

Bipartite causal inference with interference, time series data, and a random network

Zhaoyan Song Georgia Papadogeorgou

Department of Statistics, University of Florida

Abstract

In bipartite causal inference with interference, two distinct sets of units exist: interventional units, which receive treatment, and outcome units, where outcomes are measured. Which interventional units' treatment can drive which outcome units' outcomes is often depicted in a bipartite network. We study bipartite causal inference with interference from observational data across time and with a changing bipartite network. Under an exposure mapping framework, we define causal effects specific to each outcome unit, representing average contrasts of potential outcomes across time. We establish unconfoundedness of the exposure received by outcome units based on unconfoundedness assumptions on the interventional units' treatment assignment and the random graph, hence respecting the bipartite structure of the problem. Harvesting the time component of our setting, causal effects are estimable while controlling only for temporal trends and time-varying confounders. Our results hold for binary, continuous, and multivariate exposure mappings. For binary exposure, we propose three matching algorithms to estimate the causal effect by matching exposed to unexposed time periods for the same outcome unit. We show that the bias of resulting estimators is bounded. We illustrate our approach through simulation studies and a study on the effect of wildfire smoke on transportation by bicycle.

1 Introduction

Causal inference methodology most often focuses on the scenario where units are assigned to treatment or control, and an outcome is measured on the same set of units. However, in some cases, the units that receive the treatment are distinct from the units that experience the outcome. We refer to the former as interventional units, and the latter as outcome units. The outcome units, which do not get treatment themselves, are exposed to the treatment only through their connections to potentially treated

interventional units. The causal dependencies across units can be described in a bipartite network, and, as a result, this setting has been termed bipartite interference [[Zigler and Papadogeorgou, 2021](#)].

In this manuscript, we focus on bipartite causal inference with interference from observational data measured over time and a time-varying bipartite network. Each interventional unit receives a treatment level which can change over time according to an unknown assignment mechanism that depends on covariates of the interventional units, the outcome units, and the network. At each time, the bipartite network of causal dependencies across units is specified by a random process which itself can depend on all the covariates. By considering an evolving network and temporal variations in the treatment assignment, we provide a comprehensive framework for analyzing bipartite interference with temporal data.

Existing work in causal inference with bipartite interference is cross-sectional and mostly considers a fixed and known bipartite network. [Zigler and Papadogeorgou \[2021\]](#) introduced causal estimands for bipartite causal inference, and developed weighting-based estimators under clustered interference. The notion of an exposure mapping introduced in unipartite causal inference [e.g., [Aronow and Samii, 2017](#), [Forastiere et al., 2021](#)] has been extended to the bipartite setting, stating that potential outcomes depend on the interventional units' treatment through known functions of the treatment and the bipartite network. In [Zigler et al. \[2020\]](#), the bipartite network describes complex atmospheric and geographic dependencies among units. In the experimental setting, under a linear exposure mapping, [Harshaw et al. \[2023\]](#) designed estimators and inferential techniques for the effect of assigning all or none of the interventional units to treatment. [Pouget-Abadie et al. \[2019\]](#) developed experimentation techniques that improve the efficiency of estimators, and [Brennan et al. \[2022\]](#) focused on avoiding inferential bias due to interference. [Doudchenko et al. \[2020\]](#) proposed using propensity scores to account for confounding due to the network structure. All these studies have fixed causal networks, except [Wikle and Zigler \[2023\]](#), which considers a probabilistic bipartite network in a cross-sectional design.

For time series data, most of the causal inference literature focuses on the case without interference. The literature on this topic is large, and it is out of scope to review here [see [Abadie and Cattaneo, 2018](#), [Imbens, 2024](#), for surveys on the topic]. Relevant to our work are extensions of panel data methodology to the case with unit-to-unit interference. In [Cao and Dowd \[2019\]](#), [Grossi et al. \[2020\]](#), [Di Stefano and Mellace \[2020\]](#) and [Menchetti and Bojinov \[2020\]](#), a subset of units receives the treatment at some point in time and remains treated thereafter, and in [Clark and Handcock \[2021\]](#) and [Agarwal et al. \[2023\]](#) the units' treatment assignments changes over time by endogenous nature or by design. None of these methods has been considered in the bipartite setting.

This work develops a causal inference framework for bipartite interference with time series observational data and a random bipartite network, which allows researchers to capture the dynamic nature of causal relationships in real-world settings where networks and treatment assignments change over time. Our contributions are the following: (a) Under an exposure mapping framework, we define causal estimands of interest for each outcome unit in the bipartite temporal setting as contrasts of the unit’s potential outcomes under different exposure levels averaged over time (Section 2). (b) We introduce the causal assumptions of unconfoundedness for the treatment and network processes conditional on variables of the interventional units, outcome units and the network. We establish the unconfoundedness for the outcome unit’s exposure, which implies that we can estimate the outcome-unit-specific effects while conditioning *only* on temporally-varying information (Section 2). In our work, the unconfoundedness assumptions are at the level on which the randomness occurs (treatment, network) in contrast with existing work that places assumptions on the exposure directly [Zigler et al., 2020, Doudchenko et al., 2020]. These results hold for binary, continuous, or multivariate exposures. (c) Focusing on binary exposures, we develop matching procedures to estimate causal effects for each outcome unit (Section 3). The proposed algorithms match an outcome unit’s exposed time periods to unexposed time periods, one-to-one, one-to-two, or one-to-one-or-two, while specifying that matched time periods occur close in time, and satisfy balance constraints for the time-varying covariates. We show that the bias of the matching estimators is bounded, and can be made arbitrarily small by imposing stricter tolerance parameters (Section 3). (d) Our approach infers the causal effect of the interventional units’ treatment on each outcome unit separately. We establish how results on multiple outcome units can be combined to test a global null hypothesis of no treatment effect (Section 3). (e) In an extensive simulation study, we showcase that our approach performs well for estimating the outcome-unit effects and results in appropriate inference (Section 4). (f) We use our methodology to study whether smoke from wildfires affects outdoor physical activity in the San Francisco Bay area (Section 5). We find that exposure to smoke from wildfires leads to a decrease in the number of bike rental hours in the city of San Francisco, but does not significantly alter bike use in nearby regions. We conclude with a discussion (Section 6).

2 Bipartite interference with time series observational data and a random network

2.1 The setup

Let $\mathcal{N} = \{n_1, n_2, \dots, n_N\}$ denote interventional units followed over time $t = 1, 2, \dots, T$. Each of them has a collection of characteristics. We use the superscript $*$ for time-invariant variables, and the subscript t for time-specific variables. Let $\mathbf{X}_i^* = (X_{i1}, X_{i2}, \dots, X_{ip_X^*})^\top$ and $\mathbf{X}_{ti} = (X_{ti1}, X_{ti2}, \dots, X_{tip_X})^\top$ denote p_X^* time-invariant and p_X time-varying covariates, respectively, for the interventional unit n_i , and $A_{it} \in \mathcal{A}$ denote its treatment level at time t . Let $\mathbf{X}^* = (\mathbf{X}_1^* \mathbf{X}_2^* \dots \mathbf{X}_N^*)^\top$ be the $N \times p_X^*$ matrix of time-invariant covariates, $\mathbf{X}_t = (\mathbf{X}_{t1} \mathbf{X}_{t2} \dots \mathbf{X}_{tN})^\top$ the $N \times p_X$ matrix of time-varying covariates at time t , and $\mathbf{A}_t = (A_{t1}, A_{t2}, \dots, A_{tN})^\top$ the treatment vector at time t , across all interventional units, with $\mathbf{A}_t \in \mathcal{A}^N$. The interventional units do not experience the outcome.

The set of outcome units is denoted by $\mathcal{M} = \{m_1, m_2, \dots, m_M\}$. Each unit m_j has p_W^* time-invariant and p_W time-varying covariates, $\mathbf{W}_j^* = (W_{j1}, W_{j2}, \dots, W_{jp_W^*})^\top$ and $\mathbf{W}_{tj} = (W_{tj1}, W_{tj2}, \dots, W_{tjp_W})^\top$, respectively. We use $\mathbf{W}^* = (\mathbf{W}_1^* \mathbf{W}_2^* \dots \mathbf{W}_M^*)^\top$ and $\mathbf{W}_t = (\mathbf{W}_{t1} \mathbf{W}_{t2} \dots \mathbf{W}_{tN})^\top$ for the covariate matrices of dimension $M \times p_W^*$ and $M \times p_W$ across the outcome units. For each unit m_j , we measure an outcome over time, denoted by Y_{tj} .

We also consider network covariates that describe the relationship between interventional and outcome units. We use $\mathbf{P}^* = \{P_{ijs}^*\}$ for the $N \times M \times p_P^*$ array of time-invariant network covariates, and $\mathbf{P}_t = \{P_{tjjs}\}$ for the $N \times M \times p_P$ array of time-varying network covariates at time t . Notation like $\mathbf{P}_{t,j}$ is used for the $N \times p_P$ sub-matrix of \mathbf{P}_t corresponding to outcome unit m_j .

The random bipartite network can vary over time. Let \mathbf{G}_t denote the $N \times M$ matrix of bipartite connections. The vector $\mathbf{G}_{t,j} = (G_{t1j}, \dots, G_{tNj})^\top$ denotes the connectivity status of outcome unit m_j with all interventional units, taking values in \mathcal{G}_j . We consider a binary bipartite network where $G_{tij} = 1$, if n_i and m_j are connected at time t , and $G_{tij} = 0$ otherwise, but alternative, non-binary specifications of \mathbf{G} can be easily incorporated.

The outcome units do not receive an intervention themselves, rather than experience the treatment of the interventional units through the bipartite network. We formalize this using exposure mappings. For the outcome unit m_j , the function $h_{tj} : \mathcal{A}^N \times \mathcal{G}_j \rightarrow \mathcal{E}_{tj}$ maps the interventional units' treatment assignment and the outcome unit's bipartite connection vector to the outcome unit's exposure value at time t , where \mathcal{E}_{tj} denotes the set of all possible exposure values. Then, $E_{tj} = h_{tj}(\mathbf{A}_t, \mathbf{G}_{t,j})$ is the realized exposure of outcome unit m_j at time t . The function $h_{tj}(\cdot)$ might return a scalar such

as the proportion of interventional units with which m_j is connected that are treated. It can also be completely general, specified to return the vector of treatment levels for all connected interventional units, or extended to depend on covariates.

In our study, the interventional units are all forest locations across North America. For these units, potential time-invariant covariates \mathbf{X}^* include local species of vegetation. The $M = 3$ outcome units are areas in the San Francisco Bay Area, with potential time-invariant covariates \mathbf{W}^* such as demographic information. Time-varying covariates for both sets of units, $\mathbf{X}_t, \mathbf{W}_t$, are weather-related factors, such as temperature and precipitation. Network covariates \mathbf{P}^* include the geographic distance of forest-area pairs. The transformation from wildfire to smoke follows complex chemical and atmospheric reactions, and the resulting smoke can travel long distances, as described in the bipartite network \mathbf{G}_t . The Hazard Mapping System for smoke (HMS) monitors smoke plumes resulting from fires, and it combines information on the presence of wildfires, \mathbf{A}_t , with information on smoke transport, \mathbf{G}_t , to deduct the potential smoke exposure in each region in day t , E_{tj} . The outcome Y_{tj} is daily bike riding hours using Lyft’s Bay Wheels program in each region.

2.2 Potential outcomes and causal estimands

Let $Y_{tj}(\mathbf{a}_t, \mathbf{g}_{t,j})$ denote the potential outcome for unit m_j at time t had the treatment of the N interventional units been \mathbf{a}_t , and under bipartite connection $\mathbf{g}_{t,j}$ for unit m_j . We include the bipartite graph in the notation for potential outcomes to establish that the graph plays a role in how the potential outcomes vary by treatment, but we do *not* assume that the network is manipulable. The observed outcome corresponds to the potential outcome under the observed treatment and network, as $Y_{tj} = Y_{tj}(\mathbf{A}_t, \mathbf{G}_{t,j})$. This notation implicitly states that previous treatments of the interventional units do not drive the contemporaneous outcome of unit m_j . We discuss this within the context of our study in Section 5.

Each treatment vector \mathbf{a}_t might lead to a different potential outcome for unit m_j . The following assumption codifies that the interventional units’ treatment drives potential outcomes only through the resulting outcome unit’s exposure based on the bipartite mapping.

Assumption 1. For all $\mathbf{a}_t, \mathbf{a}'_t \in \mathcal{A}^N$, and $\mathbf{g}_{t,j}, \mathbf{g}'_{t,j} \in \mathcal{G}_{t,j}$, if $h_{tj}(\mathbf{a}_t, \mathbf{g}_{t,j}) = h_{tj}(\mathbf{a}'_t, \mathbf{g}'_{t,j}) = e_{tj}$, then $Y_{tj}(\mathbf{a}_t, \mathbf{g}_{t,j}) = Y_{tj}(\mathbf{a}'_t, \mathbf{g}'_{t,j})$, and the potential outcome can be denoted as $Y_{tj}(e_{tj})$.

Under Assumption 1, the collection of all potential outcomes for outcome unit m_j at time t is the set $\mathcal{Y}_{tj}(\cdot) = \{Y_{tj}(e_{tj}), \text{ for } e_{tj} \in \mathcal{E}_{tj}\}$. Similar constructions of exposure mappings and assumptions on the potential outcomes have been discussed in the unipartite [Aronow and Samii, 2017, Forastiere et al., 2021] and bipartite [Zigler et al., 2020, Harshaw et al., 2023, Doudchenko et al., 2020] interference literature. Sävje [2024] discusses the implications of using exposure mappings in the definition of

estimands and as assumptions on potential outcomes, providing interesting distinctions between the two.

We consider estimands that are specific to each outcome unit. The contrast $\tau_{tj}(e, e') = Y_{tj}(e) - Y_{tj}(e')$ represents the fundamental effect of a change in exposure from e' to e for the outcome unit m_j at time t . This unit- and time-specific estimand cannot be estimated without parametric assumptions. Instead, we consider target estimands that represent temporally-averaged causal effects for unit m_j for a change in its exposure value. Specifically,

$$\tau_j(e, e') = \frac{1}{T} \sum_{t=1}^T \tau_{tj}(e, e') = \frac{1}{T} \sum_{t=1}^T [Y_{tj}(e) - Y_{tj}(e')],$$

represents the average effect of a change in exposure for unit m_j across time, and

$$\tilde{\tau}_j(e, e') = \frac{1}{\sum_t I(E_{tj} = e)} \sum_t (Y_{tj}(e) - Y_{tj}(e')) I(E_{tj} = e)$$

over only those time periods with realized exposure equal to e . Therefore, the estimand $\tilde{\tau}_j$ resembles a temporal version of the average treatment effect on the treated. If exposures e and e' are not both possible for all time periods, the estimands $\tau_j(e, e')$ and $\tilde{\tau}_j(e, e')$ should only average over time periods where both exposure values under investigation are possible.

In Section 3.5, we discuss how focusing on temporally-average estimands might lead to weaker confounding adjustment requirements compared to estimands that average across units.

2.3 Ignorable assignments: assumptions and results

We establish unconfoundedness assumptions that allow us to estimate the estimands of Section 2. Our assumptions pertain to the interventional units' treatment assignment and the random bipartite network. Since our estimands are specific to each outcome unit, we focus on unit m_j throughout.

Assumption 2. (Unconfoundedness of the treatment assignment). The interventional units' treatment assignment at time t is independent of the potential outcomes of outcome unit m_j , conditional on a function of time $f(t)$, time-invariant and time-varying covariates of the interventional units, outcome unit m_j , and their connections, i.e., $P(\mathbf{A}_t \mid \mathcal{Y}_{tj}(\cdot), f(t), \mathbf{X}^*, \mathbf{X}_t, \mathbf{W}^*, \mathbf{W}_{tj}, \mathbf{P}^*, \mathbf{P}_{t.j}) = P(\mathbf{A}_t \mid f(t), \mathbf{X}^*, \mathbf{X}_t, \mathbf{W}^*, \mathbf{W}_{tj}, \mathbf{P}^*, \mathbf{P}_{t.j})$.

Under Assumption 2, the treatment level of interventional units can be driven by their individual characteristics, characteristics of the outcome units, and general temporal trends such as those that alter the overall prevalence of treatment. Therefore, the permitted treatment assignment mechanisms allow

for complex bipartite dependencies that relate units from the two ends of the graph. Specifically, it clarifies that confounding can arise not only from the covariates of the interventional units themselves, but also from the features of the network and outcome units that contribute to the assignment of the interventional units' treatment.

Assumption 3. (Unconfoundedness of the bipartite network). The bipartite connection vector for unit m_j is independent of the unit's potential outcomes given the treatment assignment, temporal trends $f(t)$, and characteristics of all interventional units, the outcome unit m_j , and their connections, i.e., $P(\mathbf{G}_{t,j} | \mathcal{Y}_{tj}(\cdot), \mathbf{A}_t, f(t), \mathbf{X}^*, \mathbf{X}_t, \mathbf{W}^*, \mathbf{W}_{tj}, \mathbf{P}^*, \mathbf{P}_{t,j.}) = P(\mathbf{G}_{t,j} | \mathbf{A}_t, f(t), \mathbf{X}^*, \mathbf{X}_t, \mathbf{W}^*, \mathbf{W}_{tj}, \mathbf{P}^*, \mathbf{P}_{t,j.})$.

This assumption states that how outcome unit m_j forms connections with interventional units might depend on temporal trends, and characteristics of the units. It can also depend on the realized treatment level, which is relevant in applications where the overall treatment prevalence might lead to higher or lower outreach of the interventional units. The probabilistic formalization of unconfoundedness in Assumption 3 implicitly assumes a random network generation. That said, confounding can also arise in scenarios with a known and fixed network as illustrated in [Doudchenko et al. \[2020\]](#), therefore our results are also applicable in that case.

Assumptions 2 and 3 allow for complex dependencies of the treatment and network processes on a very large class of covariates. We establish that, under these assumptions, the exposure that outcome unit m_j receives from the interventional units' treatment through the bipartite network is unconfounded. This result holds for exposure mappings that are arbitrarily complex. (The proof is in Supplement A.)

Proposition 1. (Exposure unconfoundedness). If Assumptions 2 and 3 hold, then outcome unit m_j 's exposure assignment is independent of its potential outcomes given the temporal trend $f(t)$, and covariate information on the units and the network, i.e., $P(E_{tj} | \mathcal{Y}_{tj}(\cdot), f(t), \mathbf{X}^*, \mathbf{X}_t, \mathbf{W}^*, \mathbf{W}_{tj}, \mathbf{P}^*, \mathbf{P}_{t,j.}) = P(E_{tj} | f(t), \mathbf{X}^*, \mathbf{X}_t, \mathbf{W}^*, \mathbf{W}_{tj}, \mathbf{P}^*, \mathbf{P}_{t,j.})$.

