Journal of the Royal Statistical Society Series B: Statistical Methodology, 2024, **86**, 487–511
https://doi.org/10.1093/jrsssb/qkad132
Advance access publication 15 December 2023
Original Article

ROYAL STATISTICAL SOCIETY
DATA | EVIDENCE | DECISIONS
Journal of the Statistics Society
Series B
Statistical Methodology

Identification and estimation of causal peer effects using double negative controls for unmeasured network confounding

Naoki Egami[1] **and Eric J. Tchetgen Tchetgen**[2]

[1]Department of Political Science, Columbia University, New York, NY, USA
[2]Department of Statistics and Data Science and Department of Biostatistics, Epidemiology and Informatics, University of Pennsylvania, Philadelphia, PA, USA

Address for correspondence: Naoki Egami, Department of Political Science, Columbia University, 420 West 118th Street, New York, NY 10027, USA. Email: naoki.egami@columbia.edu

Abstract

Identification and estimation of causal peer effects are challenging in observational studies for two reasons. The first is the identification challenge due to unmeasured network confounding, for example, homophily bias and contextual confounding. The second is network dependence of observations. We establish a framework that leverages a pair of negative control outcome and exposure variables (double negative controls) to non-parametrically identify causal peer effects in the presence of unmeasured network confounding. We then propose a generalised method of moments estimator and establish its consistency and asymptotic normality under an assumption about ψ-network dependence. Finally, we provide a consistent variance estimator.

Keywords: causal inference, causal inference with networks, interference, negative control, peer effect, proximal causal learning

1 Introduction

Social and biomedical scientists have long been interested in how people's behaviours are affected by peer behaviours. For example, social scientists have studied peer effects on voting behaviours, educational outcomes, criminal behaviours, and job opportunities. Epidemiologists and researchers in public health have studied related concepts such as 'contagion' effects of infectious disease and health behaviours.

Despite its importance, identification and estimation of causal peer effects have been challenging for two reasons. The first issue is that it is often difficult to identify causal peer effects in observational studies due to unmeasured network confounding, such as *homophily bias* and *contextual confounding* (Manski, 1993; Ogburn, 2018; VanderWeele & An, 2013). Homophily bias arises when people become connected due to unobserved characteristics. Contextual confounding exists when peers share some unobserved contextual factors. Highlighting concerns about these potential biases, influential papers across disciplines have criticised prior peer effect analyses from observational studies (e.g. Angrist, 2014; Lyons, 2011). Shalizi and Thomas (2011) argue that it is nearly impossible to credibly estimate causal peer effects in observational studies using direct confounding adjustment methods (e.g. regression-based adjustment) due to pervasive concerns about unmeasured network confounding.

In addition to unmeasured network confounding, another important challenge is that one needs to account for network dependence of observations in order to obtain valid statistical inference. When such network dependence is ignored, standard errors may be underestimated, and confidence intervals may be anti-conservative (Lee & Ogburn, 2021). While recent studies have allowed for some extent of network dependence across units in observational causal inference

Received: October 3, 2021. **Revised:** September 27, 2023. **Accepted:** October 6, 2023

(Forastiere et al., 2020; Leung, 2021; Ogburn et al., 2017; Tchetgen Tchetgen, Fulcher, et al., 2021; van der Laan, 2014), they have largely relied on an assumption of no uncontrolled network confounding.

In this paper, we propose to resolve these challenges by using a pair of negative control outcome (NCO) and negative control exposure (NCE) variables [*double negative controls (DNCs)*]. An NCO (also known as a placebo outcome) is an outcome variable that is known not to be causally affected by the treatment of interest. Likewise, an NCE (also known as a placebo treatment) is a treatment variable that does not causally affect the outcome of interest (Lipsitch et al., 2010). There is a long-standing tradition in biomedical and social sciences of using negative controls to detect unmeasured confounding. A non-null effect of the treatment on the NCO or a non-null effect of the NCE on the outcome of interest amounts to compelling evidence of unmeasured confounding. In the literature of causal peer effects, Egami (2018) exploits an NCO, and Liu and Tchetgen Tchetgen (2020) use an NCE to address unmeasured network confounding. While they require relatively weak assumptions to detect unmeasured network confounding, both works require much stronger assumptions for identification of causal peer effects as they each use only one type of negative control variable but not both. Recently, a series of papers (Deaner, 2018; Kallus et al., 2021; Kuroki & Pearl, 2014; Miao, Geng, et al., 2018; Miao, Shi, et al., 2018; Shi et al., 2020; Tchetgen Tchetgen, Ying, et al., 2020) propose to use DNCs for identification of causal effects (also known as proximal causal learning), but they have focused on i.i.d. or panel data settings, and to date none of these papers has considered network data.

Our contribution is to propose a general framework for using DNCs for identification and estimation of causal peer effects in the presence of uncontrolled network confounding while taking into account network dependence. We first derive non-parametric identification of causal peer effects in the presence of unmeasured network confounding by exploiting DNCs that are associated with unmeasured confounders. In particular, we incorporate DNC variables via a network outcome confounding bridge function, a network version of the outcome confounding bridge function studied in Miao, Shi, et al. (2018), Shi et al. (2020), and Tchetgen Tchetgen, Ying, et al. (2020). We discuss general approaches for selecting negative controls from network data in practice. In particular, one can use auxiliary variables that (a) do not affect network relationships and (b) do not affect variables of other units. These auxiliary variables can be seen as proxies that have no causal network effects but are associated with unmeasured network confounders. We then propose a semi-parametric estimation framework using a generalised method of moments (GMM) estimator (Hansen, 1982) for the causal peer effect, and we establish consistency and asymptotic normality of the resulting estimator under correct specification of the network confounding bridge function, and an assumption about ψ-network dependence (Kojevnikov et al., 2021), which expresses the degree of stochastic dependence between variables in terms of network distance. This assumption of ψ-network dependence restricts the speed by which network dependence decays as network distance increases, and the speed by which the density of the network changes as sample size increases. Finally, we provide a network heteroskedasticity and autocorrelation consistent (network HAC) variance estimator, with which researchers can construct asymptotic confidence intervals of causal peer effects.

The paper is organised as follows. In Section 2, we consider dyadic data to introduce the use of DNCs for identification of causal peer effects. In Section 3, we examine a more general setting where one observes data from a single network. We study identification as well as estimation and inference by accounting for ψ-network dependence. In Section 4, we illustrate our methods by applying them to the Add Health network data to infer the extent of causal peer effects in education. Section 5 concludes the paper. In online supplementary Appendix, we assess the finite sample performance of our proposed estimators via extensive simulations and provide proofs of all the results described in the main text.

1.1 Related literature

This article builds on a growing literature on causal peer effects. Various methods have been proposed to address concerns about uncontrolled network confounding in observational studies. A popular approach for identification of causal peer effects is the so-called instrumental variable (IV) method. Bramoullé et al. (2009) use IVs to deal with simultaneity in the linear-in-mean models

(Goldsmith-Pinkham & Imbens, 2013; Manski, 1993). O'Malley et al. (2014) propose to use genes as instruments to study causal peer effects of body mass index among friends. In addition to the well-known exclusion restriction, both methods assume (conditional) exogeneity of IVs. However, this assumption may be untenable in a wide range of applications because it is violated as long as IVs are associated with unmeasured variables at the source of homophily or contextual confounding. In contrast, our methods allow for and in fact leverage such association between negative controls and unobserved confounders. McFowland and Shalizi (2023) propose a consistent estimator of causal peer effects, which adjusts for estimated latent homophilous attributes in settings where the data generating process is linear and the network grows according to either a stochastic block model or a continuous space model. In contrast, we establish non-parametric identification of causal peer effects by using DNCs, and our results accommodate both latent homophily and contextual confounding. Also, our asymptotic results make an alternative assumption of ψ-network dependence (Kojevnikov et al., 2021) instead of assuming specific network models, and we allow for asymptotic normality and construction of asymptotic confidence intervals in addition to consistent estimation of causal peer effects. The main distinction is that our proposed method and McFowland and Shalizi (2023) rely on different causal assumptions. While McFowland and Shalizi (2023) assume that researchers can model the network formation process and that estimated latent variables are sufficient for controlling for homophily bias, our approach relies on two negative controls (NCO and NCE) and gives corresponding causal assumptions sufficient for identification. We view these methods as complementary to each other, and researchers can use domain knowledge to evaluate which identification strategies are most credible in a given application.

2 DNCs for dyadic data

In this section, to ground ideas, we focus on identification of causal peer effects from dyadic data. In general, even in this simple setting, causal peer effects are not identified based on standard covariate adjustment in the presence of unmeasured network confounding, such as latent homophily and contextual confounding. We propose an alternative approach based on negative controls. We consider a more general network setting in Section 3.

2.1 Notation and definitions

We consider data on dyads, i.e. pairs of two individuals. For each dyad i, let $S_i = 1$ when the two units are connected and $S_i = 0$ when the two units have no tie between them. For example, in dyadic data based on a students' friendship survey, $S_i = 1$ encodes the two students being friends with each other and $S_i = 0$ otherwise.

Suppose one has observed n independent and identically distributed samples of connected dyads ($S_i = 1$) where each dyad is labelled $i \in \{1, \ldots, n\}$. For each unit within dyads, we observe focal behaviour Y at two time points, baseline and a single follow-up. Define Y_{kt} to be the focal behaviour of unit $k \in \{1, 2\}$ at time $t \in \{1, 2\}$, where $t = 1, 2$ denotes baseline and follow-up, respectively. Without loss of generality, define $k = 1$ to be the *ego*—a unit on whom we estimate a causal peer effect—and define $k = 2$ to be the *peer*—a unit whose effect on the ego we estimate.

The outcome of interest is ego's behaviour at follow-up Y_{12}. The treatment variable of interest is peer's behaviour at baseline Y_{21}. Using the potential outcome framework (Neyman, 1923; Robins, 1986; Rubin, 1974), define $Y_{12}(y_{21})$ to be the potential outcome had possibly contrary to fact, the treatment variable been set to $Y_{21} = y_{21}$, which we assume to exist and to be well defined. Throughout the paper, we make the standard consistency assumption linking observed and potential outcomes:

$$Y_{12} = Y_{12}(Y_{21}), \quad \text{a.s.} \tag{1}$$

Our goal is to estimate the average causal peer effect (ACPE), defined as

$$\tau(y_{21}, y'_{21}) := \mathbb{E}\{Y_{12}(y_{21}) - Y_{12}(y'_{21}) \mid S = 1\},$$

where $y_{21}, y'_{21} \in \mathcal{Y}_{21}$ and $\mathcal{Y}_{21}$ is the support of Y_{21}. We condition on $S = 1$ because we only consider connected dyads.

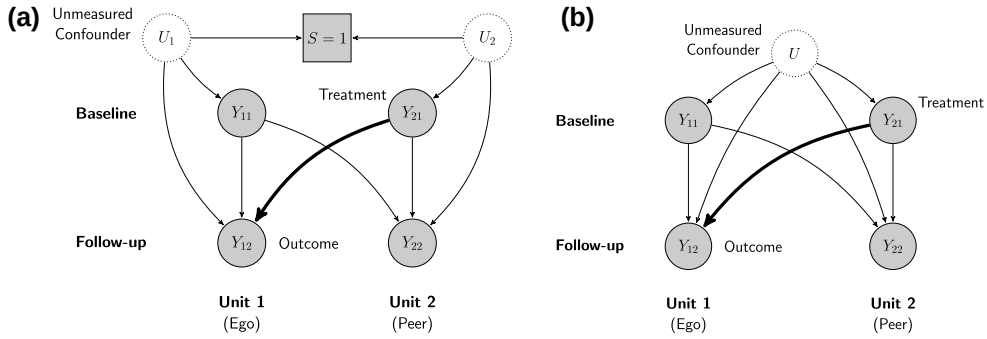

Figure 1. DAGs for dyadic data in the presence of unmeasured network confounding. (a) Unmeasured homophily. (b) Contextual confounding. The thick arrows from Y_{21} to Y_{12} indicate the causal peer effects of interest. We use shaded (dotted) nodes to denote observed (unobserved) variables. For simplicity, we suppressed observed covariates X. In panel (a), the square box around $S = 1$ represents that we observe dyads conditional on $S = 1$. DAGs = directed acyclic graphs.

