Journal of the Royal Statistical Society Series B: Statistical Methodology, 2023, **00**, 1–22
https://doi.org/10.1093/jrsssb/qkad103

Journal of the Statistics Society
Series B
Statistical Methodology

Original Article

Dynamic synthetic control method for evaluating treatment effects in auto-regressive processes

Xiangyu Zheng[1] **and Song Xi Chen**[1,2]

[1]Guanghua School of Management, Peking University, Beijing, P.R. China
[2]Center for Statistical Science, Peking University, Beijing, P.R. China

Address for correspondence: Song Xi Chen, Center for Statistical Science and Guanghua School of Management, Peking University, Beijing, 100871, P.R. China. Email: songxichen@pku.edu.cn

Abstract

Motivated by evaluating the effects of air pollution alerts on air quality, we propose the dynamic synthetic control method for micro-level data with time-varying confounders and spatial dependence under an auto-regressive model setting. We employ the empirical likelihood to define the synthetic control weights, which ensures a unique solution and permits theoretical analysis. The dynamic matching increases the feasibility of matching and enables us to assess the unconfoundedness assumption using pre-treatment data. For statistical inference, we develop a normalised placebo test to address the asymmetry issue. The method is illustrated and evaluated on numerical simulations and a case study on air pollution alerts.

Keywords: empirical likelihood, policy intervention, spatial dependence, time-varying confounding, treatment effect

1 Introduction

We consider evaluating the treatment effects caused by a policy intervention on a target group in the spatial and temporal panel data setting. In the absence of random assignment, simple comparisons of the treatment and control groups cannot demonstrate the causal link between the treatment intervention and the observed outcomes due to the unbalanced distributions of confounding variables. One common strategy is to match each of the treated units with a set of control units that share similar characteristics on the confounding variables (Rubin, 2006; Stuart, 2010). However, in the panel data setting with time-varying observations, it is difficult to find a control unit that matches the whole trajectory of the confounding variables. The synthetic control (SC) method, proposed by Abadie and Gardeazabal (2003) and further developed in Abadie et al. (2010), offered an approach to tackle this issue by creating SCs that provide a better comparison in the application setting with a single treated unit. Currie et al. (2020) conducted a comprehensive review on the SC-related papers published in the National Bureau of Economic Research papers and the top-five economic journals, and found a rapid growth in the use of the SC method since 2010. The SC method had been regarded as one of the most important innovations in the policy evaluation studies and has been widely applied since its inception (Athey & Imbens, 2017). Similar idea of adaptive weighting is also adopted in the literature of statistical inference on adaptively collected data (Deshpande et al., 2018; Hadad et al., 2021; Zhan et al., 2021; Zhang et al., 2021), where the weights are designed to account for the variance part for valid inference.

The SC method uses the outcomes and covariates in the pre-treatment period to construct a vector of weights such that the weighted averages of the control units match both the trajectory of pre-treatment outcomes and the covariates of the treated unit. Then, the counterfactual outcomes

Received: June 1, 2021. **Revised:** July 27, 2023. **Accepted:** August 21, 2023

of the treated unit in the post-treatment periods are estimated by weighted averages of the outcomes in the control group. Abadie et al. (2010) proves that the SC estimator of the causal effect is asymptotically unbiased given that the length of pre-treatment periods T_0 goes to infinity and the matching equations can be attained. Arkhangelsky et al. (2019) proposed a synthetic difference in difference (SDID) method that combined the SC and the difference in difference method to offer the double-robustness in the same factor model context as the SC, which also contained aspects useful for further theoretical analysis of the SC method.

As the SC approach is not applicable for the settings with time-varying confounders, we extend the SC to dynamic panel data setting as encountered in the air pollution alert study. We develop a dynamic synthetic control (DSC) approach which constructs dynamic weights to match the level of time-varying confounders via the empirical likelihood (EL).

The dynamic property of the weights makes it easier to find an exact match as the matching conditions do not involve the whole pre-treatment periods. We prove that under a lenient assumption, the exact match can be attained with probability approaching one.

The proposed DSC method facilitates consistent estimation of the average treatment effect on the treated (ATT) in situations with micro-level data including multiple treated units and time-varying covariates. We propose to construct a series of time-varying weights by maximising the EL (Chen & Van Keilegom, 2009; Owen, 1988; Qin & Lawless, 1994) of the weighted distribution, under the constraints that match the current state of the time-varying confounders and the lagged values of the potential outcomes under control. This is also the first work that provides the asymptotic orders of w_{it}^* for spatially weakly dependent data and under both cases with and without confounding bias, while the SC method used the simple range $0 \leq w_{it}^* \leq 1$ for deriving an upper bound of the expected bias in the error of the estimated treatment effect for independent data. The use of the EL ensures that the solution of the optimal weights is unique and their order of the magnitude can be controlled such that the consistency of the treatment effect estimates can be established. The DSC formulation via the EL is connected to the entropy balancing (EB) method (Zhao & Percival, 2016), where the weights are obtained by minimising the relative entropy between the observed and the weighted distributions. If we exchange the positions of the observed and the weighted distributions in the EB, it would yield the same result as the EL formulation. The manner of dynamic matching improves over the SC method, as it is generally hard to derive the SC weights for matching both the pre-treatment outcomes and covariates (Ferman & Pinto, 2016).

The need for the proposed DSC emerges when evaluating the effects of air pollution alerts, where the confounding variables are time-varying meteorological variables and the air-quality observations from the monitoring sites are spatially dependent. Air pollution alerts and associated mandatory emission-reduction measures are interventions by management authorities to temporarily reduce the emission levels under foreseen unfavourable meteorological conditions. Due to the social and economic costs associated with the alerts, there is an enduring question on the effectiveness of these alerts. A common approach for evaluating their effectiveness is based on the air-quality models, for instance the Comprehensive Air Quality Model with Extensions in conjunction with Weather Research Forecast model via the scenario analysis (Cheng et al., 2017; Wang et al., 2018). A limitation of the approach is its reliance on the emission inventory, which are subject to measurement errors and usually 3–4 years reporting delay. To the best of our knowledge, there is a lack of studies which examine effects of air pollution alerts based on observational data. Although motivated by the air-quality alert study, the application scenario for DSC is quite wide, as time-varying confounders appear in many other fields, for instance in marketing and medical research, with data collected from wearable devices.

Statistical inference procedures associated with the DSC method are proposed, including a test for the model specification based on the estimated effects in the pre-treatment period, and a normalised placebo test to evaluate the significance of the estimated treatment effects.

In addition to the simulation experiments to validate the theoretical results, we applied the DSC method to evaluate two air pollution alerts issued in Beijing. The estimated pseudo-treatment effects in the pre-treatment period supported the unconfoundedness assumption and the model specification. The hypothesis testing shows that the emission-reduction measure led to significant pollution reduction about 7–8 h after the start of the alerts.

The rest of the paper proceeds as follows. Section 2 introduces the DSC method and gives the theoretical results on the existence of SCs, the order of the EL weights, and the consistency of the estimation. Section 3 shows how to assess the unconfoundedness assumption using the pre-treatment data and introduces the normalised placebo test to examine the significance of the estimated treatment effects. Section 4 reports simulation results, while Section 5 demonstrates the DSC by evaluating the effects of two air pollution alerts in Beijing.

2 DSC method

Suppose that there are N_{tr} units from the treatment group and N_{co} units from the control group, which add up to a total number of $N = N_{\text{tr}} + N_{\text{co}}$ units. The units may be pollution monitoring sites in the air pollution study. All units are observed for T time periods. Let $\mathcal{T}$ and $\mathcal{C}$ be the sets of unit indexes in the treatment and the control groups, respectively. Units in $\mathcal{T}$ are exposed to a policy intervention after time T_0 and subsequently observed for $T - T_0$ time units, which are called the treated units. Units in $\mathcal{C}$ are unaffected by the intervention in the observed time span. Let D_{it} be the treatment indicator so that $D_{it} = 1$ if unit i is affected by the intervention at time t and $D_{it} = 0$ otherwise; and $D_i = D_{i,T_0+1}$ be the unit treatment indicator that denotes if unit i is exposed to the intervention (treatment) after T_0. Let Y_{it} be the outcome of interest, $\mathbf{X}_{it}$ be the p-dimensional vectors of covariates.

2.1 Average treatment effects on the treated

Adopting the notion of potential outcomes (Rubin, 1974), let $Y_{it}(0)$ and $Y_{it}(1)$ be the potential outcomes which would be observed for unit i at time t when $D_{it} = 0$ and 1, respectively. The observed outcome is $Y_{it} = Y_{it}(D_{it})$ depending on whether unit i is affected by the intervention at time t or not. We focus on the ATT group at time t over the post-intervention period:

$$\tau_t = \mathrm{E}\big(Y_{it}(1) - Y_{it}(0) | D_i = 1\big) \quad \text{for } t \in \{T_0 + 1, , \ldots, T\}. \tag{2.1}$$

Note τ_t in definition (2.1) is irrelevant to the unit index i, which implies that we assume the potential outcomes $\{Y_{it}(1), Y_{it}(0)\}$ for the units with $D_i = 1$ in the treatment group share an identical distribution.

The proposed DSC method estimates τ_t by the difference between the simple average of the treatment group and a dynamic weighted average of the control group $\sum_{j\in\mathcal{T}} \frac{Y_{it}}{N_{\text{tr}}} - \sum_{i\in\mathcal{C}} w_{it}^* Y_{it}$, where $w_{it}^* \in \sigma(\{Y_{i,s-1}, \mathbf{X}_{is}\}_{s\le t})$. By constructing the weights as dynamic, the DSC extends the setting of the SC to the case with time-varying confounders while allowing both outcomes Y_{it} and covariates $\mathbf{X}_{it}$ to be spatially dependent across units. The DSC method is also *the first* to employ the EL into the construction of $\{w_{it}^*\}$ and provide theoretical analysis on the asymptotic orders of the SC weights w_{it}^*, which is crucial to prove the consistency of $\widehat{\tau}_t$.

Before introducing the DSC method, we introduce the following common assumptions on the potential outcomes.

Assumption 1 The consistency assumption: $Y_{it} = D_{it}Y_{it}(1) + (1 - D_{it})Y_{it}(0)$ for $1 \le t \le T$, $1 \le i \le N$.

Assumption 2 The unconfoundedness assumption:
$\mathrm{E}(Y_{it}(0)|Y_{i,t-1}(0), \mathbf{X}_{it}, D_i = 1) = \mathrm{E}(Y_{it}(0)|Y_{i,t-1}(0), \mathbf{X}_{it}, D_i = 0)$.

Assumption 3 The confounding variables are unaffected by the intervention such that $\mathbf{X}_{it} = \mathbf{X}_{it}(1) = \mathbf{X}_{it}(0)$ for $1 \le t \le T$.

Assumption 4 $\mathrm{P}(D_i = 1) > 0$; $\mathrm{P}(D_i = 1|\mathbf{X}_{it}, Y_{i,t-1}(0)) < 1$ with probability one.

Assumption 1 ensures that one of the potential outcomes is observed for each unit at each period, which connects the potential outcomes with the observed outcomes. Under Assumption 1, $\mathrm{E}(Y_{it}(1)|D_i = 1) = \mathrm{E}(Y_{it}|D_i = 1)$ for $t > T_0$. Therefore, for $t > T_0$,

$$\tau_t = \mathrm{E}(Y_{it}|D_i = 1) - \mathrm{E}\big(Y_{it}(0)|D_i = 1\big), \tag{2.2}$$

where the first term is estimated by the sample average in the treatment group and the second term $\mathrm{E}(Y_{it}(0)|D_i = 1)$ will be estimated by the proposed DSC.

It should be noted that the potential outcomes at time t actually rely on the historical path of treatment indicators before t, $Y_{it}(1)$ implicitly denotes the potential outcome if unit i has been subject to the treatment after T_0, and $Y_{it}(0)$ denotes the potential outcome if unit i has been under control from the start. In the notations from Robins et al. (2000) and Boruvka et al. (2018), we may express the potential outcomes as

$$Y_{it}(1) = Y_{it}(\underbrace{0\cdots0}_{T_0}, \underbrace{1\cdots1}_{t-T_0}) \quad \text{and} \quad Y_{it}(0) = Y_{it}(\underbrace{0\cdots0}_{T_0}, \underbrace{0\cdots0}_{t-T_0}).$$

In this way, it is clearer why the expectation in Assumption 2 is conditioning on the potential outcome under control $Y_{i,t-1}(0)$, instead of the observed outcome $Y_{i,t-1}$. In our case, since the intervention is only administered after $t = T_0$ or never in our scenario, we may leave out the past treatment indicators when expressing the potential outcomes.

Following Imbens and Wooldridge (2009), we call Assumption 2 as the unconfoundedness assumption, which is for the identification of τ_t and means that given the same covariates and lagged outcomes, the conditional expectations of $Y_{it}(0)$, the potential outcomes under control, are identical between the treatment and the control groups. Assumption 3 implies that the time-varying covariates $\mathbf{X}_{it}$ are not affected by the treatment.

Under Assumptions 1–3, the conditional expectation of $Y_{it}(0)$ in the treatment group can be expressed in terms of the observed data (identifiable),

$$\mathrm{E}\big(Y_{it}(0)|\mathbf{X}_{it}, Y_{i,t-1}(0), D_i = 1\big) = \mathrm{E}\big(Y_{it}(0)|\mathbf{X}_{it}, Y_{i,t-1}(0), D_i = 0\big) = \mathrm{E}\big(Y_{it}|\mathbf{X}_{it}, Y_{i,t-1}, D_i = 0\big),$$

which suggests that τ_t can be estimable from observed data. Note that since $Y_{it} = Y_{it}(0)$ for all units no matter $D_i = 0$ or $D_i = 1$ when $t \le T_0$, the validity of Assumption 2 for $t \le T_0$ can be assessed using the pre-treatment data. A test for that purpose is given in Section 3.1. Assumption 4 ensure that τ_t is well defined and for each given value of $\mathbf{X}_{it}$ and $Y_{i,t-1}(0)$ and there is always a fraction of untreated population that can be used as controls.