Statements like the exposure unconfoundedness in Proposition 1 have been evoked as assumptions in previous work on bipartite interference [[Zigler et al., 2020](#), [Doudchenko et al., 2020](#)]. However, here, exposure unconfoundedness is established while acknowledging that, in a bipartite interference context, the exposure experienced by an outcome unit is governed by mechanisms operating at the treatment and network levels. This has crucial implications for practice. Assumptions 2 and 3 yield fruitful and practical insights for identifying confounders that exist in the treatment-outcome

or network-outcome relationships. This is particularly relevant in bipartite interference contexts for which we have a clear grasp of the physical or mechanistic processes driving the network structure. Therefore, these assumptions provide guidance that renders confounding adjustment more tangible, nuanced and actionable within the bipartite setting.

The unconfoundedness result in Proposition 1 means that we can acquire an unbiased estimator of the temporally-averaged causal effect on unit m_j , $\tau_j(e, \tilde{e})$, by comparing outcomes of time periods with similar values of the covariates in the conditioning set and different values of their exposure, and averaging over the covariate distribution across time [see Forastiere et al., 2021, for a related discussion]. We can estimate the causal effect $\tilde{\tau}_j(e, \tilde{e})$ similarly, by altering which distribution we average over, to reflect the distribution of the covariates among time periods with $E_{tj} = e$ [Abadie and Imbens, 2006]. Importantly, the covariates \mathbf{X}^* , \mathbf{W}^* and \mathbf{P}^* are constant across time. Therefore, they are implicitly *always* conditioned on when studying the same outcome unit across time. This implies that the time-invariant covariates that create differences in the assignment mechanism of treatment *across* interventional units, or the assignment mechanism of bipartite connections *across* pairs, *need not be measured* when focusing on estimands that average over time. Instead, we need to control for time-varying information only to estimate causal effects. We discuss this further in Section 3.5.

3 Estimation via matching exposed to unexposed time periods

Temporally-averaged causal effects can be estimated by controlling for time-varying information, for exposure mappings that return binary, continuous, or multivariate exposures. In our study of the effects of smoke from wildfires, a region’s exposure can be specified as binary, indicating the presence or absence of smoke exposure for the population residing in the area. Therefore, from here onwards, we focus on the estimation of causal effects under binary exposures.

Viewing the time periods as the elementary unit of observation, in Section 3.1 we introduce three matching algorithms that match exposed time periods to unexposed time periods under constraints that balance time-varying information. In Section 3.2, we define the corresponding causal effect estimators and we show that, under reasonable assumptions on the outcome model, the estimators’ bias is bounded and can be controlled by the algorithms’ tuning parameters. We discuss an inferential approach for one outcome unit in Section 3.3, and for multiple outcome units in Section 3.4. Lastly, in Section 3.5, we discuss the advantages in estimation of estimands that are specific to each outcome unit and average across time, over estimands that average across units.

3.1 Algorithms for matching exposed to unexposed time periods

We propose three novel matching algorithms for estimating causal effects in bipartite interference settings with time series observational data and a binary exposure. Since we focus on *outcome unit specific* estimands, our matching algorithms match time periods with and without exposure for the *same outcome unit*. For notational simplicity, we drop notation pertaining to the unit, and we refer to a time period as exposed if $E_t = 1$, and unexposed otherwise. We focus on the average change in the unit’s outcome among the exposed time periods, had the unit been exposed versus unexposed, defined in Section 2 as $\tilde{\tau}(1, 0)$. In the presence of multiple outcome units, the algorithms would be applied to each unit separately.

The three matching approaches, ‘Matching 1-1’, ‘Matching 1-2’, and ‘Matching 1-1/2’, match an exposed time period to one, two, or either one or two unexposed time periods, respectively. The algorithms are constructed as integer programming optimization problems with the objective of maximizing the number of matches under a set of constraints. Integer programming optimization algorithms have been previously used in the causal inference literature [Zubizarreta, 2012, Zubizarreta et al., 2013, Keele et al., 2014]. Our formalization is different since the fundamental unit of observation is time (rather than physical units) and the confounders correspond to time-varying information.

3.1.1 Matching 1-1.

We use $t_e \in \mathcal{T}_e = \{t : E_t = 1\}$ and $t_u \in \mathcal{T}_u = \{t : E_t = 0\}$ to denote a time period during which the outcome unit is exposed and unexposed, respectively. We introduce binary indicators $a_{t_e t_u}$ for each pair of exposed and unexposed time periods that describe whether the exposed time period t_e is matched to the unexposed time period t_u ($a_{t_e t_u} = 1$), or not ($a_{t_e t_u} = 0$). The objective of the optimization problem is to maximize the number of matches over all possible matching indicators $\mathbf{a} \in \{0, 1\}^{|\mathcal{T}_e| \times |\mathcal{T}_u|}$, as

$$\max_{\mathbf{a}} \sum_{t_e, t_u} a_{t_e t_u}, \tag{A}$$

where we use \sum_{t_e, t_u} to denote the summation over both sets of indices, $\sum_{t_e \in \mathcal{T}_e} \sum_{t_u \in \mathcal{T}_u}$. The optimization problem is performed under a number of constraints. Firstly, we impose that each time point, exposed or unexposed, can be matched at most once,

$$\sum_{t_u} a_{t_e t_u} \leq 1, \quad \forall t_e \in \mathcal{T}_e, \quad \text{and} \quad \sum_{t_e} a_{t_e t_u} \leq 1, \quad \forall t_u \in \mathcal{T}_u. \tag{A.1}$$

If, on the contrary, a certain time period was used in multiple matches, we have found in practice that it could have a disproportionately large influence in the causal estimator, resulting to inaccurate variance estimation and inference.

We impose additional constraints that target the balance of time-varying information including temporal trends and time-varying covariates. In reality, little (if any) information is given about the temporal trends $f(t)$. We balance temporal trends indirectly through balancing the average time of matches. Specifically, the average time difference of matched exposed and unexposed time points is at most $\delta \geq 0$, as

$$\left| \sum_{t_e, t_u} a_{t_e t_u} (t_e - t_u) \right| \leq \delta \sum_{t_e, t_u} a_{t_e t_u}. \quad (\text{A.2})$$

This constraint does not necessarily imply that *each* of the matched pairs is close in time, rather than they are close *on average*. The constant δ can be set arbitrarily small, even to $\delta = 0$. In order to improve the balance of local temporal trends and to reduce the computationally intensive search for possible match combinations, we also impose that each matched pair of time periods is at most ϵ apart in time,

$$|a_{t_e t_u} (t_e - t_u)| \leq \epsilon, \quad \forall t_e \in \mathcal{T}_e, \forall t_u \in \mathcal{T}_u. \quad (\text{A.3})$$

Finally, we balance the time-varying covariates between exposed and unexposed matched time periods. Since we exclude the outcome unit index, \mathbf{W}_t is the vector of outcome unit m_j 's covariates, and \mathbf{P}_t is the $N \times p_P$ matrix including the network covariates for m_j only. We impose that

$$\left| \sum_{t_e, t_u} a_{t_e t_u} (\mathbf{W}_{t_e} - \mathbf{W}_{t_u}) \right| \leq \mathbf{1}_{p_W} \cdot \delta' \sum_{t_e, t_u} a_{t_e t_u}, \quad (\text{A.4})$$

which states that the temporal covariates are on average balanced in matched exposed and unexposed time periods. The constant δ' is usually chosen to be 0.05 or 0.1 standard deviations of the corresponding covariate, though its choice should be part of the design phase of the study and should be decided upon without looking at outcomes [Zubizarreta, 2015].

The time-varying covariates, \mathbf{X}_t and \mathbf{P}_t are of dimension $N \times p_X$ and $N \times p_P$, respectively. In theory, balance constraints could be imposed so that matched time periods are similar with respect to all $N(p_X + p_P)$ variables. However, these balance constraints would likely be high-dimensional, and incorporating them could drastically reduce the number of matches, especially when the number of interventional units is large. Instead, we propose matching summaries of these covariates across interventional units. For a vector $\mathbf{q} = (q_1, q_2, \dots, q_N)^\top$, let $\widetilde{\mathbf{X}}_t = \mathbf{q}^\top \mathbf{X}_t = (\widetilde{X}_{t1}, \widetilde{X}_{t2}, \dots, \widetilde{X}_{tp_X})^\top$

denote the \mathbf{q} -summaries of the interventional units' p_X covariates, and similarly for $\tilde{\mathbf{P}}_t = \mathbf{q}^T \mathbf{P}_t$. We impose that

$$\begin{aligned} \left| \sum_{t_e, t_u} a_{t_e t_u} (\tilde{\mathbf{X}}_{t_e} - \tilde{\mathbf{X}}_{t_u}) \right| &\leq \mathbf{1}_{p_X} \cdot \delta' \sum_{t_e, t_u} a_{t_e t_u}, \quad \text{and} \\ \left| \sum_{t_e, t_u} a_{t_e t_u} (\tilde{\mathbf{P}}_{t_e} - \tilde{\mathbf{P}}_{t_u}) \right| &\leq \mathbf{1}_{p_P} \cdot \delta' \sum_{t_e, t_u} a_{t_e t_u}. \end{aligned} \tag{A.5}$$

The vector \mathbf{q} controls which covariate summary should be balanced, and its choice will be driven by the problem at hand. For example, by setting $q_i = \frac{1}{n}$ for all i , the algorithm balances the average covariate value across interventional units for matched time periods. Alternatively, q_i could give different weights to the covariate value of interventional units based on their geographic proximity to outcome unit m_j , or the frequency with which they are connected. For ease of exposition, we used the same vector \mathbf{q} in the balance constraints for all covariates in (A.5). Different vectors \mathbf{q} can be used for different covariates, and multiple summaries of the same covariate under different vectors \mathbf{q} could be balanced.

By investigating the problem through its true bipartite lens, rather than its unipartite counterpart, the discussion on the needs for confounding adjustment due to interventional unit, outcome unit, or network covariates becomes transparent. For example, in certain scenarios such as the study of Section 5, it is reasonable to argue that time-varying confounding due to interventional unit and network covariates does not exist, in which case constraints (A.5) need not be used. Such understanding of the needs of confounding adjustment would not be obvious if addressing the same study question by projecting it on a unipartite framework. We expand on this in Section 5.

3.1.2 Matching 1-2.

We propose an alternate matching algorithm to match an exposed time point to *two* unexposed time points, one occurring temporally *before* and one *after*. Consider matching indicators $a_{t_e t_{u_1} t_{u_2}}$ for whether the exposed time $t_e \in \mathcal{T}_e$ is matched to the unexposed timestamps $t_{u_1}, t_{u_2} \in \mathcal{T}_u$. The objective function of the optimization algorithm is to maximize the number of matches, which now are of the form (t_e, t_{u_1}, t_{u_2}) ,

$$\max_a \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}}. \tag{B}$$

The constraints are similar in spirit to the ones for algorithm (A), but they are adapted to accommodate matching of one exposed time period to two unexposed ones. Exposed and unexposed time periods can be used in a match at most once, though if an exposed time period is matched, it is matched to two

unexposed ones:

$$\begin{aligned} \sum_{t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}} &\leq 1, \quad \forall t_e \in \mathcal{T}_e \\ \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}} I(t_{u_1} = t_u) + \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}} I(t_{u_2} = t_u) &\leq 1, \quad \forall t_u \in \mathcal{T}_u \end{aligned} \quad (\text{B.1})$$

We impose constraints that balance time and time-varying covariates. Specifically, Constraint (B.2) balances, *on average*, the time of the exposed time period compared to the average time of its unexposed matches,

$$\left| \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}} \left(t_e - \frac{t_{u_1} + t_{u_2}}{2} \right) \right| \leq \delta \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}}, \quad (\text{B.2})$$

but it does not restrict the time difference for each match. Constraint (B.3) restricts the temporal difference and order of time periods for each individual match, as

$$\begin{aligned} |a_{t_e t_{u_1} t_{u_2}} (t_e - t_{u_i})| &\leq \epsilon, \quad \text{for } i = 1, 2 \\ a_{t_e t_{u_1} t_{u_2}} (t_e - t_{u_1}) &\geq 0, \quad \text{and} \quad a_{t_e t_{u_1} t_{u_2}} (t_{u_2} - t_e) \geq 0, \end{aligned} \quad (\text{B.3})$$

for all $t_e \in \mathcal{T}_e$, and $t_{u_1}, t_{u_2} \in \mathcal{T}_u$. The first line imposes that the time difference between an exposed time period and each of its matched unexposed ones cannot exceed ϵ . The second line forces a temporal sequence within each match, requiring the exposed period to fall within the unexposed ones. In addition, the average value of time-varying covariates is balanced when comparing the exposed time periods with the average of their matches. For the outcome-unit time-varying covariates this constraint is

$$\left| \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}} \left(\mathbf{W}_{t_e} - \frac{\mathbf{W}_{t_{u_1}} + \mathbf{W}_{t_{u_2}}}{2} \right) \right| \leq \mathbf{1}_{p_W} \cdot \delta' \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}}, \quad (\text{B.4})$$

and for the interventional unit and network covariates, it takes the form

$$\begin{aligned} \left| \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}} \left(\widetilde{\mathbf{X}}_{t_e} - \frac{\widetilde{\mathbf{X}}_{t_{u_1}} + \widetilde{\mathbf{X}}_{t_{u_2}}}{2} \right) \right| &\leq \mathbf{1}_{p_X} \cdot \delta' \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}}, \quad \text{and} \\ \left| \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}} \left(\widetilde{\mathbf{P}}_{t_e} - \frac{\widetilde{\mathbf{P}}_{t_{u_1}} + \widetilde{\mathbf{P}}_{t_{u_2}}}{2} \right) \right| &\leq \mathbf{1}_{p_P} \cdot \delta' \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}}. \end{aligned} \quad (\text{B.5})$$

Matching seeks unexposed periods that resemble the exposed ones in order to predict what would have happened during the exposed time periods, had they been unexposed. Therefore, matching one exposed unit to two unexposed ones can improve accuracy in predicting the missing potential outcome, over matching one to one. Ensuring that the exposed period falls between matched unexposed ones (B.3) enhances balance in temporal trends. For instance, monotonic temporal trends would be more effectively balanced when matching an exposed period with unexposed periods both before and after. Therefore, Matching 1-2 can be more accurate in imputing missing potential outcomes for exposed time periods, and more efficient in estimating causal effects, compared to Matching 1-1.

3.1.3 Matching 1-1/2.

While Matching 1-2 might provide more accurate predictions of missing potential outcomes in some cases, it might lead to fewer matched exposed units than Matching 1-1, especially in scenarios with a relatively high proportion of exposed time periods. We propose an approach that combines 1-1 and 1-2 matching and enjoys the advantages of both. This approach matches one exposed period to one or two unexposed time periods. Specifically, we consider binary matching indicators of the form $a_{t_e t_u}$ and $a_{t_e t_{u_1} t_{u_2}}$ where $a_{t_e t_u} = 1$ denotes that exposed time period t_e is matched to unexposed time period t_u , whereas $a_{t_e t_{u_1} t_{u_2}} = 1$ denotes that t_e is matched to two unexposed time periods, t_{u_1} and t_{u_2} . The target is to maximize the number of matched exposed time periods as

$$\max_{\mathbf{a}} \left(\sum_{t_e, t_u} a_{t_e t_u} + \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}} \right). \quad (\text{C})$$

The constraints we impose are similar in spirit to those in optimization problems (A) and (B), but they are re-designed to account for the presence of two types of matches. Each time period can be involved in either a match of type 1-1 or 1-2, and at most once,

$$\begin{aligned} \sum_{t_u} a_{t_e t_u} + \sum_{t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}} &\leq 1, \\ \sum_{t_e, t_u} a_{t_e t_u} I(t_u = \tilde{t}_u) + \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}} I(t_{u_1} = \tilde{t}_u) + \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}} I(t_{u_2} = \tilde{t}_u) &\leq 1, \end{aligned} \quad (\text{C.1})$$

for all $t_e \in \mathcal{T}_e$ and $\tilde{t}_u \in \mathcal{T}_u$. Moreover, the average time difference between an exposed time period and its one or two matches is bounded by the constant $\delta \geq 0$, as

$$\left| \sum_{t_e, t_u} a_{t_e t_u} (t_e - t_u) + \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}} \left(t_e - \frac{t_{u_1} + t_{u_2}}{2} \right) \right| \leq \delta \left(\sum_{t_e, t_u} a_{t_e t_u} + \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}} \right). \quad (\text{C.2})$$

Furthermore, the time difference between matched exposed and unexposed time periods is bounded by ϵ in 1-1 or 1-2 matches, and in 1-2 matches the exposed time period lies temporally between the two matched unexposed ones, as

$$\begin{aligned} |a_{t_e t_u} (t_e - t_u)| &\leq \epsilon, & |a_{t_e t_{u_1} t_{u_2}} (t_e - t_{u_i})| &\leq \epsilon, \text{ for } i = 1, 2 \\ a_{t_e t_{u_1} t_{u_2}} (t_e - t_{u_1}) &\geq 0, & \text{and } a_{t_e t_{u_1} t_{u_2}} (t_{u_2} - t_e) &\geq 0. \end{aligned} \quad (\text{C.3})$$

Next, on average, matches are balanced in terms of the outcome unit characteristics

$$\begin{aligned} \left| \sum_{t_e, t_u} a_{t_e t_u} (\mathbf{W}_{t_e} - \mathbf{W}_{t_u}) + \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}} \left(\mathbf{W}_{t_e} - \frac{\mathbf{W}_{t_{u_1}} + \mathbf{W}_{t_{u_2}}}{2} \right) \right| \\ \leq \mathbf{1}_{p_W} \cdot \delta' \left(\sum_{t_e, t_u} a_{t_e t_u} + \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}} \right) \end{aligned} \quad (\text{C.4})$$

and the time-varying covariates of the interventional units and the network,

$$\begin{aligned} \left| \sum_{t_e, t_u} a_{t_e t_u} (\tilde{\mathbf{X}}_{t_e} - \tilde{\mathbf{X}}_{t_u}) + \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}} \left(\tilde{\mathbf{X}}_{t_e} - \frac{\tilde{\mathbf{X}}_{t_{u_1}} + \tilde{\mathbf{X}}_{t_{u_2}}}{2} \right) \right| \\ \leq \mathbf{1}_{p_X} \cdot \delta' \left(\sum_{t_e, t_u} a_{t_e t_u} + \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}} \right) \\ \left| \sum_{t_e, t_u} a_{t_e t_u} (\tilde{\mathbf{P}}_{t_e} - \tilde{\mathbf{P}}_{t_u}) + \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}} \left(\tilde{\mathbf{P}}_{t_e} - \frac{\tilde{\mathbf{P}}_{t_{u_1}} + \tilde{\mathbf{P}}_{t_{u_2}}}{2} \right) \right| \\ \leq \mathbf{1}_{p_P} \cdot \delta' \left(\sum_{t_e, t_u} a_{t_e t_u} + \sum_{t_e, t_{u_1}, t_{u_2}} a_{t_e t_{u_1} t_{u_2}} \right) \end{aligned} \quad (\text{C.5})$$

Matching 1-1/2 is expected to yield more matches than Matching 1-2 since it allows some exposed time points to be matched to a single unexposed one. At the same time, when possible, it allows for exposed time periods to be matched to two unexposed ones, which can improve the balance of temporal trends and improve accuracy in imputing the missing potential outcomes for exposed time periods.

