstationary, whereas the policy variable is stationary.[4] We consider the following three general models allowing nonstationary outcomes when a policy variable does not influence y_{it} (i.e., under the null hypothesis):

$$y_{it} = \begin{cases} a_i + t + \xi_{it},\ \xi_{it} = \xi_{it-1} + u_{it} & \text{for M1}, \\ a_i + b_i t + \xi_{it},\ \xi_{it} = \xi_{it-1} + u_{it} & \text{for M2}, \\ a_i + b_{it} t + \xi_{it},\ \xi_{it} = \xi_{it-1} + u_{it} & \text{for M3}, \end{cases} \tag{4}$$

where the nonstationarity includes heterogeneous deterministic trend functions as well as a stochastic trend. Specifically: in M1, y_{it} has a homogeneous linear trend; in M2, y_{it} has a heterogeneous linear trend with time invariant individual coefficients b_i; and in M3, the individual trend coefficients b_{it} are time varying, which nests M1 and M2. A policy variable, x_{it}, is assumed to be stationary. One may introduce a staggered or continuous trend to x_{it}[5], but such modifications do not alter the main results of the paper.

Under this framework there are two mechanisms by which the stationary policy change can affect the nonstationary outcome under the alternative: (i) x_{it} can affect the level of y_{it} (designated C1, as above); and (ii) x_{it} can affect the trend coefficient b_{it} (designated C2). Under the null of no policy effectiveness, the outcome variables are given by (4). In this case the parameter β, which measures the policy impact in the following equations, is zero. The two alternative cases are written as follows.

$$y_{it} = a_i + \theta_t + \beta x_{it} + \xi_{it} \ \text{ for C1}, \tag{5}$$

$$b_{it} = b_o + \beta x_{it} + e_{it} \ \text{ for C2}, \tag{6}$$

where θ_t in (5) may be a linear or more general deterministic trend and b_{it} are the trend coefficients in M3. We consider C1 in the next subsection, showing that if the true generating process is either

[4] As shown later in the empirics, key policy variables often tend to change in a staggered manner over time. After removing the steps these variables are typically flat or stationary about certain levels. As such, they impact outcome variables only temporarily. On the other hand a policy that switched on and off at different levels of intensity could produce nonstationarity in mean.

[5] As shown in Figure 3 there were staggered increases in the Federal vaccine mandate, whereas the vaccine lottery had been in decline continuously after July in 2021.

M1 or M2, a first difference regression can account for the nonstationary process in C1. C2, however, is a more realistic formulation that has been used by labor and health economists in situations a stationary policy affects the long run behavior of outcome variables. For example, Abadie et al. (2010) showed that California's tobacco control program changed the slope of the trend in tobacco consumption. In such situations first difference regression fails, as shown in Section 2.2 below. An alternative solution is provided in Section 3.

2.1 Policy influence only on levels: C1

This model has been popularly considered in the treatment effects literature. Without a homogeneous trend assumption, difference-in-difference estimation fails to identify the treatment effects. When a policy variable x_{it} is stationary in a model such as (5), it is not a long run determinant of the outcome variable and impacts the outcome only temporarily.

Consider the following TWFE regression frequently used in applied research:

$$y_{it} = a_i + \theta_t + \beta x_{it} + v_{it}, \tag{7}$$

where θ_t is a time fixed effect and the regression error v_{it} in (7) is the stochastic trend ξ_{it}. Let $\hat{\beta}_l$ be the TWFE estimator of β in (7). To simplify notation, rewrite (7) in differenced form as

$$\dot{y}_{it} = \beta \dot{x}_{it} + \dot{v}_{it}, \tag{8}$$

where $\dot{y}_{it} = y_{it} - \frac{1}{T}\sum_{t=1}^{T} y_{it} - \frac{1}{n}\sum_{i=1}^{n} y_{it} + \frac{1}{nT}\sum_{i=1}^{n}\sum_{t=1}^{T} y_{it}$, with similar definitions for the other variables. This within-group transformation does not alter the nature of nonstationarity but it does eliminate heterogeneous individual fixed effects and any homogeneous (over individuals) trend. The consequences of using this TWFE regression when the policy variable x_{it} is stationary arise from the fact that the regression is ill-balanced because of the stochastic trend effects that remain in $\dot{v}_{it}$ and are transmitted to $\dot{y}_{it}$. This lack of balance induces a reduction in the signal to noise ratio that affects the accuracy of the TWFE estimator, $\hat{\beta}_l$. The properties of $\hat{\beta}_l$ depend in this case primarily on the cross section sample size n, so that increases in the number of time series observations do not help to shrink the variance of $\hat{\beta}_l$.[6] More formally, as $n, T \to \infty$, when x_{it} affects y_{it} only in levels as in (5), the limit distribution of $\hat{\beta}_l$ is given by

$$\sqrt{n}(\hat{\beta}_l - \beta) \to \mathcal{N}(0, V_\beta), \tag{9}$$

with asymptotic variance V_β shown in the Online Supplement (Theorem 1). A panel robust variance

[6] In this regression, the signal from the sample variance of $\dot{x}_{it}$ is dominated by the sample covariance of $\dot{x}_{it}$ and $\dot{v}_{it}$ because the regression error is nonstationary. See the Online Supplement for further analysis.

estimator, clustering by time, can be used to consistently estimate V_β.

Next, consider the case of heterogeneous trends. Many dependent variables used in TWFE regressions can be expected to involve heterogeneous trends, a complication that affects asymptotic behavior and finite sample performance even when some of the regressors themselves include trends. As long as time-homogeneous coefficients are assumed in estimation, the regression error carries the effects of heterogeneous trends, leading to failure in standard limit theory and inference. To illustrate, suppose y_{it} is generated from

$$y_{it} = a_i + b_i t + \beta x_{it} + \xi_{it}, \tag{10}$$

where ξ_{it} is nonstationary.[7] Then the induced TWFE residual in (8) includes a trend term under the null $\beta = 0$, i.e.,

$$\dot{v}_{it} = (b_i - \frac{1}{n}\sum_{i=1}^{n} b_i)(t - \frac{1}{T}\sum_{t=1}^{T} t) + \dot{\xi}_{it}, \tag{11}$$

so that cross section variation of $\dot{v}_{it}$ grows at an $O(t^2)$ rate even when ξ_{it} is stationary. When x_{it} affects y_{it} in levels under heterogeneous trends, the asymptotic behavior as $(T, n) \to \infty$ of $\hat{\beta}_l$ are shown in Theorem 3 of the Online Supplement to have the following form

$$\sqrt{n}(\hat{\beta}_l - \beta) = O_p(T^{1/2}) + \mathcal{N}(0, V). \tag{12}$$

To examine the empirical implications of these asymptotic findings, we use log COVID-19 vaccination data across 50 US states – see the empirical section for a detailed explanation of the data. We also ran a Monte Carlo simulation with 50,000 replications to explore test size, giving rejection frequencies of the null hypothesis when the null is true. For each replication, a random policy variable x^*_{it} was generated from a standard normal distribution. The outcome variable is constructed as follows. First, we computed the cross section average $\hat{\mu}_t = \frac{1}{n}\sum_{i=1}^{n} y_{it}$ of the log vaccination rates across states. Next, we generated pseudo vaccination rates as

$$y^*_{it} = \begin{cases} a_i + \hat{\mu}_t + \xi_{it}, & \text{for M1}, \\ a_i + b_i\hat{\mu}_t + \xi_{it}, & \text{for M2}, \end{cases} \tag{13}$$

where $a_i \sim_{iid} \mathcal{N}(0,1)$, $b_i \sim_{iid} \mathcal{N}(1,1)$, $\xi_{it} = \xi_{it-1} + u_{it}$ with $u_{it} \sim_{iid} \mathcal{N}(0,1)$ and $\xi_{i1} = u_{i1}$. Test power was computed from the rejection frequencies of the null when the null is false using the

[7] The asymptotic results do not change even when ξ_{it} is stationary under heterogeneous trends

Table 1: Monte Carlo Simulation Results of Level Effects with Nonstationary Errors

		$\hat{\beta}_l$		$\hat{\beta}_{\mathrm{fd}}$	
Size (10%)	T	Hom	Het	Hom	Het
	10	0.112	0.113	0.116	0.114
	15	0.113	0.113	0.113	0.113
	20	0.112	0.111	0.114	0.113
	25	0.114	0.111	0.110	0.113
	40	0.113	0.111	0.114	0.114
Power	10	0.492	0.190	0.909	0.911
	15	0.519	0.174	0.981	0.981
	20	0.526	0.160	0.996	0.997
	25	0.533	0.154	1.000	0.999
	40	0.543	0.137	1.000	1.000

following data generating process:

$$y_{it}^* = \begin{cases} a_i + \hat{\mu}_t + \beta x_{it}^* + \xi_{it}, & \text{for M1}, \\ a_i + b_i \hat{\mu}_t + \beta x_{it}^* + \xi_{it}, & \text{for M2}, \end{cases} \tag{14}$$

where $\beta = 0.1$. In all replications $n = 50$ and $T \in \{10, 15, 20, 25, 40\}$.

Table 1 shows simulation findings based on (13) and (14). Under M1 (homogeneous trend) and M2 (heterogeneous trend), the size of the test seems satisfactory even for $T = 10$. The main difference is found in test power, as asymptotic results predict. As T increases, the power of the test based on $\hat{\beta}_l$ decreases under M2, but fluctuates under M1.[8]

The following first difference (FD) regression can be used to restore test power:

$$\Delta y_{it} = a_i + \Delta\theta_t + \beta \Delta x_{it} + e_{it}, \tag{15}$$

or the two-way transformed form given by

$$\Delta \dot{y}_{it} = \beta \Delta \dot{x}_{it} + \dot{e}_{it}. \tag{16}$$

[8]Since our simulations replicate actual vaccination rates using a random ensemble, the sample sizes n and T cannot be raised greater than their actual empirical sizes ($n = 50, T = 40$). However, as shown in the supplementary appendix, results from standard Monte Carlo simulations mirror the asymptotic theory. As n increases, size becomes more accurate and power improves; but as T increases, size changes little. For M1, test power based on $\hat{\beta}_l$ fluctuates with T for a given n. For M2, test power actually decreases as T increases for a given n.

