The Annals of Applied Statistics
2017, Vol. 11, No. 2, 919–948
DOI: 10.1214/17-AOAS1023

VACCINES, CONTAGION, AND SOCIAL NETWORKS

BY ELIZABETH L. OGBURN[1] AND TYLER J. VANDERWEELE[2]

Johns Hopkins School of Public Health and Harvard School of Public Health

Consider the causal effect that one individual's treatment may have on another individual's outcome when the outcome is contagious, with specific application to the effect of vaccination on an infectious disease outcome. The effect of one individual's vaccination on another's outcome can be decomposed into two different causal effects, called the "infectiousness" and "contagion" effects. We present identifying assumptions and estimation or testing procedures for infectiousness and contagion effects in two different settings: (1) using data sampled from independent groups of observations, and (2) using data collected from a single interdependent social network. The methods that we propose for social network data require fitting generalized linear models (GLMs). GLMs and other statistical models that require independence across subjects have been used widely to estimate causal effects in social network data, but because the subjects in networks are presumably not independent, the use of such models is generally invalid, resulting in inference that is expected to be anticonservative. We describe a subsampling scheme that ensures that GLM errors are uncorrelated across subjects despite the fact that outcomes are nonindependent. This simultaneously demonstrates the possibility of using GLMs and related statistical models for network data and highlights their limitations.

1. Introduction. We are concerned here with the effect that one individual's treatment may have on another individual's outcome, when the outcome is contagious. In the infectious disease literature, this is often called an *indirect effect* of treatment [Halloran and Struchiner (1991)], while the effect of an individual's treatment on his own outcome is a *direct effect*. Indirect effects of infectious disease interventions are of significant importance for understanding infectious disease dynamics and for designing public health interventions. For example, the goal of many vaccination programs is to achieve herd immunity, whereby a large enough subset of a population is vaccinated that even those individuals who remain unvaccinated are protected against infection. This is one type of an indirect effect of a vaccination program; it has been extensively studied in the infectious disease literature [Anderson and May (1985), Fine (1993), John and Samuel (2000), O'Brien and Dagan (2003)]. Recently, interest has turned toward the identification and estimation of average individual-level indirect effects

Received November 2014; revised February 2017.
[1]Supported by ONR Grant N00014-15-1-2343.
[2]Supported by NIH Grant ES017876.
Key words and phrases. Causal inference, social networks, contagion, peer effects.

[Halloran and Struchiner (1991, 1995), Halloran and Hudgens (2012), VanderWeele and Tchetgen Tchetgen (2011a, 2011b), VanderWeele, Tchetgen Tchetgen and Halloran (2012)], such as the effect on a single member of a community of two different vaccination programs implemented on the rest of the community [Halloran and Struchiner (1995)].

VanderWeele, Tchetgen Tchetgen and Halloran (2012) demonstrated that the individual-level indirect effect of vaccination in communities of size two can be decomposed into two different effects, called the "infectiousness" and "contagion" effects. These two effects represent distinct causal pathways by which one person's vaccination may affect another's disease status. The contagion effect is the indirect effect that vaccinating one individual may have on another by preventing the vaccinated individual from getting the disease and thereby from passing it on. If an intervention to prevent infectious disease operates by reducing the susceptibility of treated individuals to the disease, thereby preventing them from becoming infected, it is operating via contagion effects. Examples of such interventions are vaccines for tetanus, hepatitis A and B, rabies, and measles [Keller and Stiehm (2000)]. The infectiousness effect is the indirect effect that vaccination might have if, instead of preventing the vaccinated individual from getting the disease, it renders the disease less infectious, thereby reducing the probability that the vaccinated infected individual transmits the disease, even if infected. The malaria transmission-blocking vaccine is designed to prevent mosquitos from acquiring, and thereby from transmitting, malaria parasites upon biting infected individuals [Halloran and Struchiner (1992)]. This vaccine has no protective effect for the vaccinated individual, but it renders vaccinated individuals less likely to transmit the disease. Therefore, any indirect effect of the malaria transmission-blocking vaccine is due entirely to an infectiousness effect. Many interventions have indirect effects that operate via both contagion and infectiousness effects, for example, any intervention that reduces susceptibility and also shortens the duration of disease among individuals who get the disease despite being treated would have both a contagion effect and an infectiousness effect.

VanderWeele, Tchetgen Tchetgen and Halloran (2012) only considered estimation of the infectiousness and contagion effects in a sample comprised of independent households of size two with one member of each household assumed to be homebound. The assumption that one individual is homebound and the assumption of independent households are restrictive, the latter because it requires that the households be sampled from distinct communities and geographic areas. Ogburn and VanderWeele (2014) considered the setting in which households are independent but both individuals may be exposed outside the household. Here, we relax the requirement of independent households of size two and provide extensions to independent groups of arbitrary size and to social networks.

Increasingly, data are available on the spread of contagious outcomes through social networks. Christakis and Fowler (2010) collected data on the evolution of a seasonal flu epidemic in the student body of a college, including information about

friendship ties among the students and about their vaccination history. This setting is considerably more complex than that of independent groups or households. Although our proposed methods are not adequately powered to detect contagion and infectiousness effects in the Harvard flu study data, we use this study as a motivating example for the development of methods to assess these effects using data from a single interconnected network. There is a growing literature on the possibility of testing for the presence of different causal mechanisms using observational data from social networks and a consensus that more rigorous methods are needed.

An emerging body of work reports results from generalized linear models (GLMs) and, for longitudinal data, generalized estimating equations (GEEs) as estimates of peer effects, or the causal effect that one individual's outcome may have on his or her social contacts' outcomes [Ali and Dwyer (2009), Cacioppo, Fowler and Christakis (2009), Fowler and Christakis (2008), Lazer et al. (2010), Rosenquist et al. (2010), Christakis and Fowler (2007, 2008, 2013)]. This work has come under criticism that can largely be summarized into two overarching themes. First, much of the criticism focuses on the ability to control for confounding when estimating peer effects, and specifically on the identifying assumptions that are required in order to tell the difference between homophily (the phenomenon by which individuals with more similar traits are more likely to form social ties with one another) and peer influence [Cohen-Cole and Fletcher (2008), Lyons (2011), Noel and Nyhan (2011), Shalizi and Thomas (2011), VanderWeele (2011)]. Homophily may not be an issue in many infectious disease settings, as an illness like the seasonal flu is unlikely to change the nature of social ties, but adequate control for confounding is still crucial (and, even here, individuals may be homophilous on unobserved traits relevant to infection risk). We assume throughout that all potential confounders of the causal effects of interest are observed. This assumption should be assessed in any application of these methods and it may not hold in many real data settings; however, we do not focus on this assumption in the remainder of this paper.

The second class of criticisms addresses the use of statistical models for independent observations in this dependent data setting. Lyons (2011) and VanderWeele, Ogburn and Tchetgen Tchetgen (2012) demonstrated the importance of ensuring that models are internally consistent when an observation can be both an outcome and a predictor (of social contacts' outcomes); this is easily accomplished by using the observations at one time point as predictors and the observations at a subsequent time point as outcomes, a solution that was implemented in many of applications of GLMs and GEEs to social network data referenced above. More challenging is the fact that, when an analysis assumes independence but observations are in fact positively correlated, as we would expect them to be for contagious outcomes in a social network, the resulting standard errors and statistical inference will generally be anticonservative. In some cases, the assumption of independent outcomes may hold under the null hypothesis [VanderWeele, Ogburn and Tchetgen Tchetgen (2012)], but it is unknown whether tests that rely on

this fact have any power to detect the presence of the causal effects of interest [Shalizi (2012)]. There is a vast literature on how to deal with dependence among observations when performing statistical inference, but this literature generally assumes that the dependence is related to a Euclidean distance metric, rendering it inapplicable to the network setting.