Even in the case where the exposure is not binary, matching algorithms designed for binary expo-

tures can still be useful for estimating interpretable effects in more complex scenarios. For instance, in the setting of [Pouget-Abadie et al. \[2019\]](#), [Doudchenko et al. \[2020\]](#), [Brennan et al. \[2022\]](#), and [Harshaw et al. \[2023\]](#) where an outcome unit’s exposure ranges from 0 to 1, one can categorize the continuous exposure to different levels, and adapt the matching algorithms to match units with different values of the categorized exposure.

3.2 The matching estimators and theoretical guarantees

The matches produced by the three algorithms are the basis for estimating the causal effect $\tilde{\tau}(1, 0)$ since they are used to impute an exposed time period’s counterfactual outcome, had it been unexposed. If the exposed time period t_e is matched one-to-one with the unexposed t_u , its counterfactual outcome is imputed as $Y_{t_e}^{\text{imp}}(0) = Y_{t_u}$. Alternatively, when an exposed time period t_e is matched with two unexposed ones, t_{u_1} and t_{u_2} , its counterfactual outcome is imputed as the average of the two unexposed outcomes, $Y_{t_e}^{\text{imp}}(0) = (Y_{t_{u_1}} + Y_{t_{u_2}})/2$.

The set of matched exposed time periods and the imputed outcomes might differ across the matching algorithms, resulting in three causal estimators, one for each algorithm. The causal estimators have the same form, defined as the average contrast of observed and imputed outcomes for matched, exposed time periods in \mathcal{T}_e^* , as

$$\hat{\tau} = \frac{1}{|\mathcal{T}_e^*|} \sum_{t \in \mathcal{T}_e^*} [Y_{t_e} - Y_{t_e}^{\text{imp}}(0)]. \quad (1)$$

Specifically, for Matching 1-1, $\mathcal{T}_e^* = \{t : E_t = 1 \text{ and } a_{t_e t_u} = 1 \text{ for some } t_u\}$, and similarly for Matching 1-2 or Matching 1-1/2. We denote these estimators as $\hat{\tau}_1$, $\hat{\tau}_2$ and $\hat{\tau}_{1/2}$. We implicitly assume that \mathcal{T}_e^* is non-empty; otherwise, matching while satisfying balance constraints is not possible and drawing causal inferences from such data might not be trustworthy.

We show that the bias of the proposed matching estimators is bounded. We consider cases where the outcome is a linear or non-linear function of the exposure, time, and time-varying covariates, in line with [Zubizarreta \[2015\]](#). We first consider the linear case.

Theorem 1. If $Y_t(e) = \beta_0 + \beta_1 e + \beta_2 t + \beta_3^\top \mathbf{W}_t + \beta_4^\top \tilde{\mathbf{X}}_t + \beta_5^\top \tilde{\mathbf{P}}_t + \epsilon_t(e)$ for all $t = 1, 2, \dots, T$, with $E(\epsilon_t(e) | E_t, t, \mathbf{W}_t, \mathbf{X}_t, \mathbf{P}_t) = 0$, then $|E(\hat{\tau} - \tau)| \leq \delta |\beta_2| + \delta' (\|\beta_3\|_1 + \|\beta_4\|_1 + \|\beta_5\|_1)$ for all matching estimators, where δ and δ' are the balance constraints tuning parameters.

The proof is in Supplement A. According to [Theorem 1](#), the bias of the matching estimators is bounded by algorithmic parameters controlling how well time-varying information is balanced, and the strength of time-varying confounding in the outcome structure. Since δ, δ' can be set arbitrarily

small, the bias of the matching estimators can, in principle, be guaranteed to be small. In practice, using small values for δ, δ' might return a small number of matches and an estimated effect that is not representative of all exposed time periods.

These results extend to the more realistic case where the outcome model is non-linear in the confounding structure. In this case, we extend the matching algorithms to impose balance constraints for auxiliary variables targeting higher order and localized versions of time and the measured covariates. For example, we define localized versions of the s^{th} outcome unit covariate, W_{ts} , by breaking its support $[a, b]$ into $(b-a)/\ell$ intervals of an arbitrary small length ℓ . The midpoint of the r^{th} interval for W_{ts} is denoted by ξ_{sr} . We construct the auxiliary variables $W_{tsr}^\dagger = (W_{ts} - \xi_{sr})I(W_{ts} \in [\xi_{sr} - \ell/2, \xi_{sr} + \ell/2])$ for covariate W_{ts} , as well as higher orders W_{tsr}^k , for $k = 1, 2, \dots, K-1$. We include balance constraints (similarly to the ones in (A.4)) for these auxiliary variables in all three matching algorithms.

We show that the bias of the causal effect estimators is still bounded, where the bound is driven by algorithmic parameters and the smoothness of functions in the outcome model.

Theorem 2. Suppose $Y_t(e) = \theta + \beta e + h_0(t) + \sum_{s=1}^{p_W} h_s(W_{ts}) + \sum_{s=1}^{p_X} h_{p_W+s}(\tilde{X}_{ts}) + \sum_{s=1}^{p_P} h_{p_W+p_X+s}(\tilde{P}_{ts}) + \epsilon_t(e)$, with $E(\epsilon_t(e)|E_t, t, \mathbf{W}_t, \mathbf{X}_t, \mathbf{P}_t) = 0$ and functions $h_0, h_1, \dots, h_{p_W+p_X+p_P}$ that are K -times differentiable on their support. If $h_s^{(k)}$ represents the k^{th} derivative of h_s , and $|h_s^{(k)}(x)| \leq c$ for some $c > 0$ for all s, x in the function's support, and $k = 1, 2, \dots, K$, then $|E(\hat{\tau} - \tau)| \leq C_T \delta + C_{WXP} \delta' + C_{TWXP} \ell^{K-1}$, where C_T, C_{WXP} and C_{TWXP} are constants proportional to c that depend on the smoothness of the functions with the corresponding indices.

The exact form of C_T, C_{WXP} and C_{TWXP} is shown in Supplement A.3. Theorem 2 establishes that, by setting the algorithms' tuning parameters δ, δ' and ℓ to be small enough, the bias of the corresponding causal effect estimators can be guaranteed to be negligible. Since the exposure is binary, the form βe in the outcome model suffices. Extending our results to allow for interactions among the covariates would be theoretically straightforward. However, practically, the matching algorithms would need to impose additional balancing constraints, which might hinder our ability to find adequate matches.

This is particularly relevant since, as with all matching procedures, the estimated effect is representative of the population of only those exposed time periods that are matched, and the targeted estimand is, in fact, $\tilde{\tau}(1, 0) = \sum_{t_e \in \mathcal{T}_e^*} [Y_{t_e}(1) - Y_{t_e}(0)] / |\mathcal{T}_e^*|$. Therefore, its interpretation might be complicated when the proportion of unmatched exposed time periods is large. In these cases, and if the causal effect is heterogeneous across time, the estimated effect might differ from the effect on all the exposed time periods $\tilde{\tau}(1, 0)$. We investigate the performance of our estimators with heterogeneous effects in the

simulations of Section 4.

3.3 Inference

Our inferential approach is formulated in a unified manner for the three matching estimators. By viewing matching as part of the design phase, we construct confidence intervals conditional on the matched data [Ho et al., 2007]. Since our estimators are averages of differences between a time period’s observed and imputed outcome, we construct Wald-type confidence intervals. Specifically, if \mathcal{T}_e^* is the set of matched exposed time periods, and $\hat{s}^2 = \sum_{t_e \in \mathcal{T}_e^*} (Y_{t_e} - Y_{t_e}^{\text{imp}}(0) - \hat{\tau})^2 / (|\mathcal{T}_e^*| - 1)$, we construct an α -level confidence interval as $\left[\hat{\tau} - z_{1-\alpha} \hat{s} / \sqrt{|\mathcal{T}_e^*|}, \hat{\tau} + z_{1-\alpha} \hat{s} / \sqrt{|\mathcal{T}_e^*|} \right]$, where $z_{1-\alpha}$ is the $1 - \alpha$ quantile of the standard normal. We similarly acquire a p-value for testing the null hypothesis of no causal effect on outcome unit m_j , $H_{j0} : \tilde{\tau}_j = 0$ v.s. $H_{jA} : \tilde{\tau}_j \neq 0$ as $p_j = P\left(|Z| > \left| \frac{\hat{\tau}_j}{\hat{s} / \sqrt{|\mathcal{T}_e^*|}} \right| \right)$, where $Z \sim N(0, 1)$. P-values in one-sided hypothesis tests can be obtained similarly.

In practice, time series data may display temporal correlation beyond what can be explained by measured covariates. Despite that, the simulations in Section 4 demonstrate that the proposed inferential approach yields valid inferences even with temporally correlated outcomes. Setting aside temporal correlation, the inferential approach is expected to be conservative when time-varying confounders are present, in that α -level confidence intervals cover the true value more than $100\alpha\%$ of the time. That is because the proposed matching algorithms do not balance time-varying covariates for each match separately. As a result, the unexposed time periods that are used to impute the potential outcomes might have substantially different values for the temporal covariates compared to the corresponding exposed time periods. Therefore, the differences of observed and imputed outcomes include fluctuations in temporal predictors, leading to an estimated variance that is larger than the truth. We illustrate this slight over-coverage in the simulations of Section 4, where we observe that balancing covariates within every match could alleviate this issue at the cost of returning fewer matches.

3.4 Testing a null hypothesis of no causal effect with multiple outcome units

Our matching algorithms and estimators are designed to evaluate the effect of the treatment on each outcome unit separately. In the presence of multiple outcome units, and when making general policy evaluations, we might be interested in studying whether the exposure has an effect on *any* of them. The hypothesis we wish to test is

$$H_0 : \tau_j = 0, \forall j = 1, 2, \dots, M \quad \text{v.s.} \quad H_A : \text{There exist at least one } j \text{ such that } \tau_j \neq 0.$$

We acquire a p-value for testing the null hypothesis of no effect on unit m_j , according to Section 3.3, for all outcome units. We adjust these p-values by performing a false discovery rate (FDR) correction for multiple comparisons [Benjamini and Hochberg, 1995]. Then, we compare the adjusted p-values to the pre-specified α -level. If all adjusted p-values are greater than α , we fail to reject the null hypothesis H_0 ; otherwise, we reject the null hypothesis and identify the affected units as those with adjusted p-values less than α .

3.5 The potential advantages of temporal analyses in bipartite settings

Alternative estimands to the ones in Section 2.2 represent cross-sectional contrasts such as the average causal effect across units for each time period, defined as $\gamma_t(e, \tilde{e}) = \frac{1}{M} \sum_{j=1}^M [Y_{tj}(e) - Y_{tj}(\tilde{e})]$. Estimation of these unit-average estimands requires that we measure and adjust for all meaningful differences across units that confound the exposure-outcome relationship, which can be complex and high-dimensional. For example, in the study of Zigler and Papadogeorgou [2021], the treatment assignment of power plants and population health can vary across the United States in intricate ways, all of which need to be adjusted for estimating unit-average effects.

In contrast, estimation of the temporally-average causal effects for each outcome unit requires that we account for time-varying confounding only, which might be simpler to understand and measure, an observation that was also noted in the interrupted time series literature [Rockers et al., 2015]. For example, our results show that if the treatment of interventional units is constant over time, any variation in the exposure for an outcome unit is due to the varying bipartite network, and therefore confounding is only due to covariates that predict the network and the outcome (Assumption 3). In this case, if the random network depends only on the units’ time-invariant characteristics like their geographic distance, no confounding adjustment would be necessary to estimate interpretable and policy-relevant estimands. Alternatively, if the network is driven by naturally-occurring processes with temporal variation such as meteorology, one would only need to account for those for causal effect estimation, which are simpler to understand and measure. Furthermore, if confounding variables show relatively smooth temporal trends during the time window under study, such as weather variables, collecting them is unnecessary since they are indirectly balanced in our matching algorithms (Theorems 1 and 2).

The inherent bipartite nature of the data suggests that temporal confounding is likely to display smoother trends compared to unipartite scenarios. In bipartite settings, the separation of physical units implies that decisions affecting one set of units may not immediately manifest and impact the other. Consequently, if an interventional unit variable influences the outcome units, it might be due to its overall trend over time, such as its average over preceding time periods. In that case, this ‘moving

average’ covariate value will be relatively smooth across time. If left unmeasured, the bias occurring due to its non-temporal component is expected to be small. We illustrate this in the simulations of Section 4.

4 Simulation Study

We perform simulations to investigate the performance of our matching estimators and the properties of our inferential procedures.

4.1 Simulation setup

We consider a setting with $N = 50$ interventional units and $M = 200$ outcome units at randomly generated locations over the $[0, 1] \times [0, 1]$ square, followed over $T = 400$ time periods.

We consider covariates for the interventional units, the outcome units, and the network: (a) (Smooth temporal trends) Covariates X_{ti1} and W_{tj1} are generated independently from Gaussian processes with the same smooth function of time as the mean, and an exponential decay kernel for the covariance matrix. Therefore, these covariates represent similar but not identical smooth temporal trends. (b) (Location-varying covariates) Covariates X_{ti2} and W_{tj2} are constant across time, $X_{ti2} = X_{i2}$ and $W_{tj2} = W_{j2}$, and they are drawn independently from a scaled beta distribution with parameters that depend on the unit’s location. Therefore, these covariates have similar structure across space. (c) (Location- and time-varying covariates) Covariates X_{ti3} , W_{tj3} and P_{tij} are independent across units, they have temporal trends, but they also include *non-smooth* temporal variation. (d) (Bipartite covariates) We define covariates for one set of units based on the covariates of the other set. For interventional units, we define location-varying covariate X_{i4} , and time-varying covariate X_{ti5} , that are weighted averages of covariates W_{i2} and W_{tj3} of neighboring outcome units, respectively. Covariates W_{i4} and W_{ti5} for outcome units are similarly defined based on covariates X_{i2} and X_{ti3} of interventional units. (e) (Non-smooth time-varying covariate) Covariates X_{t6} , W_{t6} are equal to each other, and represent *non-smooth* temporal trends. Draws from these covariates are depicted in Supplement B.

4.1.1 Data generative mechanisms.

We consider five data generative models corresponding to different confounding structures. Across these scenarios, the treatment assignment for the interventional units depends on smooth temporal trends through X_{ti1} , on interventional unit covariates through X_{ti2} , X_{ti3} , on outcome unit covariates through X_{ti4} , X_{ti5} , on non-smooth temporal trends through X_{ti3} , X_{ti5} , X_{ti6} , and on network covariates through P_{tij} . The entries of the bipartite network are generated independently from Bernoulli distributions with probability that, under the different scenarios, might depend on time and units’ spatial

proximity. The exposure of unit m_j at time t is specified as $E_{tj} = I(\sum_{i=1}^N A_{ti}G_{tij} \geq d)$. The outcome is generated based on the exposure and the covariates $W_{ti1}, W_{ti2}, \dots, W_{ti6}$ and P_{tij} which include smooth and non-smooth temporal trends, outcome and interventional unit covariates, and network covariates.

Table 1 shows the variables that are used in each data generative model component across the different scenarios. These five scenarios correspond to five different confounding structures for the exposure-outcome relationship: (a) no confounding, (b) confounding by smooth temporal variables, (c) confounding by location-varying variables, (d) confounding by all time-varying information, and (e) all types of confounding. The simulation scenarios are detailed in Supplement B. Confounding arises if a predictor of the outcome is correlated with a predictor of the treatment assignment, the bipartite network, or both. For example, in scenario (b), the role of X_{ti1} and W_{tj1} induces confounding due to the variables' common smooth temporal trend, and in scenario (d), the role of W_{tj3} and X_{ti5} induce confounding due to the outcome unit covariate. Therefore, these data generative models allow for complex confounding structures in accordance to Assumptions 2 and 3.

For each scenario, we consider three sparsity levels for the exposure by tuning d , dense, medium, and sparse, corresponding to about 150-200, 80-120, and 30-60 exposed time periods, respectively. Therefore, in total, we consider 15 simulation scenarios, and generate 500 data sets for each one of them.

Table 1: Table of five confounding scenarios, in which treatment \mathbf{A} , network graph \mathbf{G} , and observed outcome \mathbf{Y} are associated with corresponding confounding covariates.