Importantly, the FD regression includes two way fixed effects. The individual fixed effect a_i captures potential heterogeneous trend coefficients, while the time fixed effect $\Delta\theta_t$ captures common time changes. The regression error e_{it} is now stationary. Let $\hat{\beta}_{\text{fd}}$ be the TWFE estimator in (16). It is straightforward to show that as $T, n \to \infty$,

$$\sqrt{nT}(\hat{\beta}_{\text{fd}} - \beta) \to \mathcal{N}(0, V_{\beta_{\text{fd}}}), \tag{17}$$

where $V_{\beta_{\text{fd}}}$ is the asymptotic variance of $\hat{\beta}_{\text{fd}}$, which is smaller than V_{β_l} in (9). This limit theory is only valid it is known that x_{it} affects only the level of y_{it}. In practice, it may be challenging to assess whether x_{it} affects the level or the trend behavior of y_{it} or both. For when the outcome variable itself has a trend, any level change is temporary and permanent effects or long run impacts are inevitably influenced by the trend or trend coefficients. It may therefore appear reasonable to assess how a policy change affects growth rates rather than level outcomes, suggesting a TWFE regression in differences as in (2). However, if the outcome variable has nonlinear or time-variable trend behavior then TWFE regression in differences also fails to capture the impacts of a policy change. This issue is now discussed.[9]

2.2 Policy influence on trends: C2

For further analysis of potential pitfalls in TWFE estimation consider the generating process (DGP)

$$y_{it} = a_i + b_{it}t + \xi_{it}, \text{ with } \xi_{it} = \rho\xi_{it-1} + e_{it}, \tag{18}$$

where intercepts a_i affect the level of y_{it}, the b_{it} are time varying trend coefficients, and the ξ_{it} are time series[10] with a unit root when $\rho = 1$. Now suppose that a policy variable x_{it} influences b_{it}. In this event, regressions either in first differences or growth rates of y_{it} fail to deliver a satisfactory proxy for the temporal impact b_{it}. For instance, taking differences of (18) gives

$$\Delta y_{it} = b_{it}t - b_{it-1}(t-1) + \Delta\xi_{it} = b_{it} + \Delta b_{it}(t-1) + \Delta\xi_{it}, \tag{19}$$

so that Δy_{it} differs from b_{it} due to potential trend effects from both $\Delta b_{it}(t-1)$ and $\Delta\xi_{it}$. Hence, regression of Δy_{it} on fixed effects (both individual and time specific) and x_{it} will produce misleading findings about the impact of policy because of the missing trend and policy effects in the component $\Delta b_{it}(t-1)$.

For a specific example consider a case where there are two different sub-groups $\mathcal{G}_1$ and $\mathcal{G}_2$ in the

[9] A more technical discussion is given in the Online Supplement.

[10] If the DGP (10) changes to (18), the asymptotic effects are similar but the reported simulation results deteriorate and these now depend on the variance of b_{it}.

panel. In $\mathcal{G}_1$, y_{it} has a constant growth rate b_1 over time; in $\mathcal{G}_2$, y_{it} initially has a lower growth rate but after some threshold point (τ), y_{it} begins to catch up with or diverge from $y_{it} \in \mathcal{G}_1$. Formally,

$$b_{it} = \begin{cases} b_1 & \text{if } i \in \mathcal{G}_1, \\ b_2 & \text{if } i \in \mathcal{G}_2, \ \& \ t \leq \tau_i \\ b_3 + d(t-\tau_i)^{-\alpha} & \text{if } i \in \mathcal{G}_2, \ \& \ t > \tau_i \end{cases}. \tag{20}$$

Introduce a policy variable x_{it} defined by

$$x_{it} = \begin{cases} 1 & \text{if } i \in \mathcal{G}_1, \\ 0 & \text{if } i \in \mathcal{G}_2, \ \& \ t \leq \tau_i \\ 1 & \text{if } i \in \mathcal{G}_2, \ \& \ t > \tau_i \end{cases} \tag{21}$$

Key parameters in this specification are α and b_3. If $\alpha < 0$, regardless of the value of b_3, $b_{it} \in \mathcal{G}_2$ diverges, and $y_{it} \in \mathcal{G}_2$ diverges also. But when $\alpha > 0$, $b_{it} \in \mathcal{G}_2$ converges to b_3 irrespective of the initial value b_2 and, as $t \to \infty$,

$$b_{it} \to b_3 \text{ if } i \in \ \mathcal{G}_2. \tag{22}$$

When $b_3 = b_1$ the trend coefficients are eventually homogeneous.

Figure 1 gives illustrations. Ignoring random innovations in the DGPs, we generate $y_{1t} = 0.5t$ in $\mathcal{G}_1$, and consider two versions of $\mathcal{G}_2$: the first is a convergent panel with $y_{2t} = [0.5 + 0.2(t-50)^{-0.5}]t$; and the second is a divergent panel with $y_{3t} = [0.4 - 0.2(t-50)^{-0.5}]t$. Observe also that relative convergence also applies, viz., $y_{2t}/y_{1t} = 1 + \frac{0.2}{(t-50)^{0.5}} \to 1$ as $t \to \infty$ and $b_{2t}/b_{1t} = 1 + \frac{0.4}{(t-50)^{0.5}} \to 1$ as $t \to \infty$.

As shown in Panel (a) of Figure 1, the time path of y_{2t} initially diverges from y_{1t} but after $t = 50$ begins to catch up with y_{1t}. Meanwhile, y_{3t} diverges except at $t = 51$. The problem occurs when the first difference is taken and the dependent variable is Δy_{it}. Panel (b) shows the time path of Δy_{it}. At $t = 51$, both Δy_{2t} and Δy_{3t} have high spikes. For $t > 51$, both series decrease. Due to these spikes, the values of the coefficients are positive and slowly converge to constants that depend on the value of b_3.

Now assume that the following TWFE regression is employed to analyze the policy impact

$$\Delta y_{it} = a_i + \theta_t + \beta x_{it} + u_{it}, \tag{23}$$

where x_{it} is given in (21). Note that the independent variable here is not Δx_{it} but x_{it} because the policy variable directly affects the trend coefficient. Since Δy_{it} is not a straightforward linear regression because of the dummy variable temporal shift in x_{it}, a conventional TWFE regression is unsuited to capture the policy effects of x_{it}. In particular, the relationship between Δy_{it} and the

Figure 1: Time paths of y_{it} and Δy_{it}

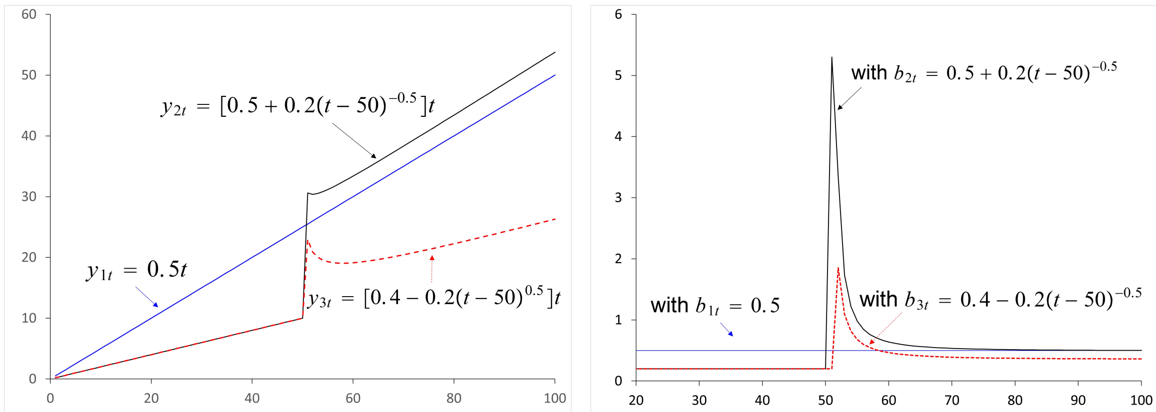

(a) y_{it} (b) Δy_{it}

Notes: Here $y_{1t} = 0.5t$, $y_{2t} = [0.5 + 0.2(t-50)^{-0.5}]t$, and $y_{3t} = [0.4 - 0.2(t-50)^{-0.5}]t$. As $t \to \infty$, y_{1t} and y_{2t} converge to each other, but y_{1t} diverges from y_{3t}.

impact of x_{it} is necessarily time varying with a regression coefficient that can change over time. Figure 1 provides a visualization of this problem, where y_{it} is generated deterministically as detailed in the figure and (23) is fitted by linear regression.[11]

Figure 2 displays how the true coefficients change over time. Here, with time series data Δy_{1t} and Δy_{2t}, the coefficient $\beta > 0$ for all t and seems to converge to 0.33. Meanwhile, with Δy_{1t} and Δy_{3t}, β seems to converge to 0.2. These examples show the limitation of TWFE regressions with differenced data Δy_{it}. Since the TWFE regressions with first difference outcomes are estimating the effectiveness of a policy variable in the short run, it is natural that $\beta > 0$ with (y_{1t} and y_{3t}) in this particular example. However, the issue is that when b_3 in (20) takes on a small value,[12] the true value of β becomes negative in the long run. Thus, such TWFE regressions are unable to produce meaningful values of the true policy effect β.

In sum, TWFE regressions formulated with first differenced outcomes are not suited to evaluate the effectiveness of policy changes if y_{it} involves non-linear trend effects. In such circumstances the problem of how to evaluate the effectiveness of policy changes is of considerable empirical interest. Two solutions are discussed in the next section.

[11] The data y_{it} in this case are generated deterministically as in the example of Figure 1 with no random components and least squares regression is used to fit the coefficients.

[12] For example, let $b_3 = 0.15$. The coefficient β is initially positive until $t < 65$, but then becomes negative and seems to converge to -0.045.