2.2 Identification challenge

Covariate adjustment based on conditional ignorability is the most common approach to identification of causal effects in observational studies (Robins, 1986; Rosenbaum & Rubin, 1983). In causal peer analysis, conditional ignorability entails assuming

$$Y_{12}(y_{21}) \perp\!\!\!\perp Y_{21} \mid Y_{11}, X, S = 1, \tag{2}$$

where X represent observed pre-treatment covariates for the dyad. However, such approach is, in general, not plausible when estimating the ACPE due to unmeasured homophily (Shalizi & Thomas, 2011) and contextual confounding (VanderWeele & An, 2013). Homophily bias arises when people become connected due to unobserved characteristics. Contextual confounding exists when peers share some unobserved contextual factors. In this paper, we use the term *unmeasured network confounding* to refer to confounding that violates conditional ignorability in causal peer analysis, and thus it contains unmeasured homophily and contextual confounding as special cases.

Figure 1a represents a causal-directed acyclic graph (DAG) illustrating latent homophily. We use U_k with $k \in \{1, 2\}$ to represent unobserved characteristics that affect a tie relationship S as well as focal behaviour Y. Because the tie relationship S is affected by U_1 and U_2, variable S is a collider in the terminology of graph theory. To estimate the ACPE, we condition on $S = 1$, and therefore, there is an unblocked back-door path $Y_{21} \leftarrow U_2 \rightarrow \boxed{S = 1} \leftarrow U_1 \rightarrow Y_{12}$, where U_1 and U_2 are unobserved, and the square box around $S = 1$ means that we observe dyads conditional on $S = 1$ (Shalizi & Thomas, 2011). Thus, even in the simple setting of dyadic data, identification of the ACPE is not possible without additional assumptions.

Figure 1b represents a causal DAG illustrating contextual confounding. Here, we use U to represent an unmeasured shared context that affects both the ego and peer. Due to the unblocked back-door path $Y_{21} \leftarrow U \rightarrow Y_{12}$, conditional ignorability is violated.

2.3 Identification with DNCs

In this section, we consider an alternative approach for identification and estimation by exploiting auxiliary variables called negative controls. In particular, we will use NCO and NCE variables, which we define below.

We first make the latent ignorability assumption, which states that conditional ignorability holds if we could measure all factors at the source of network confounding.

Assumption 1.1 (Latent ignorability). For all $y_{21} \in \mathcal{Y}_{21}$,

$$Y_{12}(y_{21}) \perp\!\!\!\perp Y_{21} \mid U, X, S = 1.$$

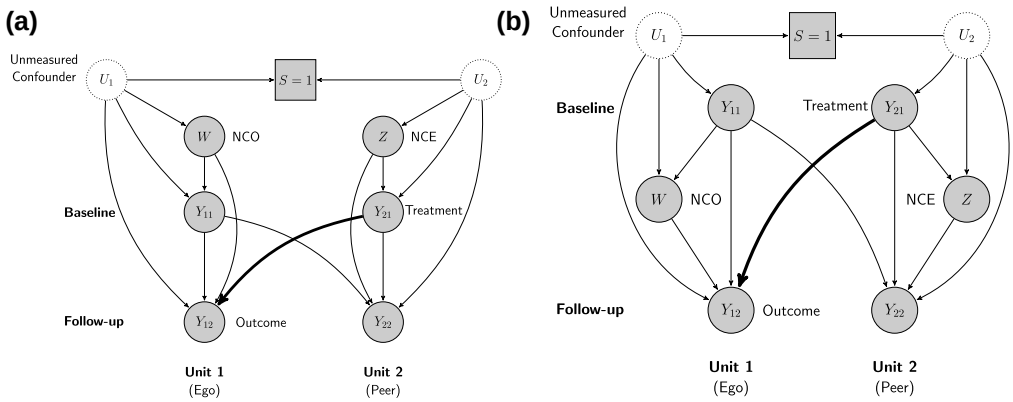

Figure 2. Examples of causal DAGs with double negative controls. Auxiliary variables W and Z are added to Figure 1a. The causal relationships between focal behaviours Y_{kt}, unmeasured confounder U_k, and the tie relationship S are the same as those in Figure 1a. In Figure 2a, W and Z are pre-treatment variables, while in Figure 2b, W and Z are intermediate variables. DAGs = directed acyclic graphs.

Assumption 1.1 states that U and X suffice to account for confounding of the relationship between Y_{21} and $Y_{12}(y_{21})$ conditional on $S = 1$, whereas X alone may not. This assumption is often plausible as there is no direct restriction on the nature of the latent characteristic U. However, this assumption alone is not sufficient for identification given that we do not observe the latent characteristic U.

The key to the proposed approach is to suppose that one can measure two auxiliary variables, NCO W and NCE Z that satisfy the following conditions.

Assumption 1.2 (Negative controls).

1. Negative control outcome (NCO):

$$W \perp\!\!\!\perp Y_{21} \mid U, X, S = 1. \tag{3}$$

2. Negative control exposure (NCE):

$$Z \perp\!\!\!\perp Y_{12} \mid Y_{21}, U, X, S = 1 \quad \text{and} \quad Z \perp\!\!\!\perp W \mid Y_{21}, U, X, S = 1. \tag{4}$$

Assumption 1.2(1) states that W is an auxiliary variable that is conditionally independent of the treatment Y_{21} given the latent confounder U, observed pre-treatment covariates X, and the dyadic type S. Assumption 1.2(2) means that Z is an auxiliary variable that is conditionally independent of the outcome Y_{12} and NCO W given the treatment Y_{21}, the latent confounder U, observed pre-treatment covariates X, and the dyadic type $S = 1$.

In practice, plausible candidates for negative controls are auxiliary variables that (a) do not affect network relationships and (b) do not affect variables of other units. Figure 2 represents examples of causal DAGs that extend Figure 1a and encode Assumption 1.2. In Figure 2a and b, W and Z are auxiliary variables that do not affect the dyadic relationship S or do not affect variables of the other unit. A variety of relationships between the focal behaviours and negative controls can be accommodated. In Figure 2a, W and Z are pre-treatment variables affecting the focal behaviours, while in Figure 2b, W and Z are intermediate variables between focal behaviour at baseline Y_{k1} and focal behaviour at follow-up Y_{k2}. In addition, in both Figure 2a and b, Y_{11} may also serve as the NCO.

Example (Negative controls). In the context of Add Health data, Cohen-Cole and Fletcher (2008) estimated peer effects on three health outcomes—acne, height, and headaches—known ex ante not to exhibit peer effects. The purpose of Cohen-Cole & Fletcher's (2008) analysis was to investigate the validity of

popular statistical approaches that assume the absence of unmeasured homophily and contextual confounding (e.g. Christakis & Fowler, 2007). By finding implausible peer effects on the three health outcomes mentioned above, Cohen-Cole and Fletcher (2008) warn that latent homophily and contextual confounding may be operating in studies of peer effects based on Add Health data.

In the proposed DNC approach, the three health outcomes can be used not only for detecting confounding but also as negative controls to potentially correct for such confounding. For example, whether a student has headaches is unlikely (a) to causally affect whether students are friends to each other and (b) to causally affect peers' headaches. In our empirical application (Section 4), unobserved study habits are important potential unmeasured network confounders and are likely to be correlated with headaches. If these conditions are plausible in applied contexts, an ego's headache and a peer's headache can be used as the NCO and NCE, respectively. In general, the choice of negative controls should be justified by domain knowledge in each application by evaluating Assumption 1.2.

Upon selecting valid NCO and NCE variables, one approach supposes that there exists an outcome confounding bridge (Miao, Shi, et al., 2018) that relates the confounders' effects on negative control outcome W to the confounders' effects on outcome of interest Y_{12}.

Assumption 1.3 (Outcome confounding bridge function). There exists a function $h(W, Y_{21}, X)$ such that for all $y_{21} \in \mathcal{Y}_{21}$,

$$\begin{aligned} &\mathbb{E}(Y_{12} \mid Y_{21} = y_{21}, U, X, S = 1) \\ &= \mathbb{E}\{h(W, y_{21}, X) \mid Y_{21} = y_{21}, U, X, S = 1\}. \end{aligned} \tag{5}$$

Assumption 1.3 states that the confounding effect of U on outcome Y_{12} is equal to the confounding effect of U on $h(W, y_{21}, X)$, a transformation of W. One simple yet important implication of this assumption is that W should be associated with U conditional on the treatment, observed covariates, and the dyadic relationship.

This assumption formally connects the confounding effect on the outcome Y_{12} and the confounding effect on the NCO W. Instead of assuming complete knowledge of its relationship, the proposed DNC approach will use NCE Z to identify it, as described below in Theorem 1.

Formally, equation (5) is a Fredholm integral equation of the first kind (Kress, 1989). Existence of a solution to this equation can be established under regularity conditions regarding the negative control relevance and a certain singular value decomposition of the operator defining the integral equation. These conditions are somewhat technical and we reserve their details to online supplementary Appendix A.1. In practice, the existence of an outcome confounding bridge can be established when the number of NCOs W is at least as large as the number of unobserved confounders U (the number of levels in W should be at least as large as the number of levels in U when both are categorical variables).

To further understand the meaning of the outcome confounding bridge assumption, we provide two examples below. The first is a linear confounding bridge. While it is restrictive, it contains a popular difference-in-differences (DID) method (Sofer et al., 2016) as a special case, and it provides clear intuition. The second is a case when unmeasured confounder U and NCO W are categorical (Shi et al., 2020). In this second example, we can simplify sufficient conditions for the existence of the outcome confounding bridge.

Example (Linear confounding bridge). Suppose researchers posit a linear confounding bridge $h(W, Y_{21}, X; \gamma) = (1, W, Y_{21}, X)^\top \gamma$. Then, Assumption 1.3 implies that the confounding effect of U on the primary outcome Y_{12} is proportional to the confounding effect of U on the negative control outcome W. With unknown

value of γ_W, we have

$$\underbrace{\mathbb{E}(Y_{12} \mid Y_{21} = y_{21}, U = u, X = x) - \mathbb{E}(Y_{12} \mid Y_{21} = y_{21}, U = u', X = x)}_{\text{Confounding effect of } U \text{ on } Y_{12}}$$
$$= \gamma_W \times \underbrace{\mathbb{E}(W \mid Y_{21} = y_{21}, U = u, X = x) - \mathbb{E}(W \mid Y_{21} = y_{21}, U = u', X = x)}_{\text{Confounding effect of } U \text{ on } W}. \tag{6}$$

While this assumption might seem unfamiliar, the parallel trends assumption necessary for the widely used DID method is a special case of this assumption (i.e. strictly stronger than Assumption 1.3). This is because the DID method uses the baseline outcome Y_{11} as the NCO and a priori sets $\gamma_W = 1$ (i.e. the confounding effect of U on the primary outcome Y_{12} is exactly equal to the confounding effect of U on NCO W). In contrast to the DID method, we only assume that the confounding effect on the primary outcome is proportional to (not necessarily equal to) the confounding effect on the NCO. Importantly, we do not need to assume the value of γ. Rather, as shown below, one can use an appropriate choice of NCE Z to identify and estimate γ. In Section 2.3.1, we provide additional discussions about the connection between the DID and the DNC approach.