Scenario	Component	Smooth time			Location-varying				Time-varying						
		t	X_1	W_1	dist	X_2	W_2	X_4	W_4	X_3	W_3	X_5	W_5	$X_6(=W_6)$	P
(a) No confounders	\mathbf{A}														
	\mathbf{G}														
	\mathbf{Y}														×
(b) Time-smooth confounders	\mathbf{A}		×												
	\mathbf{G}				×										
	\mathbf{Y}					×									
(c) Location-varying confounders	\mathbf{A}					×		×							
	\mathbf{G}														
	\mathbf{Y}					×			×						
(d) Time-varying confounders	\mathbf{A}									×		×		×	×
	\mathbf{G}	×			×										
	\mathbf{Y}										×		×	×	×
(e) All confounders	\mathbf{A}		×			×		×		×		×		×	×
	\mathbf{G}	×			×									×	×
	\mathbf{Y}					×		×		×		×		×	×

4.1.2 Estimation and inference evaluation.

We estimate the temporally-averaged causal effect specific to an outcome unit. We fit our matching algorithms employing balance constraints on the time-varying covariates W_3, W_5, W_6 and P only, shown in the last wide column of Table 1. Since the covariates X_1, W_1 represent smooth temporal trends (denoted by $f(t)$ in Proposition 1), we do *not* consider balance constraints on them, illustrating that the constraints on time suffice. For a consistent choice of tuning parameter δ' across covariates, we standardize each of the time-varying covariates W_3, W_5, W_6 and \tilde{P} using the pooled standard deviation of exposed and unexposed time periods [Rosenbaum and Rubin, 1985]. For example, the entries of W_3 are divided by $\sqrt{(\text{Var}(W_{t3} \text{ for exposed } t) + \text{Var}(W_{t3} \text{ for unexposed } t))/2}$, and similarly for the rest of the covariates. Therefore, δ' denotes the allowed covariate imbalance as the proportion of the covariate's standard deviation. We consider three sets of tuning parameters $(\delta, \delta', \epsilon)$. The results shown here correspond to values $(2, 0.05, 6)$. Alternative choices for the tuning parameters are discussed in Section 4.2.3 and shown in the Supplement. We estimate the causal effect using (1), and acquire 95% confidence intervals as detailed in Section 3.3.

Since there do not exist alternative approaches in the literature for estimating causal effects in bipartite time series settings, we implement three naïve approaches. Naïve- t uses temporal information for the single outcome unit and estimates an effect as the difference of mean outcomes between exposed and unexposed time periods. Naïve- j uses information across outcome units for a single time period and estimates an effect as the difference of mean outcomes between exposed and unexposed outcome units. Lastly, Naïve-all estimates an effect as the overall difference of mean outcomes in exposed and unexposed time periods across units. Additional details on the naïve approaches are included in Supplement C.

Finally, for the scenarios with temporal confounding (b, d, and e), we consider simulations where all treatment effects are set to zero, and evaluate the properties of the inferential technique of Section 3.4 for testing the global null hypothesis at the 0.05 level.

4.2 Simulation results

4.2.1 Estimation and inference on a single unit.

Table 2 shows the estimation and inferential results for estimating the effect for one outcome unit using the three naïve approaches, and the three matching estimators. For each estimator we report bias, mean squared error and coverage of 95% intervals. For the matching estimators, we also report the proportion of exposed time periods that were matched.

Table 2: Simulation results of single-unit estimations. Bias, mean squared error (MSE), coverage of 95% intervals (%), and proportion of exposed time points being matched (%). We show simulation results for the 3 naïve approaches and 3 matching estimators over 5 confounding scenarios, and for 3 exposure levels. ‘N’ stands for ‘naïve’. Bold values correspond to the minimum MSE across the matching methods for each simulation scenario.

		Dense				Medium				Sparse			
	Method	Bias	MSE	Cover	Prop	Bias	MSE	Cover	Prop	Bias	MSE	Cover	Prop
(a) ¹	N- <i>t</i>	0.00	0.010	95.4	-	-0.01	0.013	95.0	-	-0.01	0.022	95.0	-
	N- <i>j</i>	0.00	0.207	94.0	-	0.00	0.316	94.8	-	0.03	0.716	94.6	-
	N-all	0.00	0.000	94.8	-	0.00	0.001	95.2	-	0.00	0.001	95.8	-
	1-1	0.00	0.012	95.0	97.8	-0.01	0.020	95.4	100.0	-0.01	0.040	95.6	100.0
	1-1/2	0.00	0.011	95.8	97.8	-0.01	0.017	96.2	100.0	-0.01	0.031	94.6	100.0
	1-2	0.00	0.013	97.2	62.3	-0.01	0.016	95.4	90.8	-0.01	0.028	95.6	98.7
(b)	N- <i>t</i>	-0.93	1.22	11.4	-	-0.95	1.294	11.8	-	-0.96	1.382	15.8	-
	N- <i>j</i>	0.01	0.097	95.8	-	0.00	0.054	93.4	-	0.00	0.047	94.4	-
	N-all	-0.90	0.810	0.0	-	-0.94	0.886	0.00	-	-0.98	0.962	0.0	-
	1-1	0.00	0.018	96.0	62.5	0.00	0.024	92.8	85.1	0.010	0.037	94.8	97.8
	1-1/2	0.00	0.018	94.8	62.5	0.00	0.022	94.2	85.1	0.010	0.035	93.8	97.8
	1-2	0.00	0.023	94.0	41.0	0.00	0.024	95.0	59.9	0.01	0.034	94.8	80.6
(c)	N- <i>t</i>	0.00	0.012	95.2	-	0.00	0.023	94.6	-	0.00	0.048	95.0	-
	N- <i>j</i>	-3.58	23.305	77.6	-	-3.82	28.191	81.6	-	-4.27	45.544	88.0	-
	N-all	-3.32	11.44	0.0	-	-3.55	13.042	0.00	-	-3.89	15.654	0.0	-
	1-1	0.01	0.020	94.4	99.7	0.00	0.037	95.0	100.0	-0.01	0.080	94.8	100.0
	1-1/2	0.01	0.017	95.2	99.7	0.01	0.033	94.2	100.0	0.00	0.073	92.6	100.0
	1-2	0.00	0.016	94.4	85.5	0.00	0.032	94.0	97.4	0.00	0.071	92.0	99.1
(d)	N- <i>t</i>	-1.19	1.479	0.0	-	-1.23	1.572	0.00	-	-1.36	1.943	4.4	-
	N- <i>j</i>	0.01	0.076	96.3	-	0.00	0.042	94.4	-	0.00	0.032	94.8	-
	N-all	-1.20	1.438	0.00	-	-1.24	1.551	0.0	-	-1.37	1.872	0.0	-
	1-1	-0.05	0.023	98.0	67.6	-0.05	0.024	98.4	85.0	-0.03	0.062	95.2	99.7
	1-1/2	-0.05	0.022	98.4	67.6	-0.05	0.024	98.2	85.0	-0.04	0.054	96.6	99.7
	1-2	-0.04	0.026	98.8	43.5	-0.04	0.025	97.6	59.4	-0.04	0.053	96.0	87.5
(e)	N- <i>t</i>	2.48	6.417	0.0	-	-2.6	7.092	0.0	-	-2.78	8.114	0.0	-
	N- <i>j</i>	-2.20	11.39	79.3	-	-1.88	6.566	78.6	-	-1.84	6.457	78.1	-
	N-all	-3.46	12.048	0.0	-	-3.72	13.869	0.0	-	-4.06	16.583	0.0	-
	1-1	-0.06	0.027	98.4	75.0	-0.06	0.033	98	94.1	-0.03	0.045	98.8	99.6
	1-1/2	-0.07	0.028	98.0	75.0	-0.06	0.032	97.2	94.1	-0.04	0.041	98.4	99.6
	1-2	-0.07	0.028	99.2	48.9	-0.06	0.032	98.6	71	-0.04	0.042	97.8	88.3

¹The scenarios shown in Table 1 correspond to (a) No confounders, (b) Time-smooth confounders, (c) Location-varying confounders, (d) Time-varying confounders, and (e) All confounders.

In the case of no confounding, all estimators are unbiased. In the presence of location-varying confounding (scenario (c)), the Naïve-*j* estimator that is based on comparing outcomes across units is biased. In contrast, under temporal confounding (scenarios (b) and (d)), the Naïve-*t* estimator is biased. Naïve-all is biased under both confounding structures. In contrast, the matching estimators have minimal bias in all cases.

In terms of MSE, Matching 1-1 performs best or close to best in the dense scenarios, and Matching

Table 3: Simulation results for using Naïve- t and 3 adjusted matching methods to test the global null hypothesis for scenarios (b), (d) and (e), and under a medium exposure level.

		Average estimator mean	Average rate of p-value ≤ 0.05	Rate of min(p-value) ≤ 0.05	Rate of FDR min (adj.p-value) ≤ 0.05
(b)	Naïve- t	-0.95	0.894	1.000	1.000
	1-1	-0.01	0.053	1.000	0.088
	1-1/2	-0.01	0.054	1.000	0.090
	1-2	0.00	0.056	1.000	0.110
(d)	Naïve- t	-1.26	0.999	1.000	1.000
	1-1	-0.04	0.018	0.966	0.014
	1-1/2	-0.04	0.018	0.970	0.016
	1-2	-0.04	0.014	0.910	0.016
(e)	Naïve- t	-2.64	1.000	1.000	1.000
	1-1	-0.06	0.015	0.946	0.004
	1-1/2	-0.06	0.016	0.954	0.004
	1-2	-0.06	0.011	0.848	0.002

1-2 in the sparse scenarios. Matching 1-1/2 performs as well, or close to as-well as the best estimator across all scenarios considered, while maintaining a high proportion of matched exposed time periods. Matching 1-1 and Matching 1-1/2 yield the same proportion of matched exposed time periods. This alignment is logical as the matching algorithms aim to maximize matches, with the matches generated under Matching 1-1 also possible under Matching 1-1/2. For all three matching algorithms, the matching rate varies by the sparsity level of the exposure, with higher rates under sparser exposures. As expected, Matching 1-2 returns the smallest proportion of matched exposed time periods. Combined with the fact that it has the lowest MSE in sparse scenarios, this illustrates that Matching 1-2 returns more accurate imputed potential outcomes than Matching 1-1 or 1-1/2. The coverage of 95% intervals for the matching estimators is close to nominal across all scenarios.

4.2.2 Testing the global null hypothesis.

We focus on the scenarios with temporal, or all types of confounding (scenarios b, d, and e), and under a medium frequency for the exposure. We alter the simulations to impose that the global null holds, and impose that $\tau_j = 0$ for all outcome units. We generate 500 data sets for each one of the three scenarios. We acquire point estimates and p-values for the exposure effect on each of the 200 outcome units, and adjust the p-values using the FDR correction detailed in Section 3.4.

The optimizer returned a solution for approximately 85% of the outcome units for each matching method. The results are in Table 3. We report the average estimated effect across outcome units with matches and across data sets, the proportion of the available p-values across the outcome units and data sets that are below 0.05, the proportion of data sets where any of the available p-values is below 0.05, and the proportion of data sets where any of the FDR-adjusted p-values is below 0.05.

Since Naïve- t is biased in the presence of temporal confounding, its inferential performance suffers,

and using this estimator would mistakenly reject the global null hypothesis every time, using FDR-corrected p-values or not. For the matching estimators, up to 6 % of the p-values across outcome units and data sets are below 0.05. Hence, it is not surprising that the minimum (unadjusted) p-value across outcome units is below 0.05 in most data sets, emphasizing the necessity of controlling for multiple comparisons to maintain the level of the test, especially with numerous outcome units. With FDR-adjusted p-values, the rate of rejection of the null hypothesis is much closer to the target level of the test. Matching methods in cases (d) and (e) are conservative in rejecting the global null hypothesis, consistent with the conservative nature of individual outcome unit hypothesis tests.

4.2.3 Additional simulations.

Additional simulations are detailed in Supplement D. In Supplement D.1, we show that the matching estimators *without* any adjustment for measured covariates are unbiased under temporally-smooth confounding. These results illustrate that smooth temporal trends such as those in X_1, W_1 are implicitly adjusted through the time constraints, and need *not* be measured. In Supplement D.2, we assess the performance of matching methods with alternative tuning parameter values. Results are robust to the choice of δ and ϵ , though 1-2 often returns a substantially lower number of matches under strict conditions. Despite some residual bias under larger δ' values in some scenarios, interval coverage is close to nominal across all scenarios. In Supplement D.3, we illustrate that applying covariate constraints to *each* match alleviates overcoverage of 95% intervals, supporting the discussion in Section 3.3 regarding the conservativeness of our inferential approach. In Supplement D.4, we illustrate that confounding can be induced by predictors of the network and the outcome, even if they are not predictors of the treatment assignment, and matching estimators perform accurately in this scenario as well. All the simulations are under a homogeneous treatment effect to separate the evaluation of estimation efficiency from the discussion on targeted estimand, and ease comparison of estimators. In Supplement D.5, we show that under treatment effect heterogeneity, results from matching estimators are representative of the matched population of exposed time periods. Lastly, in Supplement D.6, we find that temporal correlation in the outcome variable has minimal impact on the performance of our inferential procedure.

5 The Impact of Wildfire Smoke Exposure on Bikeshare Hours

Wildfires contribute to increased levels of ozone and fine particulate matter. Smoke from these fires is carried by the wind to populated regions, potentially causing reductions in the population's outdoor activity [Doubleday et al., 2021]. We quantify the effect of wildfire smoke on the use of bikeshare

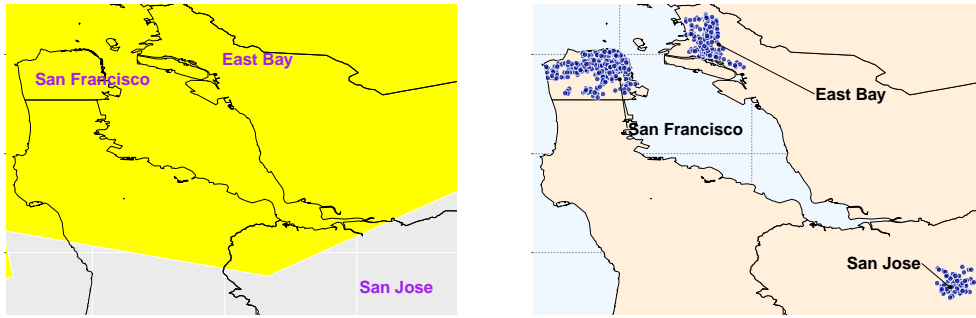


Figure 1: Left: The wildfire smoke coverage on August 31, 2021, where yellow represents exposure to wildfire smoke. Right: The bike share locations in the three outcome units.

services from January 2021 to September 2023. In this study, it is reasonable to assume that the individuals’ decision to ride a bicycle might be driven by the exposure value on the same day only, and any temporal carry-over effects are minimal. We acquire daily smoke exposure from the National Oceanic and Atmospheric Administration’s Hazard Mapping System (HMS). An area is considered exposed under light or heavier smoke thickness according to HMS, and unexposed otherwise. The outcome represents the daily total number of bikeshare hours in the San Francisco, East Bay, and San Jose areas in California, US, as measured through the publicly-available Bay Wheels data provided by Lyft. Figure 1 shows smoke exposure for a single day during our time period and the locations of the stations in the three areas. Additional information on the data is available in Supplement E.1.

Traditional unit-to-unit cross-sectional analyses would be infeasible in our study as it would be impossible to control for attributes of interventional and outcome units, with only three outcome units (see discussion in Section 3.5). However, given daily data on 1,003 days, approximately 140 of which are exposed across the three areas, our approach can be applied to estimate the effect of smoke exposure on each outcome area, without the need to consider location-varying covariates, and controlling for time-varying confounding only.

The matching algorithms balance daily temperature, humidity, precipitation, wind speed, and wind direction as potential time-varying confounders, and smooth seasonal trends are balanced implicitly. We find it plausible that no further covariates are necessary beyond weather-related data to meet the unconfoundedness assumptions. This reasoning stems from our understanding of the bipartite structure of the problem: factors (other than weather conditions) influencing wildfire occurrence and smoke dispersion patterns are unlikely to impact biking activity in distant locations, and economic indices fluctuating over time that might affect biking activity are likely unrelated to wildfire presence in North American forests. Therefore, while estimating the causal effect of exposure using data at the outcome unit level may resemble a unipartite setting, the investigation of unconfoundedness based on measured

Table 4: The effect of wildfire smoke on bikeshare hours in San Francisco, East Bay, and San Jose for Naïve- t and the matching estimators. For each region, the three columns correspond to the estimate, p-value (bold font if below 0.05), and number of matches.

	San Francisco			East Bay			San Jose		
Naïve- t	0.973	(1.000)		0.110	(1.000)		0.022	(0.810)	
Maching 1-1	-0.601	(0.014)	101	-0.043	(0.190)	103	-0.002	(0.235)	100
Matching 1-1/2	-0.521	(0.021)	101	-0.031	(0.265)	103	-0.002	(0.196)	100
Matching 1-2	-0.100	(0.356)	56	-0.001	(0.490)	59	-0.001	(0.279)	59

covariates stems by viewing the question at hand from a bipartite perspective, and explicitly investigating the treatment assignment at the level of the interventional units. Consequently, even though implementation does not necessarily differ, there exist subtle, yet important differences for how confounding might be investigated under a bipartite lens, compared to a unipartite one.

Table 4 shows the estimates and p-values from Naïve- t and the three matching estimators under the tuning parameters $(\delta, \delta', \epsilon) = (2, 0.1, 6)$. The estimate from the naïve approach is positive, implying that smoke increases biking activity. This unreasonable result is likely due to temporal confounding because most exposed time periods occur during the late summer and fall months when biking might be more prevalent. Instead, the matching estimators estimate that smoke exposure reduces biking activity in San Francisco, while riding behavior in the East Bay and San Jose is not influenced by wildfire smoke exposure. The matching algorithms use unexposed time periods during the summer and fall months only as matches (see Supplement E.2 for an illustration). These results are statistically significant at the 0.05 level for San Francisco under Matching 1-1, or 1-1/2. Since exposure is relatively dense during the summer and fall months, Matching 1-2 returns approximately half the number of matches compared to the other two matching algorithms, which might explain larger p-values for this estimator.

We conducted a sensitivity analysis to the definition of exposure, and performed the same analysis when an area is categorized as exposed during a give day under medium or high smoke thickness, and unexposed under no or light smoke thickness (see Supplement E.3). We find that all effect estimates are negative and similar or larger in magnitude in that case, even though most results are not-statistically significant, likely due to the small number of exposed days (40 or less) and as a result the small number of matches.

6 Discussion

We developed a causal inference framework for time series data with bipartite interference and a random network. Focusing on time-averaged estimands, we showed that controlling for time-varying information allows us to attribute outcome differences to exposure’s causal effects. We introduce three algorithms for matching exposed and unexposed time periods and corresponding estimators, which perform well across various scenarios. In principle, weighting or outcome modeling estimators could be employed in our context. However, in the bipartite setting, the exposure arises as a function of the treatment and network, whose mechanisms are often unknown and would need to be modeled, adding complexity to a weighting-based estimation procedure. At the same time, outcome modeling approaches would require correct specification of confounding adjustment, which is particularly complicated in bipartite settings where confounding relates to both the treatment and the network, as we show in Section 2. Matching methods bypass such issues.