2.2 Dynamic synthetic controls

To put forward the DSC method, we assume that $\mathrm{E}(Y_{it}(0)|Y_{i,t-1}(0), \mathbf{X}_{it})$ is linear as given by the following dynamic linear model:

$$Y_{it}(0) = \delta_t + \boldsymbol{\beta}_t'\mathbf{X}_{it} + \rho_t Y_{it-1}(0) + \varepsilon_{it}, \quad \text{for } 1 \le i \le N,\ 1 \le t \le T, \tag{2.3}$$

where δ_t is a unknown common factor that is time-varying. It is assumed that ε_{it} is independent of $\{X_{it}, Y_{i,t-1}(0)\}$ with zero mean at each time t. Spatial dependence among the errors $\{\varepsilon_{it}\}_{i=1}^N$ for each t is allowed provided that $\sum_{0\le|i-j|<N} \mathrm{E}(\varepsilon_{it}\varepsilon_{jt}) = O(N)$, where the latter regulates the degree of the spatial dependence.

For each post-treatment period $t > T_0$, we aim at constructing a set of SCs that match the counterfactual situation of the treatment group in the absence of treatment to estimate $\mu_t(0) = \mathrm{E}(Y_t(0)|D = 1)$, which is the key in estimating τ_t from equation (2.2). We consider constructing the SCs by weighted averages of the outcomes from units in the control group $\mathcal{C}$ to

estimate $\mu_t(0)$, namely

$$\widehat{\mu}_t(0) = \sum_{i \in \mathcal{C}} w_{it}^* Y_{it}, \tag{2.4}$$

where $\mathbf{w}_t^* = (w_{1t}^*, \ldots, w_{N_{co},t}^*)$ is a vector of non-negative weights summing to 1. The weights determine the contribution of the control units at time t in the DSC estimates $\widehat{\mu}_t(0)$.

According to Model (2.3), the weighted average

$$\sum_{i \in \mathcal{C}} w_{it}^* Y_{it} = \sum_{i \in \mathcal{C}} w_{it}^* \delta_t + \boldsymbol{\beta}_t \sum_{i \in \mathcal{C}} w_{it}^* X_{it} + \rho_t \sum_{i \in \mathcal{C}} w_{it}^* Y_{i,t-1}(0) + \sum_{i \in \mathcal{C}} w_{it}^* \varepsilon_{it}. \tag{2.5}$$

At the same time, we note the following natural but infeasible estimator $\widetilde{\mu}_t(0)$ of $\mu_t(0)$,

$$\widetilde{\mu}_t(0) =: \bar{Y}_{\cdot t}(0) = \sum_{i \in \mathcal{T}} \frac{Y_{it}(0)}{N_{tr}} = \delta_t + \boldsymbol{\beta}_t \sum_{i \in \mathcal{T}} \frac{X_{it}}{N_{tr}} + \rho_t \sum_{i \in \mathcal{T}} \frac{Y_{i,t-1}(0)}{N_{tr}} + \sum_{i \in \mathcal{T}} \frac{\varepsilon_{it}}{N_{tr}}. \tag{2.6}$$

Despite being infeasible as $Y_{it}(0)$ are unobserved in the treatment units for $t > T_0$, $\widetilde{\mu}_t(0)$ motivates two matching conditions (2.8) and (2.9) for the synthetic weights $\mathbf{w}_t^*$.

We use the EL method (Chen & Van Keilegom, 2009; Owen, 1988; Qin & Lawless, 1994) to construct the synthetic weights $\{w_{it}^*\}_{i=1}^{N_{co}}$ which can be regarded as a non-parametric distribution on the control group that put probability mass w_{it}^* on unit i at time point t, and $\prod_{i \in C} w_{it}^*$ is the EL. We would like to maximise the likelihood under the structural constraints in equations (2.8) and (2.9).

Notably, the use of the EL is crucial as it leads to a unique solution of the weights which facilitates analyses of the asymptotic properties of w_{it}^*.

Specifically, the EL SC weights $\mathbf{w}_t^* = (w_{1t}^*, \ldots, w_{N_{co},t}^*)'$ are defined as

$$\mathbf{w}_t^* = \underset{\mathbf{w}_t \in \mathcal{W}_t}{\operatorname{argmax}} \prod_{i \in C} w_{it}, \tag{2.7}$$

where $\mathcal{W}_t = \{\mathbf{w}_t : \sum_{i=1}^{N_{co}} w_{it} = 1, w_{it} \geq 0,$ and satisfy the following equations (2.8) and (2.9)$\}$,

$$\sum_{i \in \mathcal{C}} w_{it} \mathbf{X}_{it} = \sum_{j \in \mathcal{T}} \frac{\mathbf{X}_{jt}}{N_{tr}} \quad \text{and} \tag{2.8}$$

$$\sum_{i \in \mathcal{C}} w_{it} Y_{i,t-1} = \begin{cases} \sum_{j \in \mathcal{T}} Y_{j,t-1}/N_{tr}, & \text{if } t \leq T_0 + 1; \quad (2.9a) \\ \hat{\mu}_{t-1}(0) = \sum_{i \in \mathcal{C}} w_{i,t-1}^* Y_{i,t-1}, & \text{if } t > T_0 + 1. \quad (2.9b) \end{cases}$$

Then, the treatment effect τ_t is estimated by

$$\widehat{\tau}_t = \frac{\sum_{j \in \mathcal{T}} Y_{jt}}{N_{tr}} - \widehat{\mu}_t(0) = \sum_{j \in \mathcal{T}} \frac{1}{N_{tr}} Y_{jt} - \sum_{i \in \mathcal{C}} w_{it}^* Y_{it}, \quad t > T_0. \tag{2.10}$$

The structural constraints (2.8) and (2.9) match terms in equation (2.5) with the corresponding terms in equation (2.6). In particular, equation (2.8) offers matching on the covariates between the treatment and the control groups based on the second term of Model (2.3), while equations (2.9a) and (2.9b) match the lagged values of the potential outcomes $Y_{i,t-1}(0)$ before and after the time point $T_0 + 1$. The dynamic matching in equations (2.9a) and (2.9b) requires some explanations. We expect to find weights $\{w_{it}^*\}$ such that $\sum_{i \in \mathcal{C}} w_{it}^* Y_{i,t-1}(0)$ in (2.5) matches $\sum_{i \in \mathcal{T}} Y_{i,t-1}(0)/N_{tr}$ in equation (2.6). When $t \leq T_0 + 1$, as $Y_{i,t-1} = Y_{i,t-1}(0)$ for $i \in \mathcal{T}$, the observed lagged outcome is unaffected by the intervention, the matching is readily equation (2.9a). For $t > T_0 + 1$, as $Y_{i,t-1}(0)$ is unobserved for units in the treatment group $\mathcal{T}$, we replace $\sum_{i \in \mathcal{T}} Y_{i,t-1}(0)/N_{tr}$ by $\widehat{\mu}_{t-1}(0)$, the SC estimate for $\mu_{t-1}(0)$ in the last time point, leading to equation (2.9b). While equation (2.3) specifies an

AR−1 model, an extension to AR−k model can be made, which requires replace the lag−1 matching constraint in equation (2.9) to the lag−k version by requiring $\widehat{\mu}_{t-j}(0) = \sum_{i\in\mathcal{C}} w^*_{i,t-j} Y_{i,t-j}$ for $1 \le j \le k$; see Section 1.5 of the online supplementary material for details.

The matching as rendered in equations (2.9a) and (2.9b) are critical in making the dynamic synthetic control estimates $\widehat{\mu}_t(0)$ in (2.4) consistent to $\widetilde{\mu}_t(0)$. To appreciate this, let $\eta_t = \sum_{i\in\mathcal{C}} w^*_{it}\varepsilon_{it} - \sum_{j\in\mathcal{T}} \frac{1}{N_{tr}}\varepsilon_{jt}$, and $\rho^t_s = \prod_{l=s+1}^{t} \rho_l$ for $s < t$. When $t \le T_0 + 1$, according to the structural constraints (2.8) and (2.9a), $\widehat{\mu}_t(0) - \widetilde{\mu}_t(0) = \eta_t$. As for $t > T_0 + 1$, there is an accumulation of error terms as we replaced $\widetilde{\mu}_{t-1}(0)$ with $\widehat{\mu}_{t-1}(0)$ in equation (2.9b), which leads to a recursive formula $\widehat{\mu}_t(0) - \widetilde{\mu}_t(0) = \eta_t + \rho_t(\widehat{\mu}_{t-1}(0) - \widetilde{\mu}_{t-1}(0))$. Therefore, we have

$$\widehat{\mu}_t(0) - \widetilde{\mu}_t(0) = \begin{cases} \eta_t, & \text{if } t \le T_0 + 1; \\ \eta_t + \sum_{s=T_0+1}^{t-1} \rho^t_s \eta_s, & \text{if } t > T_0 + 1. \end{cases} \tag{2.11}$$

Under certain mild assumptions, the order of $\{w^*_{it}\}$ can be controlled and equation (2.11) is stochastically small as will be shown in Theorem 1.

2.3 Comparison with SC method

The proposed DSC method is motivated by the SC method in Abadie et al. (2010) and the setting of the air pollution alert study. It differs from the SC method in several aspects. First, the vector of weights $\mathbf{w}^*_t$ is time-varying and is defined for each time point to accommodate time-varying confounding variables. This is different from the SC method as the SC weights are constant in time due to the time-invariant nature of the confounders. Abadie et al. (2010) considered the following factor model:

$$Y_{it}(0) = \delta_t + \boldsymbol{\beta}_t \mathbf{X}_i + \boldsymbol{\lambda}'_t \boldsymbol{\xi}_i + \varepsilon_{it}, \tag{2.12}$$

where $\boldsymbol{\lambda}_t$ and $\boldsymbol{\xi}_i$ are unobserved factors. The weights used in constructing the counterfactual average are obtained by matching $(p + T_0)$-dimensional equations for the time-invariant covariates and the whole trajectory of T_0 pre-treatment outcomes:

$$\sum_{i\in\mathcal{C}} w^*_i \mathbf{X}_i = \sum_{j\in\mathcal{T}} \frac{\mathbf{X}_j}{N_{\text{tr}}} \quad \text{and} \quad \sum_{i\in\mathcal{C}} w^*_i Y_{i,t} = \sum_{j\in\mathcal{T}} \frac{Y_{j,t}}{N_{\text{tr}}} \quad \text{for } t = 1, \ldots, T_0. \tag{2.13}$$

The structural constraints (2.8) and (2.9) used by the proposed DSC are counterparts of equation (2.13). The essence of equation (2.13) is to match on the trajectory up to the time T_0 before intervention. This is a key difference from the formulation in DSC where we conduct repeated matching at each time period t, which is largely determined by the different models used. We consider the auto-regressive model (2.3) that suits for time-varying confounders, instead of having the factor component $\lambda_t \xi_i$ in equation (2.12).

From the practical perspective, the estimated effects in the post-treatment period $\{\widehat{\tau}_t\}$ do not require a large T_0, and only rely on the pre-treatment outcomes $\{Y_{iT_0}\}$ at time T_0. To obtain $\mathbf{w}^*_t$, we only need to match in the DSCs on $p + 1$ dimensions, which is less demanding than the $(p + T_0)$-dimensional match required in Abadie et al. (2010), and increases the feasibility of finding the weights for a match. A price for the proposed DSC method is the increased computation, while the SC method only solves the optimisation problem once, the proposed DSC needs to repeat the procedure for each post-treatment period.

An advantage of the proposed DSC method is that the weights $\{w^*_{it}\}_{i=1}^{N_{\text{co}}}$ are obtained by maximising the EL under the constraints that reflect the matched equations. This ensures that the solution is unique and also allows us to provide a complete characterisation on the order of $\{w^*_{it}\}$ in terms of N_{co}. In contrast, the SC method cannot generally ensure exact matches for large T_0. As results, there are proposals for attaining approximate matches. It is noted that the dimensions of matching equations of the DSC does not increase with T_0 and, as shown in Lemma 1, the probability of

having the exact matches tends to be 1 as the sample size the increases. In fact, the EL can be used to obtain the SC weights with equation (2.13) as the EL constraints. It may be shown that the EL solution of the synthetic weights is uniqueness. Additionally, since the DSC does not match the whole trajectory of the pre-treatment outcomes, the difference between the synthetic controls and the observed outcomes in the pre-treatment periods can be used to assess the Assumption 2 for unconfoundedness and the linear specification (2.3). To appreciate this aspect, note that $\widehat{\mu}_t(0)$ can be obtained for each pre-treatment period. If the method is valid, $\widehat{\mu}_t(0)$ should be close to the simple average on the treatment group. Otherwise, the difference between $\widehat{\mu}_t(0)$ and the average on the treatment group would be large in absolute value.

The detailed testing procedures will be introduced in Section 3.1. In contrast, note that the weights in the SC method are obtained by matching the pre-treatment outcomes, so the differences in the pre-treatment periods do not provide evidence to assess the factor model assumption.

Regarding the analysis settings of the two methods, the SC requires T_0 to be sufficiently large for the estimation to be asymptotically unbiased, while the DSC allows T_0 to be fixed but requires both N_{tr} and N_{co} to be large when establishing the estimation consistency. These are reflected in the forms of the theoretical results as we analyse the asymptotic probability of finding the exact match and prove the consistency of the DSC estimate under $N_{\text{co}}, N_{\text{tr}} \to \infty$, while Abadie et al. (2010) proved the asymptotically unbiasedness of the SC estimate when $T_0 \to \infty$ and $N_{\text{tr}} = 1$. More importantly, the two methods operate on different model settings. The DSC method may be viewed as a complement for the SC in the case with time-varying confounders and auto-regressive model structure.