Example (Categorical variables). Suppose W and U are categorical variables. We define possible values for W and U as w_i and u_j for $i = 1, \ldots, |W|$, $j = 1, \ldots, |U|$, where $|\cdot|$ denotes the cardinality of a categorical variable. Define $\Pr(W \mid y_{21}, U, x, S = 1)$ to be a $|W| \times |U|$ matrix with $\Pr(W \mid y_{21}, U, x, S = 1)_{ij} = \Pr(W = w_i \mid y_{21}, U = u_j, x, S = 1)$. In this case, sufficient conditions for the existence of the confounding bridge function can be simplified. As long as $\Pr(W \mid y_{21}, U, x, S = 1)$ have full column rank (which implies that $|W| \geq |U|$), the confounding bridge function $h(W, y_{21}, x)$ exists as a solution to $\mathbb{E}(Y_{12} \mid y_{21}, U, x, S = 1) = h(W, y_{21}, x)\Pr(W \mid y_{21}, U, x, S = 1)$, where $h(W, y_{21}, x)$ is a $1 \times |W|$ vector.

Finally, we make the following assumption about negative control relevance.

Assumption 1.4 (Negative control relevance). For any square integrable function f and any y_{21} and x, if $\mathbb{E}\{f(W) \mid Z = z, Y_{21} = y_{21}, X = x, S = 1\} = 0$ for almost all z, then $f(W) = 0$ almost surely.

This assumption states that Z is sufficiently informative about W, which is essential to ensure identification of the outcome confounding bridge function h. This condition is formally known as a completeness condition, a well-known technical condition central to the study of sufficiency in statistical inference. Many commonly used parametric and semi-parametric models, such as semi-parametric exponential family and semi-parametric location-scale family, satisfy the completeness condition (Newey & Powell, 2003).

The completeness condition has been widely used to achieve identification in non-parametric IV models (e.g. Newey & Powell, 2003). In the IV literature, completeness formalises the IV relevance condition and generalises the rank condition of linear IV models. Thus, Assumption 1.4 essentially means that Z is relevant for W conditional on (Y_{21}, X, S).

Remark To gain intuition, we consider implications of the completeness condition.

(Categorical variables) Consider a special case of categorical NCO and NCE. In this case, Assumption 1.4 requires that the number of levels in NCE must be at least as large as the number of levels in NCO.

(Continuous variables) When both NCO and NCE are continuous, Assumption 1.4 requires that the number of NCEs must be at least as large as the number of NCOs.

It is important to note that the negative control relevance assumption can be violated in practice. In online supplementary Appendix E, we conduct simulation studies where this assumption is violated and demonstrate that the proposed DNC estimator has small bias and corresponding confidence intervals maintain reasonable coverage when the violation is small. However, as one might expect, the larger is the violation, the bias is larger and coverage performance of corresponding confidence intervals becomes poorer. Simulation studies also show that, when association between negative controls and unmeasured confounders is weak, finite sample performance of the proposed methods could be poor, even though the ACPE is in principle identified. This highlights the practical importance of selecting negative controls that are strongly associated with unmeasured confounders.

To understand the potential impact of the violation of the negative control relevance assumption, we also provide a sensitivity analysis method in online supplementary Appendix C. We extend the first sensitivity analysis for negative control methods in a linear confounding bridge case (Cobzaru et al., 2022). In particular, we generalise their result to accommodate two important aspects of the causal peer effect analysis: continuous treatments and network-dependent errors. Therefore, the sensitivity analysis method provided in online supplementary Appendix C can accommodate the dyadic data (Section 2) and the network data (Section 3).

Remark Importantly, the assumption of negative control relevance with respect to W and Z is not testable for unrestricted models in the sense that no test will have power greater than size against any alternative: this is a general impossibility theorem established in Canay et al. (2013) for unrestricted models. However, for restricted models, the violation of the negative control relevance [Assumptions 1.4 and 2(4)] can be analysed as the identification problem in the GMM (Wright, 2003). For example, in a linear confounding bridge case, researchers can use a weak instrument test discussed in Andrews et al. (2019) to detect the violation of the negative control relevance assumption. For the network data (Section 3), these tests need to be generalised to take into account network-dependent errors, which requires additional substantial theoretical analysis, and is an important topic of future studies.

Remark Finally, we note that alternative completeness conditions are also sufficient for identification of the ACPE. We establish non-parametric identification of the ACPE under these alternative completeness conditions in online supplementary Appendix A.6. We further discuss completeness condition in online supplementary Appendix A.2.

The following theorem establishes non-parametric identification of the ACPE under the stated assumptions.

Theorem 1 Under Assumptions 1.1–1.4, the confounding bridge function is identified as the unique solution to the following equation:

$$\mathbb{E}(Y_{12} \mid Z, Y_{21}, X, S = 1) = \mathbb{E}\{h(W, Y_{21}, X) \mid Z, Y_{21}, X, S = 1\}.$$

Using the identified confounding bridge function $h(W, Y_{21}, X)$, the ACPE is identified by

$$\tau(y_{21}, y'_{21}) = \mathbb{E}\{h(W, y_{21}, X) - h(W, y'_{21}, X) \mid S = 1\}.$$

Importantly, under Assumptions 1.1–1.4, we have identified the ACPE without imposing any parametric restriction on the confounding bridge function. We provide a proof in online supplementary Appendix A.5.

Example (Identification under linear confounding bridge). While Theorem 1 establishes non-parametric identification of the ACPE, here we provide intuition behind identification using a simple case with a binary treatment and a binary NCE without any pre-treatment covariates, which admits a closed-form solution. Because this part of discussions is general about the DNC approach, we use A to denote the binary treatment and Y to denote the outcome variable.

Consider a linear confounding bridge function: $h(W, A; \gamma) = \gamma_0 + \gamma_1 W + \gamma_2 A$. Then, the ACPE is identified as

$$\begin{aligned}\tau(1,0) = &\mathbb{E}\{\mathbb{E}(Y \mid A=1, Z) - \mathbb{E}(Y \mid A=0, Z)\} \\ &- \mathbb{E}\{\mathbb{E}(W \mid A=1, Z) - \mathbb{E}(W \mid A=0, Z)\} \\ &\times \frac{\mathbb{E}\{\mathbb{E}(Y \mid A, Z=1) - \mathbb{E}(Y \mid A, Z=0)\}}{\mathbb{E}\{\mathbb{E}(W \mid A, Z=1) - \mathbb{E}(W \mid A, Z=0)\}}.\end{aligned}$$

We provide the proof in online supplementary Appendix A.8. The first term $\mathbb{E}\{\mathbb{E}(Y \mid A=1, Z) - \mathbb{E}(Y \mid A=0, Z)\}$ corresponds to a biased estimator of the ACPE; a regression of outcome Y on treatment A conditional on Z, which is equal to the ACPE only in the absence of unmeasured confounder. The second term $\mathbb{E}\{\mathbb{E}(W \mid A=1, Z) - \mathbb{E}(W \mid A=0, Z)\}$ captures the confounding effect on W, which would be zero in the absence of unmeasured confounding. This captures the amount of confounding to be corrected. Finally, the third term $\mathbb{E}\{\mathbb{E}(Y \mid A, Z=1) - \mathbb{E}(Y \mid A, Z=0)\}/\mathbb{E}\{\mathbb{E}(W \mid A, Z=1) - \mathbb{E}(W \mid A, Z=0)\}$ represents a ratio of the confounding effects on the outcome Y and on NCO W, which is identified by using NCE Z. It is clear here that the negative control relevance (Assumption 1.4) is essential to guarantee $\mathbb{E}\{\mathbb{E}(W \mid A, Z=1) - \mathbb{E}(W \mid A, Z=0)\} \neq 0$. Intuitively, the DNC approach subtracts an estimated bias (the second term) scaled by the differential confounding effects on the outcome and on the NCO (the third term) from the original biased estimator (the first term).

This explicit form contains two important special cases: (a) conditional ignorability (i.e. no unmeasured network confounding) and (b) the well-known DID design. First, under conditional ignorability, $\mathbb{E}\{\mathbb{E}(W \mid A=1, Z) - \mathbb{E}(W \mid A=0, Z)\} = 0$, and thus, $\tau(1,0) = \mathbb{E}\{\mathbb{E}(Y \mid A=1, Z) - \mathbb{E}(Y \mid A=0, Z)\}$, which reduces to the standard identification formula under conditional ignorability. Second, the formula reduces to DID when the confounding effect on the outcome is equal to the confounding effect on the NCO. Sofer et al. (2016) argue that DID uses a pre-treatment outcome (Y_{11} in our dyadic data setting) as W and assumes the entire bridge function is known, i.e. not only a functional form but also the value of coefficient γ, without using any NCE. In particular, the widely used assumption of parallel trends states that the confounding effects on the outcome Y_{12} and on pre-treatment outcome Y_{11} (used as the NCO) are equal, and thus, the third term is equal to one.

In contrast, the DNC approach can use any valid W (including pre-treatment outcome Y_{11} as a special case). Most importantly, we use NCE Z to estimate the confounding bridge function—the differential confounding effects on the outcome Y and on NCO W—as the ratio $\mathbb{E}\{\mathbb{E}(Y \mid A, Z=1) - \mathbb{E}(Y \mid A, Z=0)\}/\mathbb{E}\{\mathbb{E}(W \mid A, Z=1) - \mathbb{E}(W \mid A, Z=0)\}$ in the third term. It is important to emphasise that, while we use this closed-form solution in the linear confounding bridge case to provide intuition behind the DNC approach, our identification results do not impose any parametric restriction on the confounding bridge or negative controls.

2.3.1 Connection to other identification strategies

The proposed DNC estimator is connected to a wide range of existing identification strategies used for causal peer effect analysis. (a) Egami (2018) examines identification of the causal peer effect

when only one NCO is available. In this case, to identify the ACPE, researchers need to assume that the confounding effect of U on the primary outcome is the same as the confounding effect of U on the NCO. This is equivalent to the parallel trends assumption required in the DID method and the time-invariant confounder assumption required in a fixed effect estimator (Sofer et al., 2016). This is a special case of the confounding bridge assumption (Assumption 1.3). (b) Liu and Tchetgen Tchetgen (2020) examine identification of the causal peer effect when only one NCE is available by focusing on the dyadic data with homophily bias. To identify the ACPE, Liu and Tchetgen Tchetgen (2020) assume the correct specification of network formation process. The DNC approach does not require a model of network formation as we have an additional NCO. (c) The DNC approach is also connected to the IV method (Bramoullé et al., 2009; O'Malley et al., 2014). In particular, the IV exclusion restriction is similar to the NCE condition. However, the key differences are other causal assumptions required for causal identification. While the IV method requires the IV exogeneity, the IV relevance as well as monotonicity or the no-interaction assumption, the DNC approach instead requires the latent ignorability, the NCO condition, the outcome confounding bridge function, and the NC relevance. Thus, the DNC method and the IV method allow for causal identification in different scenarios and they are complementary to each other.

2.4 Estimation and inference

We now propose a strategy for estimation and inference of the ACPE. Suppose that an analyst has specified a parametric or semi-parametric model for the confounding bridge $h(W, Y_{21}, X; \gamma)$ with parameter γ. Then, based on Theorem 1, one can estimate γ by solving the following empirical moment equations:

$$\frac{1}{n}\sum_{i=1}^{n}\{Y_{i12} - h(W_i, Y_{i21}, X_i; \gamma)\} \times \eta(Z_i, Y_{i21}, X_i) = 0,$$

where η is a user-specified vector function with dimension equal to that of γ. For example, under a linear confounding bridge function, i.e. $h(W, Y_{21}, X; \gamma) = (1, W, Y_{21}, X)^\top \gamma$, one can use $\eta(Z, Y_{21}, X) = (1, Z, Y_{21}, X)^\top$.