Figure 2: Time paths of the true coefficient values of β in the two panels

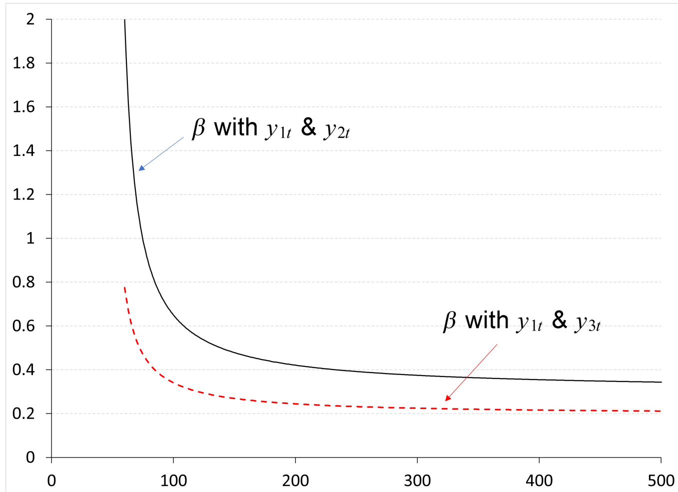

Notes: Data are generated according to the deterministic relations $y_{1t} = 0.5t$, $y_{2t} = [0.5 + 0.2(t-50)^{-0.5}]t$, and $y_{3t} = [0.4 - 0.2(t-50)^{-0.5}]t$ for $t = 1, \cdots T$, and then first differences are taken. Next, the fixed effects are eliminated by removing time series averages of Δy_{it}. The common time effects are eliminated by removing cross-sectional averages in each sub-panel. Then, the least squares is used to calculate (not estimate) the value of β.

3 A Dynamic Clustering Approach

This section describes a clustering approach to the evaluation of policy impacts in the presence of trending outcomes. The first part introduces the relative convergence test proposed by P-S, which can be used to test for a common nonlinear or stochastic trend in outcome data. If a panel of interest shares a common nonlinear or stochastic trend, then the TWFE regression in (16) is valid since x_{it} cannot influence the homogeneous trend coefficient heterogeneously.[13] The second part provides a methodology for transforming nonstationary panel data into stable multinomial club membership using a recursive clustering algorithm.

3.1 Testing for Homogeneous Trends

As in P-S, the starting point is to represent trending multidimensional data in terms of a panel components model as

$$y_{it} = b_{it}\theta_t, \tag{24}$$

where θ_t is an unknown trend which may be either a stochastic process or a nonlinear deterministic time trend. The representation in (24) is general and typically has unidentified multiple components. For instance, if the DGP were $y_{it} = a_i + b_i t + \xi_{it}$ where $\xi_{it} = \xi_{it-1} + e_{it}$, then we could

[13]Identification of the long run determinant then becomes of interest because in this case x_{it} does not affect the long run trend of y_{it}. Analysis of this case is a separate matter and left for future research.

rewrite the model in the form of (24) as $y_{it} = (a_i t^{-1} + b_i + \xi_{it} t^{-1})t = b_{it} t$.

If $y_{it}/y_{jt} \to 1$ as $t \to \infty$, then we say that y_{it} relatively converges to y_{jt} over time. Let $\hat{\mu}_t$ be the sample cross section average of the y_{it}. If $y_{it}/\hat{\mu}_t \to 1$ as $t \to \infty$ for all i, then the panel y_{it} is said to be relatively convergent to its cross section average. The ratio $h_{it} := y_{it}/\hat{\mu}_t$ traces out a *transition path* over time that manifests convergence when $h_{it} \to 1$. In this case, y_{it} shares the same (stochastic) trend, which is factored out in the ratio, thereby enabling analysis and inference about convergence. The test for relative convergence in P-S relies on the following (so-called $\log t$) regression

$$\log \frac{H_1}{H_t} - 2\log\left(L\left(t\right)\right) = a + b \log t + e_t, \tag{25}$$

which is estimated by ordinary least squares and where

$$H_t = \frac{1}{n}\sum_{i=1}^{n}\left(h_{it} - 1\right)^2, \text{ for } h_{it} = \frac{y_{it}}{\hat{\mu}_t}, \tag{26}$$

$L(t) = \log t$, $t = p+1, \ldots, T$, $p = \lfloor r \times T \rfloor$ with $r = 1/3$, and $\lfloor \cdot \rfloor$ is the integer floor function.

Under the null of relative convergence, H_t is asymptotically convergent to zero over time since $h_{it} \to 1$ as $t \to \infty$. Hence, $\log(H_1/H_t)$ is increasing over time. If the t-value for $\hat{b}$ exceeds -1.65, then the null of relative convergence is not rejected in the test at the 5% level. Note that in finite samples the term involving $2\log\left(L\left(t\right)\right)$ serves as a penalty function in the regression (25), as explained in P-S. Under relative divergence, H_t and $\log(H_1/H_t)$ should increase and decrease over time, respectively. Under fluctuations over time, H_t simply fluctuates, but in view of the penalty function of $-2\log(L(t))$, the dependent variable in (25) decreases over time. Hence, the fitted OLS coefficient $\hat{b}$ becomes significantly less than zero in this case.

For present purposes in the empirical evaluation of policy effects under trending outcomes, if the null of convergence is not rejected, then the TWFE regression with Δy_{it} is well justified since in the long run the panel y_{it} is identified as having a homogeneous (stochastic or nonlinear) trend. If the null is rejected, then data analysis may be conducted as described in the next section.

3.2 Dynamic Clustering Mechanism and Panel Logit Regression

One possible outcome is that there are few sub-group convergences, but each of these clubs diverges from the others over time, in which case a null of overall club convergence would be rejected. P-S suggested how to find convergent sub-groups by using an *convergence clustering mechanism* (CCM). This mechanism transforms the full $(n \times T)$ panel dataset into a club membership structure that features each individual member $(n \times 1)$. The CCM requires finding a core convergence club within the panel. Once a core club is identified, each individual time series is compared with the core group and is added to the convergence group if it relatively converges. Otherwise, the individual is classified to another group. Successive repetition of this procedure identifies members of the first convergent club. The clustering algorithm is then repeated with non-members of the first convergent club. The approach allows empirical researchers to explore the underlying determinants of club membership through multinomial logit regression of club membership on driver variables, as suggested in Phillips and Sul (2007b) and Phillips and Sul (2009).[14]

The present paper utilizes this approach to design a robust method of clustering club membership over time. The proposed method is straightforward and involves recursive implementation of the CCM algorithm over time to identify the clusters and cluster evolution over time. As we will show in the next section, various patterns of dynamic evolution over time can be identified by recursively estimating club memberships in this way. This dynamic version of the CCM approach employs some modifications of the original algorithm including a fixed rule for initialization in the recursive regressions[15] and a fixed rule for core member detection.[16]

The asymptotic justification for the clustering method is given in Appendix C of P-S. As long as the number of core members is not large relative to the time series sample, consistency of the clustering mechanism is easily achieved. Consistency of the dynamic version of the CCM approach can be achieved in the same way. The Online Supplement provides details of this method and reports findings of its finite sample performance from Monte Carlo simulations.

If we knew the specific functional form of dependence and values of b_{it} it would be straightforward to evaluate its determinants by running a regression of b_{it} on the relevant function $f(z_i, \theta_{x,t}, x_{it})$. In practice, finding a specific functional form is challenging and consistent estimation of all of the idiosyncratic trend coefficients b_{it} is not possible using only the data y_{it}.[17] To avoid these issues, we utilize the clustering method in P-S, as now explained.

Define J as the number of convergent sub-groups: $j = 1, \cdots, J$. The original algorithm was designed to provide club memberships based on descending order of the final observation values, y_{iT}. Hence, the first convergent sub-group always dominates the remaining convergent sub-groups. Define $\hat{C}_{it}$ as the estimated membership emerging from the application of dynamic CCM for the ith individual from 1 to t. This transformation changes nonstationary outcomes y_{it} to stationary club memberships, $\hat{C}_{it}$. If the membership does not change over time, one does not need to run panel ordered logit, but simply run ordered logit with the final club memberships at time T. Otherwise,

[14] See Sul (2019) for more detailed discussion.

[15] The $\log t$ test in P-S requires initialization of the regressions, eliminating some early observations. The discard rule in P-S removed the first 1/3 observations. This rule is problematic in the present implementation because the sample size changes in recursive regression. Instead, a fixed rule is used here in which the first 5 or 6 observations are discarded. The Online Supplement provides further discussion.

[16] The CCM algorithm estimates the initial core members based on the sample observations. To maintain the core membership in the recursive approach, the core members are fixed in the recursion by employing the entire sample in their initial detection.

[17] P-S simplifies this challenge by approximating b_{it} using a relative transition curve h_{it} as explained earlier in (26). However, in this approach the sample cross-sectional mean is often not a robust measure and can be sensitive to outliers, so it is not used here.

the next panel ordered logit regression can be run. The j-th ordered logit model is given by

$$\text{Prob}(\hat{C}_{it} > j) = \frac{\exp(a_j + z_i'\gamma_j + \theta_{x,t}'\lambda_j + x_{it}'\beta_j)}{1 + \left[\exp(a_j + z_i'\gamma_j + \theta_{x,t}'\lambda_j + x_{it}'\beta_j)\right]}, \text{ for } j = 1, \cdots, J-1, \tag{27}$$

where $\theta_{x,t}$ is a vector of known common policy variables, such as macro factors including market interest and inflation rates, and federal policy changes. Once the known common factors are included in the regression, panel fixed effects cannot be identified.[18] Note that all variables on the right hand side influence the trend coefficients in y_{it}. Variables not affecting the trend behaviors of y_{it} must not be significantly different from zero.