Our contribution to methodology for social network analysis is to adapt GLMs to ensure that the models can be correctly specified, with uncorrelated errors, even when the outcome is contagious. GLMs could be misspecified either parametrically (i.e., the marginal distribution specified for the outcome conditional on covariates is incorrect) or due to a violation of the assumption of uncorrelated errors. It is the latter that differentiates social networks from other settings and, therefore, the latter that we focus on. The former is equally an issue in this setting as in any setting in which GLMs are appropriate; researchers should be careful to specify flexible parametric models and to test the assumptions of their models with sensitivity analyses. Because it may be more robust to parametric model misspecification, we focus on hypothesis testing rather than point estimation throughout. In simulations, we demonstrate the possibility of testing for the presence of contagion and infectiousness effects using social network data and GLMs. We discuss the paradigmatic example of the effect of a vaccination on an infectious disease outcome, but effects like contagion and infectiousness are of interest in other settings as well. Our general approach to correctly specifying GLMs for a contagious outcome using network data could potentially be applied to any estimand for which GLMs are appropriate under independence. The tests that we propose have important limitations, most notably low power to detect effects unless networks are large and/or sparse. However, this work represents a proof of concept in the ongoing endeavor to develop methods for valid inference using data collected from a single network. Furthermore, we hope that in proposing a work-around to the problem of dependent observations, our work highlights and elucidates the issues of model misspecification and invalid standard errors raised by previous proposals for using GLMs and GEEs to assess peer effects using network data.

2. Social networks and contagion. Formally, a social network is a collection of individuals and the ties between them. The presence of a tie between two individuals indicates that the individuals share some kind of a relationship; what types of relationships are encoded by network ties depends on the context. For example, we might define a network tie to include familial relatedness, friendship, and shared place of work. Some types of relationships are mutual, for example, familial relatedness and shared place of work. Others, like friendship, may go in only one direction: Tom may consider Sue to be his friend, while Sue does not consider Tom to be her friend. We will assume that all ties in our network are mutual or undirected, but the principles of our method extend to directed ties. A node whose characteristics we wish to explain is called an *ego*; nodes that share ties with the

ego are its *alters* or *contacts*. If an ego's outcome may be affected by his contacts' outcomes, then we say that the outcome exhibits *induction* or *contagion*.

Social networks are crucial to understanding many features of infectious disease dynamics, and, increasingly, infectious disease researchers draw on social network data to refine their understanding of transmission patterns and treatment effects. For example, many mathematical models of infectious disease now incorporate social network structure, whereas they previously generally assumed uniform mixing among members of a community [Eubank et al. (2004), Keeling and Eames (2005), Klovdahl (1985), Klovdahl et al. (1994)]. Researchers studying sexually transmitted diseases often collect data on sexual contact networks, in part because properties of these networks can inform strategies for controlling sexually transmitted diseases [Eames and Keeling (2002, 2004), Latora et al. (2006)].

Existing methods for assessing causal effects using network data are limited. Some recent proposals give methods for assessing indirect effects when treatment can be randomized [Aronow and Samii (2013), Bowers, Fredrickson and Panagopoulos (2013), Choi (2014), Eckles, Karrer and Ugander (2014), Rosenbaum (2007), Toulis and Kao (2013), Ugander et al. (2013)], but these methods are of limited use in observational settings or for teasing apart specific types of indirect effects like the infectiousness and contagion effects. A new and powerful approach by van der Laan (2012) relies on the independence of all outcome observations at any given time point, conditional on the past. But much of the extant literature relies on GLMs and GEEs, despite the fact that the key assumption of independent outcomes across subjects is unlikely to hold in social network settings [Lyons (2011)]. In this paper, we introduce a way to ensure that GLM errors are uncorrelated across subjects despite the fact that outcomes may be nonindependent; this facilitates the use of GLMs to assess infectiousness and contagion effects in social network contexts like the Harvard College seasonal flu dataset. We demonstrate through simulations that our methods do have some power to detect the presence of contagion and infectiousness effects; however, in order to ensure that errors are uncorrelated, we make several adaptations to naive GLMs, and unfortunately these can result in low power. The applications that we discuss in this paper do not require the use of GEEs to account for within-subject dependence over time, but the general principles that we use to adapt GLMs to the network setting apply to GEEs as well.

2.1. *Harvard college flu study.* From September 1 to December 31, 2009, researchers monitored the seasonal flu among a sample of 744 undergraduates living in dorms at Harvard University [Christakis and Fowler (2010)]. A random sample of 319 students was first selected from the total undergraduate population of 6650. Each selected student was asked to name up to three friends, and, of the named friends, a sample of 425 was selected to enroll in the study. These 425 subjects were also asked to name up to three friends. Two subjects are deemed to share a network tie if either named the other as a friend. For each student, data were

collected on flu vaccination status, dorm of residence, and participation in various sports and activities. At the time of enrollment on October 23, 2009, students were asked to report on any flu symptoms or flu vaccination since September 1, 2009, and to provide the date for any reported symptoms or vaccination. From October 23 to December 31, enrolled students were emailed a survey every Monday and Thursday in which they were asked to report any new flu symptoms or vaccination since the previous emailed survey, and these data allowed daily precision regarding flu status going back to September 1. These records contained information on vaccination status and flu diagnosis for students who sought treatment on campus. For more details on the data collection and monitoring, see Christakis and Fowler (2010). In the supplementary material [Ogburn, VanderWeele and Christakis (2017)], we describe an (inconclusive) analysis of the Harvard flu network data using our methods.

3. Preliminaries and review of previous work on infectiousness and contagion. VanderWeele, Tchetgen Tchetgen and Halloran (2012) first introduced infectiousness and contagion effects and described the identification and estimation of these effects using data comprised of independent households of size 2 in which one person is assumed to be homebound. Because it entails that any case of the disease in the homebound individual was acquired from the other household member, the assumption that one individual in each pair is homebound simplifies the definition and the identification of the infectiousness and contagion effects considerably. However, it also severely limits the applicability of the estimands introduced in VanderWeele, Tchetgen Tchetgen and Halloran (2012), which are only interpretable under this assumption. In this section, we review the identification results of Ogburn and VanderWeele (2014), which do not require one member of each pair to be homebound, and provide a corresponding estimation procedure. This estimating procedure is based on but corrects a mistake in the estimation procedure of VanderWeele, Tchetgen Tchetgen and Halloran (2012).

3.1. *Notation, assumptions, and definitions.* Consider K pairs of individuals who have regular contact with one another, such that if one gets the flu the other is at risk of catching it. Borrowing terminology from the social network literature, we will refer to one individual as the alter, denoted a, and the other as the ego, denoted e. Contagion and infectiousness effects are analogous to causal mediation effects of the alter's vaccination on the ego's outcome, mediated by the alter's disease status [VanderWeele, Tchetgen Tchetgen and Halloran (2012)]. We formally define these effects below after first introducing some key notation and identifying assumptions.

For individual i in pair k, $i = a, e$, $k = 1, \ldots, K$, let $Y_{i_k}^t$ be the outcome at time t and C_{i_k} be a vector of baseline covariates. Let V_{a_k} be an indicator of vaccination for the alter in pair k. For now, we assume that in each pair the ego is unvaccinated, and that all vaccination occurs before the start of follow-up. Define $Y_{i_k}^t(v)$

to be the counterfactual outcome we would have observed for individual i in pair k at time t, if, possibly contrary to fact, the alter had received treatment v. Let M_k be a variable that lies on a causal pathway from V_{a_k} to $Y^t_{e_k}$. Let $Y^t_{e_k}(v,m)$ be the counterfactual outcome for the ego at time t that we would have observed if V_{a_k} had been set to v and M_k to m. Throughout, we make the consistency assumptions that $M_k(v) = M_k$ when $V_{a_k} = v$, that $Y^t_{e_k}(v,m) = Y^t_{e_k}$ when $V_k = v$ and $M_k = m$, and that $Y^t_{e_k}(v, M_k(v)) = Y^t_{e_k}(v)$. Let $Y^t_{e_k}(v, M_k(v'))$ be the counterfactual disease status for the ego in pair k that we would have observed at time t if V_{a_k} had been set to v and M_k to its counterfactual value under $V_{a_k} = v'$. To ensure that this counterfactual is well-defined, we assume that it is hypothetically possible to intervene on the mediator without intervening on V_{a_k}. Let $\boldsymbol{C}_k = (C_{a_k}, C_{e_k})$. Below we omit the subscript k when context allows. The expected *natural direct effect* of V_a on Y^t_e, mediated by M, is given by $E[Y^t_e(v, M(v))] - E[Y^t_e(v', M(v))]$ and the expected *natural indirect effect* is $E[Y^t_e(v', M(v))] - E[Y^t_e(v', M(v'))]$. The expected total effect of V_a on Y^t_e is the sum of these two causal mediation effects: $E[Y^t_e(v)] - E[Y^t_e(v')] = E[Y^t_e(v, M(v))] - E[Y^t_e(v', M(v'))] = E[Y^t_e(v, M(v))] - E[Y^t_e(v', M(v)] + E[Y^t_e(v', M(v))] - E[Y^t_e(v', M(v'))]$.