Once the bridge function h is estimated, one can estimate the ACPE by

$$\frac{1}{n}\sum_{i=1}^{n}\{h(W_i, y_{21}, X_i; \widehat{\gamma}) - h(W_i, y'_{21}, X_i; \widehat{\gamma})\}.$$

To appropriately account for uncertainty of the estimated bridge function and for the possibility that dimension of η might be larger than that of γ, one can combine both moments into GMM with parameter $\theta = (\tau, \gamma)$ (Hansen, 1982; Miao, Shi, et al., 2018). Define a moment for dyad i to be

$$m(Y_{i12}, Y_{i21}, W_i, Z_i, X_i; \theta) = \left\{ \begin{array}{c} \tau - \{h(W_i, y_{21}, X_i; \gamma) - h(W_i, y'_{21}, X_i; \gamma)\} \\ \{Y_{i12} - h(W_i, Y_{i21}, X_i; \gamma)\} \times \eta(Z_i, Y_{i21}, X_i) \end{array} \right\}.$$

Then, the GMM estimator is given by

$$\widehat{\theta} = \underset{\theta}{\operatorname{argmin}}\, \overline{m}(\theta)^\top \Omega \overline{m}(\theta), \tag{7}$$

where $\overline{m}(\theta) = \frac{1}{n}\sum_{i=1}^{n} m(Y_{i12}, Y_{i21}, W_i, Z_i, X_i; \theta)$ and Ω is a user-specified positive-definite weight matrix. Asymptotic properties described below hold for any positive-definite weight matrix Ω and for alternative GMM estimators, such as the two-step GMM and continuously updating GMM.

The proposed DNC estimator $\widehat{\tau}(y_{21}, y'_{21})$ for $\tau(y_{21}, y'_{21})$ is the first element of $\widehat{\theta}$ defined in equation (7). Because we consider i.i.d. samples of dyads in this section, the moment

$m(Y_{i12}, Y_{i21}, W_i, Z_i, X_i; \theta)$ is also i.i.d., and thus, under the standard regularity conditions for GMM (Hansen, 1982; Newey & McFadden, 1994), the DNC estimator is consistent and asymptotically normal. We provide expressions for asymptotic variance and its consistent estimator in online supplementary Appendix B.

3 DNCs for network data

In this section, we consider a sample of interconnected units in a single network. Extending results in Section 2, we propose the DNC approach to identification of the ACPE in the presence of unmeasured network confounding. We then examine estimation and inference while accounting for both unmeasured network confounding and network dependence of observations.

3.1 Notation and definitions

Suppose one has observed data on a population of n units interconnected by a network. We let $N_n = \{1, \ldots, n\}$ be the set of unit indices. We consider an undirected network $\mathcal{G}_n$ where ties or links between units are mutual, and connected units can affect each other. While we focus on an undirected network to ease the presentation, the proposed approach equally applies to directed networks. Formally, define $\mathcal{G}_n = (N_n, \mathbf{G})$, i.e. a set of units N_n connected by mutual ties, represented by a network adjacency matrix $\mathbf{G}$. The network adjacency matrix $\mathbf{G}$ depends on sample size n, but the index will be suppressed in the following discussion as it eases the exposition without confusion. The entry G_{ij} takes the value of one if unit i and j are connected and takes the value of zero otherwise. We follow the convention that $G_{ii} = 0$ for $i \in N_n$, and we call units i and j *peers* if $G_{ij} = 1$. We also define two network notations useful throughout the paper. We define network distance $d_n(i, j)$ to be the length of the shortest path between nodes i and j on network $\mathcal{G}_n$. We define $\mathcal{N}_n(i; s)$ to be a set of nodes that are at distance s from node i:

$$\mathcal{N}_n(i; s) = \{j \in N_n : d_n(i, j) = s\}.$$

For each $i \in N_n$, one observes (Y_{i1}, Y_{i2}, X_i), where Y_{it} denotes the focal behaviour of unit i at time $t \in \{1, 2\}$, and X_i are covariates of unit i measured prior to Y_{i1}. X_i can include network-characteristics, such as the network degree of unit i, as well as covariates of peers, e.g. the average age of peers. We call $t = 1$ baseline and $t = 2$ follow-up.

For the sake of clarity in the exposition, we restrict presentation of all main results to the causal effect of peers. It is important to emphasise that results in this section, however, do not assume null causal effects from higher order peers (e.g. peers-of-peers); we only consider such higher order peer effects as nuisance. In online supplementary Appendix F, we discuss similar results for a general case where higher order peer effects (e.g. the causal effect from peers-of-peers) are the main causal estimands of interest.

As the outcome variable, we focus on focal behaviour at follow-up Y_{i2}. In principle, it is possible to perform causal inference by defining a multivariate treatment variable based on focal behaviours of peers at baseline $\{Y_{j1}: j \in \mathcal{N}(i; 1)\}$. However, in practice, researchers may need to make a dimension-reducing assumption, known as an exposure mapping (Aronow & Samii, 2017), to define the treatment variable, $A_i = \phi(\{Y_{j1}: j \in \mathcal{N}(i; 1)\}) \in \mathbb{R}$, where function ϕ is specified by a researcher based on subject matter knowledge. For example, the most common choice is $A_i = \sum_{j=1}^{n} G_{ij} Y_{j1} / \sum_{j=1}^{n} G_{ij}$, while our results can accommodate any choice of ϕ. The potential outcome $Y_{i2}(a)$ is defined as the outcome that would realise when the treatment variable is set to $A_i = a$. We make the standard consistency assumption linking observed and potential outcomes, $Y_{i2} = Y_{i2}(A_i)$, throughout the paper. Our goal is to estimate the ACPE, defined as

$$\tau(a, a') := \frac{1}{n} \sum_{i=1}^{n} \mathbb{E}\big\{ Y_{i2}(a) - Y_{i2}(a') \big\}, \tag{8}$$

where $a, a' \in \mathcal{A}$ and $\mathcal{A}$ is the support of A. We define the expectation conditional on the observed network $\mathcal{G}_n$, while we omit this conditioning event for notational simplicity. Unlike typical causal parameters in i.i.d. settings, units' potential outcome may not share a common expectation, i.e.

$\mathbb{E}\{Y_{i2}(a) - Y_{i2}(a')\} \neq \mathbb{E}\{Y_{j2}(a) - Y_{j2}(a')\}$ for $i \neq j$. Thus, the causal estimand is explicitly written as the network average of individual level causal effects.

To identify the ACPE, existing works rely upon the assumption that the observed variables are sufficient to account for confounding of the relationship between A_i and $Y_{i2}(a)$, i.e.

$$Y_{i2}(a) \perp\!\!\!\perp A_i \mid Y_{i1}, X_i,$$

where observed pre-treatment covariates X_i can include observed covariates of peers of unit i or covariates of other units who are indirectly connected to unit i. Even though some recent methods allow for network dependence across units (e.g. Ogburn et al., 2017; Tchetgen Tchetgen, Fulcher, et al., 2021), they assume such latent network dependence does not confound the main outcome-treatment relationship. However, as discussed in Section 2.2, such assumption may be untenable in many applications due to unmeasured network confounding, including unmeasured homophily and contextual confounding.

Remark We note that an exposure mapping ϕ can be misspecified in practice. Because it is impossible to identify the ACPE while accounting for all types of misspecification of the exposure mapping, we consider one common type of misspecification in online supplementary Appendix F.2: the observed network might miss some network ties between units. By connecting it to the general theory of the GMM, we provide a framework to test this type of misspecification. Examining more general forms of misspecification is an important topic of future studies, given that, even in randomised experiments settings without unobserved network confounding, only few papers consider the possibility of misspecification (e.g. Aronow & Samii, 2017; Egami, 2021; Sävje, 2021) and, as far as we know, there is no prior work that handles misspecification in the observational network data settings even in the absence of unmeasured network confounding.

3.2 Identification assumptions with DNCs

We propose an alternative approach based on DNCs. We generalise Assumptions 1.1–1.4 in Section 2 to the network setting.

Assumption 2 1. (Latent ignorability) For all $a \in \mathcal{A}$ and all $i \in N_n$,

$$Y_{i2}(a) \perp\!\!\!\perp A_i \mid U_i, X_i.$$

2. (Negative controls). For all $i \in N_n$,

$$\begin{aligned} &\text{(Negative control outcome)} && W_i \perp\!\!\!\perp A_i \mid U_i, X_i, \\ &\text{(Negative control exposure)} && Z_i \perp\!\!\!\perp Y_{i2} \mid A_i, U_i, X_i, \quad \text{and} \\ & && Z_i \perp\!\!\!\perp W_i \mid A_i, U_i, X_i. \end{aligned}$$

3. (Outcome confounding bridge). There exists an outcome confounding bridge function $h(W_i, A_i, X_i)$ such that for all $a \in \mathcal{A}$, and all $i \in N_n$,

$$\mathbb{E}(Y_{i2} \mid A_i = a, U_i, X_i) = \mathbb{E}\{h(W_i, a, X_i) \mid A_i = a, U_i, X_i\}. \tag{9}$$

4. (Negative control relevance). For any square integrable function f and any a and x, if $\mathbb{E}\{f(W_i) \mid Z_i = z, A_i = a, X_i = x\} = 0$ for almost all z, then $f(W_i) = 0$ almost surely.

We emphasise that X_i can include network-characteristics, such as the network degree of unit i, as well as covariates of peers, e.g. the average age of peers.

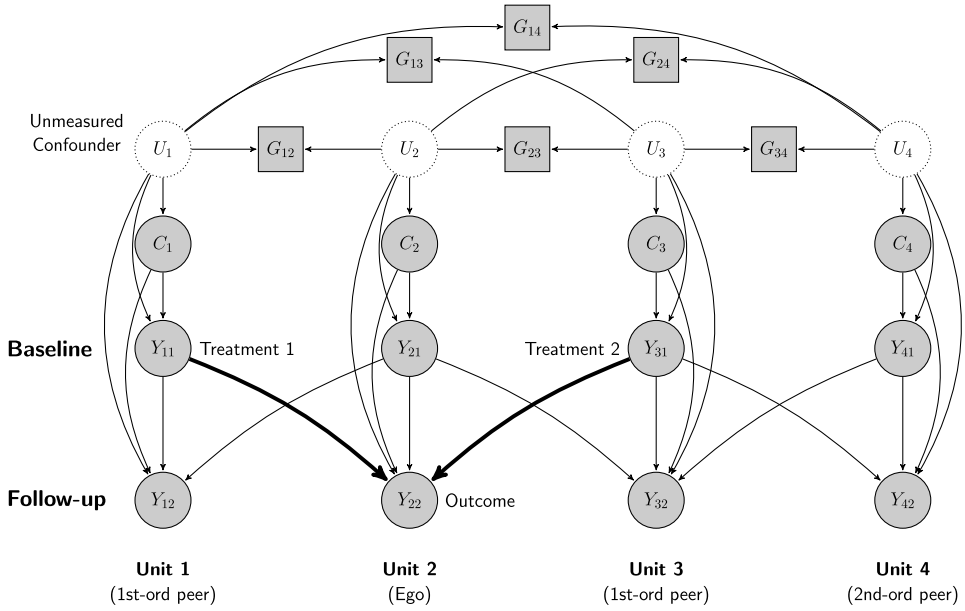

Figure 3. Example of a DAG with unobserved homophily. For concreteness, we show Unit 2 as the ego. The thick arrows from Y_{11} to Y_{22} and from Y_{31} to Y_{22} indicate the causal peer effects of interest. We use shaded (dotted) nodes to denote observed (unobserved) variables. For simplicity, we suppressed observed covariates X. Here, the square box around G represents that we observe variables conditional on network $\mathcal{G}$. DAG = directed acyclic graph.

3.3 Selecting negative controls in network settings

In the proposed DNC approach, selection of NCO and NCE is essential in practice. While in principle any negative controls satisfying Assumption 2 can be used, we discuss two convenient strategies to select negative controls in the network setting.

3.3.1 Using auxiliary variables

As in Section 2, a plausible candidate is an auxiliary variable C_i that (a) is a priori known not to affect network relationships $\mathcal{G}_n$ and (b) not to affect other units. For example, in the context of Add Health data, whether a student experiences headaches is likely to satisfy this condition (see 'Example (Negative Controls)' in Section 2.3). More generally, these auxiliary variables can be seen as proxies that have no causal network effects but are associated with unmeasured network confounders.