In the case of two sub-convergent clubs, as is the case in the empirical study of the next section, instead of a panel ordered logit regression one needs to run a panel logit regression with random effects. In this case, $\hat{C}_{it} = 1$ or 0, and the ordered logit regression becomes

$$\hat{C}_{it} = 1\{a + z_i'\gamma + \theta_{x,t}'\lambda + x_{it}'\beta + e_{it} \geq 0\}, \tag{28}$$

Note that neither conditional logit nor ordered logit regressions can identify λ since $\theta_{x,t}$ is common across individuals so that the conditional likelihood function eliminates $\theta_{x,t}'\lambda$ automatically.[18]

The economic interpretation of β in (27) and (28) is different from that of conventional TWFE regression. Since the dependent variable is club membership, the marginal effect of x_{it} becomes of interest, which indicates the change in probability when x_{it} increases by one unit. The unconditional logit regression in (28) also provides potentially complex causal effects to explain club membership. To see this, rewrite (28) with a single variable, ignoring x_{it} and e_{it}. Further, let $\gamma = \lambda = 1$ and $a = 0$. If the i-th state is in the first convergent sub-group at time t, then assign $\hat{C}_{it} = 1$; if it is in the second convergent sub-group, then assign $\hat{C}_{it} = 0$. Assume that $\theta_{x,t}$ is a federal policy which implements after $t \geq \tau$ (so $\theta_{x,t} = 0$ if $t \leq \tau$, otherwise $\theta_{x,t} = 1$), and z_i is a particular state characteristic variable. Then, depending on the values of z_i, the federal policy may influence the club membership differently.

Now consider a simple example where there are two types of individuals: the first makes a personal choice with little attention to others or any group behavior. The second type does consider other behavior before making a decision. Let $\zeta_{j,i,t}$ be a choice or outcome made by the $j-$th individual in state i at time t and allow for three time periods (1, 2, 3).[19]

[18] Consider, for example, the following conditional logit model with a single common factor and a single policy variable with two individuals for notational convenience : $\hat{C}_{it} = 1\{a_i + \lambda\theta_{x,t} + \beta x_{it} + e_{it} \geq 0\}$ with $i = 1, 2$. The conditional probability at time $t = 1$ becomes $\frac{\exp(\lambda\theta_{x,1} + \beta x_{11})}{\exp(\lambda\theta_{x,1} + \beta x_{11}) + \exp(\lambda\theta_{x,1} + \beta x_{21})} = \frac{\exp(\beta x_{11})}{\exp(\beta x_{11}) + \exp(\beta x_{21})}$. So λ cannot be identified with observed $\theta_{x,t}$.

[19] This example is purely illustrative because in practice it is unrealistic to estimate club memberships with such a small T.

Consider the choice function for the first type of individual

$$\zeta_{j,i,t} = 1\{z_{j,i} + \theta_t \geq 0\} \tag{29}$$

With no federal intervention (viz., $\theta_t = 0$ for all t), individual j's choice depends on the time invariant variable $z_{j,i}$. Assume that in state i the $z_{j,i} \sim_{iid} \mathcal{N}(z_i, 1)$. Define $Y_{it} = (\sum_{j=1}^{J_i} \zeta_{j,i,t})(J_i)^{-1}$ and $y_{it} = \log Y_{it}$, where J_i is the population of the i-th state. Then outcome y_{it} depends on z_i with higher z_i leading to higher y_{it}. Now introduce a federal intervention in period 2 so that $\theta_1 = 0$ and $\theta_2 = \theta_3 = 1$; and the outcome variable jumps at $t = 2$ but does not increase at $t = 3$.[20] Club membership will depend on the z_i and θ_t. If $z_i < 0$ and $z_i + \theta_2 < 0$, then $C_{i,2} = 0$. But if $z_i > 0$, then θ_t does not alter club membership.

Consider next the second type of individual with choice function

$$\zeta_{j,i,t} = 1\{z_{j,i} + \theta_t + \vartheta_j y_{i,t-1} \geq 0\}, \tag{30}$$

where ϑ_j is an individual reaction parameter measuring the response to a group decision in the previous period. Depending on the fraction of the second type of individual and the values of the ϑ_j, each state's reaction changes over time. A greater number of individuals of the second type and a higher ϑ_j both lead to a larger increase in outcome at $t = 3$, with the precise response and dynamic path of y_{it} dependent on multiple factors including the determinants of the choice function of this type of individual.[21] While it may be hard to statistically model the dynamic path of y_{it}, club membership itself is well determined and free of the precise nature of the individual response functions. When club membership can be empirically determined, club membership and club convergence are also estimable, opening up the use of the techniques of the present paper.

Importantly, the proposed methodology is designed to identify determinants of club memberships rather than estimate overall treatment effects. If y_{it} includes a nonlinear trend and exogenous policy variables cause changes in the trend behavior, the proposed method can identify the relevant exogenous variables but not estimate the overall treatment effects on y_{it}. But it is possible to deduce relative information about the implied treatment effects. For instance, take the case of two convergence sub-groups and define $n_{\mathcal{G}_{a,t}}$ as the number of individuals in $\mathcal{G}_a$ at time t for $a \in \{1, 2\}$.

[20] However, if θ_t were to increase over time, reflecting increased Federal intervention, then y_{it} would increase.

[21] Without correct specification of the determinants of the choice functional form, statistical modeling of the dynamic path of y_{it} is difficult and subject to misspecification, thereby affecting the properties of TWFE estimation including consistency.

Then the average outcomes for each sub-group can be estimated directly from the data by

$$\hat{\mu}_{\mathcal{G}_1} = \frac{1}{n_{\mathcal{G}_{1,t}}} \sum_{i \in \mathcal{G}_1} \frac{1}{T - \tau_i - 1} \sum_{t \geq \tau_i, \hat{C}_{it}=1}^{T} y_{it}, \tag{31}$$

$$\hat{\mu}_{\mathcal{G}_2} = \frac{1}{n_{\mathcal{G}_{2,t}}} \sum_{i \in \mathcal{G}_2} \frac{1}{T - \tau_i - 1} \sum_{t \geq \tau_i, \hat{C}_{it}=2}^{T} y_{it}. \tag{32}$$

The difference between these two averages provides an estimate of the difference between the two average overall outcomes, thereby giving information about the relative impact of the two treatments. If these two convergence sub-groups eventually merge into a single convergence group from some point $t \geq \tau$, then it is known that the treatment outcomes effectively become the same from this point.

We note finally that in any convergence sub-group empirical analysis some individuals may end up outside any of the identified groups. Such individuals are treated in P-S as outliers or divergent members of the population. In this case the relevant individuals display behavior outside the framework of the convergence analysis and these outcomes may need separate empirical study to explain their behavior.

4 COVID-19 Vaccination in the US

4.1 Data Preparation: some practical considerations

Let $\{Y_{it}\}$ be raw panel observations. If the Y_{it} appear to grow exponentially over time at an approximately constant growth rate it is a common convention to use logarithms of the raw data. In practical work a few observations may take zero or negative values, preventing the use of this transformation. If no measurement errors are suspected then these observations may be important and ignoring them may be consequential. In this case the following modification can be useful.

First, suppose that $Y_{it} \geq 0$. Define $Y_{it}^{+} = Y_{it} \times 10^{\alpha}$ where α is a large constant, noting that

$$\log(Y_{it}^{+} + 1) \simeq \log Y_{it} + \alpha \log 10 \;\; \text{if } Y_{it} > 0, \tag{33}$$

whereas

$$\log(Y_{it}^{+} + 1) = 0 \;\; \text{if } Y_{it} = 0, \tag{34}$$

This transformation does not alter the nature of the data for regression purposes as long as either time or individual fixed effects are included in the regression. To see this, let $Y_{it} = \exp(a_i + y_{it}^{*})$ if

$Y_{it} > 0$ and set $y_{it} = \log(Y_{it}^{+} + 1)$. Then

$$y_{it} - \frac{1}{T}\sum_{t=1}^{T} y_{it} = y_{it}^{*} - \frac{1}{T}\sum_{t=1}^{T} y_{it}^{*},$$

so that any fixed effects a_i are effectively eliminated in regression with the transformed data.[22]

If some Y_{it} take small negative values, then setting these observations to zero enables use of the above transformation and removes the difficulty. An alternative approach is to define the following transform

$$Y_{\min}^{+} = \min_{1\le t\le T, 1\le i\le n} Y_{it}^{+}, \tag{35}$$

and instead of adding unity to Y_{it}^{+}, add $Y_{\min}^{+}+1$. Then, it is easy to see that $\log(Y_{it}^{+}+Y_{\min}^{+}+1) = 0$ if $Y_{it}^{+} = Y_{\min}$. Otherwise, $\log(Y_{it}^{+} + Y_{\min}^{+} + 1) \simeq 0$ if $Y_{it} \le 0$ but $\log(Y_{it}^{+} + Y_{\min}^{+} + 1) > 0$ if $Y_{it} > 0$. If $Y_{\min}$ is not a big number, this modification does not materially change the nature of the data for regression purposes.

Ratio variables may similarly benefit from logarithmic transforms in eliminating individual fixed effects by regression. For instance, if $Y_{it} = W_{it}/V_{it}$ where $W_{it} = \exp(a_i^w + w_{it})$ and $V_{it} = \exp(a_i^v + v_{it})$, then

$$Y_{it} - \frac{1}{T}\sum_{t=1}^{T} Y_{it} = \exp(a_i^w - a_i^v)[\exp(w_{it} - v_{it}) - \frac{1}{T}\sum_{t=1}^{T}\exp(w_{it} - v_{it})], \tag{36}$$

and the within group transformation – subtracting the time series averages in (36) – does not eliminate fixed effects. On the other hand, fixed effects in the logarithmic transforms of these variables are eliminated by subtracting out the time series averages.

4.2 Empirical Results

As an empirical illustration, our recursive club clustering methodology was applied to state-level COVID-19 vaccination rates in the US and panel logit regressions were employed to explore the impact of vaccination policies on actual vaccination rates. By late spring of 2021 COVID-19 vaccinations were widely available in the US, but vaccination rates began to plateau even though

[22] When some Y_{it} are equal to zero, it is common empirical practice for researchers to transform Y_{it} as $\log(Y_{it}+1)$ or $\operatorname{arcsinh}(Y) = \log(\sqrt{1+Y^2}+Y)$. As Chen and Roth (2023) point out, direct comparison with transformed data can cause difficulties. Let Y_{it}^{*} be a latent (or unknown) variable and Y_{it} be the observed variable such that $Y_{it} = Y_{it}^{*} \times 10^{-\gamma}$, where γ represents a unit of account. Units of Y_{it} could be a thousand, million, or billion, and the value of $\log(Y_{it}+1)$ would be different depending on the unit of Y_{it}. But in TWFE regressions, these difficulties can usually be avoided if a large constant α is employed. If α is set to zero, then $\log(Y_{it}^{+}+1) \ne \log Y_{it}$, in contrast to (34). Furthermore, depending on the units of measurement of Y_{it}, the simple transformation $\log(Y_{it}+1)$ can lead to inconsistent estimation particularly if the overall mean of Y_{it} is close to zero.

only roughly 45% of the targeted US population were fully vaccinated by mid-May 2021.[23] There was also substantial variation in state vaccination rates: Maine had the highest vaccination rate at 49% in mid-May 2021; and Mississippi had the lowest at the time, with only 26% of residents fully vaccinated.