The auto-regression was briefly discussed in Abadie et al. (2010) under the dual dynamic linear models for outcomes Y_{it} and the covariates $\mathbf{X}_{it}$:

$$\begin{aligned} Y_{it}(0) &= \boldsymbol{\beta}_t' \mathbf{X}_{it} + \rho_t Y_{it-1}(0) + \epsilon_{it}, \\ X_{it} &= \Pi_t' \mathbf{X}_{i,t-1} + \gamma_t Y_{i,t-1}(0) + \upsilon_{it}, \end{aligned} \tag{2.14}$$

where ϵ_{it} and υ_{it} are residuals. The authors constructed the SC estimates as $\widehat{\mu}_0(t) = \sum_{i \in \mathcal{C}} w_i^* Y_{it}$, where the weights $\{w_i^*\}$ were still time-invariant and obtained by matching the covariates and outcomes at T_0 only:

$$\sum_{i \in \mathcal{C}} w_i^* \mathbf{X}_{iT_0} = \sum_{j \in \mathcal{T}} \frac{\mathbf{X}_{jT_0}}{N_{\text{tr}}} \quad \text{and} \quad \sum_{i \in \mathcal{C}} w_i^* Y_{iT_0} = \sum_{j \in \mathcal{T}} \frac{Y_{jT_0}}{N_{\text{tr}}}. \tag{2.15}$$

Note that Model (2.14) requires that the confounders $\mathbf{X}_{it}$ also satisfy an auto-regressive model, while Model (2.3) for the DSC does not specify a model on $\mathbf{X}_{it}$ as we match the time-varying confounders for each period. Besides, the variance of the SC estimates would be large as it only controls the initial status of the covariates and outcomes at the time point T_0.

One may wonder how the SC would perform under the dynamic Model (2.3). According to the SC matching equation (2.15),

$$\widehat{\mu}^{\text{SC}}_{T_0+1} - \widetilde{\mu}_{T_0+1}(0) = \eta_t + \beta' \left[\sum_{i \in \mathcal{C}} w_i^* (\boldsymbol{X}_{i,T_0+1} - \boldsymbol{X}_{i,T_0}) - \sum_{i \in \mathcal{T}} \frac{(\boldsymbol{X}_{i,T_0+1} - \boldsymbol{X}_{i,T_0})}{N_{\text{tr}}} \right],$$

which implies that for $t > T_0 + 1$,

$$\widehat{\mu}^{\text{SC}}_t - \widetilde{\mu}_t(0) = \eta_t + \rho(\widehat{\mu}^{\text{SC}}_{t-1} - \widetilde{\mu}_{t-1}(0)) + \sum_{s=T_0+1}^{t} \beta' \left[\sum_{i \in \mathcal{C}} w_i^* (\boldsymbol{X}_{is} - \boldsymbol{X}_{i,s-1}) - \sum_{i \in \mathcal{T}} \frac{(\boldsymbol{X}_{is} - \boldsymbol{X}_{i,s-1})}{N_{\text{tr}}} \right].$$

Let $\Delta_t = \beta'\left(\sum_{i\in\mathcal{C}} w_i^*(\boldsymbol{X}_{it} - \boldsymbol{X}_{i,t-1}) - \sum_{i\in\mathcal{T}} \frac{(\boldsymbol{X}_{it}-\boldsymbol{X}_{i,t-1})}{N_{\text{tr}}}\right)$, then

$$\widehat{\mu}_t^{\text{SC}}(0) - \widetilde{\mu}_t(0) = \begin{cases} \eta_t + \Delta_{T_0+1} & \text{if } t = T_0+1; \\ \eta_t + \sum_{s=T_0+1}^{t-1} \rho_s^t \eta_t + \sum_{s=T_0+1}^{t-1} \rho_s^t \sum_{l=T_0+1}^{s} \Delta_l & \text{if } t > T_0+1. \end{cases} \tag{2.16}$$

Comparing equation (2.16) with equation (2.11), it can be seen that the SC estimates would incur larger variance due to the extra terms related to Δ_t. The latter reflects the variation in $\boldsymbol{X}_{it}$, which is not accounted for by the SC formulation. Thus, under the dual dynamic linear models (2.14), which is a more restricted version of equation (2.3), both estimates are unbiased but the SC estimates will have a larger variance than the DSC estimates. Furthermore, if $\{\boldsymbol{X}_{it}\}$ does not satisfy the auto-regressive specification prescribed in the second equation of equation (2.14), the SC estimates might be biased with a larger variance, as shown by a simulation experiment in Section 4.

2.4 Theoretical analysis of weights and consistency result

The EL SC weights $\{w_{it}^*\}$ that maximise equation (2.7) can be obtained by the Lagrange multiplier method. Let $\mathbf{Z}_{it}^0 = (\mathbf{X}_{it}', Y_{it-1})'$ for $1 \le i \le N_{\text{co}}$, $1 \le t \le T$, and $\bar{\mathbf{Z}}_t^1 = (\sum_{j\in\mathcal{T}} \mathbf{X}_{jt}'/N_{\text{tr}}, \sum_{j\in\mathcal{T}} Y_{j,t-1}/N_{\text{tr}})'$ for $t \le T_0+1$ and $(\sum_{j\in\mathcal{T}} \mathbf{X}_{jt}'/N_{\text{tr}}, \widehat{\mu}_{t-1}(0))'$ for $t > T_0+1$. Then, the structural constraints (2.8) and (2.9) can be written as matching the $(p+1)$-dimensional vector $\sum_{i\in\mathcal{C}} w_{it}^* \mathbf{Z}_{it}^0 = \bar{\mathbf{Z}}_t^1$. The EL weights admit the following expression:

$$w_{it}^* = {N_{\text{co}}}^{-1}[1 + \boldsymbol{\lambda}'(\mathbf{Z}_{it}^0 - \overline{\mathbf{Z}}_t^1)]^{-1} \quad \text{for } i = 1, \ldots, N_{\text{co}}, \tag{2.17}$$

where $\boldsymbol{\lambda}$ is the Lagrange multiplier that satisfies

$$\sum_{i=1}^{N} (\mathbf{Z}_{it}^0 - \overline{\mathbf{Z}}_t^1)[1 + \boldsymbol{\lambda}'(\mathbf{Z}_{it}^0 - \overline{\mathbf{Z}}_t^1)]^{-1} = 0 \tag{2.18}$$

and can be found by numerical methods, for instance the Newton–Raphson method.

Before introducing Theorem 1 on the consistency of $\widehat{\tau}_t$, we firstly show in Lemma 1 that the feasible region $\mathcal{W}_t$ for the optimisation problem (2.7) would be non-empty with probability approaching 1 under a mild condition, such that the optimisation problem admits a finite and unique solution for $\mathbf{w}_t^*$. In addition, we provide a characterisation on the order of $\{w_{it}^*\}$ in Lemma 2, which allows us to establish the consistency result in Theorem 1. The problem (2.7) is feasible when the region $\mathcal{W}_t \neq \emptyset$, which is equivalent to $\bar{\mathbf{Z}}_t^1$ falling into the convex hull of $\{\mathbf{Z}_{it}^0\}_{i=1}^{N_{\text{co}}}$, denoted by $\mathcal{H}_t$. Under the following Assumption 5, we can prove that $\bar{\mathbf{Z}}_t^1 \in \mathcal{H}$ with probability tending to 1 as the sample size is increased. Let $\mathbf{Z}_{it} = (\mathbf{X}_{it}', Y_{it-1}(0))'$ for $1 \le i \le N$, $1 \le t \le T$.

Assumption 5 Suppose that there exists a binary variable $S_i \in \{0, 1\}$ such that $\mathbb{P}(S_i = 1 | D_i = 0) > 0$, $\mathrm{E}(\mathbf{Z}_{it} | D_i = 0, S_i = 1) = \mathrm{E}(\mathbf{Z}_{it} | D_i = 1)$, and $\mathrm{Var}(\mathbf{Z}_{it} | D_i = 1, S_i = 1)$ is of full rank.

Note that $\mathbf{Z}_{it}$ denotes the vector of confounding variables which was unaffected by the treatment but influence the potential outcome $Y_{it}(0)$. When there are systematic differences on $\mathbf{Z}_{it}$, i.e. $\mathrm{E}(\mathbf{Z}_{it} | D_i = 0) \neq \mathrm{E}(\mathbf{Z}_{it} | D_i = 1)$, we cannot compare the treatment group with the control group directly even with sufficiently large sample size. Assumption 5 requires that there exists a comparable sub-population in the control group ($D_i = 0$) that has the same expectation on the confounders $\mathbf{Z}_{it}$ in the treatment group ($D_i = 1$). To appreciate it, let us illustrate in the context of air pollution alerts example, where $\mathbf{Z}_{it}$ includes the meteorological confounders. While the distribution of meteorological variables may not be balanced across the treated region (Beijing, subject to pollution alerts) and the surrounding control region, Assumption 5 states that it suffices to have a part of the control region which shares the same expectation of meteorological variables as that in Beijing. Under Assumption 5, we can conclude that the probability of $\bar{\mathbf{Z}}_t^1 \in \mathcal{H}_t$ would converge to

1 as shown in Lemma 1 such that Problem (2.7) is feasible. Let $\Theta = \{\theta : \|\theta\| = 1\}$ be the set of unit vectors in $\mathbb{R}^{p+1}$, where $\|\cdot\|$ is the Euclidean distance.

Lemma 1 Let $\mathbb{P}_{t,N_{\text{co}}}$ be the empirical distribution of $\{\mathbf{Z}_{it}^0\}_{i=1}^{N_{\text{co}}}$ and $\mathbf{Z}_t^0$ be the random variable with distribution $\mathbb{P}_{t,N_{\text{co}}}$, then under Assumption 5,

$$\inf_{\theta\in\Theta} \mathbb{P}_{t,N_{\text{co}}}\left(\left\{\theta'\left(\mathbf{Z}_t^0 - \overline{\mathbf{Z}}_t^1\right) > 0\right\}\right) > 0 \quad \text{as} \quad N_{\text{co}}, N_{\text{tr}} \to \infty. \tag{2.19}$$

Note that $\bar{\mathbf{Z}}_t^1 \in \mathcal{H}_t$ is equivalent to that there does not exist $\theta \in \Theta$ such that $\theta'(\mathbf{Z}_{it}^0 - \bar{\mathbf{Z}}_t^1) \le 0$ for all $i = 1, \ldots, N_{\text{co}}$. Combining this fact and Lemma 1, we may conclude that as $N_{\text{co}}, N_{\text{tr}} \to \infty$, $\bar{\mathbf{Z}}_t^1$ falls into $\mathcal{H}_t$ with probability approaching 1. Therefore, the feasible region $\mathcal{W}_t$ is not empty and hence Problem (2.7) is well defined. Furthermore, note that $\mathcal{W}_t$ is convex and bounded, a finite solution of $\{w_{it}^*\}_{i=1}^{N_{\text{co}}}$ for Problem (2.7) uniquely exists by the theory of convex optimisation (see online supplementary material for details).

From now on, we assume the existence of $\{w_{it}^*\}_{i=1}^{N_{\text{co}}}$. As for the order of w_{it}^*, when $\mathrm{E}(\mathbf{Z}_{it}|D_i = 0) = \mathrm{E}(\mathbf{Z}_{it}|D_i = 1)$, we may apply the standard result of EL (Qin & Lawless, 1994) to obtain that $w_{it}^* = O_p(N_{\text{co}}^{-1})$ uniformly for $i = 1, \ldots, N_{\text{co}}$ as $N_{\text{co}}, N_{\text{tr}} \to \infty$. However, in the context of SCs, it is often the case that the means of the confounders $\mathbf{Z}_{it}$ can be quite different between the treatment and control groups, which prevents attaining the uniform order of $O_p(N_{\text{co}}^{-1})$ for the weights on control units. We need develop results on the order of the upper bound of $\{w_{it}^*\}$ to establish the consistency result of $\widehat{\tau}_t$. In the following, we analyse $\{w_{it}^*\}_{i=1}^{N_{\text{co}}}$ at a fixed t by generalising (Ghosh & Chaudhuri, 2019) to the case of multivariate covariates and spatially dependent data.

For a unit vector $\theta \in \Theta$, define $\xi_i(\theta) = \theta'[\mathbf{Z}_{it}^0 - \mathrm{E}(\mathbf{Z}_{it}^0)]$ for $1 \le i \le N_{\text{co}}$. Let $\xi_{\min}(\theta) = \min_{i\in\mathcal{C}}\{\xi_i(\theta)\}$, $\xi_{\min,2}(\theta) = \min_{i\in\mathcal{C}}\{\{\xi_i(\theta)\}\backslash\{\xi_{\min}(\theta)\}\}$. The following assumptions on $\xi_i(\theta)$ are required to derive the asymptotic order of w_{it}^*.

Assumption 6 For any $\theta \in \Theta$ with $\|\theta\| = 1$, $\{\xi_i(\theta)\}$ has positive variance (non-degenerated) and $\sum_{i-j=-\infty}^{\infty} \mathrm{cov}(\xi_i(\theta), \xi_j(\theta)) < \infty$.

Assumption 7 There exists a non-random sequence $b_n \to \infty$, $\gamma > 0$, and $\delta > 0$, such that (a) $b_n^{2\gamma+2} = o_p(n)$ and $|\xi_{\min}(\theta)| = O(b_{N_{\text{co}}})(1 + o_p(1))$; (b) $|\xi_{\min}(\theta) - \xi_{\min,2}(\theta)| \ge b_{N_{\text{co}}}^{-\gamma}$ with probability converging to 1 as $N_{\text{co}} \to \infty$; and (c) $b_n^{\gamma+2}\mathrm{P}(|\xi_i(\theta)| > b_n^{1-\delta}) \to 0$ as $n \to \infty$.