Figure 3 represents an example of a DAG with unobserved homophily where Assumption 2.2 holds. We view Unit 2 as ego, who has two peers (Units 1 and 3) and one peer-of-peers (Unit 4). In this graph, C_2 satisfies NCO conditions and three variables $\{C_1, C_3, C_4\}$ satisfy NCE conditions. Importantly, not only auxiliary variables of peers but also those of peers-of-peers also satisfy NCE conditions. While we use a DAG with unobserved homophily as an example here, the same argument also applies to contextual confounding.

Finally, we provide primitive conditions that imply negative control conditions. In particular, the following conditions capture a general approach for using auxiliary variables as negative controls.

$$C_i \perp\!\!\!\perp \{(C_j, Y_{j1}) : j \neq i\} \mid U_i, X_i, \tag{10}$$

$$Y_{i2} \perp\!\!\!\perp \{C_j : j \neq i\} \mid A_i, U_i, X_i. \tag{11}$$

Equations (10) and (11) formalise the notion that C_i should not affect network relationships and should not affect peers' variables. Lemma 1 below shows that auxiliary variable C_i can serve as a

valid negative control if it satisfies the stated conditions (Figure 3 is an example).

Lemma 1 Suppose auxiliary variable C_i satisfies the two conditions (equations (10) and (11)). Then, Assumption 2(2) holds with $W_i = C_i$ and $Z_i = \{C_j : j \neq i\}$.

Several points are worth noting. First, these are sufficient conditions, not necessary conditions for Assumption 2(2). Therefore, any negative controls that satisfy Assumption 2 can in principle be used for identification. Second, when this assumption is violated, researchers can consider an additional unobserved variable on which they can restore the conditional independence after conditioning. In this case, by including the additional unobserved confounder, Assumption 2(2) may be more credible, while Assumption 2(4) (the negative control relevance) may be less plausible, and thus, researchers may use a sensitivity analysis to investigate robustness of causal findings (see online supplementary Appendix C for details on such a sensitivity analysis). Third, Lemma 1 suggests that in practice, there may be multiple NCEs because one can use auxiliary variables of all other units $\{C_j : j \neq i\}$. Therefore, it may be possible to enhance identification and increase estimation efficiency by exploiting a large number of NCEs. We plan to examine optimal selection and specification of NCEs in future work. Finally, it is important to emphasise that, while the DNC approach requires an additional negative control compared to previous negative control methods that use only one type of negative control (e.g. Egami, 2018; Liu & Tchetgen Tchetgen, 2020), researchers can use one auxiliary variable C_i for two different purposes (NCO and NCE), and thus, the DNC approach can be used without additional practical complications.

3.3.2 *Using focal behaviours*

Second, in certain settings, we may also use focal behaviours Y_{it} of peers and those measured at baseline as plausible candidates for negative controls. In particular, ego's focal behaviour at baseline may serve as valid NCO and focal behaviours of peers-of-peers $\{Y_{jt} : j \in \mathcal{N}(i; 2), t \in \{1, 2\}\}$ may constitute valid NCEs. In Figure 3, Y_{21} also qualifies as NCO and variables $\{Y_{41}, Y_{42}\}$ qualify as NCEs. More generally, when focal behaviours of peers-of-peers constitute valid NCEs, focal behaviours of units at least of network distance 2 from node i may be credible NCEs. A hybrid approach might entail combining the auxiliary variables and focal behaviours as negative controls. In Figure 3, one may also use focal behaviour measured at baseline Y_{21} as NCO (instead of C_2) and auxiliary variables of peers $\{C_1, C_3, C_4\}$ as NCEs.

This selection of negative controls is particularly plausible when focal behaviours do not have direct causal relationships with peers' focal behaviours measured concurrently. In Figure 3, while peers' focal behaviours measured at baseline affect egos' focal behaviours measured at follow-up (e.g. Y_{11} and Y_{31} affect Y_{22}), peers' focal behaviours cannot causally affect egos' focal behaviours measured concurrently (e.g. Y_{11} and Y_{31} do not affect Y_{21}; Y_{12} and Y_{32} do not affect Y_{22}). This absence of causal simultaneity has previously been assumed in the literature of causal peer effects (Egami, 2018; Liu & Tchetgen Tchetgen, 2020; McFowland & Shalizi, 2023; Ogburn & VanderWeele, 2014; Shalizi & Thomas, 2011).

In practice, this assumption is most credible when researchers know a priori that the focal behaviours of units are indeed measured concurrently. For example, in Add Health data, students' grade-point average (GPA) are likely to be measured at the same time for students within a school, and thus, a student's GPA cannot be affected by peers' GPA in the same semester. Importantly, a student's GPA can be affected by peers' GPA in the prior semester, and a student's study habit might be affected by peer's study habits within the same semester. These, however, do not invalidate the use of GPA of peers-of-peers as NCEs as long as students' GPA within the same semester do not have direct causal relationships with each other. We further discuss practical considerations related to selection of negative controls in online supplementary Appendix A.7.

As in Section 3.3.1, we provide primitive sufficient conditions for valid negative controls. In particular, the following conditions provide a general approach for leveraging focal behaviours as negative controls.

$$Y_{i1} \perp\!\!\!\perp \{Y_{j1} : j \neq i\} \mid U_i, X_i, \tag{12}$$

$$Y_{i2} \perp\!\!\!\perp \{Y_{j1} : j \in \mathcal{N}(i; s), s \geq \tilde{s}\} \mid A_i, U_i, X_i, \quad \text{with user specified } \tilde{s}. \tag{13}$$

Equation (12) formalises the notion that focal behaviours do not have direct causal relationships with the peers' focal behaviours measured concurrently. Equation (13) captures the assumption that only peers closer than distance $\tilde{s}$ can have causal peer effects on an ego. While we focus on the causal effect from peers as the causal estimand, we do not necessarily need to assume the absence of higher order peer effects.

Lemma 2 below establishes that focal behaviours can serve as valid negative controls when they satisfy the stated conditions.

Lemma 2 Suppose that focal behaviours satisfy conditions (12) and (13). Then, Assumption 2(2) holds with $W_i = Y_{i1}$ and $Z_i = \{Y_{j1} : j \in \mathcal{N}(i; s), s \geq \tilde{s}\}$.

Again, we emphasise that these conditions are sufficient for Assumption 2(2), and thus, there may be other strategies to justify negative control selection. Any negative controls that satisfy Assumption 2 may be used for identification of the ACPE. Similarly to Lemma 1, Lemma 2 also shows that there exist multiple NCEs in practice. In our empirical application (Section 4), we use the average GPA of peers-of-peers at baseline as the NCE.

3.4 Non-parametric identification

Analogous to Theorem 1, we now establish non-parametric identification of the ACPE under Assumption 2.

Theorem 2 Under Assumption 2, the confounding bridge function is identified as the unique solution to the following equation:

$$\mathbb{E}(Y_{i2} \mid Z_i, A_i, X_i) = \mathbb{E}\{h(W_i, A_i, X_i) \mid Z_i, A_i, X_i\},$$

and the ACPE is identified by

$$\tau(a, a') = \frac{1}{n}\sum_{i=1}^{n} \mathbb{E}\{h(W_i, a, X_i) - h(W_i, a', X_i)\}.$$

We provide a proof in online supplementary Appendix A.5. Despite the complexity of network data, the identification assumptions and formula of the ACPE are remarkably similar to those in the dyadic study design. An important difference from the dyadic case emerges in estimation and inference where one must appropriately account for network dependence, a challenging task we consider next.

3.5 Estimation and inference

In this section, we consider estimation and inference for the ACPE while allowing for network-dependent observations.

We define a triangular array of $\mathbb{R}^{\nu}$-valued random vector, $L_{n,i} = (Y_{i2}, A_i, W_i, Z_i, X_i)$ for $i \in N_n$, adapted to a network $\mathcal{G}_n$, where ν is the length of the vector $L_{n,i}$. Similar to Section 2, we define a moment estimating function with parameter $\theta = (\tau, \gamma)$.

$$m(L_{n,i}; \theta) = \begin{Bmatrix} \tau - \{h(W_i, a, X_i; \gamma) - h(W_i, a', X_i; \gamma)\} \\ \{Y_{i2} - h(W_i, A_i, X_i; \gamma)\} \times \eta(A_i, Z_i, X_i) \end{Bmatrix}.$$

The GMM estimator for θ is

$$\widehat{\theta} = \underset{\theta}{\operatorname{argmin}}\, \overline{m}(\theta)^{\top} \Omega\, \overline{m}(\theta), \tag{14}$$

where $\overline{m}(\theta) = \frac{1}{n}\sum_{i=1}^{n} m(L_{n,i}; \theta)$ and Ω is a user-specified positive-definite weight matrix. Therefore,

the proposed DNC estimator $\widehat{\tau}(a, a')$ for $\tau(a, a')$ is the first element of $\widehat{\theta}$ defined in equation (14). Asymptotic results we derive below hold for the two-step GMM or other alternative GMM estimators, too.

Because we consider a sample of arbitrarily interconnected units in a network, the assumption that $m(L_{n,i}; \theta)$ is independently and identically distributed is unrealistic. Below, we consider assumptions on the observed data law that are considerably weaker, but still allow for valid inferences about the ACPE in network settings. For ease of exposition, we consider a setting in which the expected causal peer effect, $\mathbb{E}\{Y_{i2}(a) - Y_{i2}(a')\}$, is constant across units, while otherwise allowing for network-dependent (i.e. non-independent) errors. In online supplementary Appendix D.5, we extend our results to more general settings of heterogeneous ACPE.

We define ψ-network dependence (Kojevnikov et al., 2021), which encodes the degree of stochastic dependence between variables in terms of network distance.

Definition 1 (ψ-Network dependence—Kojevnikov et al., 2021). The triangular array $\{V_{n,i}\}_{i\in N_n}$, $n \geq 1$, $V_{n,i} \in \mathbf{R}^v$, is called conditionally ψ-weakly dependent given $\mathcal{G}_n$, if for each $n \in \mathbb{N}$, there exist a $\mathcal{G}_n$-measurable sequence $\beta_n = \{\beta_{n,s}\}$, $s \geq 0$, $\beta_{n,0} = 1$, and a collection of non-random functions $(\psi_{q_1,q_2})_{q_1,q_2\in\mathbb{N}}$, $\psi_{q_1,q_2}:\mathcal{L}_{v,q_1} \times \mathcal{L}_{v,q_2} \to [0, \infty)$, such that for all $(Q_1, Q_2) \in \mathcal{P}_n(q_1, q_2; s)$ with $s > 0$ and all $f_1 \in \mathcal{L}_{v,q_1}$ and $f_2 \in \mathcal{L}_{v,q_2}$,[1]

$$|\mathrm{Cov}(f_1(V_{n,Q_1}), f_2(V_{n,Q_2}) \mid \mathcal{G}_n)| \leq \psi_{q_1,q_2}(f_1, f_2)\beta_{n,s} \quad \text{a.s.} \tag{15}$$

We call the sequence $\beta_n = \{\beta_{n,s}\}_{s=1}^{\infty}$ the weak-dependent coefficients of $\{V_{n,i}\}_{i\in N_n}$.

Coefficient $\beta_{n,s}$ captures network dependence between units that are at network distance greater than or equal to s in network $\mathcal{G}_n$ in terms of the covariance of non-linearly transformed variables. Thus, a sequence of coefficients $\beta_n = \{\beta_{n,s}\}_{s=1}^{\infty}$ captures how fast network dependence between units decays as network distance s increases. Assumption 3, which we will introduce next, restricts the rate by which this network dependence decays. Importantly, ψ-network dependence permits dependence between units i and j that are only indirectly connected in the network, and thus, any two units can be dependent as long as there is a network path between them. This is in contrast to two other popular approaches; (a) dependency graphs (e.g. Stein, 1972), which can allow units to be dependent only when they are adjacent in a given network and (b) Markov random fields (e.g. Besag, 1974; Tchetgen Tchetgen, Fulcher, et al., 2021), which impose conditional independence restrictions based on the network structure (e.g. a given unit's observed data are independent of data observed for all units in the network conditional on observed data for its first-order network peers).