In order to identify the natural direct and indirect effects, we require the following four assumptions [Pearl (2001)]:

$$Y^t_e(v,m) \perp V_a | \boldsymbol{C}, \tag{3.1}$$

$$Y^t_e(v,m) \perp M | V_a, \boldsymbol{C}, \tag{3.2}$$

$$M(v) \perp V_a | \boldsymbol{C}, \tag{3.3}$$

and

$$Y^t_e(v,m) \perp M(v') | \boldsymbol{C}, \tag{3.4}$$

where $A \perp B \mid C$ denotes that A is independent of B conditional on C. Assumptions (3.1), (3.2), and (3.3) correspond to the absence of unmeasured confounders for the effects of the exposure on the outcome (V_a on Y^t_e), of the mediator on the outcome (M on Y^t_e), and of the exposure on the mediator (V_a on M), respectively. Assumption (3.4) requires that no confounder of the effect of M on Y^t_e is affected by V_a. Discussion of these assumptions in the context of mediation analysis can be found in Pearl (2001). Discussion and extension of these assumptions to settings similar to the one we consider in this paper can be found in Ogburn and VanderWeele (2014), including the discussion of how to determine which covariates must be included in $\boldsymbol{C}$.

Let T_k be the time of the first case of the disease in pair k. If neither individual in pair k is ever infected, then we define T_k to be the end of follow-up. Now $Y^{T_k}_{a_k}$ is an indicator of whether the alter is symptomatic at time T_k, that is, an indicator of whether the alter is the first individual in the pair to become infectious; if neither individual becomes infectious, then it will be 0. This indicator

will play the role of the mediator in what follows. Let $T_k(v)$ be the time at which the first infection in pair k would have occurred if the alter had, possibly contrary to fact, had vaccine status v. Let $Y_{a_k}^{T_k(v)}(v)$ be the counterfactual disease status of the alter at time $T_k(v)$ had he had vaccine status v. Define the outcome $Y_{e_k}^{T_k+s} = I(\text{individual } e_k \text{ became infectious after time } T_k \text{ and on or before time } T_k+s)$. The constant s should be chosen to be the sum of the infectious period (f) and the incubation period (b) of the disease under study. The infectious period is the length of time during which an infected individual is infectious, and the incubation period is the length of time between being infected and becoming infectious. If the alter becomes infectious at time T, then he can infect the ego until time $T + f$. If infected at time $T + f$, the ego will become infectious at time $T_k + f + b = T_k + s$. Therefore, if the alter infects the ego, the ego must be infectious by time $T_k + s$. We assume that the probability of the ego contracting the disease within a fixed follow-up interval, if exposed at time t, is constant in t. This ensures that the time of the first infection T is not a confounder of the mediator-outcome relationship, which would constitute a violation of assumption (3.4) because T is affected by V_a. Note that $Y_{e_k}^{T_k+s} = 0$ whenever $Y_{a_k}^{T_k} = 0$: the latter occurs either when the ego is the first to become infectious, at time T_k, in which case the indicator condition defining $Y_{e_k}^{T_k+s}$ is false (it requires individual e_k to become infectious *after* T_k), or when neither the alter nor the ego is observed to become infectious.

Throughout, we define a disease case to begin when an individual becomes infectious. If infectiousness does not coincide with the appearance of disease symptoms, then we may not observe the timing of disease cases directly, but we could infer the time based on when symptoms appear and on known disease dynamics. For example, an individual with the flu will generally be infectious one day before he is symptomatic [Earn, Dushoff and Levin (2002)]. Therefore, if flu is the disease under study we would classify an individual as having the disease beginning one day before he reported having flu symptoms. We also assume throughout that there are no asymptomatic carriers of the disease.

The contagion effect is given by a contrast of counterfactuals of the form $Y_e^{T(v')+s}(v, Y_a^{T(v')}(v'))$ where, unlike in the mediation framework we described above, the variables $Y_e^{T(v')+s}$ and $Y_a^{T(v')}$ that play the roles of outcome and mediator may be a different random variable in the two terms in the contrast. Specifically, the population average contagion effect is $E[Y_e^{T(1)+s}(0, Y_a^{T(1)}(1))] - E[Y_e^{T(0)+s}(0, Y_a^{T(0)}(0))]$, and $Y_a^{T(0)}$ and $Y_a^{T(1)}$ will be different random variables whenever $T(0) \neq T(1)$, as will $Y_e^{T(0)+s}$ and $Y_e^{T(1)+s}$ [Ogburn and VanderWeele (2014)]. This contrast is the difference in expected counterfactual outcomes for the ego s days after the time of the first infection when the vaccine status of the alter is held constant at 0 but his infection status is set to that under vaccination in the first term and to that under no vaccination in the second term of the contrast. It captures the effect that vaccination might have had on the disease status of the ego

by preventing the alter from contracting the disease. The nested counterfactuals are well-defined because we can imagine a (purely hypothetical) intervention on Y_a^T that would not require intervening on V_a, for example, by administering immune boosters to prevent the alter from being infected or by exposing the alter to a high dose of flu virus in a laboratory setting to cause infection. This thought experiment supports the idea that Y_a^T is governed by mechanisms in addition to those involved in V_a; if it were not then it would be impossible to conceive of a counterfactual in which V_a is set to, for example, 0 while Y_a^T is set to its counterfactual value under $V_a = 1$. [See Robins and Richardson (2010) for further discussion of when these nested counterfactuals are well-defined.]

The population average infectiousness effect is $E[Y_e^{T(1)+s}(1, Y_a^{T(1)}(1))] - E[Y_e^{T(1)+s}(0, Y_a^{T(1)}(1))]$ [Ogburn and VanderWeele (2014)]. This is the effect of one individual's treatment on another's disease status, not mediated through the first individual's disease status. This effect operates if treatment renders cases of disease among treated individuals less likely to be transmitted. Suppose that the alter in group k would get the flu first if vaccinated, that is, $Y_{a_k}^{T_k(1)}(1) = 1$. Then the infectiousness effect is the difference in counterfactual outcomes for the ego comparing the scenario in which the alter is vaccinated and infected first with the scenario in which to the alter is unvaccinated and infected first. If the alter in group k would not get the flu first under vaccination, then the infectiousness effect for group k is null.