Assumption 6 is to ensure the distribution of $\xi_i(\theta)$ is non-degenerate with positive variance and ensure the law of large numbers $\lim_{N_{\text{co}}\to\infty} \sum_{i=1}^{N_{\text{co}}} \xi_i/N_{\text{co}} = 0$ holds (Anderson, 1971). Assumption 7 restricts the growth rate of the extreme order statistics $|\xi_{\min}(\theta)|$, $|\xi_{\min}(\theta) - \xi_{\min,2}(\theta)|$, and the decay rate of the tail probability $\mathrm{P}(|\xi_i(\theta)| > b_n^{1-\delta})$. By the standard results of extreme value theory (Leadbetter et al., 1983), many distributions of $\xi_i(\theta)$ would allow $\{\xi_i(\theta)\}_{i=1}^{N_{\text{co}}}$ satisfying Assumption 7 given that (a) the distribution of $\xi_i(\theta)$ is in the domain of attraction of some extreme value distribution and (b) $\{\xi_{s_1}(\theta), \xi_{s_2}(\theta), \ldots\}$ is a stationary random field with weak dependence, where $\boldsymbol{s}_i \in \mathbb{R}^q$ denotes the 'spatial' location of unit i and $\xi_{\boldsymbol{s}_i}(\theta) = \xi_i(\theta)$ for $q \in \mathbb{Z}^+$. For example, in the case study of air pollution alert $\boldsymbol{s}_i \in \mathbb{R}^2$ denotes the geographical location of site i. Specifically, (a) means that there exist normalising constants a_n, b_n such that $F_{\xi_i(\theta)}^n(a_n x + b_n) \to G(x)$ where $F_{\xi_i(\theta)}$ denotes the distribution function of $\xi_i(\theta)$ and $G(x)$ is in the form of Gumbel, Fréchet or Weibull distributions. The weak dependence assumed in (b) is to allow the theory for independent random variables to go through in the dependent case. Leadbetter and Rootzén (1998) proved that the required weakly dependence can be formulated by a coordinate-wise mixing condition, which extends the typical 'strong mixing' condition in each coordinate direction. Doukhan and Gómez (2017) considered the upper bound of covariance $|\mathrm{Cov}(f(X_{\mathbf{i}_1}, \ldots, X_{\mathbf{i}_u}), g(X_{\mathbf{j}_1}, \ldots, X_{\mathbf{j}_v}))|$ of any pair of l−distant finite sequences $(\mathbf{i}_1, \ldots, \mathbf{i}_u)$, $(\mathbf{j}_1, \ldots, \mathbf{j}_v)$ and any pair of bounded function f, g

with finite Lipschitz modules. Specially, when $\{\xi_i(\theta)\}$ are Gaussian distributed, we may take $b_n = \sqrt{\log n}$, $\gamma > \frac{1}{2}$, and $0 < \delta < \frac{1}{2}$ such that $\{\xi_i(\theta)\}$ would satisfy Assumption 7.

Under Assumptions 6 and 7, we may establish the order of w_{it}^* for the general case of $\mathrm{E}(\mathbf{Z}_{it}|D_i = 0) - \mathrm{E}(\mathbf{Z}_{it}|D_i = 1) \neq 0$.

Lemma 2 When N_{co} and $N_{\mathrm{tr}} \to \infty$,

(i) if $\mathrm{E}(\mathbf{Z}_{it}|D_i = 0) - \mathrm{E}(\mathbf{Z}_{it}|D_i = 1) = 0$, then $w_{it}^* = O_p(N_{\mathrm{co}}^{-1})$ uniformly for $i = 1, \ldots, N_{\mathrm{co}}$;

(ii) if $\mathrm{E}(\mathbf{Z}_{it}|D_i = 0) - \mathrm{E}(\mathbf{Z}_{it}|D_i = 1) \neq 0$, then $\max\{w_{it}^* : i \neq i_{\min}\} = O_p(b_{N_{\mathrm{co}}}^{\gamma+1} N_{\mathrm{co}}^{-1})$ and $w_{i_{\min},t}^* = O_p(b_{N_{\mathrm{co}}}^{-1})$ for $i = i_{\min}$ under Assumption 7, where $i_{\min}$ is the index of $\xi_{\min}(\theta^*)$.

Lemma 2 characterises the asymptotic order of $\{w_{it}^*\}$. In the case when $\mathrm{E}(\mathbf{Z}_{it}|D_i = 0) - \mathrm{E}(\mathbf{Z}_{it}|D_i = 1) \neq 0$, the weights other than $w_{i_{\min},t}^*$ are bounded by $O_p(b_{N_{\mathrm{co}}}^{\gamma+1} N_{\mathrm{co}}^{-1}) = o_p(N_{\mathrm{co}}^{-1/2})$. Lemma 2 ensures that the error term η_t in equation (2.11) is stochastically small and leads to the consistency of $\widehat{\tau}_t$ as in the next theorem.

Theorem 1 Under Assumptions 1–7 and Model (2.3), the matching equations (2.8) and (2.9) can be attained with probability converging to 1 and there exists a unique solution for the EL SC weights $\{w_{it}^*\}_{i=1}^{N_{\mathrm{co}}}$. Given that the error terms $\{\varepsilon_{it}\}_{i=1}^{N}$ satisfy $\sum_{0 \leq |i-j| < N} \mathrm{E}(\varepsilon_{it}\varepsilon_{jt}) = O(N)$, the SC estimate $\widehat{\tau}_t \to \tau_t$ in probability as both N_{tr} and $N_{\mathrm{co}} \to \infty$.

Theorem 1 establishes the consistency of $\widehat{\tau}_t$ under the asymptotic regime where N_{tr} and $N_{\mathrm{co}} \to \infty$, while T is fixed, which is motivated by the setting of the air pollution alert study. Under a suggestion of a referee, we further examine the consistency of the DSC estimator $\widehat{\tau}_T$ for large T by imposing conditions on the relative rates of divergence between T and N_{co} as shown in the next proposition. Specifically, we need to assume the following conditions.

Assumption 8 In the case of $E[Z_{it}|D_i = 1] \equiv E[Z_{it}|D_i = 0]$, (a) if $\overline{\lim}_{T\to\infty}|\rho_T| > 1$, then $T = o_P(\log N_{\mathrm{co}}) = o_P(\log N_{\mathrm{tr}})$; (b) if $\overline{\lim}_{T\to\infty}|\rho_T| = 1$, then $O(T\prod_{t=1}^{T}|\rho_t|) = o_P(N_{\mathrm{tr}}^{1/2}) = o_P(N_{\mathrm{co}}^{1/2})$. In the case of $E[Z_{it}|D_i = 1] \neq E[Z_{it}|D_i = 0]$, (i) if $\overline{\lim}_{T\to\infty}|\rho_T| > 1$, then $T = o_P(\log b_{N_{\mathrm{co}}}) = o_P(\log N_{\mathrm{tr}})$; (ii) if $\overline{\lim}_{T\to\infty}|\rho_T| = 1$, then $O(T\prod_{t=1}^{T}|\rho_t|) = o_P(N_{\mathrm{tr}}^{1/2}) = o_P(b_{N_{\mathrm{co}}})$.

The conditions in Assumption 8 are provided under different situations depending on whether $\mathrm{E}[\mathbf{Z}_{it}|D_i = 1] = \mathrm{E}[\mathbf{Z}_{it}|D_i = 0]$ and the upper limits $\{\rho_t\}$ being larger than 1 or equal to 1, and does not require $Y_{it}(1)$ or $Y_{it}(0)$ as stationary. This is because the order of $(\widehat{\tau}_T - \tau_T)$ involves the auto-regressive coefficients $\{\rho_t\}$ and the synthetic weights $\{w_{it}^*\}$, while the order of $\{w_{it}^*\}$ is related to whether $\mathrm{E}[\mathbf{Z}_{it}|D_i = 1]$ is equal to $\mathrm{E}[\mathbf{Z}_{it}|D_i = 0]$ as described in Lemma 2. It is noted that if $\overline{\lim}_{T\to\infty}|\rho_T| < 1$, $\widehat{\tau}_T$ is consistent without requiring the conditions on the relative rates between T and N_{tr} and N_{co}. Thus, Assumption 8 does not need to provide for this case.

Proposition 1 Under Assumptions 1–8 and Model (2.3), $\widehat{\tau}_T \to \tau_T$ in probability as N_{tr}, N_{co}, and $T \to \infty$.

In practice, $\bar{\mathbf{Z}}_t^1$ may fall outside the convex hull $\mathcal{H}_t$ due to limited sample size, then the weights $\mathbf{w}_t^*$ can be solved by minimising the distance between the weighted average and the matching object. Specifically, let $\Delta_t(\mathbf{w}_t) = \| \sum_{i\in\mathcal{C}} w_{it}\mathbf{Z}_{it}^0 - \bar{\mathbf{Z}}_t^1 \|$. Then $\mathbf{w}_t^*$ can be obtained by minimising $\Delta_t(\mathbf{w}_t)$ subject to $\sum_{i=1}^{N_{\mathrm{co}}} w_{it} = 1$ and $w_{it} \geq 0$. This is a convex optimisation problem where the feasible region under constraint is bounded, closed and convex. Thus, there exists a global optimum and a solution is assured.

3 Statistical inference

In this section, we present the inference procedures associated with the proposed DSC method on how to (a) assess the unconfoundedness assumption and model specification using the pre-treatment data and (b) assess the significance of the estimated effects $\widehat{\tau}_t$.

3.1 Assess the unconfoundedness assumption and model specification

Recall that the unconfoundedness Assumption 2 requires that

$$\mathrm{E}\big(Y_{it}(0)|Y_{i,t-1}(0), \mathbf{X}_{it}, D=1\big) = \mathrm{E}\big(Y_{it}(0)|Y_{i,t-1}(0), \mathbf{X}_{it}, D=0\big).$$

Since $Y_{it}(0)$ and $Y_{i,t-1}(0)$ are observable for both of the treated units with $D_i = 1$ and the control units with $D_i = 0$ during the pre-treatment period when $t \le T_0$, we can use the pre-treatment data to assess a portion of this assumption.

As τ_t in equation (2.1) and $\widehat{\tau}_t$ in equation (2.10) are only defined for $t > T_0$ in the post-treatment period, to extend the notations to the pre-treatment periods, we define

$$\tau_t = \mathrm{E}\big(Y_{it} - Y_{it}(0)|D_i = 1\big) = \begin{cases} 0, & \text{if } 1 \le t \le T_0; \\ \mathrm{E}\big(Y_{it}(1) - Y_{it}(0)|D_i = 1\big), & \text{if } T_0 < t \le T. \end{cases} \tag{3.1}$$

Then, the SC estimate $\widehat{\tau}_t$ can be also extended to the entire time span,

$$\widehat{\tau}_t = \sum_{i\in\mathcal{T}} \frac{1}{N_{\mathrm{tr}}} Y_{it} - \sum_{i\in\mathcal{C}} w_{it}^* Y_{it}, \quad 1 \le t \le T, \tag{3.2}$$

where w_{it}^* for $t \le T_0$ is also defined by the constrained EL optimisation problem (2.7).

Note that $Y_{it} = Y_{it}(0)$ for the treated units with $D_i = 1$ when $t \le T_0$, so the first term in equation (3.2) is the sample average $\bar{Y}_{\cdot t}(0)$ in the treatment group. Also, recall that $\sum_{i\in\mathcal{C}} w_{it}^* Y_{it}$ is an estimate of $\mathrm{E}(Y_{it}(0)|D_i = 1)$ by equation (2.4). Therefore, if the SC $\sum_{i\in\mathcal{C}} w_{it}^* Y_{it}$ is a good estimation for $\mathrm{E}(Y_{it}(0)|D_i = 1)$, $\widehat{\tau}_t$ is expected to be around zero in the pre-treatment period. Specifically, the null hypothesis at $t \le T_0$ is

$$H_0^t: Y_{it}(0) = \delta_t + \boldsymbol{\beta}_t' \mathbf{X}_{it} + \rho_t Y_{it-1}(0) + \varepsilon_{it} \quad \text{holds for } i \in \mathcal{C} \cup \mathcal{T}. \tag{3.3}$$

Under H_0^t, the SC estimate in pre-treatment periods

$$\widehat{\tau}_t = \sum_{i\in\mathcal{T}} \frac{\varepsilon_{it}}{N_{\mathrm{tr}}} - \sum_{i\in\mathcal{C}} w_{it}^* \varepsilon_{it}. \tag{3.4}$$

And note that the weights $w_{it}^* \in \sigma(\{\mathbf{X}_{it}, Y_{i,t-1}\}_{i=1}^N)$, we have $\mathrm{E}(\widehat{\tau}_t) = 0$.

To test $H_0^t: \mathrm{E}(\widehat{\tau}_t) = 0$ at a $t \le T_0$, we consider the distribution of $\widehat{\tau}_t$ conditional on $\mathcal{F}_t = \sigma\{\mathbf{X}_{i1}, \ldots, \mathbf{X}_{it}, Y_{i1}, \ldots, Y_{i,t-1}\}$, the information set of confounding covariates and lagged outcomes up to time t. Assume $\{\varepsilon_{it}\}_{i=1}^N$ are spatially stationary and isotropic, and there exists a function $\sigma_t^2(\cdot)$ such that $\mathrm{Cov}(\varepsilon_{it}, \varepsilon_{jt}) = \sigma_t^2(\|\boldsymbol{s}_i - \boldsymbol{s}_j\|)$, where $\boldsymbol{s}_i$ denotes unit i's spatial location. The conditional variance $\mathrm{Var}(\widehat{\tau}_t|\mathcal{F}_t) = N_{\mathrm{tr}}^{-2} \sum_{i,j\in\mathcal{T}} \sigma_t^2(d_{ij}) + \sum_{i,j\in\mathcal{C}} w_{it}^* w_{jt}^* \sigma_t^2(d_{ij}) + \sum_{i\in\mathcal{T},j\in\mathcal{C}} N_{\mathrm{tr}}^{-1} w_{jt}^* \sigma_t^2(d_{ij})$ for $t \le T_0$, where $d_{ij} := \|\boldsymbol{s}_i - \boldsymbol{s}_j\|$. In implementation, $\sigma^2(d)$ is estimated by the sample covariance $\widehat{\sigma}^2(d)$ from the residuals obtained from fitting the dynamic linear model (2.3) using the ordinary least square method. Let $\widehat{\sigma}_{\widehat{\tau}_t}^2$ be the corresponding estimation for $\mathrm{Var}(\widehat{\tau}_t|\mathcal{F}_t)$. Then we use $\widehat{\tau}_t/\widehat{\sigma}_{\widehat{\tau}_t}$ as the test statistic.