Using the notion of ψ-network dependence, we make the following assumptions about the observed data distribution that accommodate relatively strong network dependence, but still allow for making valid inferences about the ACPE.

Assumption 3 The triangular arrays $\{m(L_{n,i}; \theta)\}_{i\in N_n, n\geq 1}$ and $\{\frac{\partial}{\partial\theta} m(L_{n,i}; \theta)\}_{i\in N_n, n\geq 1}$ are conditionally ψ-weakly dependent given $\mathcal{G}_n$, respectively, for all $\theta \in \Theta$ with the weak dependent coefficients β_n that satisfy the following conditions, $\sup_{n\geq 1} \max_{s\geq 1} \beta_{n,s} < \infty$, and for some constant λ, $\psi_{q_1,q_2}(f_1, f_2) \leq \lambda \times q_1 q_2 (\|f_1\|_\infty + \mathrm{Lip}(f_1))(\|f_2\|_\infty + \mathrm{Lip}(f_2))$, where $\mathrm{Lip}(f)$ represents Lipschitz constant of f, and $\|\cdot\|_\infty$ denotes the sup norm, i.e. $\|f\|_\infty = \sup|f(x)|$. There exist $p > 4$ and a sequence $\xi_n \to \infty$ such that

1. $\beta_{n,\xi_n}^{(p-1)/p} = o_{a.s.}(n^{-3/2})$,

[1] $\mathcal{P}_n(q_1, q_2; s)$ denotes the collection of two sets of nodes of size q_1 and q_2 with distance between each other of at least s. Formally, $\mathcal{P}(q_1, q_2; s) = \{(Q_1, Q_2) : Q_1, Q_2 \subset N_n, |Q_1| = q_1, |Q_2| = q_2, \text{ and } d_n(Q_1, Q_2) \geq s\}$. $\mathcal{L}_{v,q_1}$ and $\mathcal{L}_{v,q_2}$ denote the collection of bounded Lipschitz real functions on $\mathbf{R}^{v\times q_1}$ and $\mathbf{R}^{v\times q_2}$, respectively.

2. for each $k \in \{1, 2\}$,

$$\frac{1}{n^{k/2}} \sum_{s \geq 0} r_n(s, \xi_n; k) \beta_{n,s}^{1-(k+2)/p} = o_{a.s.}(1),$$

where we define the neighbourhood shell:

$$\begin{aligned} r_n(s, \xi_n; k) = \inf_{a>1} \left\{ \frac{1}{n} \sum_{i \in N_n} |\mathcal{N}_n(i; s)|^{\frac{a}{a-1}} \right\}^{1-\frac{1}{a}} \\ \times \left\{ \frac{1}{n} \sum_{i \in N_n} \max_{j \in \mathcal{N}_n(i;s)} |\widetilde{\mathcal{N}}_n(i; \xi_n) \setminus \widetilde{\mathcal{N}}_n(j; s-1)|^{ka} \right\}^{\frac{1}{a}} \end{aligned} \tag{16}$$

and the within-s peers $\widetilde{\mathcal{N}}_n(i; s) = \{j \in N_n : d_n(i, j) \leq s\}$. We use $|\cdot|$ to denote the cardinality of a set.

This is an adaptation of Condition ND in Kojevnikov et al., (2021) to our set-up. Assumption 3(1) restricts the speed by which weak-dependent coefficients $\beta_{n,s}$ decay as network distance s increases. The first term of r_n captures the number of peers at the network distance s and the second term captures the extent to which edges in the network tend to cluster together. Therefore, Assumption 3(2) restricts the speed by which the density of the network changes as sample size increases. When network dependence $\beta_{n,s}$ decays faster with network distance s, it can accommodate denser networks. See Kojevnikov et al., (2021) for further discussion on these conditions.

We are now ready to state asymptotic properties of the proposed DNC estimator.

Theorem 3 Under the conditions given in Theorem 2, Assumption 3 and standard GMM regularity conditions,[2] as n goes to infinity, the DNC estimator is consistent:

$$\widehat{\tau}(a, a') \xrightarrow{p} \tau(a, a'),$$

and asymptotically normal:

$$\frac{\widehat{\tau}(a, a') - \tau(a, a')}{\sqrt{\sigma^2/n}} \xrightarrow{d} \text{Normal}(0, 1),$$

where $\xrightarrow{p}$ denotes convergence in probability and $\xrightarrow{d}$ denotes convergence in distribution. The asymptotic variance σ^2 is the (1, 1)th element of matrix Σ where

$$\Sigma = \Gamma_0 \Lambda_0 \Gamma_0^\top, \ \Lambda_0 = \text{Var}\left(\frac{1}{\sqrt{n}} \sum_{i=1}^{n} m(L_{n,i}; \theta_0) \right),$$

$$\Gamma_0 = (M_0^\top \Omega M_0)^{-1} M_0^\top \Omega, \ M_0 = \frac{1}{n} \sum_{i=1}^{n} \mathbb{E}\left\{ \frac{\partial}{\partial \theta} m(L_{n,i}; \theta_0) \right\},$$

and we define θ_0 to be the true parameter such that, for all units $i \in N_n$, $\mathbb{E}\{m(L_{n,i}; \theta)\} = 0$ only when $\theta = \theta_0$.

[2] In online supplementary Appendix D, we provide the regularity conditions widely used in the GMM framework (Hansen, 1982; Newey & McFadden, 1994).

We provide a proof in online supplementary Appendix D. We also provide simulation studies in online supplementary Appendix E.

To estimate the standard error of the DNC estimator, the key is to estimate Λ_0. To account for network-dependent errors, we rely on the network HAC variance estimator (Kojevnikov et al., 2021) adapted to our setting:

$$\widehat{\Lambda} = \sum_{s \geq 0} \omega(s/b_n) \left\{ \frac{1}{n} \sum_{i \in N_n} \sum_{j \in \mathcal{N}_n(i;s)} m(L_{n,i}; \widehat{\theta}) m(L_{n,j}; \widehat{\theta})^\top \right\}, \tag{17}$$

where a kernel function $\omega(\cdot)$ is defined as follows: $\omega : \overline{\mathbb{R}} \to [-1, 1]$ such that $\omega(0) = 1$, $\omega(c) = 0$ for $|c| \geq 1$ and $\omega(c) = \omega(-c)$ for all $c \in \overline{\mathbb{R}}$. Examples include the truncated, Parzen, and Tukey-Hanning kernels. b_n denotes a bandwidth of the network HAC variance estimator. This bandwidth determines the extent to which $\widehat{\Lambda}$ takes into account long range network dependence; kernel weight $\omega(s/b_n) > 0$ for $s < b_n$ and $\omega(s/b_n) = 0$ for $s \geq b_n$.

The variance estimator of the DNC estimator can be computed as the (1, 1)th element of matrix $\widehat{\Sigma}$ defined as

$$\widehat{\Sigma} = \widehat{\Gamma}\widehat{\Lambda}\widehat{\Gamma}^\top, \tag{18}$$

where $\widehat{\Gamma} = (\widehat{M}^\top \Omega \widehat{M})^{-1} \widehat{M}^\top \Omega$, and $\widehat{M} = \frac{1}{n} \sum_{i=1}^n \frac{\partial}{\partial \theta} m(L_{n,i}; \widehat{\theta})$.

We now establish consistency of this network HAC variance estimator. The following result formally restricts the speed by which bandwidth b_n should increase as sample size increases. When network dependence $\beta_{n,s}$ decays slower and the average number of network peers at distance s, captured by $\rho_n(s)$, increases faster with network distance, bandwidth b_n should increase faster as sample size increases.

Theorem 4 Under the conditions given in Theorem 3, suppose the choice of bandwidth and kernel satisfies the following condition with p that satisfies Assumption 3.

$$\lim_{n \to \infty} \sum_{s \geq 0} |\omega(s/b_n) - 1| \rho_n(s) \beta_{n,s}^{1-2/p} = 0 \quad \text{a.s.}, \tag{19}$$

where $\rho_n(s) = \frac{1}{n} \sum_{i=1}^n |\mathcal{N}_n(i; s)|$. Then, $\widehat{\Sigma} \xrightarrow{p} \Sigma$.

We provide a proof in online supplementary Appendix D.

In practice, we recommend using the default choice of bandwidth b_n provided by Kojevnikov et al., (2021), i.e.

$$b_n = \text{constant} \times \frac{\log(n)}{\log\{\max(\text{average degree}, 1.05)\}}. \tag{20}$$

In our simulation studies (online supplementary Appendix E), we set the constant in equation (20) to 1.0 and find this default choice performs well across various settings.

Finally, under conditions given in Theorem 4, we obtain an asymptotically valid $(1 - \alpha)$ confidence interval for $\tau(a, a')$ by $[\widehat{\tau}(a, a') - \Phi(1 - \alpha/2) \times \widehat{\sigma}/\sqrt{n},\ \widehat{\tau}(a, a') + \Phi(1 - \alpha/2) \times \widehat{\sigma}/\sqrt{n}]$ where $\widehat{\sigma}$ is the square root of the (1, 1)th element of matrix $\widehat{\Sigma}$ defined in equation (18) and $\Phi(\cdot)$ denotes the quantile function for the standard normal distribution.

3.6 Linear DNC estimator

While the theoretical results in the previous section apply to any parametric and semi-parametric confounding bridge functions, here we discuss an important special case under a linear specification for the confounding bridge function, which admits a closed-form solution. Suppose we

assume a linear confounding bridge:

$$h(W_i, A_i, X_i; \gamma) = \gamma_\alpha + \gamma_A A_i + \gamma_W W_i + \gamma_X^\top X_i.$$

Under this linear model, $\tau(a, a') = \gamma_A \times (a - a')$. We can estimate coefficients $\gamma = (\gamma_\alpha, \gamma_A, \gamma_W, \gamma_X^\top)^\top$ by fitting the linear GMM estimator:

$$\widehat{\gamma} = (\mathbf{V}_W^\top \mathbf{V}_Z \Omega \mathbf{V}_Z^\top \mathbf{V}_W)^{-1} \mathbf{V}_W^\top \mathbf{V}_Z \Omega \mathbf{V}_Z^\top \mathbf{Y}, \tag{21}$$

where $\mathbf{V}_W$ is a matrix with n rows with ith row $V_{iW} = (1, A_i, W_i, X_i^\top)^\top$, $\mathbf{V}_Z$ is a matrix with n rows with ith row $V_{iZ} = (1, A_i, Z_i^\top, X_i^\top)^\top$, and $\mathbf{Y}$ is a n-dimensional vector with ith element equal to Y_{i2}.

To account for network dependence, we adopt the network HAC variance estimator in Section 3.5. The key is to estimate

$$\Lambda_0^{\text{lin}} = \text{Var}\left(\frac{1}{\sqrt{n}} \sum_{i=1}^{n} e_{i2} V_{iZ}\right) \quad \text{and} \quad e_{i2} = Y_{i2} - \gamma^\top V_{iW}. \tag{22}$$

Using the network HAC variance estimator, we can estimate the variance of $\widehat{\gamma}$ as

$$\widehat{\text{Var}}(\widehat{\gamma}) = n(\mathbf{V}_W^\top \mathbf{V}_Z \Omega \mathbf{V}_Z^\top \mathbf{V}_W)^{-1} \mathbf{V}_W^\top \mathbf{V}_Z \Omega \widehat{\Lambda}^{\text{lin}} \Omega \mathbf{V}_Z^\top \mathbf{V}_W (\mathbf{V}_W^\top \mathbf{V}_Z \Omega \mathbf{V}_Z^\top \mathbf{V}_W)^{-1},$$

where

$$\widehat{\Lambda}^{\text{lin}} = \sum_{s \geq 0} \omega(s/b_n) \left\{ \frac{1}{n} \sum_{i \in N_n} \sum_{j \in \mathcal{N}_n(i;s)} \widehat{e}_{i2} \widehat{e}_{j2} V_{iZ} V_{jZ}^\top \right\}, \quad \text{and} \quad \widehat{e}_{i2} = Y_{i2} - \widehat{\gamma}^\top V_{iW}.$$

Remark In general settings of network-dependent errors (Section 3.5), one must estimate bandwidth b_n for the network HAC variance estimator because how far network dependence persists is a priori unknown. However, in some settings, we prove that one can analytically select the bandwidth. We provide technical details as Lemma 7 in online supplementary Appendix D.4.