Despite the merits of the proposed framework, several open questions remain. How to define and estimate effects for time series settings under more complicated exposure mappings that might involve multivariate or continuous exposures persists as an open question. Towards that front, future work could investigate matching, outcome modeling, and weighting approaches with the relative merits and drawbacks discussed above. Moreover, in certain applied contexts it would be crucial to allow for potential outcomes to depend on previous exposures [Bojinov and Shephard, 2019]. A careful consideration of necessary assumptions and an approach for estimation in this setting remains to be addressed in future research.

Acknowledgements

The authors thank Corwin Zigler, Fabrizia Mealli, Lucas Henneman, Jean Pouget-Abadie, and Jennifer Brennan for their useful comments in the preparation of this manuscript.

Funding

The authors are grateful for support from the National Science Foundation under Grant No 2124124.

References

- Alberto Abadie and Matias D Cattaneo. Econometric methods for program evaluation. *Annual Review of Economics*, 10:465–503, 2018.
- Alberto Abadie and Guido W Imbens. Large sample properties of matching estimators for average treatment effects. *Econometrica*, 74(1):235–267, 2006.

- Anish Agarwal, Sarah H Cen, Devavrat Shah, and Christina Lee Yu. Network synthetic interventions: A causal framework for panel data under network interference. 2023.
- Peter M. Aronow and Cyrus Samii. Estimating average causal effects under general interference, with application to a social network experiment. *Annals of Applied Statistics*, 11:1912–1947, 12 2017.
- Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)*, 57(1):289–300, 1995.
- Iavor Bojinov and Neil Shephard. Time series experiments and causal estimands: Exact randomization tests and trading. *Journal of the American Statistical Association*, 114:1665–1682, 10 2019.
- Jennifer Brennan, Vahab Mirrokni, and Jean Pouget-Abadie. Cluster randomized designs for one-sided bipartite experiments. *Advances in Neural Information Processing Systems*, 35:37962–37974, 2022.
- Jianfei Cao and Connor Dowd. Estimation and inference for synthetic control methods with spillover effects. *arXiv preprint arXiv:1902.07343*, 2019.
- Duncan A Clark and Mark S Handcock. An approach to causal inference over stochastic networks. *arXiv preprint arXiv:2106.14145*, 2021.
- Roberta Di Stefano and Giovanni Mellace. The inclusive synthetic control method. *Discussion Papers on Business and Economics, University of Southern Denmark*, 14, 2020.
- Annie Doubleday, Youngjun Choe, Tania M Busch Isaksen, and Nicole A Errett. Urban bike and pedestrian activity impacts from wildfire smoke events in seattle, wa. *Journal of Transport & Health*, 21:101033, 2021.
- Nick Doudchenko, Minzhengxiong Zhang, Evgeni Drynkin, Edoardo Airoldi, Vahab Mirrokni, and Jean Pouget-Abadie. Causal inference with bipartite designs. 10 2020.
- Laura Forastiere, Edoardo M Airoldi, and Fabrizia Mealli. Identification and estimation of treatment and interference effects in observational studies on networks. *Journal of the American Statistical Association*, 116(534):901–918, 2021.
- Giulio Grossi, Patrizia Lattarulo, Marco Mariani, Alessandra Mattei, and O Oner. Synthetic control group methods in the presence of interference: The direct and spillover effects of light rail on neighborhood retail activity. *arXiv preprint arXiv:2004.05027*, 2020.
- Christopher Harshaw, Fredrik Sävje, David Eisenstat, Vahab Mirrokni, and Jean Pouget-Abadie. Design and analysis of bipartite experiments under a linear exposure-response model. *Electronic Journal of Statistics*, 17(1):464–518, 2023.

- Daniel E Ho, Kosuke Imai, Gary King, and Elizabeth A Stuart. Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. *Political analysis*, 15(3):199–236, 2007.
- Guido W Imbens. Causal inference in the social sciences. *Annual Review of Statistics and Its Application*, 11, 2024.
- Luke Keele, Rocío Titiunik, and José R Zubizarreta. Enhancing a geographic regression discontinuity design through matching to estimate the effect of ballot initiatives on voter turnout. *Journal of the Royal Statistical Society, Series A*, 178:223–239, 2014.
- Fiammetta Menchetti and Iavor Bojinov. Estimating causal effects in the presence of partial interference using multivariate bayesian structural time series models. *Harvard Business School Technology & Operations Mgt. Unit Working Paper*, (21-048), 2020.
- Jean Pouget-Abadie, Kevin Aydin, Warren Schudy, Kay Brodersen, and Vahab Mirrokni. Variance reduction in bipartite experiments through correlation clustering. *Advances in Neural Information Processing Systems*, 32, 2019.
- Peter C Rockers, John-Arne Røttingen, Ian Shemilt, Peter Tugwell, and Till Bärnighausen. Inclusion of quasi-experimental studies in systematic reviews of health systems research. *Health Policy*, 119(4):511–521, 2015.
- Paul R Rosenbaum and Donald B Rubin. Constructing a control group using multivariate matched sampling methods that incorporate the propensity score. *The American Statistician*, 39(1):33–38, 1985.
- Fredrik Sävje. Causal inference with misspecified exposure mappings: separating definitions and assumptions. *Biometrika*, 111(1):1–15, 2024.
- Nathan B Wikle and Corwin M Zigler. Causal health impacts of power plant emission controls under modeled and uncertain physical process interference. *arXiv preprint arXiv:2306.05665*, 2023.
- Corwin Zigler, Vera Liu, Fabrizia Mealli, and Laura Forastiere. Bipartite interference and air pollution transport: Estimating health effects of power plant interventions. 12 2020.
- Corwin M Zigler and Georgia Papadogeorgou. Bipartite causal inference with interference. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 2021.
- José R Zubizarreta. Using mixed integer programming for matching in an observational study of kidney failure after surgery. *Journal of the American Statistical Association*, 107(500):1360–1371, 2012.
- José R Zubizarreta, Dylan S Small, Neera K Goyal, Scott Lorch, and Paul R Rosenbaum. Stronger instruments via integer programming in an observational study of late preterm birth outcomes. *The Annals of Applied Statistics*, pages 25–50, 2013.

José R. Zubizarreta. Stable weights that balance covariates for estimation with incomplete outcome data. *Journal of the American Statistical Association*, 110:910–922, 7 2015.

Supplementary materials

Table of Contents

A Theoretical results	3
A.1 Proof of exposure unconfoundedness	3
A.2 Proofs of bias bounds	3
A.3 The bounding constants in Theorem 2	6
B A Detailed Simulation Description of Section 4	6
B.1 The covariates	7
B.2 The treatment assignment, bipartite network, and outcome	10
C Description for naïve approaches	11
C.1 Estimator	11
C.2 Inference	11
D Additional Simulation Results	12
D.1 Unadjusted matching	12
D.2 Analysis of tuning parameters	13
D.3 Modification of matching methods: Covariate balancing within matches	14
D.4 Network-outcome confounding can lead to exposure-outcome confounding	17
D.5 Simulations under heterogeneous treatment effects	18
D.6 The impact of temporal correlation in the outcome variable on inference	20
E Additional study information	20
E.1 Information on data availability and creation	20

E.2 Illustrations of matched data for San Francisco	21
E.3 Results under alternative definition of exposure	21

A Theoretical results

A.1 Proof of exposure unconfoundedness

Proof of Proposition 1. First, since \mathbf{X}^* , \mathbf{W}^* , \mathbf{P}^* are time-invariant, we can treat them as constants in the condition and thus omit them for simplicity. Next, under Assumption 2:

$$\begin{aligned}
& P(E_{tj} = e | \mathcal{Y}_{tj}(\cdot), f(t), \mathbf{W}_{tj}, \mathbf{X}_t, \mathbf{P}_{t.j}) \\
&= \sum_{\forall \mathbf{g}, \mathbf{a}} P(E_{tj} = e | \mathbf{G}_{.tj} = \mathbf{g}, \mathbf{A}_t = \mathbf{a}, \mathcal{Y}_{tj}(\cdot), f(t), \mathbf{W}_{tj}, \mathbf{X}_t, \mathbf{P}_{t.j}) \cdot \\
&\quad P(\mathbf{G}_{.tj} = \mathbf{g} | \mathbf{A}_t = \mathbf{a}, \mathcal{Y}_{tj}(\cdot), f(t), \mathbf{W}_{tj}, \mathbf{X}_t, \mathbf{P}_{t.j}) \cdot \\
&\quad P(\mathbf{A}_t = \mathbf{a} | \mathcal{Y}_{tj}(\cdot), f(t), \mathbf{W}_{tj}, \mathbf{X}_t, \mathbf{P}_{t.j}) \\
&= \sum_{\forall \mathbf{g}, \mathbf{a}} I(h_{tj}(\mathbf{a}, \mathbf{g}) = e) P(\mathbf{G}_{.tj} = \mathbf{g} | \mathbf{A}_t = \mathbf{a}, f(t), \mathbf{W}_{tj}, \mathbf{X}_t, \mathbf{P}_{t.j}) P(\mathbf{A}_t = \mathbf{a} | f(t), \mathbf{W}_{tj}, \mathbf{X}_t, \mathbf{P}_{t.j})
\end{aligned}$$

The first equality holds because of the law of total probability and the conditional probability formula. The first component in the second equality is simply the indicator function of the exposure vector equivalent to a known vector or not; the second utilizes Assumption 3 and the third applies Assumption 2. \square

A.2 Proofs of bias bounds

Proof of Theorem 1. We consider the case of bounding the bias of the matching algorithms when the outcome model has a linear form in the exposure and time. Define I_1 as the set of (t_e, t_u) and I_2 as the set of (t_e, t_{u_1}, t_{u_2}) . Omitting j , set $\mathbf{E} = (E_1, E_2, \dots, E_T)^T$ for the observed exposure time series at unit m_j . Let $\mathbf{t} = (1, 2, \dots, T)^T$ be the sequence of time points.

- Matching 1-1

$$\begin{aligned}
& |E(\hat{\tau} - \tau)| = \\
&= \left| E \left(E \left(\frac{1}{|I_1|} \sum_{(t_e, t_u) \in I_1} Y_{t_e} - Y_{t_u} \right) - \left(\frac{1}{|\mathcal{T}_e|} \sum_{t_e \in \mathcal{T}_e} Y_{t_e}(1) - Y_{t_e}(0) \right) \middle| \mathbf{E}, \mathbf{t}, \mathbf{W}, \mathbf{X}, \mathbf{P} \right) \right| \\
&= \left| \frac{1}{|I_1|} \sum_{(t_e, t_u) \in I_1} \left(\beta_1 + \beta_2 t_e - \beta_2 t_u + \beta_3^\top \mathbf{W}_{t_e} - \beta_3^\top \mathbf{W}_{t_u} + \beta_4^\top \widetilde{\mathbf{X}}_{t_e} - \beta_4^\top \widetilde{\mathbf{X}}_{t_u} + \beta_5^\top \widetilde{\mathbf{P}}_{t_e} - \beta_5^\top \widetilde{\mathbf{P}}_{t_u} \right) - \beta_1 \right| \\
&\leq \frac{1}{|I_1|} |\beta_2| \left| \sum_{(t_e, t_u) \in I_1} (t_e - t_u) \right| + \frac{1}{|I_1|} \left| \sum_s \beta_{3s} \sum_{(t_e, t_u) \in I_1} (W_{t_e s} - W_{t_u s}) \right| +
\end{aligned}$$

$$\begin{aligned} & \frac{1}{|I_1|} \left| \sum_s \beta_{4s} \sum_{(t_e, t_u) \in I_1} (\tilde{X}_{t_e s} - \tilde{X}_{t_u s}) \right| + \frac{1}{|I_1|} \left| \sum_s \beta_{5s} \sum_{(t_e, t_u) \in I_1} (\tilde{P}_{t_e s} - \tilde{P}_{t_u s}) \right| \\ & < \delta |\beta_2| + \delta' (\|\beta_3\|_1 + \|\beta_4\|_1 + \|\beta_5\|_1). \end{aligned}$$

- Matching 1-2

$$\begin{aligned} & |\mathbb{E}(\hat{\tau} - \tau)| = \\ & = \left| \mathbb{E} \left(\mathbb{E} \left(\frac{1}{|I_2|} \sum_{(t_e, t_{u_1}, t_{u_2}) \in I_2} Y_{t_e} - \frac{1}{2} (Y_{t_{u_1}} + Y_{t_{u_2}}) \right) - \left(\frac{1}{|\mathcal{T}_e|} \sum_{t_e \in \mathcal{T}_e} Y_{t_e}(1) - Y_{t_e}(0) \right) \middle| \mathbf{E}, \mathbf{t}, \mathbf{W}, \mathbf{X}, \mathbf{P} \right) \right| \\ & = \sum_{(t_e, t_{u_1}, t_{u_2}) \in I_2} \left(\beta_1 + \beta_2 t_e - \frac{\beta_2}{2} (t_{u_1} + t_{u_2}) + \beta_3^\top \mathbf{W}_{t_e} - \frac{\beta_3^\top}{2} (\mathbf{W}_{t_{u_1}} + \mathbf{W}_{t_{u_2}}) + \right. \\ & \quad \left. \beta_4^\top \tilde{\mathbf{X}}_{t_e} - \frac{\beta_4^\top}{2} (\tilde{\mathbf{X}}_{t_{u_1}} + \tilde{\mathbf{X}}_{t_{u_2}}) + \beta_5^\top \tilde{\mathbf{P}}_{t_e} - \frac{\beta_5^\top}{2} (\tilde{\mathbf{P}}_{t_{u_1}} + \tilde{\mathbf{P}}_{t_{u_2}}) \right) - \beta_1 \Big| \\ & \leq \frac{1}{|I_2|} |\beta_2| \left| \sum_{(t_e, t_{u_1}, t_{u_2}) \in I_2} \left(t_e - \frac{1}{2} (t_{u_1} + t_{u_2}) \right) \right| + \frac{1}{|I_2|} \left| \sum_s \beta_{3s} \sum_{(t_e, t_{u_1}, t_{u_2}) \in I_2} W_{t_e s} - \frac{1}{2} (W_{t_{u_1} s} + W_{t_{u_2} s}) \right| + \dots \\ & < \delta |\beta_2| + \delta' (\|\beta_3\|_1 + \|\beta_4\|_1 + \|\beta_5\|_1). \end{aligned}$$

- Matching 1-1/2

$$\begin{aligned} & |\mathbb{E}(\hat{\tau} - \tau)| \\ & = \left| \mathbb{E} \left(\mathbb{E} \left(\frac{1}{|I|} \sum_{(t_e, t_u) \in I_1} (Y_{t_e} - Y_{t_u}) + \sum_{(t_e, t_{u_1}, t_{u_2}) \in I_2} \left(Y_{t_e} - \frac{1}{2} (Y_{t_{u_1}} + Y_{t_{u_2}}) \right) \right) - \right. \\ & \quad \left. - \left(\frac{1}{|I|} \sum_{j \in I} Y_{t_e}(1) - Y_{t_e}(0) \right) \middle| \mathbf{E}, \mathbf{t}, \mathbf{W}, \mathbf{X}, \mathbf{P} \right) \right| \\ & = \left| \frac{1}{|I|} \sum_{(t_e, t_u) \in I_1} \left(\beta_1 + \beta_2 t_e - \beta_2 t_u + \beta_3^\top \mathbf{W}_{t_e} - \beta_3^\top \mathbf{W}_{t_u} + \beta_4^\top \tilde{\mathbf{X}}_{t_e} - \beta_4^\top \tilde{\mathbf{X}}_{t_u} + \beta_5^\top \tilde{\mathbf{P}}_{t_e} - \beta_5^\top \tilde{\mathbf{P}}_{t_u} \right) + \right. \\ & \quad \sum_{(t_e, t_{u_1}, t_{u_2}) \in I_2} \left(\beta_1 + \beta_2 t_e - \frac{\beta_2}{2} (t_{u_1} + t_{u_2}) + \beta_3^\top \mathbf{W}_{t_e} - \frac{\beta_3^\top}{2} (\mathbf{W}_{t_{u_1}} + \mathbf{W}_{t_{u_2}}) + \right. \\ & \quad \left. \left. + \beta_4^\top \tilde{\mathbf{X}}_{t_e} - \frac{\beta_4^\top}{2} (\tilde{\mathbf{X}}_{t_{u_1}} + \tilde{\mathbf{X}}_{t_{u_2}}) + \beta_5^\top \tilde{\mathbf{P}}_{t_e} - \frac{\beta_5^\top}{2} (\tilde{\mathbf{P}}_{t_{u_1}} + \tilde{\mathbf{P}}_{t_{u_2}}) \right) - \beta_1 \Big| \\ & \leq \frac{1}{|I|} |\beta_2| \left| \sum_{(t_e, t_u) \in I_1} (t_e - t_u) + \sum_{(t_e, t_{u_1}, t_{u_2}) \in I_2} \left(t_e - \frac{1}{2} (t_{u_1} + t_{u_2}) \right) \right| + \end{aligned}$$

$$\begin{aligned} & \frac{1}{|I|} \left| \sum_s \beta_{3s} \left(\sum_{(t_e, t_u) \in I_1} W_{t_e s} - W_{t_u s} + \sum_{(t_e, t_{u_1}, t_{u_2}) \in I_2} W_{t_e} - \frac{1}{2}(W_{t_{u_1} s} + W_{t_{u_2} s}) \right) \right| + \dots \\ & < \delta |\beta_2| + \delta' (\|\beta_3\|_1 + \|\beta_4\|_1 + \|\beta_5\|_1). \end{aligned}$$

□

Proof of Theorem 2.

We study the bias of the matching estimator based on Matching 1-1. The bias bounds of estimators based on Matching 1-2 and Matching 1-1/2 are derived similarly, and therefore are omitted here.