Determinants of this variation in state vaccination rates are naturally of considerable interest to policy makers, epidemiologists, and social scientists. Some preliminary research conducted over the summer of 2021 pointed to partisanship having a strong association with vaccination rates. Specifically, it was found that the percentage of votes cast for Donald Trump in the 2020 presidential election was a primary predictor of vaccination rates: the higher the Trump vote, the lower the vaccination rate, on average. Around the same time in 2021, cities, counties, and states attempted to bolster their waning vaccination rates by implementing various vaccine incentive campaigns such as vaccine lotteries and cash for vaccination. Numerous studies have examined the efficacy of such incentives in various states and counties across the United States, and have come to differing conclusions. Some found modest increases in vaccinations resulting from vaccine lotteries or cash incentives, while others found no statistical evidence that these lotteries or cash incentives increased vaccinations, even finding small negative impacts in some cases. Table 2 provides reference details for some of these explicit findings in the literature.

Table 2: Findings of Vaccination Incentives

Small Positive Effect	Zero or Small Negative Effect
Barber and West (2022)	Chang et al. (2021)
Brehm et al. (2022)	Dave et al. (2021)
Sehgal (2021)	Lang et al. (2022)
Wong et al. (2022)	Thirumurthy et al. (2022)
	Walkey et al. (2021)

By late summer 2021, policies mandating vaccinations for particular sub-populations were being announced and implemented at the state and the federal level. To examine the impact of vaccination incentives and policies on vaccination rates, we assembled a dataset of vaccination policies and incentives at the state level, including policies that were implemented in large cities or counties within a state. Appendix A explains how the state policy dataset was constructed and appendix table 6.1 provides summary statistics for the various state-level policies.

In addition, we created a separate federal-level vaccine mandate variable. This variable includes information from a combination of vaccine mandate announcements that were national in scale.

[23]The term *fully vaccinated* was defined at the time as two doses of the Pfizer or Moderna vaccine, or a single dose of the Johnson & Johnson vaccine.

More specifically, it includes the military vaccine mandate, the various vaccine mandates announced by President Biden on September 9, 2021[24], and mandates by private employers that typically were national in scope. Details on the construction of the federal vaccine mandate variable are given in Appendix B, and appendix table 6.2 displays the relevant events and dates that were used.

Our state-level vaccination data came from the publicly available county-level data from the Centers for Disease Control and Prevention (CDC), spanning the period from December 13, 2020 to February 9, 2022. Data prior to May 12, 2021 was discarded because COVID vaccines were initially in short supply and difficult to obtain, meaning that discrepancies in vaccination rates across states during this early period may not have been voluntary but simply due to availability. By mid-May 2021 Covid-19 vaccines were easily accessible in most areas of the United States. Daily county-level data were converted to weekly state-level data and logarithms of the resulting vaccination rates were recorded. There were a few points of decreasing cumulative vaccination rates for a short period in a small number of states in the data, which was likely due to state or county reporting errors. To correct for these, we applied Stata's HP filter[25] with a smoothing parameter of 1600,[26] and then subtracted y_{min}, as suggested in Section 3.1 of the Online Supplement.

As Figure 3 shows, the implementation percentages of both federal-level mandates, and state-level mandates are very low through late July 2021, at which point they both sharply increase and then stabilize. State vaccination mandates stabilize around 0.45, meaning about 45% of states had some form of state-level vaccination mandate in place by September 2021. Federal vaccine mandates sharply increase over roughly the same time period, although the increase is more stairways than a single sharp jump, as is the case with state-level mandates. The federal-level mandates level off at an implementation rate of 100% since all states were impacted by this federal-level mandate. Vaccination lotteries also show a sharp increase, but the increase is several weeks earlier than the vaccine mandate increases. This shows that states initially tried to incentivize people to get vaccinated through positive incentives. After peaking in June of 2021, the use of lotteries to incentivize vaccination started to steadily fall away until converging at zero.

Cumulative vaccination rates are partial sum time series and therefore typically stochastically nonstationary. For the reasons explained earlier involving the effects of regression imbalance, these characteristics suggest caution in the use of random effects or fixed effects linear regression to

[24]Not all of the federal mandates announced by President Biden on September 9, 2021 were ultimately implemented. The mandate on large employers was struck down in court, and the mandate for employees of federal contractors was blocked for months before an August 31, 2022 announcement from the federal government that it would not be enforcing the mandate.

[25]The HP filter is by far the most commonly used filter in empirical studies that have employed the P-S CCM algorithm. We considered other filters in the empirical application with little changes in the results.

[26]Since the data are weekly, higher values than the quarterly smoothing parameter 1600 are sometimes preferred. In the present case our goal is to smooth the series only moderately because use of a smoothing parameter that is too large produces a filtered trend that is almost linear. Provided the smoothing parameter is neither very small nor very large, the empirical results were not sensitive to the specific choice. Details are provided in the online supplement.

Figure 3: log(Cumulative Vaccination Rate) and Select Policy Common Trends.

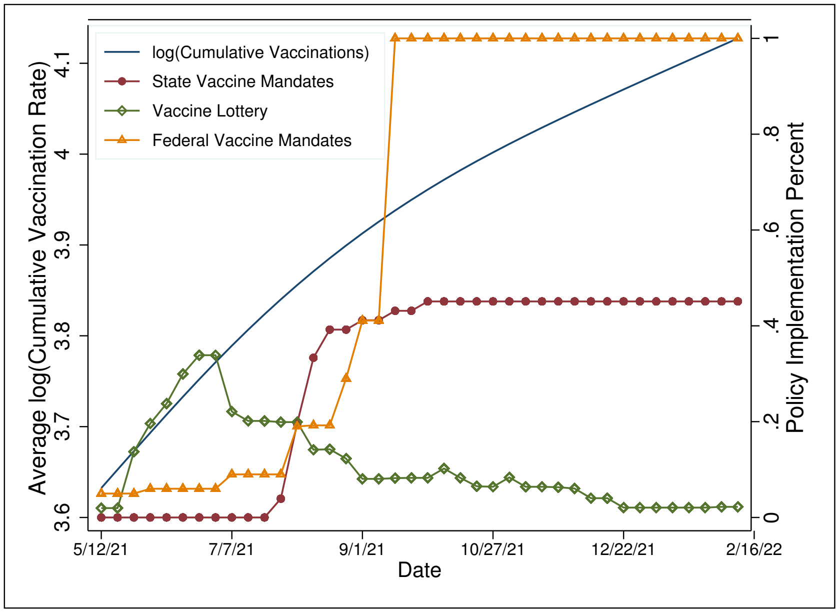

Notes: The left vertical axis scale is the logarithm of the national average cumulative vaccination rate, which corresponds with the smooth blue line. The right vertical axis scale is the adoption percentage of vaccination policies. The maroon, green, and orange lines are measured on the right axis, and each is the national average of the policy at each time t. The maroon and green lines are sample state-level vaccination policies, vaccine mandates for state employees and/or healthcare workers, and vaccination lotteries, respectively. The orange line is federal-level vaccine mandates. Log cumulative vaccinations do not show any jump or discontinuity over time and the path of this variable appears impervious to the policy variables being enacted.

assess the impact of vaccination policies on cumulative vaccination rates. Using first differences (new vaccination numbers) as the dependent variable does not resolve the imbalance when the data involve non-linear heterogeneous trends; and first difference specifications are less helpful in addressing the primary issue of modeling discrepancies in overall state vaccination rates.

To address heterogeneity and nonlinearity in the trend behavior of the data our empirical approach is to classify state vaccination rates into groups where homogeneous trends are manifest. The groupings were obtained by applying the P-S automatic clustering technique to log cumulative vaccination rates, y_{it}, producing individual club membership data $\hat{C}_{it}$, which is the estimated convergence club that state i belongs to at time t. As previously noted, $\hat{C}_{it}$ takes on the value of 0 or 1 when there are two convergence clubs. Also, as noted earlier, when vaccination rates began to plateau in each state a variety of vaccination incentive schemes and policies were announced and implemented. Examining club membership data at a single fixed point in time does not reveal the dynamic effects on club membership over time as these policies were rolled out. To explore the evolutionary relationship between club membership and individual state policy implementation we employed the dynamic club clustering technique described earlier in Section 3.

We first applied the automatic clustering mechanism on the full data set from May 12, 2021 to February 9, 2022, using $m = 6$, where m is the first sample observation used in the regression, a setting that matches findings from the simulations reported in Section 4 of the Online Supplement. The t-ratio from the initial $\log t$ regression of the entire data set was 7.24, so that the null of convergence was not rejected, implying that vaccination rates of all fifty states plus the District of Columbia were converging to the same long run national average in February, 2022. Implementation of the P-S clustering algorithm produced a core group of eight members: Connecticut, the District of Columbia, Maine, Maryland, Massachusetts, New York, Rhode Island, and Vermont. This club membership outcome is a static full sample result that is uninformative regarding the actual process of convergence and, in particular, the important empirical question of whether state convergence may have occurred without intervention or whether vaccination policies impacted state convergence over time.

To address this question dynamic clustering was employed with recursive sampling to estimate club membership evolution over time. Core membership was fixed to the aforementioned seven states plus the District of Columbia. From this core membership we applied the remaining steps of the clustering process using the first thirteen weeks of data from May 12, 2021 to August 4, 2021. Interestingly, with this shortened dataset involving vaccination rates only from the late spring to the mid-summer of 2021, there was no evidence of convergence to a single long run average. Instead, during the summer of 2021 there were two distinct clubs: one comprising thirteen members (twelve states plus the District of Columbia) with relatively high vaccination rates, and a second club consisting of 38 states with relatively low vaccination rates.