Remark When specifying the functional form of the outcome confounding bridge function, it is useful to recall that the outcome confounding bridge function is defined such that the confounding effect of U_i on the primary outcome Y_{i2} is equal to the confounding effect of U_i on $h(W_i, a, X_i)$. Therefore, especially when NCO W is the pre-treatment outcome variable (e.g. examples in Section 3.3.2 and the empirical application in Section 4), it is intuitive to specify a model for $h(\cdot)$ as they are on the same scale and $h(\cdot)$ should capture how much the confounding effect of U on the focal behaviour changes over time (i.e. effects on Y_{i2} and Y_{i1}).

Existing methods use parametric or semi-parametric estimation under an assumption of no uncontrolled network confounding (e.g. Forastiere et al., 2020; Ogburn et al., 2017; Tchetgen Tchetgen, Fulcher, et al., 2021; van der Laan, 2014). We follow the existing literature and use parametric or semi-parametric estimation in this paper in order to consider a more challenging scenario with unmeasured network confounding. However, the potential misspecification of the functional form for the outcome confounding bridge is an important practical concern. Without additional functional form assumptions, arbitrary misspecification of $h(\cdot)$ can induce unbounded bias.

In practice, researchers can mitigate concerns about functional form misspecification by considering more flexible estimation of the outcome confounding bridge function. In the i.i.d. data setting, an increasing number of papers develop non-parametric estimation and inference methods for the DNC approach by incorporating modern machine learning estimators (Ghassami et al., 2022; Kallus et al., 2021; Mastouri et al., 2021; Singh, 2020). The mathematical structure is also similar to the non-parametric IV problem (Canay et al., 2013; Newey, 2013; Newey & Powell, 2003), which also recently incorporates various machine learning estimators (e.g. Bennett et al., 2023; Lewis & Syrgkanis, 2018) in the i.i.d. data setting. Extending those non-parametric estimation to the DNC approach in the network data setting (by incorporating network-dependent errors) is a fruitful direction for future work.

Finally, when there are multiple NCEs, researchers can also conduct an overidentification test within the GMM framework to test misspecification of the functional form of $h()$. It is important to note that selected NCE variables might be invalid in practice, and thus, such a test should be seen as an omnibus test of the negative control assumption and model misspecification of the outcome bridge function. Rejection in such a test should be interpreted as global alarm to misspecification.

4 Empirical application: causal peer effects in education

We apply our method to Add Health data to evaluate causal peer effects of education outcomes in a friendships network. The study of causal peer effects in education has a long history in the social sciences, and many studies have shown moderate positive peer effects (e.g. Sacerdote, 2011). While some papers have used experimental or quasi-experimental methods where classmates or roommates are randomly assigned by schools, the vast majority of existing evidence comes from observational studies of causal peer effects. In the absence of randomisation, researchers have adjusted for a variety of observed pre-treatment covariates and a host of fixed effects (e.g. fixed effects for schools or network components). However, such approaches rely on a conditional ignorability assumption, and assume away unmeasured latent homophily or contextual confounding. For example, students who have higher education performance might become friends with other high-performing students due to unobserved characteristics. In this case, strong association between one's education outcome and her friends' outcomes cannot be interpreted as the causal peer effect. Potential bias due to such unmeasured network confounding can undermine the validity of causal conclusions. To explore the possibility of unmeasured network confounding, we use the proposed DNC approach to identify and estimate the ACPE.

4.1 Data

Add Health project collected survey data from students in grades 7–12 from a nationally representative sample of over 100 private and public schools in years 1994–1995 in the United States. There were two types of surveys conducted in years 1994–1995, both of which provide information about social and demographic characteristics of respondents, education level and occupation of their parents, and their friendship links (i.e. their best friends, up to five females and up to five males). The first type was an in-school survey, which were administered to students in schools from September 1994 until April 1995. It includes over 90,000 students across about 140 schools. Each school administration occurred on a single day within one 45- to 60-min class period. The second type was an in-home survey, in which students answered a 90-min in-home interview. It includes over 20,000 students. The in-home interview was conducted at least 90 days after the in-school survey, except for 10 students who we exclude from analysis. Our analysis focuses on 10,264 students who completed both in-school and in-home surveys and answered questions related to variables we use below. The in-school survey serves as the baseline time period and the in-home survey as the follow-up period.

We examine a network based on the friendship information collected in the in-school survey. We define friendships as symmetric relationships: the pair of students i and j in the same school

are coded as friends if either i lists j as a friend, or j lists i as a friend, or both. The average degree of the friendship network is 6.20.

4.2 Set-up

We use the GPA at follow-up period Y_{i2} as the outcome variable. GPA ranges between 1 and 4, and is computed based on the average grade-point of four subjects; English, Mathematics, History/Social Studies, and Science. The treatment variable is the average GPA of the network peers at baseline $A_i = \sum_{j \in \mathcal{N}(i;1)} Y_{j1}/|\mathcal{N}(i;1)|$, where $|\mathcal{N}(i;s)|$ denotes the number of the sth order network peers of unit i.

Following Section 3.3, we consider two sets of negative controls. First, we consider focal behaviours of peers and those measured at baseline as negative controls. In particular, we use the GPA at baseline Y_{i1} as the NCO, i.e. $W_i = Y_{i1}$. We then use the average GPA of peers-of-peers at baseline as the NCE, i.e. $Z_i = \sum_{j \in \mathcal{N}(i;2)} Y_{j1}/|\mathcal{N}(i;2)|$. These are valid negative controls when the GPAs of students are recorded concurrently within schools. Second, we also use an auxiliary variable for negative controls. To make the assumption about the confounding bridge function (Assumption 2(3)) more plausible, we again use the GPA at baseline as the NCO, i.e. $W_i = Y_{i1}$, while we use level of peers' headaches as the NCE. Specifically, we use the average level of headaches of peers as NCE, i.e. $\widetilde{Z}_i = \sum_{j \in \mathcal{N}(i;1)} C_j/|\mathcal{N}(i;1)|$, where C_j captures level of headaches by student j measured at baseline. According to Cohen-Cole and Fletcher (2008), whether a student has headaches is a plausible negative control because it is unlikely (a) to causally affect whether students are friends to each other and (b) to causally affect headaches of other units.

In accordance with prior analyses of these data in the literature, we also include a series of observed pre-treatment covariates X_i, including fixed effects for each school, the number of friends, age, gender, race, born in the United States, health status, physical fitness, school attendance, student grade, motivation in education, school attachment, homework, self-esteem, household size, living with mother, living with father, parental care, mother's education, father's education, whether both parents work, relationship with teachers, and relationship with friends.

Following Section 3.6, we use a linear DNC estimator. In particular, we specify a linear confounding bridge as,

$$h(W_i, A_i, X_i; \gamma) = \gamma_\alpha + \gamma_A A_i + \gamma_W W_i + \gamma_X^\top X_i,$$

where (A_i, W_i) are the treatment variable and the NCO, and X_i represent observed pre-treatment covariates defined above, including fixed effects for schools.

Our causal estimand is γ_A, which captures the causal effect on a given student's GPA induced by one point increase in the average GPA of her peers. We use the proposed network HAC variance estimator to compute standard errors and 95% confidence intervals. We report standard errors based on the default choice of the bandwidth, while results for the analytical bandwidth selection are similar and therefore not reported.

We compare our proposed DNC estimator to the ordinary least squares (OLS) regression estimator, which relies on the assumption of conditional ignorability where we adjust for pre-treatment outcome Y_{i1} and X_i. In absence of unmeasured network confounding, we expect DNC estimates and OLS estimates to be comparable (i.e. within sampling variability). For the OLS estimator, we apply the network HAC variance estimator (Kojevnikov et al., 2021) to residuals in order to make comparison clear. We note that the estimated standard errors for the OLS are valid only when there is no unmeasured network confounding.

4.3 Results

Table 1 reports estimates of the ACPE, standard errors, p-values, and 95% confidence intervals for each method. The OLS estimate suggests that the estimated ACPE is as large as 0.176 and statistically significant. However, our proposed DNC estimator indicates that there may be a large amount of unmeasured network confounding operating in this network which cannot be accounted for by a standard regression analysis, despite having accounted for a large number of pre-treatment covariates. The DNC estimate based on NCE Z_i is 0.033, is less than 20% of the OLS estimate, and is not statistically significant. We can empirically assess the negative control

Table 1. Estimated average causal peer effects on GPA

Method	Estimate	Stand. error	p-Value	95% CI
OLS under conditional ignorability	0.176	0.015	.000	(0.147, 0.206)
DNC estimator with NCE Z_i: the average GPA of $\mathcal{N}(i; 2)$	0.033	0.049	.497	(−0.063, 0.129)
DNC estimator with NCE $\widetilde{Z}_i$: the average level of headaches of $\mathcal{N}(i; 1)$	0.078	0.183	.668	(−0.280, 0.437)

Note. DNC = double negative control; GPA = grade-point average; NCE = negative control exposure; OLS = ordinary least squares.

relevance assumption with the following test: we first regress the NCO on the NCE, treatment, and observed pre-treatment covariates, and then test the null hypothesis that coefficients of the NCE are zero. We find that the F-statistic is 81.44 and the p-value is .00, suggesting the negative control relevance holds.

The DNC estimate based on NCE $\widetilde{Z}_i$ shows a similar pattern: a point estimate is 0.078 and is not statistically significant. We find that a test of the negative control relevance yields the F-statistic 4.76 and the p-value is .03, suggesting the association between the NCO and NCE is weaker for this second NCE. This is why the standard error based on the second NCE $\widetilde{Z}_i$ is larger than the one based on the first NCE (Miao, Shi, et al., 2018). These results show that the OLS estimator under conditional ignorability can suffer from more than 100% bias.

5 Concluding remarks

In this article, we have developed the DNC approach to identification and estimation of causal peer effects. In contrast to existing literature, we take into account both unmeasured network confounding and network dependence of observations. We then provide a GMM estimator for the ACPE and establish its consistency and asymptotic normality under conditions of ψ-network dependence. We also derive the network HAC variance estimator, with which researchers can construct asymptotic confidence intervals.

While this paper has focused on the causal peer effect, we can straightforwardly extend our results to a related but different causal quantity of interest, the spillover effect (Ogburn & VanderWeele, 2014). The causal peer effect is defined as the causal effect of peers' focal behaviours on an ego's focal behaviour (e.g. the GPA of peers serves as the treatment and ego's GPA as the outcome), whereas to define the spillover effect, one has a treatment variable defined separately from the focal behaviour of interest (e.g. scholarships received by peers as the treatment and ego's GPA as the outcome). See Ogburn and VanderWeele (2014) and online supplementary Appendix F for further discussion on the difference between causal peer effects and spillover effects. Despite this difference, by selecting appropriate negative controls, researchers can rely on the same proposed DNC approach to identify and estimate spillover effects in the presence of unmeasured confounding in observational studies.

Our findings have established a theoretical basis for future research on non-parametric estimation of causal peer effects with DNCs for unmeasured network confounding. This will be able to extend previous studies that consider more flexible semi-parametric and non-parametric estimation for DNC adjustment in i.i.d. settings (e.g. Ghassami et al., 2022; Kallus et al., 2021; Shi et al., 2020) and methods that examine network effects without unmeasured network confounding (e.g. Forastiere et al., 2020; Ogburn et al., 2017; Tchetgen Tchetgen, Fulcher, et al., 2021).

Conflict of interest: None declared.