We bound the quantity $h_s(W_{t_e s}) - h_s(W_{t_u s})$ below, and similar bounds can be acquired for the functions h_s that correspond to $\tilde{\mathbf{X}}$ and $\tilde{\mathbf{P}}$.

$$\begin{aligned} & \left| \sum_{(t_e, t_u) \in I_1} h_s(W_{t_e s}) - h_s(W_{t_u s}) \right| = \\ & = \left| \sum_{r=1}^{(b-a)/\ell} \sum_{(t_e, t_u) \in I_1} \left(\sum_{k=1}^{K-1} \gamma_k W_{t_e s r}^{k\dagger} + R_{t_e, r, K} \right) - \left(\sum_{k=1}^{K-1} \gamma_k W_{t_u s r}^{k\dagger} + R_{t_u, r, K} \right) \right| \\ & \leq \sum_{r=1}^{(b-a)/\ell} \left(\sum_{k=1}^{K-1} |\gamma_k| \left| \sum_{(t_e, t_u) \in I_1} W_{t_e s r}^{k\dagger} - W_{t_u s r}^{k\dagger} \right| + 2|I_1| \frac{|h_s^{(K)}(\xi_{sr})|}{K!} (\ell/2)^K \right) \\ & \leq |I_1| \left(\sum_{r=1}^{(b-a)/\ell} \left(\sum_{k=1}^{K-1} \delta' |\gamma_k| + 2 \frac{|h_s^{(K)}(\xi_{sr})|}{K!} (\ell/2)^K \right) \right) \\ & \leq |I_1| \left((b_s - a_s)/\ell \left(\sum_{k=1}^{K-1} \delta' c/k! + 2 \frac{c}{K!} (\ell/2)^K \right) \right), \end{aligned}$$

where $\gamma_k := \frac{h_s^{(k)}(\xi_{sr})}{k!} \leq \frac{c}{k!}$ is the coefficient of the Taylor expansion of order k around ξ_{sr} , and $R_{t,r,k}$ is the residual of Taylor expansion such that $|R_{t,r,k}| \leq \frac{|h_s^{(K)}(\xi_{sr})|}{K!} (\ell/2)^K \leq \frac{c}{K!} (\ell/2)^K$. Therefore, we can show

$$\begin{aligned} & |\mathbf{E}(\hat{\tau} - \tau)| = \\ & = \left| \mathbf{E} \left(\mathbf{E} \left(\frac{1}{|I_1|} \sum_{(t_e, t_u) \in I_1} Y_{t_e} - Y_{t_u} \right) - \left(\frac{1}{|\mathcal{T}_e|} \sum_{t_e \in \mathcal{T}_e} Y_{t_e}(1) - Y_{t_e}(0) \right) \middle| \mathbf{E}, \mathbf{t}, \mathbf{W}, \mathbf{X}, \mathbf{P} \right) \right| \\ & = \left| \frac{1}{|I_1|} \sum_{(t_e, t_u) \in I_1} \left(\theta + \beta + (h_0(t_e) - h_0(t_u)) + \sum_{s=1}^{p_W} (h_s(W_{t_e s}) - h_s(W_{t_u s})) + \right. \right. \\ & \quad \left. \left. \sum_{s=1}^{p_X} (h_{p_W+s}(\tilde{X}_{t_e s}) - h_{p_W+s}(\tilde{X}_{t_u s})) + \sum_{s=1}^{p_P} (h_{p_W+p_X+s}(\tilde{P}_{t_e s}) - h_{p_W+p_X+s}(\tilde{P}_{t_u s})) \right) - \beta_1 \right| \end{aligned}$$

$$\begin{aligned}
&\leq \left((T-1)/\ell \left(\sum_{k=1}^{K-1} \delta c/k! + 2 \frac{c}{K!} (\ell/2)^K \right) \right) + \sum_s \left((b_s - a_s)/\ell \left(\sum_{k=1}^{K-1} \delta' c/k! + 2 \frac{c}{K!} (\ell/2)^K \right) \right) \\
&\leq C_T \delta + C_{WXP} \delta' + C_{TWP} \ell^{K-1}.
\end{aligned}$$

□

A.3 The bounding constants in Theorem 2

Although we match discrete time periods, we think of the smooth temporal trend as a continuous function for time over the whole interval $[0, T]$. For the time-varying covariates, their supports are indexed by their corresponding additive function, which is $[b_s - a_s]$ for every $h_s, s = 1, 2, \dots, p_W + p_X + p_P$. The constants C_T, C_{WXP} and C_{TWP} are equal to

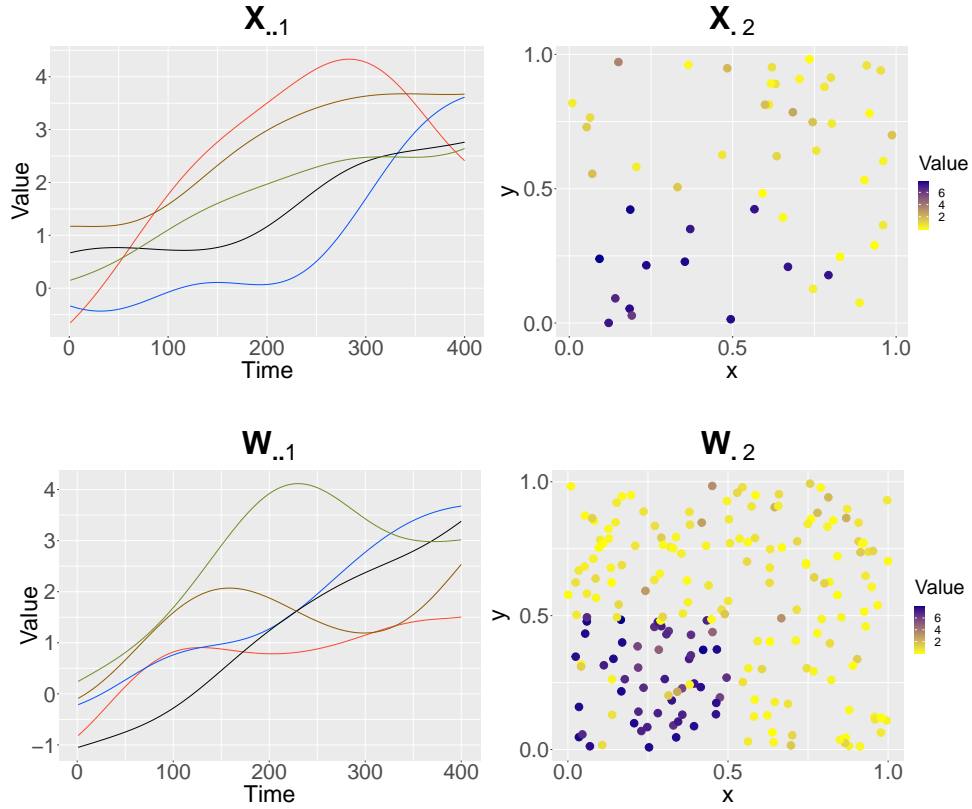
$$\begin{aligned}
C_T &= \sum_{k=1}^{K-1} \frac{(T-1)c}{\ell k!}, \\
C_{WXP} &= \sum_{s=1}^{p_W+p_X+p_P} \sum_{k=1}^{K-1} \frac{(b_s - a_s)c}{\ell k!}, \\
C_{TWP} &= \left(\frac{1}{2}\right)^{K-1} \left(\frac{(T-1)c}{K!} + \sum_{s=1}^{p_W+p_X+p_P} \frac{(b_s - a_s)c}{K!} \right).
\end{aligned}$$

B A Detailed Simulation Description of Section 4

We consider settings with $N = 50$ interventional units and $M = 200$ outcome units at randomly generated locations on the $[0, 1] \times [0, 1]$ square in an (x, y) coordinate system. The locations of the interventional and outcome units are generated in the following manner. For the interventional units, the x coordinates of units 1 to 10, 31 to 40 are generated independently from $\text{Uniform}(0, 0.5)$, while the x coordinates for the remaining units are generated from a $\text{Uniform}(0.5, 1)$ distribution. The y coordinates for the interventional units are generated from a $\text{Uniform}(0, 0.5)$ distribution for units 1 to 10 and 21 to 30, and from $\text{Uniform}(0.5, 1)$ distribution for the remaining units. Similarly, for the outcome units, the x coordinates for units 1 to 68 and 113 to 156 are drawn from a $\text{Uniform}(0, 0.5)$ distribution, while the x coordinates for the remaining outcome units are drawn from a $\text{Uniform}(0.5, 1)$ distribution. The y coordinates are drawn from $\text{Uniform}(0, 0.5)$ for units 1 to 50 and 101 to 150, and from $\text{Uniform}(0.5, 1)$ for the remaining units.

We generate covariates, treatments, graphs, exposures, and outcomes over $T = 400$ time periods. We create scenarios with or without different confounders between exposure and outcome.

Figure S.2: Visualization of interventional and outcome unit covariates over time and units. The top left panel shows 5 realizations from the Gaussian process for $\mathbf{X}_{..1}$. The top right figure is the $[0, 1] \times [0, 1]$ location square with one realization for the location of the interventional points and the corresponding location-varying covariate $\mathbf{X}_{.2}$. The bottom right figure is 5 realizations from the Gaussian process for $\mathbf{W}_{..1}$ and the bottom right figure is a realization from the locations of the outcome units colored by the values of one realization of $\mathbf{W}_{.2}$.



B.1 The covariates

We consider six covariates for the interventional units, six covariates for the outcome units, and one network covariate. Specifically, the covariates we consider are as follows:

- Interventional unit covariates:

1. Confounders equal to smooth functions of time: $\mathbf{X}_{.i1} = (X_{1i1}, X_{2i1}, \dots, X_{Ti1})$ is generated from a Gaussian process with Gaussian correlation kernel and mean representing a smooth temporal trend as $\mathbf{X}_{.i1} \sim \mathcal{N}(f(t), \Sigma)$, where $\Sigma_{t_1, t_2} = \frac{\exp(-(t_1 - t_2)^2)}{2 \cdot 100^2}$ and $f(t) = \frac{3}{400}t$. Realizations of this covariate $\mathbf{X}_{.i1}$ for five interventional units is shown on the top-left panel of Figure S.2.
2. Location-varying confounders: We consider a covariate that varies across unit but not across time, $X_{ti2} = X_{i2}$. For units i inside the $[0, 0.5] \times [0, 0.5]$ rectangle, we generate $X_{i2}/8 \sim \text{Beta}(9, 1)$;

Figure S.3: Visualization of non-smoothing interventional covariates $\mathbf{X}_{..3}$ and $\mathbf{X}_{..6}$. The left panel shows one realization for $\mathbf{X}_{..3}$ and the right panel shows one realization for $\mathbf{X}_{..6}$.


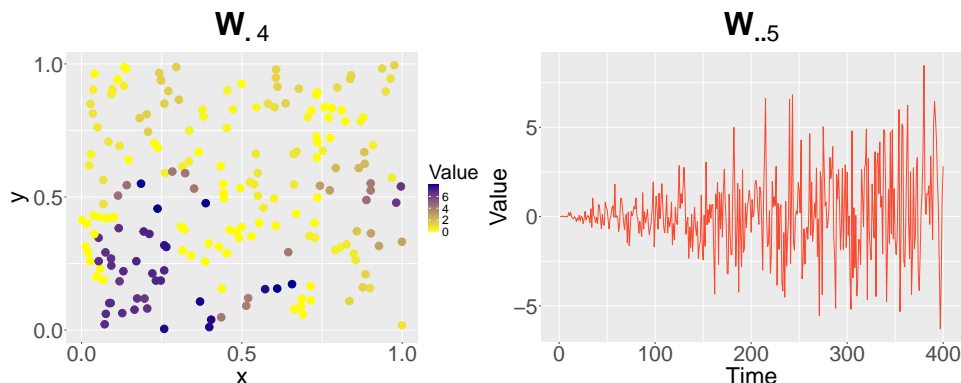


Figure S.4: Visualization of the location-varying outcome covariate $\mathbf{W}_{.4}$ and the non-smooth covariate $\mathbf{W}_{..5}$. The left panel is the $[0, 1] \times [0, 1]$ location square with one realization for the location of the outcome units and the corresponding location-varying covariate $\mathbf{W}_{.4}$. The right panel shows one realization for $\mathbf{W}_{..5}$ for the outcome units which is defined based on the interventional units' covariate $\mathbf{X}_{..3}$ in Figure S.3 and their geographical distance.



for the rest i 's, $X_{i2}/8 \sim \text{Beta}(1, 9)$. A visualization of this covariate for all intervention units is shown at the top-right panel of Figure S.2.

3. Time-varying confounders: $X_{ti3} \sim N(0, t/100), i = 1, \dots, N, t = 1, \dots, T$. This variable varies across time without a smooth pattern. A realization is drawn on the left of Figure S.3.

- Outcome unit covariates:

1. Confounders equal to smooth functions of time: $\mathbf{W}_{.j1}$ is generated by the same Gaussian process as $\mathbf{X}_{.i1}$, as $\mathbf{W}_{.j1} \sim N(f(t), \Sigma)$, where $\Sigma_{t_1, t_2} = \frac{\exp(-(t_1 - t_2)^2)}{2 \cdot 100^2}$ and $f(t) = \frac{3}{400}t$. Realizations of this covariate $\mathbf{W}_{.i1}$ for five outcome units is shown on the bottom-left panel of Figure S.2. It is evident that covariate $\mathbf{X}_{.i1}$ for the interventional units and covariate $\mathbf{W}_{.j1}$ for the outcome units represent smooth temporal trends that are similar, but different. Therefore, if $\mathbf{X}_{.i1}$ is a predictor

of the interventional units' treatment across time, and $\mathbf{W}_{.j1}$ is a predictor of the outcome units' outcome across time, then the common smooth temporal trend in $\mathbf{X}_{.i1}$ and $\mathbf{W}_{.j1}$ confounds the exposure-outcome relationship of interest.

2. Location-varying confounders: for units j inside the $[0, 0.5] \times [0, 0.5]$ square, we let their covariates $W_{tj2} = W_{j2}$, where $W_{j2}/8 \sim \text{Beta}(9, 1)$ for all t ; for the rest j 's, $W_{tj2} = W_{j2}$, where $W_{j2}/8 \sim \text{Beta}(1, 9)$ for all t . A visualization of this covariate for all outcome units is shown at the bottom-right panel of Figure S.2. It is evident that the location-varying covariates $\mathbf{X}_{.2}$ and $\mathbf{W}_{.2}$ share spatial trends. Therefore, if the former is a predictor of the interventional units' treatment, and the latter of the outcome units' outcome, then there exists location-varying confounding of the exposure-outcome relationship.
3. Time-varying confounders: $W_{tj3}/2 \sim \text{Beta}(t/100, 2), j = 1, 2, \dots, M, t = 1, 2, \dots, T$. This variable has a *non-smooth* temporal trend.

- Network covariates: We generate $p_P = 1$ network covariate, and the network covariate array \mathbf{P}_t is an $N \times M$ matrix. The entries of the matrix are generated independently as $P_{tij} \sim \text{Beta}(t/50, 10)$, $i = 1, 2, \dots, N, j = 1, 2, \dots, M, t = 0, 1, \dots, T$. Therefore, the network covariates have a *non-smooth* temporal trend.
- Bipartite covariates: We consider additional covariates for interventional units based on the weighted average of neighboring outcome unit covariates, and vice versa. We define the $N \times M$ matrix \mathbf{R} , with entries $r_{ij} = 1$ if unit i and j are within distance 0.1, and $r_{ij} = 0$ otherwise. We set X_{i4} and $\mathbf{X}_{.i5}$ to denote the i th interventional unit-related location and time-varying covariates such that:

4. $X_{i4} = \frac{\sum_j r_{ij} W_{j2}}{\sum_j r_{ij}}, i = 1, 2, \dots, N$, and
5. $X_{ti5} = \frac{\sum_j r_{ij} W_{tj3}}{\sum_j r_{ij}}, i = 1, 2, \dots, N, t = 1, 2, \dots, T$.

Similarly we set W_{j4} and $\mathbf{W}_{.j5}$ to denote the j th outcome-unit-related location and time-varying covariates. Here, it is useful to introduce the matrix \mathbf{Q} , whose columns are the normalized version of the columns of \mathbf{R} . Specifically, the (i, j) th entry is equal to $q_{ij} = r_{ij} / \sum_i r_{ij}$. Then, the covariates are defined as

4. $W_{j4} = \sum_i q_{ij} X_{i2}, j = 1, 2, \dots, N$, and
5. $W_{tj5} = \sum_i q_{ij} X_{ti3}, j = 1, 2, \dots, N, t = 1, 2, \dots, T$.

The columns of the matrix \mathbf{Q} correspond to the vectors \mathbf{q} we introduce in Section 3.1 for defining the interventional and network covariate summaries. One realization of the location-varying covariate

$W_{.4}$ for all outcome units and one realization of the non-smooth covariate $W_{.5}$ are presented on the left and right panel of Figure S.4.

- For every unit, there is an associated time-varying location-invariant confounder, termed $X_{ti6} = W_{tj6} = Z_t$, following $\mathcal{N}(t/200, \log(t)/10)$ everywhere. A realization of this covariate is shown on the right of Figure S.3.

B.2 The treatment assignment, bipartite network, and outcome

We consider five scenarios regarding the confounding structure. In terms of the treatment assignment, the five scenarios consider the following generation of treatment vectors over the interventional units across time.

(a) No confounders:

$$\begin{aligned} A_{ti} &\sim \text{Uniform}(0, 1), \\ G_{tij} &\sim \text{Ber}(\rho), \quad \text{where } \rho = 0.17, \\ Y_{tj} &= E_{tj} + W_{j2} + \epsilon_{tj}. \end{aligned}$$

(b) Only time-smooth confounders exist:

$$\begin{aligned} A_{ti} \mid X_{ti1} &\sim \text{Ber}(1/(1 + \exp(X_{ti1}/1.2))) \\ G_{tij} &\sim \text{Ber}(\rho_{ij}), \quad \text{where } \rho_{ij} = 1/(2(1 + \exp \text{dist}(i, j))), \\ Y_{tj} &= E_{tj} + W_{tj1} + \epsilon_{tj} \end{aligned}$$

(c) Only location-varying confounders exist:

$$\begin{aligned} A_{ti} \mid X_{i2}, X_{i4} &\sim \text{Ber}(1/(1 + 0.3 \exp(X_{i2} - X_{i4}/40))) \\ G_{tij} &\sim \text{Ber}(\rho_{ij}), \quad \text{where } \rho_{ij} = 1/(1.7(1 + \exp \text{dist}(i, j))) \\ Y_{tj} &= E_{tj} + 4W_{j2} + 4W_{j4} + \epsilon_{tj} \end{aligned}$$

(d) Only time-varying confounders exist:

$$\begin{aligned} A_{ti} \mid X_{ti3}, X_{ti5}, X_{ti6}, \mathbf{P}_t &\sim \text{Ber}\left(\frac{1}{1 + \exp(X_{ti3}/2 + X_{ti5}/2 + X_{ti6} + \sum_j r_{ij} P_{tij} / (10 \sum_j r_{ij}))}\right) \\ G_{tij} \mid t &\sim \text{Ber}(\rho_{tij}), \quad \text{where } \rho_{tij} = 1/(1 + 0.1 \exp(\sin(\pi t/1000) + \exp \text{dist}(i, j))) \end{aligned}$$

$$Y_{tj} = E_{tj} + W_{tj3} + W_{tj5} + \sum_i q_{ij} P_{tij} + W_{tj6} + \epsilon_{tj}$$

(e) All confounders exist:

$$A_{ti} | \mathbf{X}_t, \mathbf{P}_t \sim \text{Ber}(\eta_{tij}), \quad \text{where}$$

$$\eta_{tij} = \frac{1}{1 + 0.45 \exp(X_{ti1}/20 + X_{i2} + X_{ti3}/100 + X_{i4} + X_{ti5}/20 + X_{ti6}/1.5 + \sum_j r_{ij} P_{tij} / (10 \sum_j r_{ij}))}$$

$$G_{tij} | t \sim \text{Ber}(\rho_{tij}), \quad \text{where } \rho_{tij} = 1 / (1 + 0.1 \exp(\sin(\pi t / 1000) + \exp(\text{dist}(i, j)))$$

$$Y_{tj} = E_{tj} + W_{tj1} + 2 * W_{j2} + W_{tj3} + 0.1 * W_{tj3}^2 + 2 * W_{j4} + W_{tj5} + 1 / (1 + \exp(W_{tj5})) + \sin\left(\sum_i q_{ij} P_{tij}\right) + 2W_{tj6} + \epsilon_{tj}$$

C Description for naïve approaches

Implementing three naïve approaches requires a full dataset including exposure status E_{tj} and outcomes Y_{tj} among all units m_1, m_2, \dots, m_M and all time points $t = 1, 2, \dots, T$. While Naïve-all uses the full dataset, Naïve- t uses only the data from the first outcome unit across time, and Naïve- j uses the data across all outcome units but only for the first time point.