By the consistency assumption we made in Section 3.1 above, $E[Y_e^{T(1)+s}(1, Y_a^{T(1)}(1))] = E[Y_e^{T(1)+s}(1)]$ and $E[Y_e^{T(0)+s}(0, Y_a^{T(0)}(0))] = E[Y_e^{T(1)+s}(0)]$. Just as the natural indirect and natural direct effects decompose the total effect, the indirect effect of the alter's vaccination on the ego (analogous to the total effect) decomposes into the sum of the contagion and infectiousness effects (analogous to the natural indirect and direct effects, respectively) as follows:

$$\begin{aligned}
&E\big[Y_e^{T(1)+s}(1)\big] - E\big[Y_e^{T(0)+s}(0)\big] \\
&\quad = E\big[Y_e^{T(1)+s}\big(1, Y_a^{T(1)}(1)\big)\big] - E\big[Y_e^{T(0)+s}\big(0, Y_a^{T(0)}(0)\big)\big] \\
&\quad = E\big[Y_e^{T(1)+s}\big(1, Y_a^{T(1)}(1)\big)\big] - E\big[Y_e^{T(1)+s}\big(0, Y_a^{T(1)}(1)\big)\big] \\
&\qquad + E\big[Y_e^{T(1)+s}\big(0, Y_a^{T(1)}(1)\big)\big] - E\big[Y_e^{T(0)+s}\big(0, Y_a^{T(0)}(0)\big)\big].
\end{aligned}$$

So far, we have described all effects on the difference scale, but everything we have written applies equally to effects on the ratio and odds ratio scales. On the ratio and odds ratio scales, the indirect effect of vaccination decomposes into a product of the contagion and infectiousness effects. On the ratio scale, the average indirect effect of V_a on the disease status of the ego is $E[Y_e^{T(1)+s}(1)]/E[Y_e^{T(0)+s}(0)]$, which is a product of the average infectiousness effect, $E[Y_e^{T(1)+s}(1, Y_a^{T(1)}(1))]/E[Y_e^{T(1)+s}(0, Y_a^{T(1)}(1))]$, and the average contagion effect, $E[Y_e^{T(1)+s}(0, Y_a^{T(1)}(1))]/E[Y_e^{T(0)+s}(0, Y_a^{T(0)}(0))]$. On the odds ratio

scale for a binary outcome, the decomposition is

$$\begin{aligned}
&\frac{E[Y_e^{T(1)+s}(1)](1-E[Y_e^{T(0)+s}(0)])}{E[Y_e^{T(0)+s}(0)](1-E[Y_e^{T(1)+s}(1)])} \\
&\quad = \frac{E[Y_e^{T(1)+s}(1,Y_a^{T(1)}(1))](1-E[Y_e^{T(1)+s}(0,Y_a^{T(1)}(1))])}{E[Y_e^{T(1)+s}(0,Y_a^{T(1)}(1))](1-E[Y_e^{T(1)+s}(1,Y_a^{T(1)}(1))])} \\
&\qquad \times \frac{E[Y_e^{T(1)+s}(0,Y_a^{T(1)}(1))](1-E[Y_e^{T(0)+s}(0,Y_a^{T(0)}(0))])}{E[Y_e^{T(0)+s}(0,Y_a^{T(0)}(0))](1-E[Y_e^{T(1)+s}(0,Y_a^{T(1)}(1))])},
\end{aligned}$$

where the first line is the indirect effect, the second line is the infectiousness effect, and the third line is the contagion effect.

3.2. *Estimation of infectiousness and contagion effects in groups of size two.* The contagion and infectiousness effects are analogous to the natural indirect and direct effects, respectively, of the effect of V_a on Y_e^{T+s} with Y_a^T as the mediator. Natural indirect and direct effects have been written about extensively in the causal inference and mediation literature [see, e.g., Pearl (2001), Robins and Greenland (1992), Robins and Richardson (2010)] and it is well known how to estimate them in a variety of settings [Imai, Keele and Tingley (2010), Valeri and VanderWeele (2013)]. This setting differs from those considered by other authors because the outcome Y_e^{T+s} is, by definition, equal to 0 whenever Y_a^T is equal to 0; therefore, one must be careful to ensure that any model specified for for $E[Y_e^{T+s} \mid V_a, Y_a^T, \boldsymbol{C}]$ is consistent with this restriction. VanderWeele, Tchetgen Tchetgen and Halloran (2012) describe how to estimate the contagion and infectiousness effects on the ratio scale in households of size two when one individual is homebound, but the procedure they present overlooks this restriction and, therefore, the models they suggest are overidentified and fitting procedures may fail to converge.

We describe a procedure for estimating the contagion and infectiousness effects that is appropriate for the setting considered in VanderWeele, Tchetgen Tchetgen and Halloran (2012) and Ogburn and VanderWeele (2014). We describe estimation of the effects on the difference and ratio scales. Estimation of effects on the odds ratio scale is also possible. Suppose that assumptions (3.1) through (3.4) hold for the effect of V_a on Y_e^{T+s} with Y_a^T as the mediator and covariates $\boldsymbol{C}$, and that the following two models are correctly specified:

$$\log\{E[Y_e^{T+s} \mid V_a, Y_a^T = 1, \boldsymbol{C}]\} = \gamma_0 + \gamma_1 V_a + \gamma_2' \boldsymbol{C}, \tag{3.5}$$

$$\operatorname{logit}\{E[Y_a^T \mid V_a, \boldsymbol{C}]\} = \eta_0 + \eta_1 V_a + \eta_2' \boldsymbol{C}. \tag{3.6}$$

If the outcome is rare, then (3.5) can be replaced with a logistic model; the choices of link functions and the specifications of the right-hand sides of equations (3.5)

and (3.6) are flexible. The contagion effect conditional on covariates $\boldsymbol{C} = \boldsymbol{c}$ on the difference scale is given by

$$\begin{aligned}
&E\big[Y_e^{T(1)+s}\big(0, Y_a^{T(1)}(1)\big) \mid \boldsymbol{c}\big] - E\big[Y_e^{T(0)+s}\big(0, Y_a^{T(0)}(0)\big) \mid \boldsymbol{c}\big] \\
&\quad = 0 + E\big[Y_e^{T+s} \mid V_a = 0, Y_a^T = 1, \boldsymbol{c}\big]\big\{E\big[Y_a^T \mid V_a = 1, \boldsymbol{c}\big] \\
&\qquad - E\big[Y_a^T \mid V_a = 0, \boldsymbol{c}\big]\big\} \\
&\quad = e^{\gamma_0 + \gamma_2' \boldsymbol{c}} \left\{ \frac{e^{\eta_0 + \eta_1 + \eta_2' \boldsymbol{c}}}{1 + e^{\eta_0 + \eta_1 + \eta_2' \boldsymbol{c}}} - \frac{e^{\eta_0 + \eta_2' \boldsymbol{c}}}{1 + e^{\eta_0 + \eta_2' \boldsymbol{c}}} \right\}
\end{aligned}$$

and the infectiousness effect conditional on covariates $\boldsymbol{C} = \boldsymbol{c}$ is given by

$$\begin{aligned}
&E\big[Y_e^{T(1)+s}\big(1, Y_a^{T(1)}(1)\big) \mid \boldsymbol{c}\big] - E\big[Y_e^{T(1)+s}\big(0, Y_a^{T(1)}(1)\big) \mid \boldsymbol{c}\big] \\
&\quad = 0 + E\big[Y_a^T \mid V_a = 1, \boldsymbol{c}\big]\big\{E\big[Y_e^{T+s} \mid V_a = 1, Y_a^T = 1, \boldsymbol{c}\big] \\
&\qquad - E\big[Y_e^{T+s} \mid V_a = 0, Y_a^T = 1, \boldsymbol{c}\big]\big\} \\
&\quad = \frac{e^{\eta_0 + \eta_1 + \eta_2' \boldsymbol{c}}}{1 + e^{\eta_0 + \eta_1 + \eta_2' \boldsymbol{c}}} \big\{ e^{\gamma_0 + \gamma_1 + \gamma_2' \boldsymbol{c}} - e^{\gamma_0 + \gamma_2' \boldsymbol{c}} \big\}.
\end{aligned}$$

The contagion and infectiousness effects can be estimated by fitting models (3.5) and (3.6) and plugging the parameter estimates into the expressions above. Consistent estimators of the asymptotic variance of these estimands can be obtained via the bootstrap, or they can be derived using the delta method [similar to those derived in Valeri and VanderWeele (2013) for the natural direct and indirect effects]. Alternatively, a Monte Carlo based approach similar to Imai, Keele and Tingley (2010) can be used for estimation of the effects and their standard errors. Software packages like SAS and SPSS mediation macros [Valeri and VanderWeele (2013)] or the R mediation package [Imai, Keele and Tingley (2010)] cannot be used in this setting because instead of (3.5), which models the conditional expectation of the Y_e^{T+s} only in the $Y_a^T = 1$ stratum; these packages require fitting a model for $E[Y_e^{T+s} \mid V_a, Y_a^T, \boldsymbol{C}]$.