Note that there exists testing multiplicity as we need to test H_0^t repeatedly for $t = 1, \ldots, T_0$. To address this issue, we adopt the false discovery rate (FDR) procedure (Benjamini & Yekutieli,

2001) to control the expected rate of false 'discoveries' under arbitrary dependence assumptions. Let p_t be the p-value for testing at time point t and $p_{(1)} \leq p_{(2)} \leq \cdots \leq p_{(T_0)}$ be their ordered values. To control FDR at level α, we compare $p_{(j)}$ with $\alpha \frac{j}{T_0 \times C(T_0)}$, where $C(T_0) = \sum_{i=1}^{T_0} \frac{1}{i}$ is the harmonic number. If there exists $p_{(j)} < \alpha \frac{j}{T_0 \times C(T_0)}$, let $\bar{j} = \max\{j : p_{(j)} < \alpha \frac{j}{T_0 \times C(T_0)}\}$, the null hypothesis is rejected for the time points corresponding to $p_{(1)}, \ldots, p_{(\bar{j})}$.

In applications, we may adopt the above testing method to guide the selection of the appropriate model specification in equation (2.3). Specifically, we may select the covariates first based on the domain knowledge and include the first-order lagged outcomes at first. If the testing results suggest violation of the model specification, we may either change the covariates or include higher order lagged outcomes in the dynamic linear model, as demonstrated in an empirical sensitivity analysis in the context of the air pollution alert in Section 5. Furthermore, one may consider a non-linear extension via the time-varying generalised additive model $Y_{it}(0) = \delta_t + m_y(t, Y_{i,t-1}(0)) + \sum_{k=1}^{p} m_k(t, X_{it}[k]) + \varepsilon_{it}$, where we denote $X_{it}[k]$ as the kth covariate of $\boldsymbol{X}_{it}$ of unit i at time t. However, the non-linear case is much more involved as it requires constructing an SC object for each unit in the treatment group, in stead of one SC for the whole treatment group in the linear case, which requires further analysis.

3.2 Normalised placebo test

To evaluate the significance of the estimated treatment effects, Abadie et al. (2010) adopted an inferential technique called the placebo test, which examines whether the absolute values of the treatment effect estimated by the SC method for the treatment group is large relative to the pseudo-effect for a randomly selected group unaffected by the treatment. The placebo test is a version of the permutation test (Canay et al., 2017). In the context of synthetic control method, the validity of this test requires that the distributions of the estimated effects are identical for the treatment group and the randomly selected groups given that there is no treatment effect. In this section we need to additionally assume that $\{\varepsilon_{it}\}_{i=1}^{N}$ are normally distributed with the same variance among units, as Hahn and Shi (2017) proved that the normality of ε_{it} is necessary for such property to be correct up to a normalisation.

Before introducing the normalised placebo test, we analyse $\widehat{\tau}_t$ under the null hypothesis that there is no treatment effect and the potential outcomes $Y_{it}(1) = Y_{it}(0)$. Let $\eta_t = \sum_{i \in \mathcal{C}} w_{it}^* \varepsilon_{it} - \sum_{j \in \mathcal{T}} \frac{\varepsilon_{jt}}{N_{\text{tr}}}$, and $\sigma_{\eta_t}^2 = \text{Var}(\eta_t | \mathcal{F}_t) = N_{\text{tr}}^{-2} \sum_{i,j \in \mathcal{T}} \sigma_t^2(d_{ij}) + \sum_{i,j \in \mathcal{C}} w_{it}^* w_{jt}^* \sigma_t^2(d_{ij}) + \sum_{i \in \mathcal{T}, j \in \mathcal{C}} N_{\text{tr}}^{-1} w_{jt}^* \sigma_t^2(d_{ij})$ for $1 \leq t \leq T$. (3.4) shows that when $t \leq T_0$, $\widehat{\tau}_t = \eta_t$. For a time point $t > T_0 + 1$ in the post-intervention period, $\widehat{\tau}_t = \eta_t + \sum_{s=T_0+1}^{t-1} \rho_s^t \eta_{t-s}$ by equation (2.11). In summary, under the null hypothesis that there is no treatment effect, we have

$$\widehat{\tau}_t = \widehat{\mu}_t(0) - \widetilde{\mu}_t(0) = \begin{cases} \eta_t, & \text{for } t \leq T_0 + 1; \\ \eta_t + \sum_{s=T_0+1}^{t-1} \rho_s^t \eta_{t-s}, & \text{for } t > T_0 + 1. \end{cases} \tag{3.5}$$

The accumulation of error terms when $t > T_0 + 1$ is resulted from using $\widehat{\mu}_{t-1}(0)$ as an approximation for $\widetilde{\mu}_t(0)$ in the matching constraint (2.9b). Let σ_t be the conditional standard deviation of $\widehat{\tau}_t$, then $\sigma_t = \sigma_{\eta_t}$ for $t \leq T_0 + 1$ and $\sigma_t = (\sigma_{\eta_t}^2 + \sum_{s=T_0+1}^{t-1} (\rho_s^t)^2 \sigma_{\eta_s}^2)^{\frac{1}{2}}$ for $t > T_0 + 1$. As σ_t is related to weights $\{w_{it}^*\}_{i=1}^{N_{\text{co}}}$, and the weights would change for another pair of treatment/control groups in the placebo runs, we should normalise the placebo estimates of τ_t with respect to the corresponding deviations. Details will be introduced later.

In implementing the placebo test, a group of N_{tr} control units are randomly selected from the control group $\mathcal{C}$ to be a placebo treatment group. Let $\mathcal{T}^{(k)}$ denote the set of indexes of the placebo group in the kth round of random selection, and $\mathcal{C}^{(k)} = \mathcal{C} \setminus \mathcal{T}^{(k)}$ be the control group for $\mathcal{T}^{(k)}$. Then we apply the proposed DSC method to the placebo treatment group $\mathcal{T}^{(k)}$ and the control group $\mathcal{C}^{(k)}$ as if there was a policy intervention imposed on $\mathcal{T}^{(k)}$ from $T_0 + 1$. Let $\{w_{it}^{*(k)}\}$ be the EL SC weights

for the treatment group $\mathcal{T}^{(k)}$ and define the placebo estimate for the pseudo-effect at time t

$$\widehat{\tau}_t^{(k)} = \sum_{j\in\mathcal{T}^{(k)}} \frac{1}{N_{\text{tr}}} Y_{jt} - \sum_{i\in\mathcal{C}^{(k)}} w_{it}^{*(k)} Y_{it}.$$

Let $\eta_t^{(k)} = \sum_{i\in\mathcal{T}^{(k)}} \frac{\varepsilon_{it}}{N_{\text{tr}}} - \sum_{i\in\mathcal{C}^{(k)}} w_{it}^{*(k)}\varepsilon_{it}$. Note that $Y_{it} = Y_{it}(0)$ for $i \in \mathcal{C}$ by Assumption 1 and $\mathcal{T}^{(k)} \subset \mathcal{C}$, $\mathcal{C}^{(k)} \subset \mathcal{C}$, the placebo estimates

$$\widehat{\tau}_t^{(k)} = \eta_t^{(k)} \quad \text{for} \quad t \le T_0 + 1 \quad \text{and} \quad \widehat{\tau}_t^{(k)} = \eta_t^{(k)} + \sum_{s=T_0+1}^{t-1} \rho_s^t \eta_s^{(k)} \quad \text{for} \quad t > T_0 + 1. \tag{3.6}$$

Since in most applications N_{co} is large relative to N_{tr}, such procedures can be implemented for many rounds, say K times, to obtain a set of placebo estimates.

The rationale of the placebo test is that under the null hypothesis of no treatment effects, the empirical distribution $\{\widehat{\tau}_t^{(k)}\}_{k=1}^K$ are used to approximate the distribution $\widehat{\tau}_t$. While the variance of $\widehat{\tau}_t^{(k)}$ relies on the DSC weights $\{w_{it}^{*(k)}\}$, we need to normalise the estimated pseudo-effects $\widehat{\tau}_t^{(k)}$ to attain the same variance as $\widehat{\tau}_t$. Let

$$z_t^{(k)} = \frac{\widehat{\tau}_t^{(k)}}{\widehat{\sigma}_t^{(k)}} \times \widehat{\sigma}_t$$

be the normalised SC estimations, where $\sigma_t^{(k)}$ is the standard deviation of $\widehat{\tau}_t^{(k)}$. Then we use the empirical distribution of $\{z_t^{(k)}\}_{k=1}^K$ to approximate the distribution of $\widehat{\tau}_t$ under the null hypothesis that there is no treatment effect at time t. When reducing effects are expected as in air pollution alerts, the alternative hypothesis is $\tau_t < 0$ and the p-value is approximated empirically by random replications as $K^{-1}\sum_{k=1}^K 1(z_t^{(k)} \le \widehat{\tau}_t)$.

4 Simulation study

We report results from simulation experiments designed to validate the main theoretical results established in the previous sections that (a) the exact matches can be attained under Assumption 5, (b) the proposed DSC method provides a consistent estimation to the underlying treatment effect, and compare the DSC method with SC method. The simulations were conducted under both cases in Lemma 2 such that $\mathrm{E}(\mathbf{Z}_{it}|D_i = 0) - \mathrm{E}(\mathbf{Z}_{it}|D_i = 1)$ is zero and non-zero. We also investigated the accuracy of the DSC estimator $\widehat{\tau}_t$ and the profiles of the dynamic synthetic control weights to confirm the results of Lemma 2.

The samples of the potential outcomes were generated according to

$$\begin{aligned} Y_{it}(0) &= \delta_t + \rho Y_{i,t-1}(0) + \beta'\mathbf{X}_{it} + \varepsilon_{it}; \\ Y_{it}(1) &= Y_{it}(0) + \tau_{it}, \end{aligned} \tag{4.1}$$

where $\varepsilon_{it} \sim N(0, \sigma_\varepsilon^2)$, $\tau_{it} \sim N(\tau_t, \sigma_\tau^2)$, and $\mathbf{X}_{it}$ represented p-dimensional time-varying confounders. We set the underlying treatment effect $\tau_t = -30/[1 + \exp\{-0.2(t - T_0 - 6)\}]$, $\sigma_\varepsilon^2 = \sigma_\tau^2 = 1$, $\rho = 0.8$, $p = 4$, and $\beta = (1, 1, 1, 1)^T$. Then the samples of observed outcomes were generated by $Y_{it} = Y_{it}(0)\mathbb{I}(t \le T_0) + Y_{it}(1)\mathbb{I}(t > T_0)$ for $i \in \mathcal{T}$ and $Y_{it} = Y_{it}(0)$ for $i \in \mathcal{C}$.

We considered three settings for $\{\{\mathbf{X}_{it}\}_{i=1}^N\}_{t=1}^T$ generation: (a) spatially and temporally independent Gaussian random variables, (b) temporally independent but spatially dependent Gaussian random variables, and (c) temporally dependent but spatially independent setting with Gamma innovations. In Setting (b), the spatial dependence across units was based on an isotropic semivariogram $\gamma_k(\boldsymbol{s}_{i_1}, \boldsymbol{s}_{i_2}) = \frac{1}{2}\mathrm{Var}(X_{i_1tk} - X_{i_2tk}) = 1 - e^{\|\boldsymbol{s}_{i_1} - \boldsymbol{s}_{i_2}\|}$, where X_{itk} is the kth dimension of $\mathbf{X}_{it}$ and $\boldsymbol{s}_i$ is the spatial coordinate of the ith unit. We generated $\{\mathbf{X}_{it}\}_{i=1}^N$ according to the covariance

matrix determined by the semi-variogram at each t. Setting (c) generated the covariate $\mathbf{X}_{it}$ from an auto-regressive model $\mathbf{X}_{it} = \boldsymbol{\mu}_i + \boldsymbol{\Gamma}_{i,t} + 0.5\boldsymbol{\Gamma}_{i,t-1}$ for each i, where each dimension of $\boldsymbol{\Gamma}_{i,t}$ and $\boldsymbol{\Gamma}_{i,t-1}$ are independent Gamma(4, 2) distributed random variables centred by the expectations.

Under each of the three settings for $\mathbf{X}_{it}$, we conducted simulation experiments for the cases having the identical or non-identical means between the treatment and control groups, which corresponded to cases (i) and (ii) of Lemma 2, respectively. In the non-identical mean case, the covariates $\mathbf{X}_{it}$ in the control group were generated from a mixture distribution such that one half of the samples shared the same distribution as the treatment group and the other half had a shift $\boldsymbol{d}$ in the mean $\boldsymbol{\mu}$. Under such design, there was a sub-population of the control group that had the same expectation of $\mathbf{X}_{it}$ with the treatment group satisfying Assumption 5. Taking Setting (a) as an example, $\mathbf{X}_{it}$ were generated from $N(\boldsymbol{\mu}, I_p)$ in the treatment group and were from a mixture distribution $\frac{1}{2}N(\boldsymbol{\mu}, I_p) + \frac{1}{2}N(\boldsymbol{\mu} + \boldsymbol{d}, I_p)$ in the control group. The mean $\boldsymbol{\mu} = (1, 5, 10, 15)$ and the shift $\boldsymbol{d} = \delta \times (1, 1, 1, 1)$ for $\delta = 0$, 0.5 and 1, respectively.

We set the total time $T = 72$ and the pre-treatment duration $T_0 = 48$, the same as the case study in the next section. The number of the treatment units $N_{\text{tr}} = 100$, and the number of control units $N_{\text{co}} = 100$ and 500 for each experiment setting, respectively. The experiments were replicated for 1,000 times for each setting. We implemented the proposed DSC method in equations (2.8) and (2.9), the SC method in equation (2.15), and the naive method of taking raw differences.