Figure 4 shows tracked behavior of each of these two clubs over time based on the first subsample from from May 12, 2021 to August 4, 2021. Panel A and B show the average relative transition paths in Club 1 and Club 2, respectively. Relative transition path is defined as $h_{it} := y_{it}/\hat{\mu}_t$ where $\hat{\mu}_t$ is the cross-sectional average of y_{it}. The average relative transition in Club 1 is defined as $n_{\mathcal{G}_1}^{-1} \sum_{i \in \mathcal{G}_1}^{n_{\mathcal{G}_1}} h_{it}$ where $n_{\mathcal{G}}$ is the number of individuals in Club 1. Note that there are 13 members in Club 1 and 38 states in Club 2. The average relative paths of Club 1 and Club 2 moved away from unity until July 2021, at which point they began to converge to unity over time. These paths reveal that overall relative convergence may be expected if the respective movements are sustained over time, with the average relative transition measures approaching unity for both clubs from different directions.

Working from the given initial club membership obtained for the original (May 12 - August 4) sample, a recursive analysis was commenced by adding a further week to the original sample, giving the new sample span from May 12, 2021 to August 11, 2021. The clustering algorithm was re-applied, using the same fixed core. This resulting outcome again produced two clubs, but Club 1 had all original thirteen members in the higher vaccination group that were in the May 12, 2021 - August 4, 2021 sample plus two additional states. This process was continued adding one week

Figure 4: Cross section averages and average relative transition curves based on initial club membership

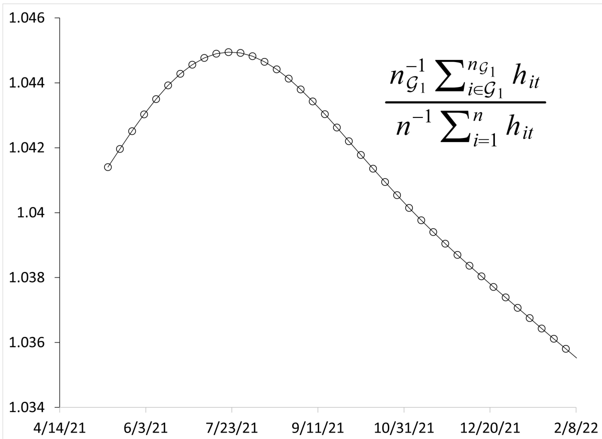

(a) Average Relative Transition in Club 1

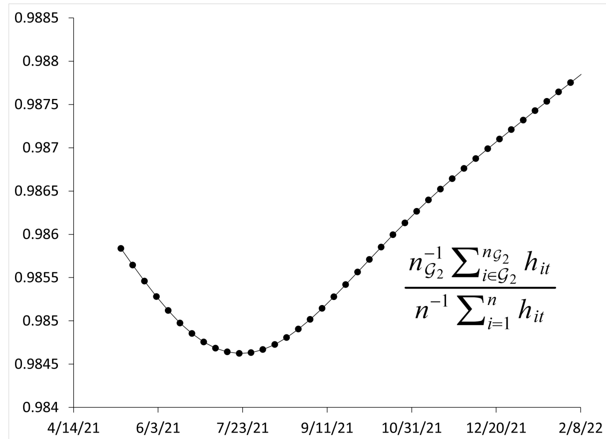

(b) Average Relative Transition in Club 2

Notes: Both Panel (a) and Panel (b) trace the average transition paths of state vaccination rates initially in Club 1 (i.e., those in Club 1 after the initial clustering process using data from May 12, 2021 to August 4, 2021). Panels (a) and (b) trace average relative transition paths for Club 1 and 2, respectively. The relative transition measure is each state's cumulative vaccination rate divided by the national average at each time t; and the average relative transition measure averages the individual measures over each club.

at a time and re-running the club clustering process to the end of the sample. With the addition of each additional week to the sample, the outcome produced more states joining the relatively high vaccination club, Club 1, each week. This pattern continued until the sample included data from May 12, 2021 - September 1, 2021. At this point, all of the states had relatively converged, forming a single convergence club. Relative convergence to a single club continued to hold for each additional week included until the recursion covered the entire sample, May 12, 2021 to February 9, 2022.

Figure 5 shows the dynamic membership evolution of Club 1 from May 12 through September 1, 2021. The top left panel shows the twelve states plus the District of Columbia in Club 1 from May 12, 2021 - August, 2021. Each subsequent panel shows the states belonging to Club 1 as the sample recursively expands. Club 1 membership evidently grows over time, as expected from Figure 4. The last panel, in the lower left position, shows that when the dataset includes weeks from May 12, 2021 through September 1, 2021, all the US states are seen to have the same club membership and full convergence applies.

Figure 6 shows federal-level vaccine mandate variables plotted alongside the evolving Club 1 membership. Two federal-level vaccine mandate variables are displayed in the graph: one includes private employer vaccine mandates and the other does not. By design the fraction of states in Club 1 held constant through to August 4, at which point it began to rise steadily for several weeks

Figure 5: Dynamic State Membership in Club 1

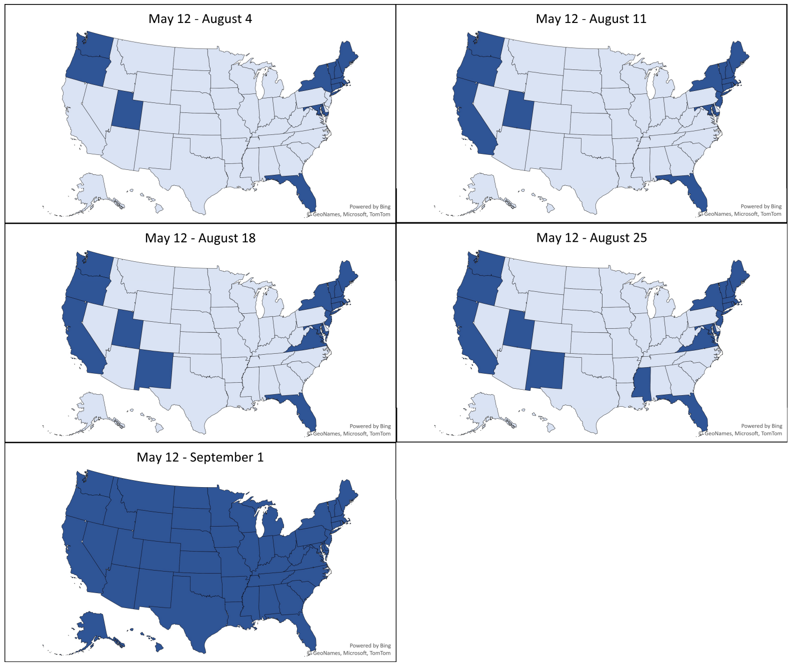

Notes: The states shaded dark blue are in Club 1 and the light blue states are in Club 2. As the sample added weeks the membership of Club 1 continued to grow until September 1, 2021, at which point Club 1 included all states giving a single convergence club.

merging into a single convergence club in September. The federal-level vaccine mandates started to increase slowly at first over May-July and subsequently rose rapidly through to September, at which time President Biden announced four federal-level mandates estimated to impact 100 million American workers. The fraction of state membership of Club 1 closely tracks the course of these federal-level mandates.

The federal mandates variable, the state policy dataset, and the dynamics of club membership offer the opportunity to explore the impact of federal-level mandate announcements and state-level policies on club membership. An unconditional panel logit (random effects) regression in (28) was used to examine some of the effects of these time-varying policies. We used the combined federal and employer mandates as a common factor θ_t, and various state specific variables as z_i including political trends, state demographic characteristics, such as population density, median household income, and education, and the percentage of people employed by industry for each state. Various state-level vaccination policy variables comprise x_{it}, including state incentives for vaccinations and state vaccination policies.

Table 3 column (5) shows the results of the preferred specification from the panel unconditional

Figure 6: National-Level Mandates and Dynamic Club 1 Membership

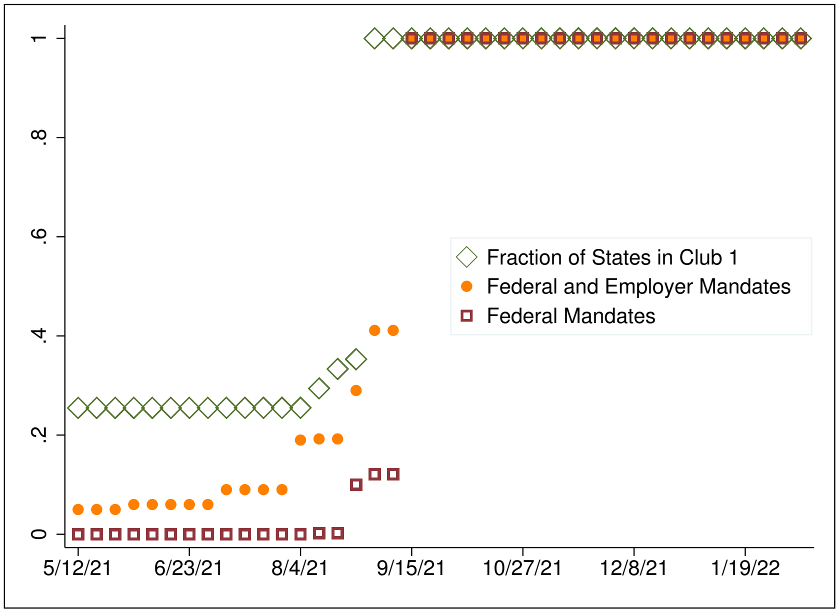

Notes: The federal mandates variable is plotted in maroon, this variable combined with employer mandated vaccinations is plotted in orange, and the fraction of states belonging to Club 1 is plotted in green. The fraction of states in Club 1 and the federal vaccine mandate variable are flat until early August 2021. The federal mandates variable, the one capturing both the federal mandates and the employer mandates, slowly increased in early summer 2021 before increasing dramatically from July to August 2021, and again from August to September 2021. All three variables end up at unity by September 9, 2021.

logit (random effects) regressions.[27] The large coefficient on the federal-level mandate variable shows that the probability of being in Club 1 given the federal-level mandate is extremely high.[28] This matches Figure 6 where transition to Club 1 moves very closely with the implementation of the federal-level mandates. Interestingly, no state-level vaccination policies or incentives had any significant impact on club membership. Also, contrary to our findings in the summer of 2021, by February 2022 our results suggest that once population density, median household income, the percentage of the population that is foreign-born, and the industry composition of a state are all controlled for, political party is not associated with club membership in a statistically significant way. States with a larger percent of foreign born individuals were more likely to be in Club 1 initially. States with higher numbers of health care and social assistance workers as well as higher percentages of people employed in the retail trade industry were also initially more likely to be in Club 1, whereas states with higher percentages of employees working in wholesale trades were less likely to be in Club 1. The McFadden pseudo R^2 value is very high for our preferred

[27] Table 9 in the online supplement reports results of other specifications used in the unconditional panel logit regressions.