If the ego can also be vaccinated, then V_e must be included in $\boldsymbol{C}$. If V_a interacts with V_e or with any other covariates, these interactions can be incorporated into the models and pose no difficulty for estimation. To test whether there is a contagion effect, we can simply test whether $\eta_1 = 0$. To test whether there is an infectiousness effect, we can simply test whether $\gamma_1 = 0$.

Using the parameters of models (3.5) and (3.6), we can also estimate the contagion and infectiousness effects on the ratio scale. The contagion effect conditional

on $\boldsymbol{C} = \boldsymbol{c}$ is given by

$$\begin{aligned}
&\frac{E[Y_e^{T(1)+s}(0, Y_a^{T(1)}(1)) \mid \boldsymbol{c}]}{E[Y_e^{T(0)+s}(0, Y_a^{T(0)}(0)) \mid \boldsymbol{c}]} \\
&\quad = \frac{0 + E[Y_e^{T+s} \mid V_a = 0, Y_a^T = 1, \boldsymbol{c}]E[Y_a^T \mid V_a = 1, \boldsymbol{c}]}{0 + E[Y_e^{T+s} \mid V_a = 0, Y_a^T = 1, \boldsymbol{c}]E[Y_a^T \mid V_a = 0, \boldsymbol{c}]} \\
&\quad = \frac{E[Y_a^T \mid V_a = 1, \boldsymbol{c}]}{E[Y_a^T \mid V_a = 0, \boldsymbol{c}]} \\
&\quad = \frac{e^{\eta_1} + e^{\eta_0+\eta_1+\eta_2'\boldsymbol{c}}}{1 + e^{\eta_0+\eta_1+\eta_2'\boldsymbol{c}}}
\end{aligned} \tag{3.7}$$

and the infectiousness effect is given by

$$\begin{aligned}
&\frac{E[Y_e^{T(1)+s}(1, Y_a^{T(1)}(1)) \mid \boldsymbol{c}]}{E[Y_e^{T(1)+s}(0, Y_a^{T(1)}(1)) \mid \boldsymbol{c}]} \\
&\quad = \frac{0 + E[Y_e^{T+s} \mid V_a = 1, Y_a^T = 1, \boldsymbol{c}]E[Y_a^T \mid V_a = 1, \boldsymbol{c}]}{0 + E[Y_e^{T+s} \mid V_a = 0, Y_a^T = 1, \boldsymbol{c}]E[Y_a^T \mid V_a = 1, \boldsymbol{c}]} \\
&\quad = \frac{E[Y_e^{T+s} \mid V_a = 1, Y_a^T = 1, \boldsymbol{c}]}{E[Y_e^{T+s} \mid V_a = 0, Y_a^T = 1, \boldsymbol{c}]} \\
&\quad = e^{\gamma_1}.
\end{aligned} \tag{3.8}$$

Under the restriction that $Y_e^{T+s} = 0$ whenever $Y_a^T = 0$, the contagion effect on the ratio scale is simply a measure of the effect of the alter's vaccination on the alter's outcome. This contagion effect is mathematically undefined if $E[Y_e^{T+s} \mid V_a = 0, Y_a^T = 1, \boldsymbol{c}] = 0$, that is, if the alters's outcome has no effect on the ego's outcome, but it is natural to define the effect to be equal to the null value of 1 in this case. The infectiousness effect on the ratio scale is simply a measure of the effect of the alter's vaccination on the ego's outcome among pairs in which the alter is infectious first, that is, in the $Y_a^T = 1$ stratum.

4. New method methods for identifying and estimating infectiousness and contagion in groups of more than two. Although allowing both individuals in a household to be infected from outside the household generalizes the results of VanderWeele, Tchetgen Tchetgen and Halloran (2012), it still requires the strong assumption, inherent in the identifying assumptions described in Section 3, that the alter and ego do not share any potentially infectious contacts. If both of the individuals in a given household could be infected from outside the household by the same mutual friend, then that friend's disease status would be a confounder of the mediator-outcome relationship; if unobserved, it would constitute a violation

of assumption (3.2). We can relax the assumption of no mutual contacts outside of the household by collecting data on any such contacts and controlling for them as covariates in our estimating procedure.

In this section, we consider identification and estimation of the contagion and infectiousness effects when independent groups of individuals are sampled. We assume that each group includes a pair of individuals who furnish the exposure, mediator, and outcome variables, plus all mutual and potentially infectious contacts of the pair. Several types of sampling procedures could give rise to this data structure, for example, one possibility would be to sample workplaces and randomly select two individuals to play the role of the alter and ego; another would be to sample household pairs first, ascertain the identities of potential mutual contacts outside of the home, and include all such contacts in the data collection moving forward. As long as it does not introduce bias into the subsample, the choice of sampling procedure does not affect the identification or estimation results described below.

Let k index the kth group, $k = 1, \ldots, K$. Let Y_{ik}^t be an indicator of whether individual i in group k has had the disease by day t. As in Section 3, we define a case of the disease to begin when the individual becomes infectious and let $s = f + b$ be the sum of the infectious and incubation periods for the disease. We assume that vaccination occurs before the start of follow-up. Given a nonrare outcome like the flu and time measured in discrete intervals like days, it is likely that we would observe multiple individuals to become infectious on the same day. We therefore do not make the assumption, made in Section 3, that no two individuals can be observed to become infectious at the same time. For group k, let e_k index the ego, whose flu status we wish to study, and let a_k index the alter, whose vaccination status may or may not have an effect on the ego's disease status. We index the other individuals in group k by $1, 2, \ldots, n_k$. Let T_k be the time of the first infection in the kth alter-ego pair. When context allows, we omit the subscript k. The definition of the mediator needs to be modified slightly to reflect the fact that the alter and the ego could become infectious at the same time: let Y_a^T be an indicator of whether the alter was infectious and the ego healthy at time T. Let Y_e^{T+s} be an indicator of whether the ego became infectious between time $T + b$, which is the first time at which the alter could have infected the ego, and time $T + s$, which is the last time at which the alter could have infected the ego. This definition preserves the interpretation of Y_a^T as an indicator that the alter was infectious before the ego; if the ego and the alter simultaneously fell ill on day T then Y_a^T will be 0, which is desirable because the ego cannot have caught the disease from the alter if they both fell ill on the same day. It also preserves the restriction, discussed in Section 3, that Y_e^{T+s} is equal to 0 whenever Y_a^T is.

$Y_e^{T(v')+s}(v, Y_a^{T(v')}(v'))$ is the counterfactual flu status of the ego at time $T(v') + s$ had the alter's vaccine status been set to v and his flu status at time $T(v')$ set to its counterfactual value under vaccine status v', where $T(v')$ is the time at which the first infection in the alter-ego pair would have occurred if V_a had been set to v'.

The effects of interest are the average contagion effect

$$Con = \frac{E[Y_e^{T(1)+s}(0, Y_a^{T(1)}(1))]}{E[Y_e^{T(0)+s}(0, Y_a^{T(0)}(0))]} \tag{4.1}$$

and the average infectiousness effect

$$Inf = \frac{E[Y_e^{T(1)+s}(1, Y_a^{T(1)}(1))]}{E[Y_e^{T(1)+s}(0, Y_a^{T(1)}(1))]}, \tag{4.2}$$

where the expectations are taken over all ego-alter pairs.

In order to identify the effects defined in (4.1) and (4.2), we must measure and control for all confounders of the relationships between Y_e^{T+s} and Y_a^T, and in particular the potential mutual infectious contacts of the alter and ego. To motivate our procedure for controlling for these confounding contacts, consider the simple case of a group of size three, comprised of a child (ego), a parent (alter), and a grandparent. In the event that the grandparent contracted the flu first and transmitted it to both the child and the parent, the grandparent's flu status would clearly be a confounder of the mediator-outcome relationship. But the grandparent's entire disease trajectory is not a potential confounder; in particular anything that happens to the grandparent after time T, that is after the first infection in the parent-child pair, occurs after the mediator and cannot possibly confound the mediator-outcome relationship. In this simple, three-person group, it suffices to control for an indicator of whether the grandparent has been infectious by time $T - b$, where T is the time of the first infection between the parent and child, and $T - b$ is the latest time at which the grandparent could have been the cause of an infection at time T.