C.1 Estimator

For Naïve- t , we split the temporal data for the first outcome unit into two groups: the exposed time periods, $t \in \mathcal{T}_e$ if $E_{jt} = 1$, and the unexposed time periods, $t \in \mathcal{T}_u$. Similarly, for Naïve- j , we split the data for the outcome units in the exposed outcome units, $j \in \mathcal{J}_e$ if $E_{jt} = 1$, and the unexposed outcome units, $j \in \mathcal{J}_u$ otherwise. For Naïve-all, the tuple $(t, j) \in \mathcal{S}_e$ if $E_{jt} = 1$ and $(t, j) \in \mathcal{S}_u$ otherwise. We use the estimators $\hat{\tau}^{\text{Naïve-}t}$, $\hat{\tau}^{\text{Naïve-}j}$ and $\hat{\tau}^{\text{Naïve-all}}$ as follows:

$$\begin{aligned} \hat{\tau}^{\text{Naïve-}t} &= \frac{1}{|\mathcal{T}_e|} \sum_{t \in \mathcal{T}_e} Y_{tj} - \frac{1}{|\mathcal{T}_u|} \sum_{t \in \mathcal{T}_u} Y_{tj}, \\ \hat{\tau}^{\text{Naïve-}j} &= \frac{1}{|\mathcal{J}_e|} \sum_{j \in \mathcal{J}_e} Y_{tj} - \frac{1}{|\mathcal{J}_u|} \sum_{j \in \mathcal{J}_u} Y_{tj}, \\ \hat{\tau}^{\text{Naïve-all}} &= \frac{1}{|\mathcal{S}_e|} \sum_{(t,j) \in \mathcal{S}_e} Y_{tj} - \frac{1}{|\mathcal{S}_u|} \sum_{(t,j) \in \mathcal{S}_u} Y_{tj}. \end{aligned}$$

C.2 Inference

We adopted the classic Wald-type confidence interval construction for two independent samples. Denote s_e as the standard deviation of outcomes with index from $\mathcal{T}_e, \mathcal{J}_e$ or \mathcal{S}_e , and s_u as the standard

deviation of outcomes with index from \mathcal{T}_u , \mathcal{J}_u or \mathcal{S}_u for each naïve approach. The confidence intervals for Naïve- t , Naïve- j , and Naïve-all are

$$\begin{aligned}\hat{\tau}^{\text{Naïve-}t} &\pm Z_\alpha \sqrt{\frac{(|\mathcal{T}_e| - 1)s_e^2 + (|\mathcal{T}_u| - 1)s_u^2}{|\mathcal{T}_e| + |\mathcal{T}_u| - 2}} \sqrt{\frac{1}{|\mathcal{T}_e|} + \frac{1}{|\mathcal{T}_u|}}, \\ \hat{\tau}^{\text{Naïve-}j} &\pm Z_\alpha \sqrt{\frac{(|\mathcal{J}_e| - 1)s_e^2 + (|\mathcal{J}_u| - 1)s_u^2}{|\mathcal{J}_e| + |\mathcal{J}_u| - 2}} \sqrt{\frac{1}{|\mathcal{J}_e|} + \frac{1}{|\mathcal{J}_u|}},\end{aligned}$$

and

$$\hat{\tau}^{\text{Naïve-all}} \pm Z_\alpha \sqrt{\frac{(|\mathcal{S}_e| - 1)s_e^2 + (|\mathcal{S}_u| - 1)s_u^2}{|\mathcal{S}_e| + |\mathcal{S}_u| - 2}} \sqrt{\frac{1}{|\mathcal{S}_e|} + \frac{1}{|\mathcal{S}_u|}},$$

respectively.

D Additional Simulation Results

D.1 Unadjusted matching

We state the methodology for unadjusted matching with time series data given exposure and outcome values at each time period. For the three matching criteria, 1-1, 1-2, and 1-1/2, we balance time but not the time-varying confounders. Specifically, we remove constraints (A.4)-(A.5) for 1-1 objective (A), (B.4)-(B.5) for 1-2 objective (B), and (C.4)-(C.5) for 1-1/2 objective (C).

We evaluate the performance of the unadjusted estimator in the simulations of Section 4. In Table S.5 we show the bias, mean squared error, coverage, and proportion of matched exposed time periods when applying the three unadjusted matching algorithms. Unadjusted matching estimators are unbiased with nominal coverage in the absence of *non-smooth* time-varying confounders, even in the presence of confounding temporal trends, in cases (a), (b) and (c). However, in the scenarios with non-smooth time-varying confounding, unadjusted estimators will be biased. The proportion of exposed time points that are matched using unadjusted approaches is only slightly higher than the proportion for adjusted approaches in some scenarios.

We also demonstrated the performance of unadjusted matching under multiple hypothesis simulations in Section 4.2.2. The results are shown in Table S.6. The unadjusted methods give reliable inferences for testing the global null in the absence of confounders or in the presence of only time-smooth confounders. In the presence of non-smooth temporal confounding, the unadjusted estimators' bias will lead to identifying statistically significant causal effects too often.

Table S.5: Simulation results of single-unit estimations. Bias, mean squared error (MSE), coverage of 95% intervals (%), and proportion of exposed time points being matched (%). We show simulation results for 3 unadjusted matching methods. ‘U’ stands for ‘unadjusted’.

		Dense				Medium				Sparse			
	Method	Bias	MSE	Cover	Prop	Bias	MSE	Cover	Prop	Bias	MSE	Cover	Prop
(a) ¹	U1-1	0.00	0.012	96.0	97.8	0.00	0.02	94.8	100	0.00	0.039	95.2	100
	U1-1/2	0.00	0.011	94.4	97.8	0.00	0.017	95.0	100	-0.01	0.035	95.0	100
	U1-2	-0.01	0.014	94.2	62.3	-0.01	0.016	95.4	90.8	-0.01	0.029	94.0	98.7
(b)	U1-1	0.00	0.019	95.4	62.5	0.00	0.024	93.6	85.1	0.01	0.036	95.6	97.8
	U1-1/2	0.00	0.020	94.6	62.5	0.00	0.022	94.6	85.1	0.00	0.033	94.2	97.8
	U1-2	0.00	0.023	95.2	41.0	0.00	0.025	93.6	59.9	0.00	0.034	95.4	80.6
(c)	U1-1	0.01	0.019	94.6	99.7	0.01	0.041	93.6	100.0	0.00	0.086	94.4	100.0
	U1-1/2	0.00	0.017	95.2	99.7	0.01	0.038	93.6	100.0	-0.01	0.079	94.4	100.0
	U1-2	0.00	0.016	94.4	85.5	0.00	0.030	92.2	97.4	0.00	0.064	93.6	99.2
(d)	U1-1	-0.41	0.234	54.4	69.3	-0.37	0.19	60.8	86.1	-0.35	0.224	80.8	99.8
	U1-1/2	-0.42	0.232	50.0	69.3	-0.37	0.19	59.8	86.1	-0.36	0.225	75.2	99.8
	U1-2	-0.45	0.287	52.8	46.7	-0.41	0.235	59.2	63.0	-0.35	0.209	77.6	91.7
(e)	U1-1	-0.65	0.486	26.4	76.0	-0.59	0.424	33.0	94.5	-0.62	0.491	53.4	99.7
	U1-1/2	-0.63	0.458	21.4	76.0	-0.59	0.414	36.8	94.5	-0.60	0.466	49.0	99.7
	U1-2	-0.67	0.538	29.4	51.0	-0.64	0.495	32.4	73.7	-0.62	0.487	46.2	91.3

¹The scenarios shown in Table S.5 correspond to (a) No confounders, (b) Time-smooth confounders, (c) Location-varying confounders, (d) Time-varying confounders, and (e) All confounders.

Table S.6: Simulation results for using 3 unadjusted matching methods to test the global null hypothesis for scenarios (b), (d) and (e), and under a medium exposure level. ‘U’ stands for ‘unadjusted’.

		Average estimator mean	Average rate of p-value ≤ 0.05	Rate of min(p-value) ≤ 0.05	Rate of FDR min (adj.p-value) ≤ 0.05
(b)	U1-1	-0.01	0.055	1.000	0.092
	U1-1/2	-0.01	0.054	1.000	0.092
	U1-2	0.00	0.055	1.000	0.104
(d)	U1-1	-0.39	0.426	1.000	0.998
	U1-1/2	-0.39	0.449	1.000	1.000
	U1-2	-0.43	0.438	1.000	1.000
(e)	U1-1	-0.63	0.728	1.000	1.000
	U1-1/2	-0.63	0.757	1.000	1.000
	U1-2	-0.68	0.743	1.000	1.000

D.2 Analysis of tuning parameters

Following the simulation study in Section 4.1.2, we consider two additional choices of the algorithmic parameters $(\delta, \delta', \epsilon)$ per simulation: $(0, 0.05, 2)$ and $(2, 0.1, 6)$. In Table S.7, we compare the performance of tuning either (δ, ϵ) that correspond to tuning parameters of time, or δ' that is a tuning parameter for the measured covariates in comparison to the default setting $(\delta, \delta', \epsilon) = (2, 0.05, 6)$ (which is shown in Table 2).

Some conclusions from this comparison are discussed in Section 4.2.3. Here, we also point out

that, even though Matching 1-2 performed best under sparse scenarios in Table 2, when (δ, ϵ) are small, Matching 1-2 is no longer best. That is because the tight constraints exclude many possible 1-2 pairs, and the proportion of exposed time points that are matched is almost half of that of 1-1 or 1-1/2 matching strategies, resulting in large variability for the matching 1-2 estimator.

D.3 Modification of matching methods: Covariate balancing within matches

In Section 4.1.2 we observe that coverage of our 95% intervals is sometimes larger than the nominal level. Here, we illustrate that this issue of over-coverage is alleviated when covariate balance is imposed for *every* match, supporting the discussion in Section 3.3. We extend the matching algorithms in (A), (B) and (C) to include additional constraints. For simplicity, we only show the constraints for the outcome unit covariates, but similar constraints are adopted for the interventional and network time-varying covariates as well. Given a small value δ'' , for matching 1-1 in (A) we impose

$$|a_{t_e t_u}(\mathbf{W}_{t_e} - \mathbf{W}_{t_u})| \leq \mathbf{1}_{p_W} \cdot \delta'', \quad \forall t_e \in \mathcal{T}_e, \forall t_u \in \mathcal{T}_u,$$

and similarly for Matching 1-2 in (B) we impose

$$\left| a_{t_e t_{u_1} t_{u_2}} \left(\mathbf{W}_{t_e} - \frac{\mathbf{W}_{t_{c_1}} + \mathbf{W}_{t_{c_2}}}{2} \right) \right| \leq \mathbf{1}_{p_W} \cdot \delta'', \quad \forall t_e \in \mathcal{T}_e, \forall t_{u_1}, t_{u_2} \in \mathcal{T}_u.$$

For Matching 1-1/2 in (C) we impose both constraints. Here, we refer to these matching algorithms and the corresponding estimators as *extended* matching algorithms and estimators. We compare them to the algorithms and estimators introduced in Section 3 (which impose only constraints on the overall matched populations, and not within each match), which we refer to here as the *standard* algorithms and estimators.

We evaluated the performance of these extended matching estimators against the standard matching estimators under confounding scenario (d) with time-varying confounder and medium sparsity level. We set $\delta'' = 0.25$, assuming the covariates have been standardized as discussed in Section 4.1.2. We set the remaining tuning parameters equal to the values $(\delta, \delta', \epsilon) = (2, 0.05, 6)$.

In Figure S.5, we see that the three extended matching estimators are close to unbiased for the true causal effect (which is equal to 1). Any residual bias is due to residual covariate imbalance. The bias of the three matching estimators based on the extended algorithms is -0.08 for Matching 1-1, -0.10 for Matching 1-1/2, and -0.09 for Matching 1/2. Since bias can affect coverage rates, we compare coverage of 95% intervals of the extended matching estimators, with the results of the standard estimators under tuning parameters $(\delta, \delta', \epsilon) = (2, 0.1, 6)$ which show similar bias (see Table S.7, bias

Table S.7: Simulation results of single-unit estimations under different tuning parameters.

Method		Dense				Medium				Sparse			
		Bias	MSE	Cover	Prop	Bias	MSE	Cover	Prop	Bias	MSE	Cover	Prop
(a) No confounders													
1-1	(1) ¹	0.00	0.013	97.2	87.2	-0.01	0.021	93.8	98.3	0.00	0.039	95.0	99.9
	(2)	0.00	0.012	95.0	97.8	-0.01	0.020	95.4	100.0	-0.01	0.040	95.6	100.0
	(3)	-0.01	0.014	94.0	96.3	-0.01	0.019	95.6	100.0	-0.01	0.040	95.8	100.0
1-1/2	(1)	0.00	0.013	94.4	87.2	0.00	0.019	93.4	98.3	-0.01	0.035	94.8	99.9
	(2)	0.00	0.011	95.8	97.8	-0.01	0.017	96.2	100.0	-0.01	0.031	94.6	100.0
	(3)	-0.01	0.014	94.6	96.3	0.00	0.018	95.6	100.0	-0.01	0.034	94.2	100.0
1-2	(1)	0.00	0.021	94.2	40.9	0.00	0.024	94.4	64.0	-0.01	0.037	93.2	82.0
	(2)	0.00	0.013	97.2	62.3	-0.01	0.016	95.4	90.8	-0.01	0.028	95.6	98.7
	(3)	-0.01	0.017	95.6	61.4	-0.01	0.016	96.0	90.8	-0.01	0.029	95.0	98.7
(b) Time-smooth confounders													
1-1	(1)	0.01	0.021	94.4	58.4	0.01	0.025	96.0	78.7	0.01	0.039	95.8	93.2
	(2)	0.00	0.018	96.0	62.5	0.00	0.024	92.8	85.1	0.01	0.037	94.8	97.8
	(3)	-0.01	0.021	94.4	63.7	0.00	0.024	93.2	85.1	0.00	0.037	94.6	97.8
1-1/2	(1)	0.00	0.020	94.4	58.4	0.01	0.025	95.8	78.7	0.01	0.036	95.8	93.2
	(2)	0.00	0.018	94.8	62.5	0.00	0.022	94.2	85.1	0.01	0.035	93.8	97.8
	(3)	-0.01	0.020	95.8	63.7	-0.01	0.023	94.4	85.1	0.01	0.034	94.4	97.8
1-2	(1)	0.01	0.034	93.6	28.5	0.01	0.036	93.8	43.0	0.02	0.048	96.2	59.6
	(2)	0.00	0.023	94.0	41.0	0.00	0.024	95.0	59.9	0.01	0.034	94.8	80.6
	(3)	-0.01	0.025	94.8	41.7	0.00	0.023	95.2	59.9	0.00	0.034	94.2	80.6
(c) Location-varying confounders													
1-1	(1)	0.00	0.023	93.0	96.3	0.00	0.040	94.2	99.6	0.01	0.082	93.0	99.8
	(2)	0.01	0.020	94.4	99.7	0.00	0.037	95.0	100.0	-0.01	0.080	94.8	100.0
	(3)	0.00	0.021	95.8	97.6	0.00	0.037	95.6	100.0	-0.01	0.081	95.0	100.0
1-1/2	(1)	0.00	0.021	91.8	96.3	0.01	0.037	93.8	99.6	0.00	0.068	94.2	99.9
	(2)	0.01	0.017	95.2	99.7	0.01	0.033	94.2	100.0	0.00	0.073	92.6	100.0
	(3)	0.00	0.020	93.8	97.6	0.01	0.033	94.8	100.0	0.00	0.071	93.6	100.0
1-2	(1)	0.01	0.025	94.6	59.7	0.01	0.037	94.4	79.1	0.00	0.063	94.6	88.6
	(2)	0.00	0.016	94.4	85.5	0.00	0.032	94.0	97.4	0.00	0.071	92.0	99.1
	(3)	-0.01	0.020	93.4	82.5	0.00	0.030	95.2	97.4	0.00	0.068	93.2	99.1
(d) Time-varying confounders													
1-1	(1)	-0.04	0.023	99.2	59.2	-0.04	0.027	98.8	74.8	-0.04	0.066	96.6	92.5
	(2)	-0.05	0.023	98.0	67.6	-0.05	0.024	98.4	85.0	-0.03	0.062	95.2	99.7
	(3)	-0.09	0.038	96.4	69.5	-0.09	0.037	95.2	85.8	-0.06	0.070	93.4	99.8
1-1/2	(1)	-0.04	0.024	99.4	59.2	-0.05	0.028	98.2	74.8	-0.04	0.067	97.0	92.7
	(2)	-0.05	0.022	98.4	67.6	-0.05	0.024	98.2	85.0	-0.04	0.054	96.6	99.7
	(3)	-0.10	0.040	96.0	69.5	-0.09	0.036	96.2	85.8	-0.07	0.062	94.2	99.8
1-2	(1)	-0.03	0.037	99.0	27.3	-0.04	0.040	99.0	38.3	-0.02	0.079	97.4	57.4
	(2)	-0.04	0.026	98.8	43.5	-0.04	0.025	97.6	59.4	-0.04	0.053	96.0	87.5
	(3)	-0.10	0.040	97.8	45.4	-0.08	0.039	97.6	61.0	-0.08	0.065	94.0	89.2
(e) All confounders													
1-1	(1)	-0.06	0.027	98.8	66.2	-0.05	0.031	99.0	85.2	-0.04	0.049	98.8	93.9
	(2)	-0.06	0.027	98.4	75.0	-0.06	0.033	98.0	94.1	-0.03	0.045	98.8	99.6
	(3)	-0.13	0.047	92.2	75.9	-0.12	0.047	94.2	94.5	-0.09	0.056	97.2	99.7
1-1/2	(1)	-0.06	0.026	98.8	66.2	-0.05	0.031	99.4	85.2	-0.04	0.046	99.4	94.1
	(2)	-0.07	0.028	98.0	75.0	-0.06	0.032	97.2	94.1	-0.04	0.041	98.4	99.6
	(3)	-0.13	0.049	91.6	75.9	-0.12	0.048	94.4	94.5	-0.10	0.054	96.2	99.7
1-2	(1)	-0.06	0.041	99.6	31.4	-0.05	0.043	99.2	47.6	-0.04	0.053	99.4	60.2
	(2)	-0.07	0.028	99.2	48.9	-0.06	0.032	98.6	71.0	-0.04	0.042	97.8	88.3
	(3)	-0.13	0.051	96.6	50.1	-0.12	0.050	96.6	72.5	-0.10	0.057	96.2	89.6

¹Tuning parameters: (1) strict: $(\delta, \delta', \epsilon) = (0, 0.05, 2)$, (2) medium: $(\delta, \delta', \epsilon) = (2, 0.05, 6)$ (the results in Table 2 correspond to this choice of tuning parameters), (3) loose: $(\delta, \delta', \epsilon) = (2, 0.1, 6)$.