To make the comparison with the SC fairer, as suggested by a referee, we considered two extensions of the SC. The first extension, which is called SC_ext1, has the weights $\{w_i^{\text{SC\_ext1}}\}$ still being static, but the covariate values $\boldsymbol{X}_{it}$ for $t > T_0$ are also added to the matching equations such that

$$\sum_{i\in\mathcal{C}} w_i^{\text{SC\_ext1}} Y_{iT_0} = \sum_{i\in\mathcal{T}} Y_{iT_0}/N_{\text{tr}} \quad \text{and} \quad \sum_{i\in\mathcal{C}} w_i^{\text{SC\_ext1}} \boldsymbol{X}_{it} = \sum_{i\in\mathcal{T}} \boldsymbol{X}_{it}/N_{\text{tr}} \quad \text{for} \quad t = T_0, \ldots, T.$$

The second extension, called SC_ext2, allows the weights $\boldsymbol{w}_t^{\text{SC\_ext2}}$ to be dynamically updated with respect to $\boldsymbol{X}_{it}$ for each $t > T_0$, but does not involve the recursive matching of lagged outcomes and the EL.

$$\sum_{i\in\mathcal{C}} w_{it}^{\text{SC\_ext2}} Y_{iT_0} = \sum_{i\in\mathcal{T}} Y_{it_0}/N_{\text{tr}} \quad \text{and} \quad \sum_{i\in\mathcal{C}} w_{it}^{\text{SC\_ext2}} \boldsymbol{X}_{it} = \sum_{i\in\mathcal{T}} \boldsymbol{X}_{it}/N_{\text{tr}}.$$

Table 1. The average root mean square errors (RMSEs) of estimated treatment efforts $\widehat{\tau}_t$ based on 1,000 replications for the DSC, the SC, the SC_ext1, the SC_ext2, and the Raw-difference

X_{it} setting	Method	$\delta = 0$		$\delta = 0.5$		$\delta = 1$	
		$N_{\text{co}} = 100$	$N_{\text{co}} = 500$	$N_{\text{co}} = 100$	$N_{\text{co}} = 500$	$N_{\text{co}} = 100$	$N_{\text{co}} = 500$
Independent normally distributed	DSC	0.24	0.18	0.59	0.47	1.46	0.69
	SC	0.54	0.41	4.46	4.38	8.98	8.66
	SC_ext1	0.51	0.39	2.78	2.66	6.82	6.35
	SC_ext2	0.51	0.40	3.72	3.38	7.62	6.95
	Raw-diff	0.53	0.41	5.03	5.01	10.01	10.12
Spatially dependent normally distributed	DSC	0.62	0.47	1.21	0.77	2.89	0.87
	SC	0.87	0.68	5.42	5.13	10.37	10.02
	SC_ext1	0.85	0.65	2.89	2.31	7.12	6.27
	SC_ext2	0.85	0.67	4.30	4.06	8.42	7.93
	Raw-diff	0.85	0.67	6.53	6.55	13.19	13.31
Temporally dependent gamma innovation	DSC	0.19	0.13	0.37	0.23	0.68	0.35
	SC	0.31	0.23	2.94	2.87	5.31	5.26
	SC_ext1	0.28	0.19	2.32	2.19	3.57	3.69
	SC_ext2	0.29	0.22	2.40	2.62	4.67	3.82
	Raw-diff	0.32	0.21	4.13	4.21	8.25	8.34

Note. The RMSEs were averaged over the post-treatment period (48 h) under three settings of the confounders X_{it} and with difference amount of disparity ($\boldsymbol{d}$) in the means between the treatment and control groups. DSC = dynamic synthetic control; SC = synthetic control.

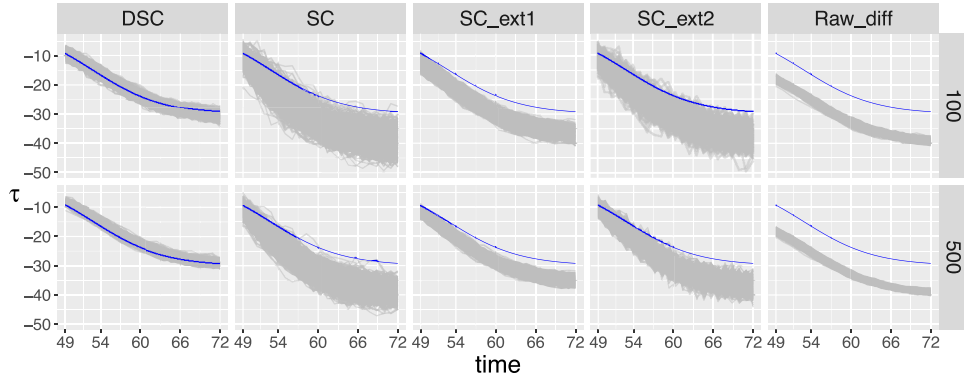

Figure 1. Estimated treatment effects $\widehat{\tau}_t$ (semitransparent grey lines) of 1,000 simulations and the underlying effects τ_t (non-transparent blue lines) by the DSC, SC, SC_ext1, SC_ext2, and Raw-difference methods under Setting (a) for the covariates with $\delta = 1$ and $N_{co} = 100$ (top row) and 500 (bottom row). DSC = dynamic synthetic control; SC = synthetic control.

Table 1 reports the averaged root mean squared errors (RMSEs) over the post-treatment period. It was observed that the exact matches were attained for all replications of the simulation at all time points for the DSC weights $\{w_{it}^*\}$ with numerical errors $\| \sum_{i \in \mathcal{C}} w_{it}^* \mathbf{Z}_{it}^0 - \bar{\mathbf{Z}}_t^1 \| < 10^{-6}$, which confirmed Theorem 1 that the matching equations (2.8) and (2.9) can be attained with probability converging to 1. In contrast, the weights $\{w_i^*\}$ of the SC method in equation (2.15) only matched $\mathbf{X}_{it}$ and Y_{it} at $t = T_0$, which could not control the bias in the post-treatment periods. The average matching errors $\sum_{t=T_0+1}^{T} \| \sum_{i \in \mathcal{C}} w_i^* \mathbf{Z}_{it}^0 - \bar{\mathbf{Z}}_t^1 \| / (T - T_0)$ with static SC weights ranged from 0.24 to 5.82 over the considered experiment settings, which explained the bias of SC estimates.

Table 1 shows that the two extended SC, SC_ext1, and SC_ext2, were able to reduce the RMSEs of the original SC method by adding the post-treatment covariates in the matching equations, especially in the experiment settings with confounding bias under $\delta > 0$. Between the two extensions, SC_ext1 had lower RMSEs, which was benefited from more matching equations. However, both SC_ext1 and SC_ext2 were unable to remove the estimation bias due to not matching the lagged outcomes.

Figure 1 shows more detailed results at each time point in all replications by the five methods when $\delta = 1$ for Setting (a) for the covariates with spatially and temporally independent Gaussian $\boldsymbol{X}_{it}$ as were introduced after equation (4.1), where the grey and blue lines correspond to the estimates $\hat{\tau}_t$ and the underlying τ_t, respectively. Comparing the results of the DSC method with those of the four methods, the DSC method had the smallest RMSEs in all settings. More importantly, as the sample size N_{co} was increased from 100 to 500 in each setting, the RMSEs of the estimated treatment effects were reduced indicating the consistency of the DSC estimation. In contrast, as N_{co} was increased, the RMSEs of the three versions of SC and the Raw-difference method were not much reduced due to the systematic bias of the two estimators. Compared among the original SC method, both SC_ext1 and SC_ext2 showed smaller bias and lower variance by adding more information to the matching equations, especially for SC_ext1. It was observed that as the disparity parameter δ in the confounders X_{it} was increased, the RMSEs of all three SC estimators got larger. In contrast, the DSC still managed to reduce the RMSEs by increasing the sample size while the SC-related methods and the Raw-difference method were still severely biased.

The bias of the SC method was due to the static matching procedure. As shown in Figure 1, as the SC only matched $\mathbf{X}_{iT_0}$ and Y_{iT_0}, the bias of $\widehat{\tau}_t^{\text{DSC}}$ originated from varying $\{\mathbf{X}_{is}\}_{s=T_0+1}^{t}$ accumulated in the post-treatment periods. The bias of the Raw-difference method originates from the distribution shift of covariates $\boldsymbol{X}_{it}$ between the treatment and control groups when $\boldsymbol{d} \neq \mathbf{0}$. The increased sample size could only reduce the variance of the estimations while was not able to eliminate the bias, which explained that the RMSEs reduced little for the SC method with increased sample size. It is noted that the DSC demonstrated a smaller variance than the SC in the simulation results reported in Figure 1, which confirmed the theoretical derivation conveyed in (2.16) that shows the SC would incur accumulation of variance as the SC does not take into account the time-

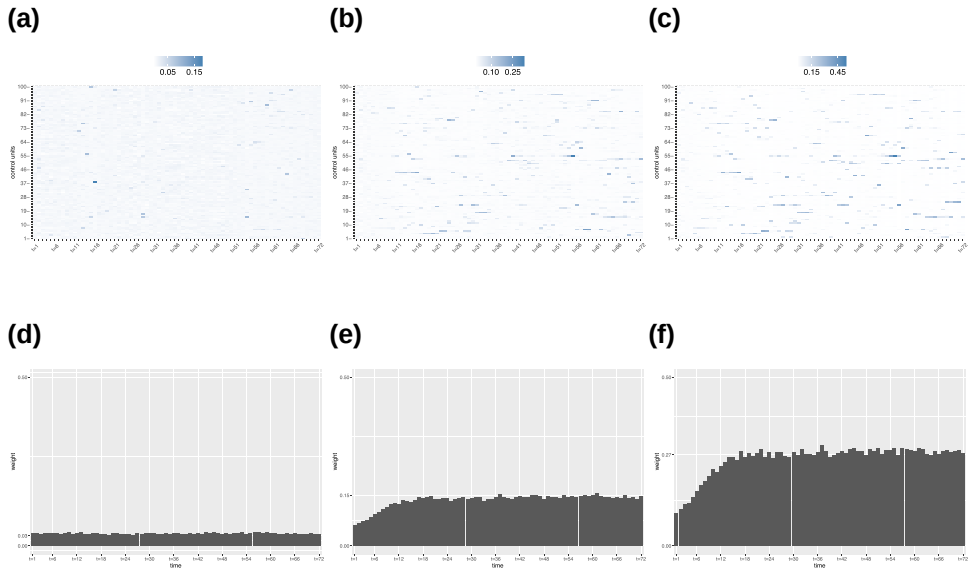

Figure 2. The dynamic synthetic control weights $\{w_{it}^*\}_{i=1\,t=1}^{N_{co}\,T}$ in one replication under Setting (a) and different disparity parameters $\boldsymbol{d}$ (a–c), and the average maximum weights for each time point t over the 1,000 replications (d–f) with the sample size $N_{tr} = N_{co} = 100$. (a) $\boldsymbol{d} = (0, 0, 0, 0)$, (b) $\boldsymbol{d} = (0.5, 0.5, 0.5, 0.5)$, (c) $\boldsymbol{d} = (1, 1, 1, 1)$, (d) $\boldsymbol{d} = (0, 0, 0, 0)$, (e) $\boldsymbol{d} = (0.5, 0.5, 0.5, 0.5)$, and (f) $\boldsymbol{d} = (1, 1, 1, 1)$.

varying confounding of the covariates. Except for the point estimation, we implemented the placebo test to examine the significance of the estimated effects by the DSC method shown in Figure 1. The results showed that the underlying τ_t is outside the confidence intervals in the 500 replications without committing the type-2 errors. To examine the coverage under the null hypothesis (3.3) , we also added a simulation scenario with $\tau_t \equiv 0$. The results showed that the coverage were around 95% in all time points, see online supplementary material for details.

To demonstrate the results on the synthetic weights, Figure 2a–c provides the heat-maps of the DSC weights $\{w_{it}^*\}_{i=1\,t=1}^{N_{co}\,T}$ in one replication for $\delta = 0$, 0.5, 1 and the sample size $N_{co} = 100$ under the basic Setting (a). To make up for the fact that the heat-maps were only based in one replication, Figure 2d–f shows the average maximum weights over the 1,000 replications at each time t. Figure 2a shows that most of the DSC weights appeared to be fairly uniformly scattered for $\boldsymbol{d} = 0$ when there was no confounding bias, which confirmed the result of Lemma 2 that $w_{it}^* = O_p(N_{co}^{-1})$ having a uniform bound when $\mathrm{E}(\mathbf{Z}_{it}|D_i = 0) - \mathrm{E}(\mathbf{Z}_{it}|D_i = 1) = 0$. In Figure 2b and c where $\boldsymbol{d} \neq \mathbf{0}$, the maximum weights at each time point were much higher than the other weights, which echoed Lemma 2 that there exists one weight in $\{w_{it}^*\}_{i=1}^{N_{co}}$ with a larger order than the other weights when $\mathrm{E}(\mathbf{Z}_{it}|D_i = 0) - \mathrm{E}(\mathbf{Z}_{it}|D_i = 1) \neq 0$. Therefore, $\sum_{i=1}^{N_{co}} (w_{it}^*)^2$ would be larger in the presence of the confounding bias, which led to a larger variance in the estimation $\widehat{\tau}_t$ according to equation (2.11), which explained the inferior performance of the DSC when $\delta \neq 0$ in Table 1. Figure 2d–f also shows that the increasing degree of confounding bias led to heavier heterogeneity among the weights $\{w_{it}^* | i = 1, \ldots, N_{co}\}$ of all control units. Note that Figure 2e and f indicates an increasing trend in the beginning, which was because $Y_{it}(0)$ at $t = 1$ was sampled from the same distribution in the treatment and control groups but differences accumulated due to the disparity of $\mathbf{X}_{it}$.