In practice, we will likely have to sample groups of a size greater than three in order to control for confounding by potential mutually infectious contacts. It may be sufficient to control for a summary measure of the infections occurring before $T - b$ in each group. If each infectious contact of an individual has an independent and identical probability of transmitting the disease to the individual, then the sum $\sum_{i=1}^{n_k} Y_{ki}^{T-b}$ of indicators of whether each mutual contact has been infectious by time $T - b$ suffices to control for confounding by potential mutual infectious contacts. Under a different transmission model, the proportions of the alter's and of the ego's contacts who were infectious by time $T - b$ could be the operative summary measure. This would be the case if the number of potentially infectious encounters of the ego (or alter) with each of his friends is proportional to the total number of friends, so that an individual with one friend is more likely to be infected by that one friend than an individual with 100 friends is to be infected by any one of them. If some of the mutual contacts may have been vaccinated, then it is unreasonable to assume that each infectious contact of an individual has an independent and identical probability of transmitting the disease to the individual; instead separate summary measures (sum or proportion infectious by time $T - b$) should be included for vaccinated and for unvaccinated contacts. In what follows, we will assume that the sum is an adequate summary measure.

4.1. *Alternative sampling schemes.* Alter-centric sampling can also be used to collect data on variables that suffice to identify the contagion and infectiousness effects. Instead of sampling an alter-ego pair and all of their mutual contacts, we can sample an individual, to serve as the alter, and all of his potentially infectious contacts. The ego is randomly selected from among the alter's contacts. Conditional on the number of the alter's contacts who have been infectious by day $T - b$, Y_a^T is independent of the number of mutual contacts who were infectious by time $T - b$. Therefore, the number of mutual contacts cannot be a confounder of the relationship between Y_a^T and Y_e^{T+s}. This obviates the need to ascertain the identity or disease status of all mutual contacts. However, the number of potentially infectious contacts of a single person can be vast, and it may be easier to identify mutual contacts of a pair of individuals, for example, by identifying the overlap in their routines and interactions with others, than to identify all of the contacts of any one individual.

5. Infectiousness and contagion in social networks. So far, we have assumed that our observations, comprised of groups of individuals, are independent of one another. This assumption will, in general, be violated when the alter-ego pairs are sampled from a single community or social network. We introduce some new notation for this context after briefly describing the example that will serve as the basis for our exposition and later for our simulations and data analysis. Consider tracking the seasonal flu in the student population of a college at which all students live in dorms on campus. Each student is a node in the network. We define a tie to exist between two nodes if the individuals regularly interact with one another in a way that could facilitate transmission of the flu, for example, if two individuals are roommates, eat together in the dining hall, or are close friends, then their nodes share a tie. We observe each individual's flu status every day over the course of the flu season, which lasts for 100 days.

The contagion and infectiousness effects *Con* and *Inf*, defined in Section 4, are not estimable from social network data using the methods that we propose below. This is because the estimation procedure that we describe requires conditioning on specific covariates that change the interpretation of the contagion and infectiousness effects. We define new conditional contagion and infectiousness effects and give assumptions under which the new estimands are estimable from network data using GLMs. In what follows, we focus our attention on correct specification of the error structure of the GLMs we postulate; correct specification of the parametric form of the outcome distribution conditional on covariates is important in practice but, because this is equally of concern in the i.i.d. setting, it is not our focus here. In part because it may be more robust to parametric model misspecification, we focus on hypothesis testing rather than point estimation below.

5.1. *Assumptions.* Along with assumptions (3.1)–(3.4), we make several additional assumptions that facilitate inference using social network data. Define

$\mathcal{A}_i = \{j : i \text{ and } j \text{ share a tie}\}$ to be the collection of indices for individual i's contacts. We assume that

$$Y_i^t \perp Y_j^r \mid \left\{ \sum_{m \in \mathcal{A}_i : V_m = v} Y_m^{t-b}, v = 0, 1 \right\} \qquad \text{for all } j \notin \mathcal{A}_i \text{ and } r \leq t. \tag{5.1}$$

The set in the conditioning event includes the number of vaccinated contacts of individual i who were infectious on or before day $t - b$ and the number of unvaccinated contacts of individual i who were infectious on or before day $t - b$. This assumption says that the outcome of individual i at time t is independent of all past outcomes for noncontacts of i, conditional on a summary measure of the flu history of the contacts of i. In other words, contacts act as a causal barrier between two nodes who do not themselves share a tie. If two individuals, i and j, do not share a tie, then they can have no effect on one another's disease status that is not through their contacts' disease statuses. Because $t - b$ is the latest time at which a disease transmission could affect Y_i^t, we do not need to condition on the contacts' outcomes past that time. This assumption implies that the total number of vaccinated and unvaccinated contacts of individual i who have been infectious by day $t - b$ are a sufficient summary measure of the complete history of all of i's contacts. It could easily be modified so that the probability of being infected at any given time depends on a different summary measure, for example, on the proportion of alters who were infectious at or before time $t - b$.

We also assume that

$$Y_i^t \perp V_j \mid \left\{ \sum_{m \in \mathcal{A}_i : V_m = v} Y_m^{t-b}, v = 0, 1 \right\} \qquad \text{for all } j \notin \mathcal{A}_i \tag{5.2}$$

and that, for any covariate C that is required for (3.1) through (3.4) to hold,

$$Y_i^t \perp C_j \mid \left\{ \sum_{m \in \mathcal{A}_i : V_m = v} Y_m^{t-b}, v = 0, 1 \right\} \qquad \text{for all } j \notin \mathcal{A}_i. \tag{5.3}$$

These assumptions state that any effect of the covariates (including vaccination) of nodes without ties to i on i's disease status would again have to be mediated by the disease statuses of i's contacts. Assumption (5.2) implies that the infectiousness effect is not transitive: whether individual j caught the flu from a vaccinated or unvaccinated person has no influence on whether individual j transmits the flu.

Embedded in assumptions (5.1)–(5.3) is the assumption that all ties are equivalent and all nonties are equivalent with respect to transmission of the outcome. This is likely to be a simplification of reality. It can be relaxed (see Section 5.3), but we make it now for heuristic purposes. It rules out the possibility that some types of ties, like roommates, are more likely to facilitate disease transmission than others, like friends who live in different dorms. It allows an individual to come into contact with and possibly infect (or be infected by) people with whom he does not share a tie, but it entails that he will come into contact with any individual in the

network who is not his contact with equal probability. This rules out, for example, the possibility that an individual is more likely to be infected by the friends of his friends than by a distant node on the network. If transmission can occur between individuals who do not share a tie in the network, then the background level of disease across the network should be included in the conditioning events in (5.1)–(5.3).

We also make the no-unmeasured-confounding assumption that, if there exists a person with whom two individuals in the network interact regularly, then that person is also in the network (with ties to both individuals). In some settings, it may be possible to satisfy this condition, for example, in full sociometric studies conducted de novo, or in studies of online data.

5.2. *Estimation and hypothesis testing*. Consider the following strategy for estimating a new contagion and new infectiousness effect, defined below:

1. Randomly select from the network K pairs of nodes such that the two nodes in each pair share a tie, but, for each pair, neither node has a tie to a node in any other pair or to the contacts of any member of any other pair. The number of possible such pairs will depend on the network size and topology. Randomly select one member of each pair to be the ego and one to be the alter.

2. Index the pairs by k, and let e_k index the ego and a_k the alter in the kth pair. For the kth pair, define a group, also indexed by k, that includes nodes a_k, e_k, $\mathcal{A}_{e_k}$, and $\mathcal{A}_{a_k}$. That is, it includes the alter-ego pair and all nodes with ties to either the alter or the ego. (Below, we suppress the index k when context allows.) Due to the way we selected pairs, none of the members of group k can belong to any other group. As in the sections above, T_k is the time of the first infection in the pair (a_k, e_k). Let $\boldsymbol{C}_k$ be a collection of covariates for group k, where the variables included in $\boldsymbol{C}$ are precisely those required for assumptions (3.1) through (3.4) to hold for outcome $Y_e^{T(1)+b}$, mediator $Y_a^{T(1)}$, and treatment V_a. Note that V_e should be included in $\boldsymbol{C}$ as it is likely to be a confounder of the mediator – outcome relationship. The number of mutual contacts of the alter and ego who were infectious by time $T - b$ must also be included.