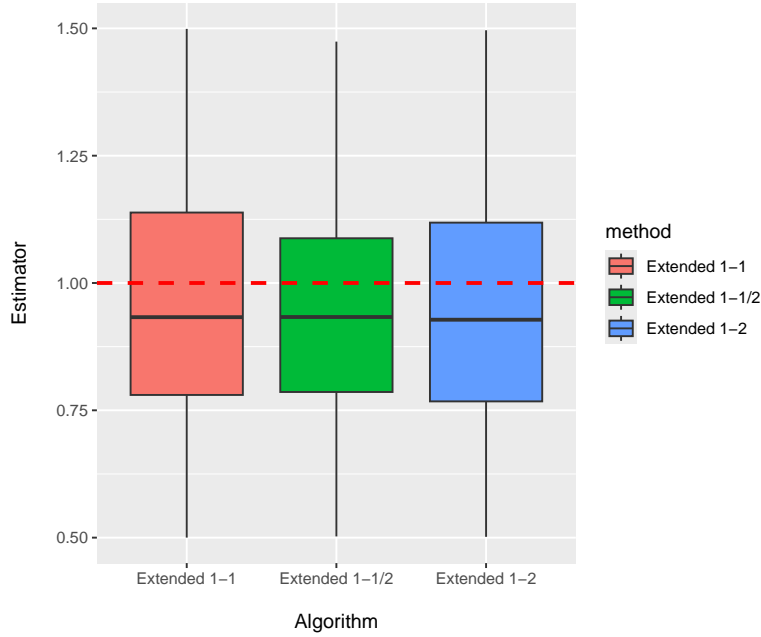


Figure S.5: Boxplots of estimated causal effects over 500 data sets based on the extended estimators.

−0.09 for Matching 1-1, −0.09 for Matching 1-1/2, and −0.08 for Matching 1-2). The coverage rate for the extended 1-1, 1-1/2, and 1/2 estimators is 95.4%, 93.0%, and 92.8%, respectively, somewhat lower than the corresponding coverage rates of the standard 1-1, 1-1/2, and 1/2 matching estimators (included in Table S.7) which are equal to 95.2%, 96.2% and 97.6%, respectively. We suspect that the extended estimators have coverage rate below 95% due to the minor amount of bias, and we expect that they will have closer to nominal coverage compared to the standard matching estimators in unbiased scenarios.

However, even though inference might be closer to the nominal level (rather than being conservative) when the extended estimators are employed, they might suffer due to small number of matches depending on the value of δ'' . Even with $\delta'' = 0.25$, the proportion of matched exposed time periods decreases significantly (compared to the algorithms that do not impose constraints within each match) to 19.7%, 30.5%, and 19.9% for Extended 1-1, Extended 1-1/2, and Extended 1-2, respectively (compared to 85.8%, 85.8% and 61% for standard 1-1, 1-1/2, and 1/2 matching, respectively). We can conclude that the extending matching that incorporates the within-match adjustment has the benefit of coverage that is closer to the nominal level, but suffers from a lower proportion of matched exposed time periods, which might lead to higher overall uncertainty.




Figure S.6: Boxplots of causal effect estimates for the naïve methods, the matching estimators without adjustment for measured covariates (discussed in Supplement D.1), and the matching estimators with balance constraints for measured covariates. This scenario corresponds to the case with exposure-outcome confounding due to network-outcome dependence on location-varying information.

D.4 Network-outcome confounding can lead to exposure-outcome confounding

We illustrate that confounding of the exposure-outcome relationship can arise due to confounding of the network-outcome relationship, even in the absence of confounding between the interventional units' treatment and the outcome. To do so, we revisit the no-confounding scenario (a) shown Table 1 and modify the generation of the network to impose network-outcome confounding. Specifically, we generate the entries of \mathbf{G} independently as $G_{tij} \sim \text{Ber}(\rho_{ij})$, for $\rho_{ij} = 1/(1.7(1 + \exp \text{dist}(i, j)))$. Under this modification, the network \mathbf{G} and the outcomes are both associated with location-varying information. We consider the medium sparsity level.

Estimated effects are shown in Figure S.6. We see that Naïve-all and Naïve-j are both biased, illustrating that the network-outcome dependence on location-varying information has introduced confounding of the exposure-outcome relationship. We also see that the matching estimators are unbiased, whether they are adjusting for measured covariates (denoted as 1-1, 1-1/2, and 1-2) or not (denoted as U1-1, U1-1/2, and U1-2).

D.5 Simulations under heterogeneous treatment effects

We performed a simulation where exposure effects are no longer homogeneous over time for an individual outcome unit. We define three new quantities, termed $\tilde{\tau}_j^{1-1}(1, 0)$, $\tilde{\tau}_j^{1-1/2}(1, 0)$ and $\tilde{\tau}_j^{1-2}(1, 0)$, as the average difference in the potential outcomes between exposures 0 and 1 over the exposed time periods *that are matched* according to 1-1, 1-1/2, and 1-2, respectively. For example, $\tilde{\tau}_j^{1-1}(1, 0)$ can be written as

$$\tilde{\tau}_j^{1-1}(1, 0) = \frac{1}{\sum_{t: \text{matched in 1-1}} I(E_{tj} = 1)} \sum_{t: \text{matched in 1-1}} (Y_{tj}(1) - Y_{tj}(0))I(E_{tj} = 1).$$

These quantities resemble the estimand $\tilde{\tau}_j(1, 0)$ defined in Section 2, but they average over the *matched* exposed time periods only. For ease of notation, we denote them as $\tilde{\tau}^{1-1}$, $\tilde{\tau}^{1-1/2}$ and $\tilde{\tau}^{1-2}$.

The data generating models follow the scenarios (b) and (d) with medium exposure, where we alter the outcome model to specify heterogeneous treatment effects. Specifically, we generate outcomes according to

$$Y_{tj} = (1 + \epsilon'_{tj})E_{tj} + 0.005(400 - t)E_{tj} + W_{tj3} + W_{tj5} + \sum_i q_{ij}P_{tij} + W_{tj6} + \epsilon_{tj},$$

where ϵ'_{tj} are generated independently from $N(0, 1)$. We simulate 500 data sets and consider the estimation of effects for the first outcome unit.

We compare the matching approaches to Naïve- t . We exclude Naïve- j and Naïve-all from simulations under heterogeneity since the estimands they target average across units, whereas Naïve- t and the matching estimators target unit-specific effects that average across time.

The bias, MSE, Coverage and proportion of exposed units that are matched are shown in Table S.8. We evaluate the performance of the three matching estimators and Naïve- t for estimating the average effect over all time points τ , and the average effect over all exposed time points $\tilde{\tau}$. In the presence of heterogeneous effects across time, all of the estimators are biased for the effect over all time periods, τ . However, the matching estimators perform relatively well for estimating the effect over the exposed time periods, $\tilde{\tau}$, and substantially better than the Naïve- t approach that suffers from confounding bias. Particularly, Matching 1-1 and 1-1/2 which match 85% of the exposed time points are close to unbiased with appropriate coverage of confidence intervals.

We also evaluate the performance of each matching estimator against the causal effect over the set of exposed time periods that are in fact matched, in that we evaluate the 1-1 matching estimator for estimating $\tilde{\tau}_{1-1}$, the 1-1/2 matching estimator for estimating $\tilde{\tau}_{1-1/2}$ and the 1-2 matching estimator for

Table S.8: Simulation results under heterogeneous effects for a single outcome unit. Bias, mean squared error (MSE), coverage of 95% intervals (%) of the Naïve- t method and matching methods for estimating the causal effect over all time periods, over the exposed time periods, and over the matched exposed time periods. The proportion of exposed time periods that are matched on average when employing each matching method is reported once.

Estimand	Method	Bias	MSE	Coverage	Proportion
(b) Time-smooth confounders					
τ	Naïve- t	-0.62	0.45	29.2	-
	1-1	0.41	0.21	58.6	85.1
	1-1/2	0.41	0.21	56.4	85.1
	1-2	0.30	0.14	83.0	59.9
$\tilde{\tau}$	Naïve- t	-1.09	1.26	1.4	
	1-1	-0.07	0.04	97.6	
	1-1/2	-0.06	0.03	97.0	
	1-2	-0.16	0.07	94.6	
$\tilde{\tau}_{1-1}$	Naïve- t	-1.02	1.12	2.0	
	1-1	0.00	0.03	98.2	
$\tilde{\tau}_{1-2}$	Naïve- t	-1.03	1.12	2.0	
	1-2	0.01	0.03	98.8	
$\tilde{\tau}_{1-1/2}$	Naïve- t	-0.92	0.92	4.8	
	1-1/2	0.01	0.03	99.4	
(d) Time-varying confounders					
τ	Naïve- t	-0.85	0.79	4.0	-
	1-1	0.25	0.10	83.0	85.1
	1-1/2	0.25	0.10	84.4	85.1
	1-2	0.13	0.06	96.2	59.4
$\tilde{\tau}$	Naïve- t	-1.22	1.55	0.0	
	1-1	-0.12	0.04	97.0	
	1-1/2	-0.12	0.04	96.8	
	1-2	-0.24	0.09	92.8	
$\tilde{\tau}_{1-1}$	Naïve- t	-1.14	1.35	0.4	
	1-1	-0.03	0.03	99.2	
$\tilde{\tau}_{1-2}$	Naïve- t	-1.14	1.35	0.4	
	1-2	-0.04	0.02	98.8	
$\tilde{\tau}_{1-1/2}$	Naïve- t	-1.01	1.10	1.0	
	1-1/2	-0.03	0.02	99.6	

estimating $\tilde{\tau}_{1-2}$. When compared against the effect over the matched exposed population, the estimators are unbiased with appropriate coverage. This suggests that even though the effect of exposure is not constant at each time point, matching approaches can still serve as a useful tool for inferring the average effect of the exposure over the population that was in fact matched.

Table S.9: Simulation results of single-unit estimation in the presence of temporal correlation $\rho = 0.2, 0.4, 0.6$ and 0.8 in the random error term. in the outcome variable. Coverage of 95% intervals (%) of matching methods.

ρ	1-1	1-1/2	1-2
0.2	96.6	96.0	95.8
0.4	96.2	96.0	96.0
0.6	96.6	94.8	94.6
0.8	94.2	93.2	94.4

D.6 The impact of temporal correlation in the outcome variable on inference

To analyze the impact of temporal correlation in the outcome variable on coverage, we alter scenario (a) under medium exposure such that the outcome error term ϵ is equal to $\epsilon_{tj} = \rho\epsilon_{(t-1)j} + \sqrt{1 - \rho^2}e_t$, where $e_t \sim N(0, 1)$, and ρ is a tuning parameter adjusting the amount of autocorrelation of the error term. We simulate 500 data sets under different levels of autocorrelation, $\rho = 0.2, 0.4, 0.6$ and 0.8 , and record the coverage of the three matching estimators. The results are shown in Table S.9. The correlation in the outcome across time has minimal impact on the coverage of the confidence intervals, which is close to 95% across all scenarios, even when auto-correlation in the error term is very high ($\rho = 0.8$).

E Additional study information

E.1 Information on data availability and creation

In our study, the interventional units correspond to the different locations of forested areas in North America, each of which might experience a wildfire or not. Hazardous smoke produced by wildfires can travel long distances, and affect population exposure and behavior across different areas. The National Oceanic and Atmospheric Administration’s Hazard Mapping System (HMS) combines data from polar and geostationary satellites in real time, enabling experts to accurately identify and track the dispersion of smoke. The HMS data can be found at <https://www.ospo.noaa.gov/Products/land/hms.html>, which report daily smoke exposure across the United States as light, medium or high based on smoke thickness.

We focus on three outcome units: San Francisco, San Jose and East Bay in northern California, which might experience exposure to smoke from wildfires or not, based on random patterns of smoke transport and dispersion. Using HMS, we consider an area as exposed if the smoke thickness is light or higher. During most of the time periods under study (January 2021 to September 2023), all three

regions are either simultaneously exposed or unexposed. Overall, the numbers of wildfire-exposed days are 147 for San Francisco, 147 for East Bay and 137 for San Jose. The scenario in Figure 1 where two of the three regions are exposed is rare, but we include it for illustration purposes.

The outcome of interest corresponds to the total daily amount of bicycle riding time using Lyft’s Bay Wheels bikeshare program, which records daily bikeshare usage since 2017 with more than 450 bikeshare stations. This database is available for public use at <https://www.lyft.com/bikes/bay-wheels/system-data>. We can obtain from it that more than 97% of rides last less than 1 hour, meaning it is unlikely to rent a bike overnight, and rides begin and end within the same area. Therefore, it is reasonable to assume there is little spillover effect across time or outcome units.

The time-varying covariates are regional temperature, dew, humidity, wind speed and wind direction among three bike locations, which are retrieved at <https://www.ncei.noaa.gov/cdo-web/datasets>.

Due to the temperature and drought season, there are more wildfires and gusty winds in late summer than in the other seasons. As a result, the exposed time periods and the corresponding matched unexposed ones, concentrate around the late summer and fall (see Supplement E.2). Therefore, the matching algorithms use information on exposed and unexposed time periods during the summer and fall. This is in contrast to the naïve approach, which uses data on all seasons, including the winter months where both exposure and bike activity are expected to be lower.

E.2 Illustrations of matched data for San Francisco

Our data set includes information from January 2021 until September 2023. Most exposed time periods, and as a result most matches, occur during the summer and fall months. As an illustration, the exposed and unexposed time periods for a *subset* of our time window and for San Francisco are shown in Figure S.7, with the red color denoting whether the time period was used in a match according to Matching 1-1, Matching 1-1/2, or Matching 1-2. Under $\epsilon = 6$, Matching 1-1 and Matching 1-1/2 have similar but not identical matching patterns, while Matching 1-2 matches fewer exposed time periods.

E.3 Results under alternative definition of exposure

As discussed in Supplement E.1, the HMS categorizes smoke exposure as no exposure, light, medium, or high exposure. In our analysis of Section 5, we considered an area at a given time period as exposed if the HMS classification was light or higher, and unexposed otherwise. Here, we evaluate the sensitivity of our conclusions when an area is considered exposed under medium or high smoke exposure, and unexposed under no smoke exposure or light smoke. Under this definition, out of the 1003 total




Figure S.7: Exposure and matching information for Matching 1-1, Matching 1-1/2, and Matching 1-2 for San Francisco during the period from June 1, 2021, to October 15, 2021. The hollow dots are time periods that are not matched while the red dots correspond to matched time periods.

number of days, San Francisco was exposed during 39 days, the East Bay during 40 days, and San Jose during 37 days.

Table S.10 shows the causal effect estimates under this alternative specification of exposure. We find that all estimates are negative, indicating that smoke exposure leads to a reduction in bikeshare hours for all three regions. Therefore, these results agree with the ones in Table 4. The effect estimates when categorizing a time period as exposed under medium or high smoke (Table S.10) are similar or larger in magnitude compared to the effect estimates when categorizing a time period as exposed under light, medium or high smoke (Table 4). Since, here, a time period is classified as exposed under heavier smoke conditions, we believe that this comparison might be because heavier smoke affects bikeshare hours more heavily than lighter smoke exposure. However, these effect estimates are not statistically significant. We believe that this is largely because of the small number of matches (ranging from 24 to 36), that is partially explained by the small number of days with medium or high smoke exposure.

Table S.10: The effect of wildfire smoke on bikeshare hours in San Francisco, East Bay, and San Jose for Naïve- t and the matching estimators. For each region, the three columns correspond to the estimate, p-value, and number of matches. In this analysis, an area is considered exposed at a given time period if it is classified to have medium or high smoke exposure according to HMS, and unexposed if it is classified to have no smoke or light smoke exposure.

	San Francisco			East Bay			San Jose		
Naïve- t	0.752	(0.983)		0.105	(0.983)		0.118	(0.995)	
Maching 1-1	-0.537	(0.152)	35	-0.050	(0.166)	36	-0.086	(0.093)	33
Matching 1-1/2	-0.615	(0.078)	35	-0.085	(0.103)	36	-0.048	(0.170)	33
Matching 1-2	-0.437	(0.199)	26	-0.161	(0.044)	27	-0.082	(0.094)	24

A similar analysis that considers time periods to be exposed under heavy smoke exposure only would not be feasible in our data set. That is because the number of days with heavy smoke thickness was small across our time window. Specifically, in the three areas there are approximately 17 days of heavy exposure, 22 days of medium exposure, 70 days of light exposure, and more than 800 days of no exposure. Therefore, analyzing this exposure as categorical would be largely infeasible.