5 Evaluating air pollution alerts

We estimated the effects of two air pollution alerts in Beijing, which motivated our proposal of the DSC method. The air pollution alerts in Beijing have four colour-coded categories, blue, yellow, orange, and red alerts in the ascending severity. Specific instructions of emergency measures are given according to each category. During the studied air pollution alerts, the emission-reduction measures were solely conducted in Beijing, which allowed us to use the monitoring stations in

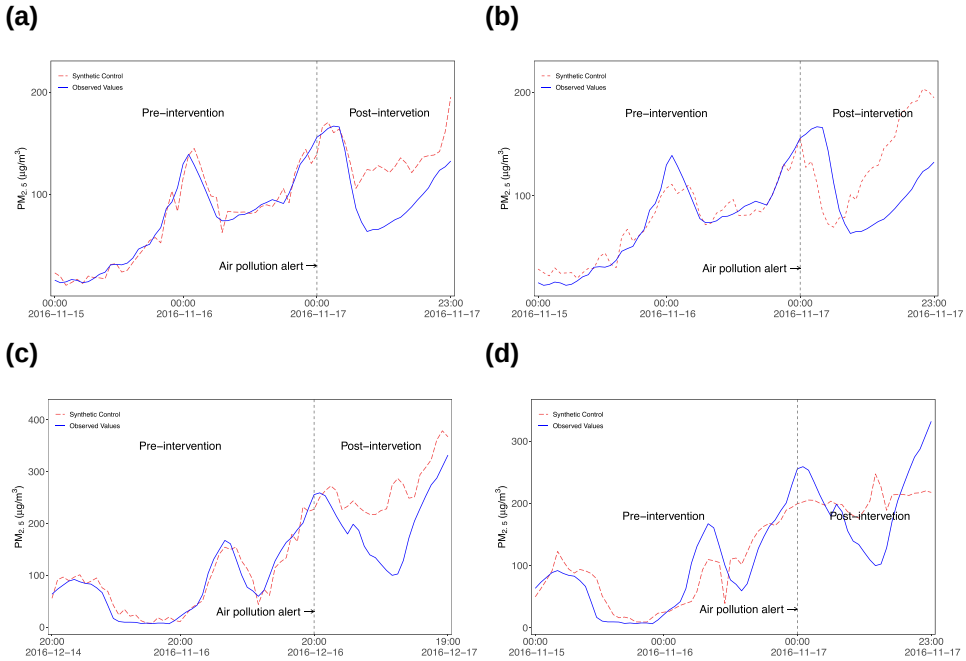

Figure 3. The estimated average potential outcomes $\widehat{\mu}_t(0)$ (dashed line) and the observed average $\sum_{i\in\mathcal{T}} \frac{Y_{it}}{N_{\text{tr}}}$ (solid line) for the two alerts by the DSC and SC methods, respectively. (a) The orange alert (DSC). (b) The orange alert (SC). (c) The red alert (DSC) and (d) The red alert (SC). DSC = dynamic synthetic control; SC = synthetic control.

the neighbouring provinces as control units to estimate the potential levels of air pollution in Beijing in the absence of the alerts.

We used hourly pollution data at the monitoring station level, which enabled the incorporation of multiple treated units with sufficient control units. Our analysis involved an orange alert starting from 00:00 on 17 November 2016 and a red alert starting from 20:00 on 26 December 2016. We aimed to estimate the effects of these alerts on the hourly $PM_{2.5}$ concentration during the 24 h after the start of the alerts.

The region under treatment was Beijing, where 20 air-quality stations located in the urban core districts were considered excluding sites in the outer suburbs to remove the potential interference from the non-treatment regions outside Beijing. To construct the synthetic controls, we considered 74 air-quality stations in Tianjin, Hebei, Shandong, and Shanxi, the municipality and provinces neighbouring Beijing, as there were no air pollution alerts in these regions and the air pollution levels were comparable to those in Beijing.

The outcome variable was the hourly $PM_{2.5}$ concentration measured in $\mu g/m^3$. Since $PM_{2.5}$ concentration is highly influenced by meteorological conditions (Chen et al., 2017; Liang et al., 2015), we considered meteorological variables as covariates. Hourly meteorological variables for each air-quality monitoring station was sourced from the nearest meteorological station of China's Central Meteorological Agency (CMA). We selected four meteorological variables, wind speed, humidity, dew point temperature, and air pressure, which are known to be influential on the $PM_{2.5}$ levels. These four meteorological variables and the lagged $PM_{2.5}$ were used to create the synthetic control weights.

There were longitudinal measurements at $N = 94$ monitoring stations, which consisted of $N_{\text{tr}} = 20$ stations in the treatment region and $N_{\text{co}} = 74$ stations in the control region. Pre-treatment periods were extended back for $T_0 = 48$ h from the start of the alerts for verification of Assumption 2 and model specification (2.3). We applied the proposed method to estimate the effect of pollution alerts on the hourly $PM_{2.5}$ during the 24 h after the intervention. So the total length of the studied time span is $T = 72$.

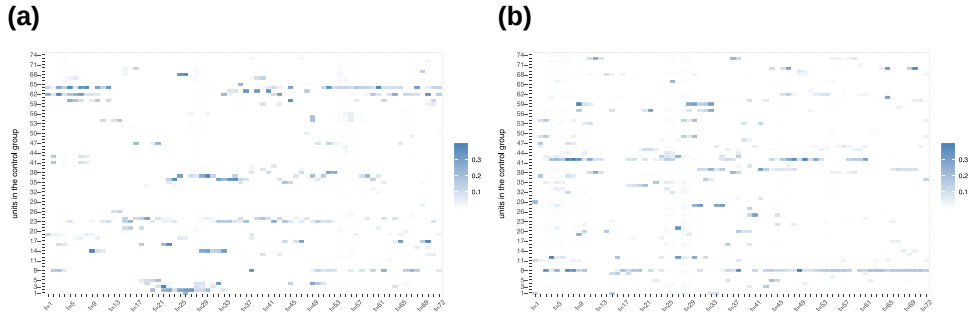

Figure 4. The dynamic synthetic control weights $\{w_{it}^*\}_{i=1\,t=1}^{N_{co}\,T}$ for the two alerts: (a) the orange alert and (b) the red alert.

Figure 3 shows the DSC estimates $\widehat{\mu}_t(0)$ and the average observed outcomes $\sum_{i\in\mathcal{T}} Y_{it}/N_{\text{tr}}$ in Beijing in periods before and after the pollution alerts by the proposed DSC method and the SC method. The synthetic Beijing (dashed line) was constructed by the weighted average of the monitoring stations in the regions under control at each time period, to estimate the counterfactual levels of $PM_{2.5}$ if the pollution alerts were not issued. The trajectories of the observed average $PM_{2.5}$ in the Beijing area exhibited a diurnal pattern and the DSC estimates also displayed similar periodic patterns both before and after the starts of the pollution alerts. During the pre-intervention periods of both alerts and the early stage of the post-intervention periods of the orange alert, the DSC estimates closely tracked the trajectory of the observed average $PM_{2.5}$ concentrations in Beijing. For the post-intervention periods, the DSC estimates showed an overall higher level of $PM_{2.5}$ concentrations and ended up with a higher peak for both alerts, which suggested that the emission-reduction measures prevented the potential deterioration of the air quality. By contrast, when we adopted the SC method, the static weights $\{w_i^*\}_{i\in\mathcal{C}}$ could not exactly match all the concentrations in the pre-treatment periods and would underestimate the post-alert effects.

The average treatment effects on the treated $\{\tau_t\}$ were estimated by the differences between the observed averages in Beijing and the SC estimates $\widehat{\mu}_t(0)$. For the orange alert, the potential average concentration of $PM_{2.5}$ would be $139.0\mu g/m^3$ in the absence of the contingency measures offered by the alert, while the observed average was $105.3\mu g/m^3$. The $PM_{2.5}$ was reduced by $33.8\mu g/m^3$ on average with standard error (SE) being $4.4\mu g/m^3$, which accounts for 24.3% of the potential level. For the red alert, the emission-reduction measures reduced the $PM_{2.5}$ concentration from a potential average of $269.2\mu g/m^3$ to $198.8\mu g/m^3$, where the decreasing effect was $70.4\mu g/m^3$ on average ($SE = 9.6\mu g/m^3$) and the reduction is 26.2%. This suggested that it was necessary to launch the red alert as otherwise the pollution level would be much higher.

Figure 4 shows the optimal weights $\{w_{it}^*\}_{i=1\,t=1}^{N_{co}\,T}$ calculated by the definitions in (2.7)–(2.9), where the range of values are signified by colour. It can be seen that at each time point, there existed only a few stations with relatively large weights (signified by deeper blue), while the other stations were endowed with small weights that are close to zeros (signified by the colours close to white). For example, for the orange alert at $t = 1$, only 3 of the 74 weights were lager than 0.1, which are 0.27, 0.13, and 0.12, while the other weights were all less then 10^{-5}. This result corresponded to the second case when $\mathrm{E}(\mathbf{Z}_{it}|D_i = 0) \neq \mathrm{E}(\mathbf{Z}_{it}|D_i = 1)$ as shown in Lemma 2.

To assess Assumption 2 and the linear model specification, we tested the null hypothesis H_0^t: $Y_{it}(0) = \delta_t + \boldsymbol{\beta}_t'\mathbf{X}_{it} + \rho_t Y_{it-1}(0) + \varepsilon_{it}$ holds for $i \in \mathcal{C} \cup \mathcal{T}$ for $t = 1, \ldots, T_0$. For both alerts, we could not reject the null hypotheses $\{\mathrm{E}(\widehat{\tau}_t) = 0\}_{t=1}^{T_0}$ when FDR was controlled at level 0.05, which supported the Assumption 2 and Model (2.3). To evaluate the effectiveness of the proposed DSC method, we adopted the normalised placebo test in Section 3.2. The placebo experiment was replicated for 500 times. Figure 5 shows the normalised $\widehat{\tau}_t^{(k)}$ in the placebo runs vs. the estimated effect $\widehat{\tau}_t$ for the air pollution episodes of the two alerts. The graphical displays indicate the scales of the estimated treatment effects $\widehat{\tau}_t$ was well in in the middle of the distribution of normalised placebo $\widehat{\tau}_t^{(k)}$ in the pre-treatment period, and then started to deviate from the body of the normalised

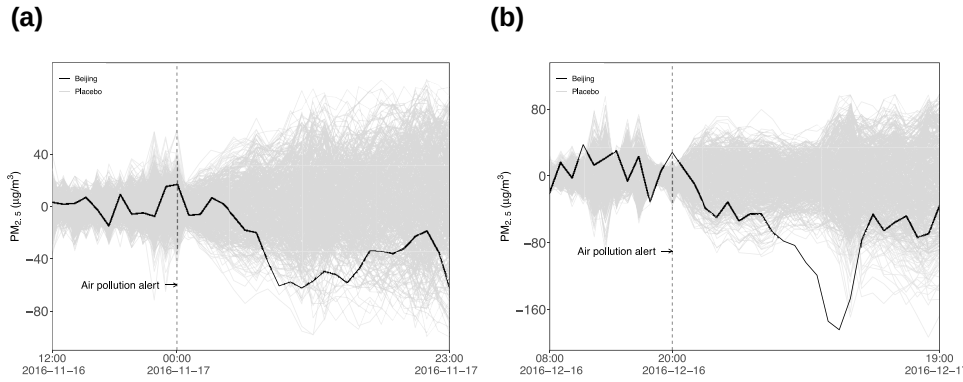

Figure 5. The estimated effects $\widehat{\tau}_t$ for Beijing (black line) and the normalised differences $\{z_t^{(k)}\}_{k=1}^{500}$ for the placebo treatment group (grey lines): (a) the orange alert and (b) the red alert.

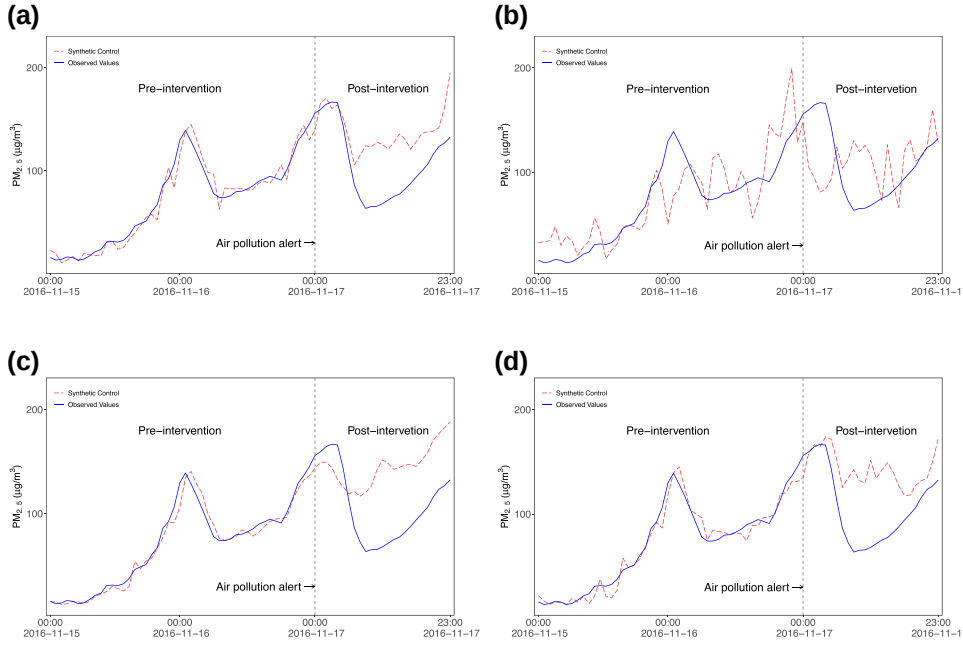

Figure 6. Estimated average potential outcomes $\widehat{\mu}_t(0)$ (dashed line) and the observed average $\sum_{i\in\mathcal{T}} \frac{Y_{it}}{N_{\text{tr}}}$ (solid line) for the orange alert under four matching strategies. Panel (a) used the original version that includes meteorological covariates and the first lagged $PM_{2.5}$, Panels (b) and (c) only matched the meteorological covariates and the first lagged $PM_{2.5}$, respectively, and Panel (d) matched the meteorological covariates and the first and second lagged $PM_{2.5}$.

placebo $\widehat{\tau}_t^{(k)}$ after the starts of the alert. The p-value was calculated by $500^{-1}\sum_{k=1}^{500} 1(z_t^{(k)} \leq \widehat{\tau}_t)$ (see online supplementary material for details), which showed that the most significant effects under the orange alert were attained in the period of 8–15 h after the start of the alert, and those under the red alert were 6–17 h after the start of the alert.