3. Let $U_{e_k}^{T_k+f}$ and $L_{e_k}^{T_k+f}$ be the number of unvaccinated and vaccinated nodes, respectively, with ties to e_k who were infectious by time $T_k + f$. Define $U_{a_k}^{T-b}$ and $L_{a_k}^{T-b}$ similarly as the number of unvaccinated and vaccinated nodes, respectively, with ties to a_k who were infectious by time $T_k - b$. Recall that f is the infectiousness period and b the incubation period, defined in Section 3.1.

4. Estimate an average conditional contagion effect

$$Con^* = \frac{E[Y_e^{T(1)+s}(0, Y_a^{T(1)}(1)) \mid U_a^{T-b}, L_a^{T-b}, U_e^{T+f}, L_e^{T+f}, \boldsymbol{C}]}{E[Y_e^{T(0)+s}(0, Y_a^{T(0)}(0)) \mid U_a^{T-b}, L_a^{T-b}, U_e^{T+f}, L_e^{T+f}, \boldsymbol{C}]}$$

and an average conditional infectiousness effect

$$Inf^* = \frac{E[Y_e^{T(1)+s}(1, Y_a^{T(1)}(1)) \mid U_a^{T-b}, L_a^{T-b}, U_e^{T+f}, L_e^{T+f}, \boldsymbol{C}]}{E[Y_e^{T(1)+s}(0, Y_a^{T(1)}(1)) \mid U_a^{T-b}, L_a^{T-b}, U_e^{T+f}, L_e^{T+f}, \boldsymbol{C}]}$$

and their standard errors.

Through Step 2, the procedure we described is nearly identical to the proposal in Section 4, the only difference being that groups are extracted from a network in Step 1 rather than being independently ascertained. Consideration for this sampling scheme becomes crucial when we estimate the parameters of GLMs like (3.5) and (3.6). The standard errors derived from these GLMs are consistent only if the errors across groups are uncorrelated. The errors are indeed uncorrelated for independent groups, but, in the network setting, they generally are not. However, the set of additional covariates introduced in Step 3 essentially blocks the flow of information between groups. Conditional on these additional covariates, the errors are uncorrelated, even in the network setting (see Section 5.2.1 for justification). Roughly, because $U_{e_k}^{T_k+f}$ and $L_{e_k}^{T_k+f}$ summarize the disease statuses of the ego's contacts b days before the outcome $Y_{e_k}^{T(1)+s}$ is assessed, conditioning on them ensures that the outcomes are uncorrelated across groups. Because $U_{a_k}^{T-b}$ and $L_{a_k}^{T-b}$ summarize the disease statuses of the alter's contacts b days before the mediator $Y_{a_k}^{T(1)}$ is assessed, conditioning on them ensures that mediators are uncorrelated across groups.

The effects defined in Step 4 differ from *Con* and *Inf* only in the conditioning set, but this changes slightly the causal effect being estimated. Conditioning on U_a^{T-b} and L_a^{T-b} is just like conditioning on an extra pair of confounders: these variables occur before the mediator and are independent of the treatment; therefore, they can be considered to be pre-treatment covariates. On the other hand, U_e^{T+f} and L_e^{T+f} occur after the mediator and lie on a possible pathway from the mediator to the outcome. Conditioning on these variables blocks the path from Y_a^T to Y_e^{T+s} that operates when the alter infects a friend of the ego, who then infects the ego. However, conditioning on these variables leaves the direct path from Y_a^T to Y_e^{T+s} open, and this path operates whenever the alter infects the ego directly. We may be interested in contagion and infectiousness effects that operate via transmission directly from the alter to the ego, in which case Con^* and Inf^* are the causal effects of interest. We assume in what follows that Con^* and Inf^* are the causal effects of interest in social network setting, but note that if *Con* and *Inf* are of interest instead, tests using Con^* and Inf^* will generally be conservative and consistent for hypothesis tests about *Con* and *Inf*.

Figure 1 depicts a causal diagram for two alter-ego pairs where $\boldsymbol{C}$ is null and $T_1 + f < T_2 - b$. The black arrows represent the causal pathways involved in the contagion and infectiousness effects, the grey arrows represent the causal pathways that are blocked by conditioning (boxes are drawn around the conditioned-on

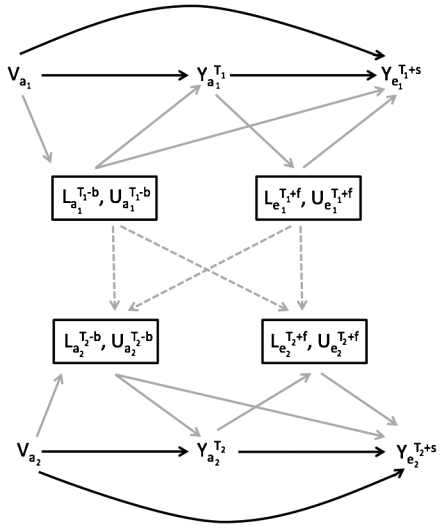

FIG. 1. *Causal Diagram for two alter-ego pairs when* $\boldsymbol{C}$ *is null and* $T_1 + f < T_2 - b$.

variables), and the dashed arrows represent possible causal paths from group 1 to group 2. Depending on how the nodes included in groups 1 and 2 are connected in the underlying network and on the magnitude of the difference of T_1 and T_2, all, some, or none of the dashed arrows could be present.

5.2.1. *Justification for the use of GLMs.* Suppose that the models

$$\begin{aligned} &g\big(E\big[Y_e^{T+s} \mid V_a, Y_a^T = 1, U_a^{T-b}, L_a^{T-b}, U_e^{T+f}, L_e^{T+f}, \boldsymbol{C}\big]\big) \\ &\quad = \beta_0 + \beta_1 V_a + \beta_2 U_a^{T-b} + \beta_3 L_a^{T-b} + \beta_4 U_e^{T+f} + \beta_5 L_e^{T+f} + \beta_6' \boldsymbol{C} \end{aligned} \tag{5.4}$$

and

$$\begin{aligned} &m\big(E\big[Y_a^T \mid V_a, U_a^{T-b}, L_a^{T-b}, U_e^{T+f}, L_e^{T+f}, \boldsymbol{C}\big]\big) \\ &\quad = \alpha_0 + \alpha_1 V_a + \alpha_2 U_a^{T-b} + \alpha_3 L_a^{T-b} + \alpha_4 U_e^{T+f} + \alpha_5 L_e^{T+f} + \alpha_6' \boldsymbol{C} \end{aligned} \tag{5.5}$$

are correctly specified for $g(\cdot)$, $m(\cdot)$ known link functions, that is, the systematic components of the two GLMs are correctly specified with known link functions. We have only to prove that the errors from model (5.4) are uncorrelated with one another and that the errors from model (5.5) are uncorrelated with one another, that is, the random components of the two GLMs are independent across observations [Breslow (1996), Gill (2001)].