[28] We also ran the model for column (5) from Table 3 with a two-week lag in each of the policy variables. Including the two-week lag did not impact the findings. See Table 10 in the Online Supplement.

specification, showing that the improvement in our model from the based intercept only specification is substantial, indicating the fitted model regression almost fully explains club membership.

For comparison with the above findings linear random effects regressions were run according to the following specifications using levels and differences as the dependent variable:

$$y_{it} = a + z_i'\gamma + \lambda\theta_t + x_{it}'\beta + \epsilon_{it}, \tag{37}$$

$$\Delta y_{it} = a + z_i'\gamma + \lambda\theta_t + x_{it}'\beta + \epsilon_{it}, \tag{38}$$

$$y_{it} = a_i + \lambda\theta_t + x_{it}'\beta + \epsilon_{it}, \tag{39}$$

$$\Delta y_{it} = a_i + \lambda\theta_t + x_{it}'\beta + \epsilon_{it}, \tag{40}$$

where y_{it} is the logarithm of state i's cumulative vaccination rate, Δy_{it} is the log of the number of new vaccinations per 10,000 people (the first difference of vaccinations), θ_t is the federal-level mandates, z_i are state fixed effects, and x_{it} is a vector of state-level policies and the number of new infections per 10,000 people in a state each week. The results are displayed in columns 1-4 of Table 3. The R^2 values show that the policies explain significantly more of the variation in cumulative vaccination rates than they do new vaccinations. The coefficients and levels of significance are very similar to the random effects model and fixed effects model for each of the dependent variables. For this reason we limit our discussion of the regression results here to the fixed effects model for both dependent variables.

When the log cumulative vaccination rate is the dependent variable, new infections and federal-level vaccine mandates are both positive and highly significant. But, state-level policies have no significant impact on cumulative vaccinations, with the exception of bans on proof of vaccination. That positive coefficient result suggests that if a state implemented a ban on proof of vaccination, cumulative vaccinations in that state would increase, which is a curious outcome that may be the spurious result of the nonlinear trend effects discussed earlier.

In the regressions with new (first differenced) vaccinations as the dependent variable there are several anomalous signs in the fitted coefficients. For instance, the signs on new infections, federal-level vaccine mandates, and state-level mandates on state employees are all negative and counter-intuitive as higher infection rates and vaccination mandates are more likely to increase than reduce new vaccinations. Further, the empirical results imply that the only state-level policy that increased new vaccinations per 10,000 people was a ban on mask mandates. It might be argued that banning mask mandates led people with high risk aversion to Covid-19 infection to get vaccinated because they felt less secure, but those people were already most likely to be vaccinated. As discussed earlier in Section 3, use of first differences does not eliminate time trend effects in the data and these counter-intuitive results are again the likely outcome of misspecification and failure

Table 3: Regression Results

	Linear Models				Logit Model
Dependent Variable:	y_{it}		Δy_{it}		$\hat{C}_{it}$
	(1)	(2)	(3)	(4)	(5)
	FE	RE	FE	RE	RE
New Infections per 10,000 People (x_{it})	0.005	0.005	-0.004	-0.004	
	(0.0001)	(0.0001)	(0.002)	(0.002)	
Federal Mandate and Employer Mandates ($\theta_{x,t}$)	0.251	0.252	-0.438	-0.442	101.0
	(0.012)	(0.012)	(0.053)	(0.053)	(22.39)
State Incentives (x_{it})					
Lottery	0.003	0.003	0.041	0.037	-1.468
	(0.012)	(0.013)	(0.047)	(0.046)	(2.038)
Cash	0.016	0.016	0.090	0.089	-0.084
	(0.011)	(0.011)	(0.049)	(0.050)	(3.405)
Community Outreach	0.028	0.029†	-0.043	-0.047	-2.440
	(0.015)	(0.014)	(0.051)	(0.045)	(3.372)
State Policies (x_{it})					
Vaccine Mandate State Employees	0.016	0.017	-0.199	-0.190	-1.919
	(0.012)	(0.011)	(0.056)	(0.055)	(2.258)
Indoor Vaccine Mandate	0.016	0.015	0.043	0.041	14.33
	(0.011)	(0.011)	(0.125)	(0.121)	(36.00)
Mask Mandate	0.007	0.007	-0.051	-0.051	2.096
	(0.012)	(0.012)	(0.094)	(0.090)	(3.567)
Ban on Proof of Vaccination	0.086	0.071	-0.221	-0.201	-3.164
	(0.020)	(0.016)	(0.129)	(0.099)	(5.166)
Mask Mandate Ban	0.013	0.006	0.193	0.129	-1.810
	(0.059)	(0.045)	(0.075)	(0.062)	(6.160)
Political (a_i)					
Percent of State House that is Republican		-0.550		-0.130	7.037
		(0.104)		(0.212)	(24.41)
Percent of Vote for Trump 2020		-0.042		-0.262	-13.66
		(0.073)		(0.104)	(10.53)
State Characteristics (a_i)					
Population Density		-0.010		-0.002	11.36
		(0.020)		(0.019)	(6.153)
Median Household Income		0.026		0.018	2.379
		(0.014)		(0.018)	(2.497)
Percent Foreign Born		-0.001		0.126	1.327
		(0.003)		(0.005)	(0.609)
Percent of People Employed by Industry (a_i)					
Health Care and Social Assistance		0.048		0.004	3.626
		(0.011)		(0.017)	(1.412)
Government and Government Enterprises		-0.006		0.002	-2.055
		(0.066)		(0.008)	(1.291)
Retail Trade		0.029		0.020	11.90
		(0.025)		(0.043)	(3.968)
Wholesale Trade		-0.015		-0.018	-13.07
		(0.036)		(0.058)	(5.901)
Transportation and Warehousing		-0.019		-0.017	-9.513
		(0.026)		(0.041)	(6.242)
n	51	51	51	51	51
T	40	40	40	40	40
R^2	0.832	0.808	0.664	0.632	
McFadden's R^2					0.936

Notes: Numbers in parentheses are standard errors. Median household income is measured in tens of thousands of dollars and population density is per 1,000 square miles. The dependent variable for the linear level model, y_{it}, is log(cumulative vaccination rate) and the dependent variable for the linear first-differenced model, Δy_{it}, is log(new vaccinations per 10,000 people). The binary club membership obtained from the P-S club clustering technique, $\hat{C}_{it}$, is the dependent variable in the logit regressions. The coefficient on the federal-level mandates in the unconditional logit model is large and all state level policies had no impact on the likelihood of being in the high vaccination club.

to capture separate group behavior in the data.

In sum, comparing the regressions results across Table 3, the panel logit regressions seem to provide the most plausible and intuitive findings. The natural explanation is that the panel logit models provide well specified formulations that take account of club membership arising from nonlinear trend effects and separate group convergence behavior that together determine cumulative vaccinations.

5 Conclusion

When outcome variables have nonlinear and possibly stochastic trends, evaluating the effectiveness of policy changes by using TWFE regressions can be problematic. This paper shows the underlying reasons for this empirical problem and proposes an alternative approach. The key idea is a simple method to transform panel nonstationary outcome data into panel multinomial data by using a dynamic clustering method based on the relative convergence test of Phillips and Sul (2007a). This approach allows researchers to use panel logit regressions to investigate how policies that are implemented can impact convergence club membership or the long-run behavior of the dependent variables over time.

The dynamic convergence clustering mechanism is applied to state-level Covid-19 vaccination rates in an empirical example of this methodology. Our findings indicate that there were initially two distinct convergence clubs, but over time all states converged to a single club. Finally, we use panel logit to show how national and state policies impacted club membership over time, and we demonstrate how the regression results from panel logit regressions appear to give more realistic results than linear models.

There are two drawbacks of the proposed method. First, the number of time series observations cannot be too small. To estimate time varying convergent club membership the time series sample size T should be large enough to capture the evolution of club membership. In the Covid-19 vaccination empirical example $T = 40$ observations were available for the time series sample and the recursive sampling procedure was initiated with sample size $T = 13$. This choice complies with the minimum sample size $T = 10$ for the clustering algorithm that was used in Phillips and Sul (2007a). When T is smaller we suggest not using the dynamic CCM but instead static CCM and running a cross-sectional logit or multinomial logit regression.

Second, to use the proposed method, the panel should be balanced since the dynamic CCM tracks each individual club membership over time. If some data are missing within time periods, then interpolation and filtering to smooth out the series can be employed. Since the proposed method is designed for analyzing long run effects, small modifications of this type typically do not affect the membership findings. If data are missing at the beginning or end of the sample, then backward or forward forecasting would be required and the accuracy of such modifications is not studied in the paper.