To gain information on the sensitivity of the proposed DSC method to changes in the model specification, we experimented the following three alternative model settings:

(i) $Y_{it}(0) = \delta_t + \beta' \mathbf{X}_{it} + \varepsilon_{it}$ and obtain the weights by matching only the meteorological covariates.
(ii) $Y_{it}(0) = \delta_t + \rho Y_{i,t-1}(0) + \varepsilon_{it}$ and obtain the weights by matching only the lag value of $PM_{2.5}$.
(iii) $Y_{it}(0) = \delta_t + \rho_1 Y_{i,t-1}(0) + \rho_2 Y_{i,t-2}(0) + \beta' \mathbf{X}_{it} + \varepsilon_{it}$ and match the second lagged $PM_{2.5}$ additionally.

Figure 6a shows the DSC results for the orange alert where the weights were calculated according to equations (2.7)–(2.9), while Panels (b)–(d) of Figure 6 provide analogous DSC results under the alternative model Settings (a)–(c), respectively. In particular, Figure 6b shows that only matching the four meteorological covariates was not sufficient to track the trend of $PM_{2.5}$, as reflected in the divergence between the blue and red lines in the pre-treatment periods. Figures 6c and d are quite similar to the results in Panel (a), which suggested that by either excluding the covariates $\boldsymbol{X}_{it}$ or including the second-order lagged value $Y_{i,t-2}$ would not affect the results much as long as the first lagged $PM_{2.5}$ was included. In summary, the analyses suggested that having the first-order lagged $PM_{2.5}$ in the dynamic matching equations was crucial to establish the unconfoundedness assumption, while the meteorological covariates were less important and having the second-order lagged $PM_{2.5}$ seemed unnecessary.

6 Discussion

We propose the DSC method for dynamic average treatment effect estimation with multiple treated units and time-varying covariates. We employ the EL to construct the DSC weights and provide theoretical analyses on the asymptotic orders of the weights and consistency results on the effect estimation. Statistical inferences for the DSC are developed which allows us to assess model specification and to attain the significance of the estimated treatment effects.

Although the DSC method is proposed to evaluate the effects of air pollution alerts, it may be applied to evaluate dynamic treatment effects in other studies. For example, assessing the effects of health interventions based on data recorded by wearable devices (Boruvka et al., 2018; Liao et al., 2020), estimating the effect of epileptic seizures on the dynamic patterns of the electroencephalogram, and evaluating the impacts of promotion strategies on consumer-generated product reviews. Future work may consider extending the DSC method to allow more flexible models. For now, the DSC method relies on a linear model over both the treatment and control groups, which is the same as the SC method. We may add non-parametric transformation on the covariate before matching such that the outcome is linear in the transformed covariates.

A question is that if the DSC can be extended for the policy evaluation. The inverse probability weighting method is a commonly used estimation of ATT for policy evaluation as demonstrated in Zhou et al. (2018). Specifically, let $\boldsymbol{Z}$ be the pre-treatment covariates and $\pi: \mathcal{Z} \mapsto [0, 1]$ denote a policy propensity score that maps $\boldsymbol{Z}$ to the probability of receiving the policy. The ATT of policy π at time t is $\tau_t = \mathrm{E}[Y_t(1) - Y_t(0) | D = 1] = \mathrm{E}[Y_t | D = 1] - \mathrm{E}[Y_t \cdot (1 - D)/(1 - \pi(\boldsymbol{Z}))]$. Hence, τ_t can be estimated by $\widehat{\tau}_t = \sum_{D_i=1} Y_{it}/N_{\mathrm{tr}} - \sum_{D_i=0} Y_{it}/(1 - \pi(\boldsymbol{Z}_i))$ via the propensity score. Comparing with the DSC estimate $\sum_{D_i=1} Y_{it}/N_{\mathrm{tr}} - \sum_{D_i=0} w_{it}^* Y_{it}$, it is seen that $1/(1 - \pi(\boldsymbol{Z}_i))$ plays a similar role as the weight w_{it}^* in the DSC. However, to extend the DSC approach for the policy evaluation, one has to extend the DSC method to non-linear models as the the propensity score of receiving the policy is a non-linear function of the covariates Z. One possible way for doing the non-linear extension is via the generalised additive model as discussed toward the end of Section 3.1, which requires further analysis.

It is noted that Brodersen et al. (2015) employed a structural time series model with time-varying covariates to infer the causal effects of a policy intervention. The study shared the same goal with the DSC. It assumed a more detailed model for the outcome under control, which was a linear function of a multi-dimensional state vector including covariates, unobserved local linear trend and a seasonality term, where the latter were modelled explicitly. In contrast, the DSC method does not specify an explicit form of the latent factor δ_t. However, Brodersen et al. (2015) can make more formal inference than the DSC as it can simulate the posterior predictive distribution of the counterfactual outcomes based on the explicit model of $Y_t(0)$ and the specification of the prior distributions, while the

DSC method can only produce a point estimation of the counterfactual average and has to rely on the placebo treatment simulation to gauge on the significance of the estimated effects.

The DSC weighting scheme outlined in equations (2.7)–(2.9) can be directly applied to estimate the average instantaneous effect in the staggered rollout design (Xiong et al., 2019), despite that the starting time of treatment may vary by unit in the staggered rollout design. However, the proposed DSC method cannot be directly applied to estimate the average lagged effect as it would requires more assumptions. The detailed formulations and the related discussion are available in Section 1.5 (extended discussions) in the online supplementary material.

Acknowledgments

The authors are grateful to five anonymous reviewers for helpful comments and suggestions which have improved the content and presentation of the paper.

Conflict of interest: None declared.

Funding

This work was supported by the National Natural Science Foundation of China, grant numbers: 92046021, 12071013 and 12026607; and the Key Laboratory of Mathematical Economics and Quantitative Finance at Peking University.

Data availability

The data on the air pollution alerts are available in the online supplementary material.

Supplementary material

Supplementary material is available online at *Journal of the Royal Statistical Society: Series B*.

References

Abadie A., Diamond A., & Hainmueller J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *Journal of the American Statistical Association*, *105*(490), 493–505. https://doi.org/10.1198/jasa.2009.ap08746

Abadie A., & Gardeazabal J. (2003). The economic costs of conflict: A case study of the basque country. *American Economic Review*, *93*(1), 113–132. https://doi.org/10.1257/000282803321455188

Anderson T. (1971). *The statistical analysis of time series*. John Wiley & Sons.

Arkhangelsky D., Athey S., Hirshberg D. A., Imbens G. W., & Wager S. (2019). *Synthetic difference in differences* (Technical Report). National Bureau of Economic Research.

Athey S., & Imbens G. W. (2017). The state of applied econometrics: Causality and policy evaluation. *Journal of Economic Perspectives*, *31*(2), 3–32. https://doi.org/10.1257/jep.31.2.3

Benjamini Y., & Yekutieli D. (2001). The control of the false discovery rate in multiple testing under dependency. *Annals of Statistics*, *29*(4), 1165–1188. https://doi.org/10.1214/aos/1013699998

Boruvka A., Almirall D., Witkiewitz K., & Murphy S. A. (2018). Assessing time-varying causal effect moderation in mobile health. *Journal of the American Statistical Association*, *113*(523), 1112–1121. https://doi.org/10.1080/01621459.2017.1305274

Brodersen K. H., Gallusser F., Koehler J., Remy N., & Scott S. L. (2015). Inferring causal impact using Bayesian structural time-series models. *The Annals of Applied Statistics*, *9*(1), 247–274. https://doi.org/10.1214/14-AOAS788

Canay I. A., Romano J. P., & Shaikh A. M. (2017). Randomization tests under an approximate symmetry assumption. *Econometrica*, *85*(3), 1013–1030. https://doi.org/10.3982/ECTA13081

Chen S. X., & Van Keilegom I. (2009). A review on empirical likelihood methods for regression. *Test*, *18*(3), 415–447. https://doi.org/10.1007/s11749-009-0159-5

Chen Z., Cai J., Gao B., Xu B., Dai S., He B., & Xie X. (2017). Detecting the causality influence of individual meteorological factors on local $PM_{2.5}$ concentration in the Jing-Jin-Ji region. *Scientific Reports*, *7*(1), 40735. https://doi.org/10.1038/srep40735

Cheng N., Zhang D., Li Y., Xie X., Chen Z., Meng F., Gao B., & He B. (2017). Spatio-temporal variations of $PM_{2.5}$ concentrations and the evaluation of emission reduction measures during two red air pollution alerts in Beijing. *Scientific Reports*, *7*(1), 1–12. https://doi.org/10.1038/s41598-017-08895-x

Currie J., Kleven H., & Zwiers E. (2020). Technology and big data are changing economics: Mining text to track methods. *AEA Papers and Proceedings*, *110*, 42–48.
Deshpande Y., Mackey L., Syrgkanis V., & Taddy M. (2018). Accurate inference for adaptive linear models. In *International Conference on Machine Learning* (pp. 1194–1203). PMLR.
Doukhan P., & Gómez J. G. (2017). On extreme values in stationary weakly dependent random fields. In *Workshop on Cyclostationary Systems and Their Applications* (pp. 92–110). Springer.
Ferman B., & Pinto C. (2016). *Revisiting the synthetic control estimator* (Technical Report). University Library of Munich, Germany.
Ghosh S., & Chaudhuri S. (2019). 'Empirical likelihood under mis-specification: Degeneracies and random critical points', arXiv, arXiv:1910.01396, preprint: not peer reviewed.
Hadad V., Hirshberg D. A., Zhan R., Wager S., & Athey S. (2021). Confidence intervals for policy evaluation in adaptive experiments. *Proceedings of the National Academy of Sciences*, *118*(15), e2014602118. https://doi.org/10.1073/pnas.2014602118
Hahn J., & Shi R. (2017). Synthetic control and inference. *Econometrics*, *5*(4), 52. https://doi.org/10.3390/econometrics5040052
Imbens G. W., & Wooldridge J. M. (2009). Recent developments in the econometrics of program evaluation. *Journal of Economic Literature*, *47*(1), 5–86. https://doi.org/10.1257/jel.47.1.5
Leadbetter M., Lindgren G., & Rootzén H. (1983). Extreme value theory and strength of materials. In *Extremes and related properties of random sequences and processes* (pp. 267–277). Springer. https://doi.org/10.1007/978-1-4612-5449-2_14
Leadbetter M., & Rootzén H. (1998). On extreme values in stationary random fields. In *Stochastic processes and related topics* (pp. 275–285). Springer. https://doi.org/10.1007/978-1-4612-2030-5_15
Liang X., Zou T., Guo B., Li S., Zhang H., Zhang S., Huang H., & Chen S. X. (2015). Assessing Beijing's $PM_{2.5}$ pollution: Severity, weather impact, APEC and winter heating. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *471*(2182), 20150257. https://doi.org/10.1098/rspa.2015.0257
Liao P., Klasnja P., & Murphy S. (2020). Off-policy estimation of long-term average outcomes with applications to mobile health. *Journal of the American Statistical Association*, *116*(533), 382–391. https://doi.org/10.1080/01621459.2020.1807993
Owen A. B. (1988). Empirical likelihood ratio confidence intervals for a single functional. *Biometrika*, *75*(2), 237–249. https://doi.org/10.1093/biomet/75.2.237
Qin J., & Lawless J. (1994). Empirical likelihood and general estimating equations. *The Annals of Statistics*, *22*(1), 300–325. https://doi.org/10.1214/aos/1176325370
Robins J., Hernán M., & Brumback B. (2000). Marginal structural models and causal inference in epidemiology. *Epidemiology (Cambridge, Mass.)*, *11*(5), 550–560. https://doi.org/10.1097/00001648-200009000-00011
Rubin D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, *66*(5), 688–701. https://doi.org/10.1037/h0037350
Rubin D. B. (2006). *Matched sampling for causal effects*. Cambridge University Press.
Stuart E. A. (2010). Matching methods for causal inference: A review and a look forward. *Statistical Science*, *25*(1), 1. https://doi.org/10.1214/09-STS313
Wang X., Wei W., Cheng S., Li J., Zhang H., & Lv Z. (2018). Characteristics and classification of $PM_{2.5}$ pollution episodes in Beijing from 2013 to 2015. *Science of the Total Environment*, *612*(1), 170–179. https://doi.org/10.1016/j.scitotenv.2017.08.206
Xiong R., Athey S., Bayati M., & Imbens G. (2019). Optimal experimental design for staggered rollouts. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3483934
Zhan R., Hadad V., Hirshberg D. A., & Athey S. (2021). Off-policy evaluation via adaptive weighting with data from contextual bandits. *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. https://doi.org/10.1145/3447548.3467456
Zhang K., Janson L., & Murphy S. (2021). Statistical inference with m-estimators on adaptively collected data. *Advances in Neural Information Processing Systems*, *34*, 7460–7471. https://doi.org/10.48550/arXiv.2104.14074
Zhao Q., & Percival D. (2016). Entropy balancing is doubly robust. *Journal of Causal Inference*, *5*(1), 20160010. https://doi.org/10.1515/jci-2016-0010
Zhou Z., Athey S., & Wager S. (2018). 'Offline multi-action policy learning: Generalization and optimization', arXiv, arXiv:1810.04778, preprint: not peer reviewed.