Funding

The research reported in this publication was supported by the National Institutes of Health (award R01AI27271, R01CA222147, R01GM139926, R01AG065276, to E.J.T.T.) and the

Department of Political Science at Columbia University (for N.E.). This research was approved by the Columbia University Institutional Review Board (#AAAT3813).

Data availability

This research uses data from Add Health, a programme project designed by J. Richard Udry, Peter S. Bearman, and Kathleen Mullan Harris, and funded by a grant P01-HD31921 from the Eunice Kennedy Shriver National Institute of Child Health and Human Development, with cooperative funding from 17 other agencies. Special acknowledgement is due to Ronald R. Rindfuss and Barbara Entwisle for assistance in the original design. Persons interested in obtaining Data Files from Add Health should contact Add Health, The University of North Carolina at Chapel Hill, Carolina Population Center, Carolina Square, Suite 210, 123 W. Franklin Street, Chapel Hill, NC 27516 (addhealth_contracts@unc.edu). No direct support was received from grant P01-HD31921 for this analysis.

Supplementary material

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

References

Andrews I., Stock J. H., & Sun L. (2019). Weak instruments in instrumental variables regression: Theory and practice. *Annual Review of Economics*, *11*(1), 727–753. https://doi.org/10.1146/economics.2019.11.issue-1

Angrist J. D. (2014). The perils of peer effects. *Labour Economics*, *30*, 98–108. https://doi.org/10.1016/j.labeco.2014.05.008

Aronow, P. M., & Samii, C. (2017). Estimating average causal effects under general interference, with application to a social network experiment. *The Annals of Applied Statistics*, *11*(4), 1912–1947. https://doi.org/10.1214/16-AOAS1005

Bennett A., Kallus N., Mao X., Newey W., Syrgkanis V., & Uehara M. (2023). 'Minimax instrumental variable regression and L_2 convergence guarantees without identification or closedness', arXiv, arXiv:2302.05404, preprint: not peer reviewed.

Besag, J. (1974). Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *36*, 192–225. https://doi.org/10.1111/j.2517-6161.1974.tb00999.x

Bramoullé Y., Djebbari H., & Fortin B. (2009). Identification of peer effects through social networks. *Journal of Econometrics*, *150*(1), 41–55. https://doi.org/10.1016/j.jeconom.2008.12.021

Canay I. A., Santos A., & Shaikh A. M. (2013). On the testability of identification in some nonparametric models with endogeneity. *Econometrica*, *81*(6), 2535–2559. https://doi.org/10.3982/ECTA10851

Christakis N. A., & Fowler J. H. (2007). The spread of obesity in a large social network over 32 years. *New England Journal of Medicine*, *357*(4), 370–379. https://doi.org/10.1056/NEJMsa066082

Cobzaru R., Welsch R., Finkelstein S., Ng K., & Shahn Z. (2022). 'Bias formulas for violations of proximal identification assumptions', arXiv, arXiv:2208.00105, preprint: not peer reviewed.

Cohen-Cole E., & Fletcher J. M. (2008). Detecting implausible social network effects in acne, height, and headaches: Longitudinal analysis. *British Medical Journal*, *337*, a2533. https://doi.org/10.1136/bmj.a2533

Deaner B. (2018). 'Proxy controls and panel data', arXiv, arXiv:1810.00283, preprint: not peer reviewed.

Egami N. (2018). 'Identification of causal diffusion effects under structural stationarity', arXiv, arXiv:1810.07858, preprint: not peer reviewed.

Egami N. (2021). Spillover effects in the presence of unobserved networks. *Political Analysis*, *29*(3), 287–316. https://doi.org/10.1017/pan.2020.28

Forastiere, L., Airoldi, E. M., & Mealli, F. (2020). Identification and estimation of treatment and interference effects in observational studies on networks. *Journal of the American Statistical Association*, *116*, 1–18. https://doi.org/10.1080/01621459.2020.1768100

Ghassami A., Ying A., Shpitser I., & Tchetgen Tchetgen E. (2022). Minimax kernel machine learning for a class of doubly robust functionals with application to proximal causal inference. In *International Conference on Artificial Intelligence and Statistics* (pp. 7210–7239). PMLR.

Goldsmith-Pinkham P., & Imbens G. W. (2013). Social networks and the identification of peer effects. *Journal of Business & Economic Statistics*, *31*(3), 253–264. https://doi.org/10.1080/07350015.2013.801251

Hansen L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica*, *50*(4), 1029–1054. https://doi.org/10.2307/1912775

Kallus N., Mao X., & Uehara M. (2021). 'Causal inference under unmeasured confounding with negative controls: A minimax learning approach', arXiv, arXiv:2103.14029, preprint: not peer reviewed.

Kojevnikov, D., Marmer, V., & Song, K. (2021). Limit theorems for network dependent random variables. *Journal of Econometrics*, *222*(2), 882–908. https://doi.org/10.1016/j.jeconom.2020.05.019
Kress R. (1989). *Linear integral equations*. (Vol. *82*). Springer.
Kuroki M., & Pearl J. (2014). Measurement bias and effect restoration in causal inference. *Biometrika*, *101*(2), 423–437. https://doi.org/10.1093/biomet/ast066
Lee Y. & Ogburn E. L. (2021) Network dependence can lead to spurious associations and invalid inference. *Journal of the American Statistical Association*, *116*(535), 1–15. https://doi.org/10.1080/01621459.2020.1782219
Leung M. P. (2021). 'Network cluster-robust inference', arXiv, arXiv:2103.01470, preprint: not peer reviewed.
Lewis G., & Syrgkanis V. (2018). 'Adversarial generalized method of moments', arXiv, arXiv:1803.07164, preprint: not peer reviewed.
Lipsitch M., Tchetgen Tchetgen E. J., & Cohen T. (2010). Negative controls: A tool for detecting confounding and bias in observational studies. *Epidemiology*, *21*(3), 383–388. https://doi.org/10.1097/EDE.0b013e3181d61eeb
Liu L., & Tchetgen Tchetgen E. (2020). 'Regression-based negative control of homophily in dyadic peer effect analysis', arXiv, arXiv:2002.06521, preprint: not peer reviewed.
Lyons R. (2011). The spread of evidence-poor medicine via flawed social-network analysis. *Statistics, Politics, and Policy*, *2*(1). https://doi.org/10.2202/2151-7509.1024
Manski C. F. (1993). Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies*, *60*(3), 531–542. https://doi.org/10.2307/2298123
Mastouri A., Zhu Y., Gultchin L., Korba A., Silva R., Kusner M., Gretton A., & Muandet K. (2021). Proximal causal learning with kernels: Two-stage estimation and moment restriction. In *International Conference on Machine Learning* (pp. 7512–7523). PMLR.
McFowland E., III & Shalizi C. R. (2023). Estimating causal peer influence in homophilous social networks by inferring latent locations. *Journal of the American Statistical Association*, *118*(541), 707–718. https://doi.org/10.1080/01621459.2021.1953506
Miao W., Geng Z., & Tchetgen Tchetgen E. J. (2018). Identifying causal effects with proxy variables of an unmeasured confounder. *Biometrika*, *105*(4), 987–993. https://doi.org/10.1093/biomet/asy038
Miao W., Shi X., & Tchetgen Tchetgen E. (2018). 'A confounding bridge approach for double negative control inference on causal effects', arXiv, arXiv:1808.04945, preprint: not peer reviewed.
Newey W. K. (2013). Nonparametric instrumental variables estimation. *American Economic Review*, *103*(3), 550–556. https://doi.org/10.1257/aer.103.3.550
Newey W. K., & McFadden D. (1994). Large sample estimation and hypothesis. In R. Engle & D. McFadden (Eds.), *Handbook of econometrics* (pp. 2112–2245). North Holland.
Newey W. K., & Powell J. L. (2003). Instrumental variable estimation of nonparametric models. *Econometrica*, *71*(5), 1565–1578. https://doi.org/10.1111/ecta.2003.71.issue-5
Neyman J. (1923). On the application of probability theory to agricultural experiments. Essay on principles (with discussion). Section 9 (translated). *Statistical Science*, *5*, 465–472. https://doi.org/10.1214/ss/1177012031
Ogburn E. L. (2018). Challenges to estimating contagion effects from observational data. In S. Lehmann & Y.-Y. Ahn (Eds.), *Complex spreading phenomena in social systems* (pp. 47–64). Springer.
Ogburn E. L., Sofrygin O., Diaz I., & van der Laan M. J. (2017). 'Causal inference for social network data', arXiv, arXiv:1705.08527, preprint: not peer reviewed.
Ogburn E. L., & VanderWeele T. J. (2014). Causal diagrams for interference. *Statistical Science*, *29*(4), 559–578. https://doi.org/10.1214/14-STS501
O'Malley A. J., Elwert F., Rosenquist J. N., Zaslavsky A. M., & Christakis N. A. (2014). Estimating peer effects in longitudinal dyadic data using instrumental variables. *Biometrics*, *70*(3), 506–515. https://doi.org/10.1111/biom.v70.3
Robins J. (1986). A new approach to causal inference in mortality studies with a sustained exposure period—Application to control of the healthy worker survivor effect. *Mathematical Modelling*, *7*(9–12), 1393–1512. https://doi.org/10.1016/0270-0255(86)90088-6
Rosenbaum P. R., & Rubin D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, *70*(1), 41–55. https://doi.org/10.1093/biomet/70.1.41
Rubin D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, *66*(5), 688–701. https://doi.org/10.1037/h0037350
Sacerdote B. (2011). Peer effects in education: How might they work, how big are they and how much do we know thus far? In E. A. Hanushek, S. Machin, & L. Woessmann (Eds.), *Handbook of the economics of education* (Vol. 3, pp. 249–277). Elsevier.
Sävje F. (2021). 'Causal inference with misspecified exposure mappings', arXiv, arXiv:2103.06471, preprint: not peer reviewed.

Shalizi C. R., & Thomas A. C. (2011). Homophily and contagion are generically confounded in observational social network studies. *Sociological Methods & Research*, *40*(2), 211–239. https://doi.org/10.1177/0049124111404820

Shi X., Miao W., Nelson J. C., & Tchetgen Tchetgen E. J. (2020). Multiply robust causal inference with double-negative control adjustment for categorical unmeasured confounding. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *82*(2), 521–540. https://doi.org/10.1111/rssb.12361

Singh R. (2020). 'Kernel methods for unobserved confounding: Negative controls, proxies, and instruments', arXiv, arXiv:2012.10315, preprint: not peer reviewed.

Sofer T., Richardson D. B., Colicino E., Schwartz J., & Tchetgen Tchetgen E. J. (2016). On negative outcome control of unobserved confounding as a generalization of difference-in-differences. *Statistical Science*, *31*(3), 348. https://doi.org/10.1214/16-STS558

Stein C. (1972). A bound for the error in the normal approximation to the distribution of a sum of dependent random variables. In *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability, Volume 2: Probability Theory* (Vol. 6, pp. 583–603). University of California Press.

Tchetgen Tchetgen E. J., Fulcher I., & Shpitser I. (2021). Auto-g-computation of causal effects on a network. *Journal of the American Statistical Association*, *116*(534), 833-844. https://doi.org/10.1080/01621459.2020.1811098

Tchetgen Tchetgen E. J., Ying A., Cui Y., Shi X., & Miao W. (2020). 'An introduction to proximal causal learning', arXiv, arXiv:2009.10982, preprint: not peer reviewed.

van der Laan M. J. (2014). Causal inference for a population of causally connected units. *Journal of Causal Inference*, *2*(1), 13–74. https://doi.org/10.1515/jci-2013-0002

VanderWeele T. J., & An W. (2013). Social networks and causal inference. In S. L. Morgan (Ed.), *Handbook of causal analysis for social research* (pp. 353–374). Springer.

Wright J. H. (2003). Detecting lack of identification in GMM. *Econometric Theory*, *19*(2), 322–330. https://doi.org/10.1017/S0266466603192055