6 Appendix

6.1 Appendix A: State Policy Variables

We created a database that tracked state (and District of Columbia) announcements of vaccination lotteries, cash for vaccination incentives, community outreach programs, vaccine mandates for state employees and or healthcare workers, indoor vaccine mandates or mandates for gatherings over a certain number of people, mask mandates, bans on proof of vaccination, and bans on mask mandates. (Only mask mandates that were re-implemented after June of 2021 were included since virtually every state had some form of mask mandate at the beginning of the pandemic.) This database is weekly and tracks policies from March 2021 to February 2022. The policies are tracked from the date of their announcement. We also included polices that were implemented by large cities or counties since it was occasionally the case that a large city or county would implement a policy which impacted many people in a state, but the policy was not implemented at the state level. Chicago, for example, gave cash incentives for vaccination, but the state of Illinois did not. The population of Chicago makes up 20.9 percent of the population of Illinois, so for the weeks that Chicago offered cash incentives, we populated the cash field for Illinois in the dataset with a value of .209 rather than 1. We went state by state and gathered information about policies that were implemented and the timing of the policies. We found policy data from AARP (formally the American Association of Retired Persons), The National Governor's Association, Becker Hospital Report, The Rockefeller Foundation, The Kaiser Family Foundation, Ballotpedia, and various other websites. We also made a list of the two largest counties and cities in each state and used Google to search for any policies implemented in those localities. Any other local policies in cities or counties other than the two largest in each state that came up in our Google searches were included (as their percentage of state population) if the locality made up at least two percent of the state population.

Appendix Table 6.1: Summary Statistics for State Vaccination Policies

	Lottery	Cash	Com Out	VMSE	IVM	MM	BPV	BMM
Mean	0.108	0.067	0.053	0.301	0.050	0.170	0.388	0.190
Cross Section Median	0.025	0.004	0.000	0.000	0.000	0.028	0.000	0.000
Overall Median	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Cross Section Variance	0.022	0.013	0.087	0.112	0.019	0.072	0.218	0.149
Overall Variance	0.094	0.050	0.127	0.211	0.040	0.122	0.238	0.154
Number of States	33	29	16	23	11	33	21	10

Notes: The number of states listed in the final row is the total number of states that implemented the policy at some point. If a large locality within the state implemented the policy, that is included in the state count. Cross sectional median is the median of the time series mean of each state. Cross section variance is the variance of the time series mean of each state. All variables were tracked from their announcement date. Some policies were challenged legally, but were still tracked from the date of their announcement. Com Out: Community Outreach, VMSE: Vaccine mandate for state employees and/or health care workers in a state, IVM: Indoor vaccine mandate or vaccine mandate for gatherings over a certain size, BPV: Ban on requiring proof of vaccination, BMM: Ban on mask mandates.

6.2 Appendix B: Federal Policy Variables

On September 9, 2021 President Biden announced vaccine mandates that would be rolled out over the next several months. The mandates applied to most federal employees, employees of federal contractors, medical workers who worked at facilities that accepted Medicare and Medicaid reimbursement, and employers with 100 or more employees. It was estimated that the mandates would apply to roughly 100 million US workers. Rather than use a binary indicator variable for federal vaccine mandates, equal to zero before and unity after September 9, we created a linearly interpolated variable that captured the increase in federal-level vaccination mandates during the time period. Prior to President Biden's September 9th announcement, a mandate on members of the military was already in place, and hundreds of employers (many of them with employees nationwide) in the United States chose privately to have their own employer vaccination requirements. In order to fully capture mandates at the federal-level we felt it important to include employer mandates that impacted workers nationwide.

To quantify how many employees were under employer mandates, we used a Gallup poll that was taken monthly from May through December 2021. (Jones, 2021) The poll asked workers to the best of their knowledge whether their employer would require vaccination against COVID-19. In May of 2021 only five percent of employees said their employers mandated vaccination. By October 2021, just five months later, that number increased to 36 percent of workers that had employer mandated vaccination. We combined the employer mandated percentages with the military and federal vaccine mandates to construct the federal-level vaccine mandate variable. We also took into account announcements made by the Secretary of Defense and the White House that signaled that vaccine mandates were likely to come in the near future. Appendix Table B shows the dates that were used to construct the federal vaccine mandate variable as well as the combined federal and employer mandates variable.

We included the signals of future mandates because of those who were eligible for vaccination by the summer of 2021, but who were still not vaccinated, there were two groups. The first group consisted of those who were vaccine hesitant and wanted to wait for full Food and Drug Administration (FDA) approval of the vaccines, or planned to get vaccinated and just hadn't gotten around to it yet. The other group was the vaccine resistant who were opposed to the vaccine at almost any cost and were willing to suffer the consequences of not being vaccinated if mandates were enacted. Vaccine mandates, whether federal or at the employer level, did not likely increase vaccines among the latter group in a significant way. The former group, however, were likely influenced by such mandates, and the mere announcement of the mandates were sufficient to nudge them into action and get vaccinated. (The Pfizer vaccine also was granted full FDA approval during this same time period, on August 23, 2021.)

Appendix Table 6.2: Federal and Employer Vaccine Mandates and Club 1 Membership Size

Date	Federal and Employer Mandates	Members in C1
May 2021	5% of employees report having employer vaccine mandate	13
June 2021	6% of employees report having employer vaccine mandate	13
July 2021	9% of employees report having employer vaccine mandate	13
Aug 9, 2021	Secretary of Defense sent message of intent to mandate COVID-19 vaccination for the military	13
Aug 2021	19% of employees report having employer vaccine mandate	15
Aug 23, 2021	President Biden's Press Secretary announces more stringent vaccine mandates coming	17
Aug 24, 2021	Secretary of Defense announces memorandom to fully vaccinate members of the military	17
Sept 2021	29% of employees report having employer vaccine mandate	18
Sept 9, 2021	President Biden announces federal vaccine mandates	51

Notes: The above dates were used to construct the common factor variables. We constructed a pure federal vaccine mandate variable, along with a federal-level vaccine mandates variable, which combined the federal vaccine mandates with employer mandates. To construct the continuous federal mandate variable we made the base of the 100 million workers that were predicted to be affected by the September 9th federal vaccine mandates and added 2,395,993, the size of the military. When a signal of upcoming mandates was made the federal mandate variable took on a value of ten percent of those who would be impacted. For example, when the Secretary of Defense sent a message about the intent to implement a military vaccine mandate, the federal mandate variable went from 0 to .0023 (10% of the people who would be impacted by the mandate (members of the military) divided by the base of 102,395,993). When the actual announcement was made, the numerator went to 100% of those impacted by the mandate. The employer mandate variable was equal to whatever percentage of employees reported having vaccine mandates at their place of work each month. The federal level variable was a combination of the two, which became equal to one on September 9, 2021 when President Biden made the vaccine mandate announcement.

References

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller, "Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program," *Journal of the American statistical Association*, 2010, *105* (490), 493–505.

Barber, Andrew and Jeremy West, "Conditional cash lotteries increase COVID-19 vaccination rates," *Journal of Health Economics*, 2022, *81*, 102578.

Brehm, Margaret E., Paul A. Brehm, and Martin Saavedra, "The Ohio Vaccine Lottery and Starting Vaccination Rates," *American Journal of Health Economics*, 2022, *8* (3), 387–411.

Chaisemartin, Clément De and Xavier d'Haultfoeuille, "Two-way fixed effects estimators with heterogeneous treatment effects," *American Economic Review*, 2020, *110* (9), 2964–96.

Chang, Tom, Mireille Jacobson, Manisha Shah, Rajiv Pramanik, and Samir B Shah, "Financial Incentives and Other Nudges Do Not Increase COVID-19 Vaccinations among the Vaccine Hesitant," Working Paper 29403, National Bureau of Economic Research October 2021.

Chen, Jiafeng and Jonathan Roth, "Logs with zeros? some problems and solutions," *arXiv preprint arXiv:2212.06080*, 2023.

Dave, Dhaval, Andrew I. Friedson, Benjamin Hansen, and Joseph J. Sabia, "Association Between Statewide COVID-19 Lottery Announcements and Vaccinations," *JAMA Health Forum*, 10 2021, *2* (10), e213117–e213117.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Jones, Jeffery M., "Worker COVID-19 Vaccine Requirements Leveling Off," *Gallup*, 2021.

Lang, David, Lief Esbenshade, and Robb Willer, "Did Ohio's Vaccine Lottery Increase Vaccination Rates? A Pre-Registered, Synthetic Control Study," *Journal of Experimental Political Science*, 2022, p. 1–19.

Phillips, Peter C B and Donggyu Sul, "Transition modeling and econometric convergence tests," *Econometrica*, 2007a, *75* (6), 1771–1855.

____ **and** ____ , "Some empirics on economic growth under heterogeneous technology," *Journal of Macroeconomics*, 2007b, *29* (3), 455–469.

——— **and** ——— , “Economic transition and growth,” *Journal of applied econometrics*, 2009, *24* (7), 1153–1185.

Rambachan, Ashesh and Jonathan Roth, “A more credible approach to parallel trends,” *Review of Economic Studies*, 2023, p. rdad018.

Roth, Jonathan, Pedro HC Sant’Anna, Alyssa Bilinski, and John Poe, “What’s trending in difference-in-differences? A synthesis of the recent econometrics literature,” *Journal of Econometrics*, 2023.

Sehgal, Neil K.R., “Impact of Vax-a-Million Lottery on COVID-19 Vaccination Rates in Ohio,” *The American Journal of Medicine*, 2021, *134* (11), 1424–1426.

Sul, Donggyu, *Panel data econometrics: Common factor analysis for empirical researchers*, Routledge, 2019.

Thirumurthy, Harsha, Katherine L. Milkman, Kevin G. Volpp, Alison M. Buttenheim, and Devin G. Pope, “Association between statewide financial incentive programs and COVID-19 vaccination rates,” *PLOS ONE*, 2022, *17* (3), e0263425.

Walkey, Allan J., Anica Law, and Nicholas A. Bosch, “Lottery-Based Incentive in Ohio and COVID-19 Vaccination Rates,” *JAMA*, 08 2021, *326* (8), 766–767.

Wong, Charlene A., William Pilkington, Irene A. Doherty, Ziliang Zhu, Hattie Gawande, Deepak Kumar, and Noel T. Brewer, “Guaranteed Financial Incentives for COVID-19 Vaccination: A Pilot Program in North Carolina,” *JAMA Internal Medicine*, 01 2022, *182* (1), 78–80.