THEOREM 5.1. *Let* $\varepsilon_{a_k} = Y_{a_k}^{T_k} - m^{-1}(\alpha_0 + \alpha_1 V_{a_k} + \alpha_2 U_a^{T_k-b} + \alpha_3 L_a^{T_k-b} + \alpha_4 U_e^{T_k+f} + \alpha_5 L_e^{T_k+f} + \alpha_6' \boldsymbol{C}_k)$. *Then* ε_{a_k} *and* ε_{a_h} *are uncorrelated.*

PROOF. Without loss of generality, assume that $T_k > T_h$. Under correct specification of (5.5), $E[\varepsilon_{a_k}] = E[\varepsilon_{a_h}] = 0$. Therefore, $\operatorname{Cov}(\varepsilon_{a_k}, \varepsilon_{a_h}) = E[\varepsilon_{a_k}\varepsilon_{a_h}]$. Letting S_k denote the set of variables $\{V_{a_k}, U_a^{T-b}, L_a^{T-b}, U_e^{T+f}, L_e^{T+f}, \boldsymbol{C}_k\}$, we have

$$
\begin{aligned}
&E[\varepsilon_{a_k}\varepsilon_{a_h}] \\
&\quad = E\big[E[\varepsilon_{a_k}\varepsilon_{a_h} \mid S_k, S_h]\big] \\
&\quad = E\big[E\big[\{Y_{a_k}^{T_k} - E[Y_{a_k}^{T_k} \mid S_k]\}\{Y_{a_h}^{T_h} - E[Y_{a_h}^{T_h} \mid S_h]\} \mid S_k, S_h\big]\big] \\
&\quad = E\big[E[Y_{a_k}^{T_k} - E[Y_{a_k}^{T_k} \mid S_k] \mid S_k, S_h] \times E[Y_{a_h}^{T_h} - E[Y_{a_h}^{T_h} \mid S_h] \mid S_k, S_h]\big] \\
&\quad = E\big[\{E[Y_{a_k}^{T_k} \mid S_k, S_h] - E[Y_{a_k}^{T_k} \mid S_k]\} \times E\{Y_{a_h}^{T_h} - E[Y_{a_h}^{T_h} \mid S_h] \mid S_k, S_h\}\big] \\
&\quad = E\big[\{E[Y_{a_k}^{T_k} \mid S_k] - E[Y_{a_k}^{T_k} \mid S_k]\} \times E\{Y_{a_h}^{T_h} - E[Y_{a_h}^{T_h} \mid S_h] \mid S_k, S_h\}\big] \\
&\quad = 0.
\end{aligned}
$$

The second equality follows from the correct specification of (5.5). The third equality holds because, by assumptions (5.1), (5.2), and (5.3), $Y_{a_k}^{T_k} \perp Y_{a_h}^{T_h} \mid S_k, S_h$. The fifth inequality holds because $Y_{a_k}^{T_k} \perp S_h \mid S_k$, again by assumptions (5.1), (5.2), and (5.3). □

THEOREM 5.2. *Let* $\varepsilon_{e_k} = Y_{e_k}^{T_k+s} - g^{-1}(\beta_0 + \beta_1 V_{a_k} + \beta_2 U_{a_k}^{T_k-b} + \beta_3 L_{a_k}^{T_k-b} + \beta_4 U_e^{T_k+f} + \beta_5 L_{e_k}^{T_k+f} + \beta_6' \boldsymbol{C}_k)$. *Then* ε_{e_k} *and* ε_{e_h} *are uncorrelated.*

The proof of Result 2 is very similar to the proof of Result 1 and we therefore omit it. It relies on the fact that, conditional on the fact that $T + f = T + s - b$ and, therefore, conditioning on $U_{e_k}^{T_k+f}$ and $L_{e_k}^{T_k+f}$ satisfies the conditions of assumptions (5.1), (5.2), and (5.3) and renders $Y_{e_k}^{T_k+s}$ independent of outcomes, vaccines, and covariates for other groups.

5.2.2. *Implementation*. Step 1 of the estimating procedure is the only departure from standard analyses that are available in many off-the-shelf software packages. The goal of this step is to select alter-ego pairs independently of their position in the network, in order to avoid possible bias due to selection related to covariates such as number of network ties, degree of connectedness, or centrality. First, randomly select one node from a list of all of the nodes in the network to be the first ego, then randomly select one node from the list of the first node's alters. This is the first alter-ego pair. Create a new list that excludes the two members of the pair and all of their alters. Repeat the selection step based on the new list. If the node selected has no alter on the list, exclude it from the list and select another node. Iterate until there is no node on the list with an alter on the list.

Because the number of groups selected in step 1 becomes the sample size of the analytical procedure, it is desirable to identify as many conditionally independent

groups as possible. Random selection of alter-ego pairs is the most conservative procedure in that it requires no additional assumptions in order to be appropriate. If there is no effect modification by correlates of position in the network (e.g., by degree of connectedness or centrality), then it might be possible to use results from the literature on graph cutting, in particular the literature on maximal independent sets, to increase the number of conditionally independent alter-ego pairs identified in step 1. Alternatively, prioritizing the selection of nodes with few network ties would be expected to increase the number of alter-ego pairs relative to random selection.

Steps 2 and 3 of the testing procedure are perfunctory. If we define $\boldsymbol{C}^* = (U_a^{T-b}, L_a^{T-b}, U_e^{T+f}, L_e^{T+f}, \boldsymbol{C})$ to be a new collection of covariates, then step 4 proceeds as in Sections 3 and 4. Interactions between components of $\boldsymbol{C}^*$ and the other predictors in the model can easily be accommodated. We estimate 95% confidence intervals for Con^* and Inf^* based on the estimates and standard errors calculated in Step 4. We reject the hypothesis of no contagion effect if our confidence interval for Con^* does not include the null value and we reject the hypothesis of no infectiousness effect if our confidence interval for Inf^* does not include the null value.

5.3. *Relaxing some assumptions*. Although we specified models for the conditional expectations of $Y_{e_k}^{T_k+s}$ and $Y_{a_k}^{T_k}$ that are linear in covariates above, any identified functional form for $E[Y_{e_k}^{T_k+s} \mid V_{a_k}, Y_{a_k}^{T_k} = 1, U_a^{T-b}, L_a^{T-b}, U_e^{T+f}, L_e^{T+f}, \boldsymbol{C}_k]$ and $E[Y_{a_k}^{T_k} \mid V_{a_k}, U_a^{T-b}, L_a^{T-b}, U_e^{T+f}, L_e^{T+f}, \boldsymbol{C}_k]$ can be used.

We assumed throughout that vaccination occurs before the start of follow-up, but this is not necessary for our methods. If vaccination can occur during follow-up, define V_i^t to be an indicator of having been vaccinated by time t. Assume that the effect of vaccination, including any infectiousness effect, is immediate. If an individual becomes infectious on day T, he would have been infected on day $T-b$. If he was vaccinated by time $T-b$, then the vaccine would have been in full effect at the time of infection. Then V_a^{T-b} can replace V_a as the "treatment" in the contagion, infectiousness, and indirect effects. We similarly redefine the summary measures for vaccinated and unvaccinated contacts of the alter and ego that appear in assumptions (5.1) through (5.3) and that are included in $\boldsymbol{C}$. Include V_e^{T-b} in the set of confounders because the mediator occurs at time T and, therefore, the ego's vaccination status at time $T-b$ suffices to control for any confounding.

We assumed throughout that the infectious and incubation periods (f and b) are constant across individuals. These assumptions, along with the assumption that the effect of vaccination is immediate, could be relaxed if the determinants of time to efficacy of vaccine, length of infectious period, and length of incubation period were observed covariates. In this case, we could, for example, infer effective time of vaccination, incubation period, and infectious period for each individual based on their covariates.

We assumed in Section 5.2 that the probability of disease transmission between two connected nodes does not depend on the type of tie. This assumption can be avoided with the addition of several covariates to models (5.4) and (5.5): we would condition on the type of tie that exists between the alter and the ego, and also include separate U and L terms for each type of tie. We also assumed in Section 5.2 that an individual will come into contact with any individual in the network who is not his contact with equal probability. This can be relaxed by expanding the k groups we define in Step 1 of the estimation procedure to include nodes within several degrees of separation from the alter and ego.

Perhaps the most challenging assumption we have made is that of conditional independence of the alter-ego pairs. To meet this assumption requires subsampling large nonoverlapping groups of nodes from the network, resulting in an sample size that is much smaller than the number of nodes in the network, and even then the assumption will likely be violated in some applications. Writing about considerably different target estimands, a handful of researchers have explored the tradeoff between subsetting a social network into large independent clusters and subsetting the network into smaller clusters that may be weakly correlated with one another. The former approach reduces or eliminates bias while the latter approach introduces bias but reduces variance by increasing the sample size [Eckles, Karrer and Ugander (2014), Toulis and Kao (2013), Ugander et al. (2013)]. Our goal in this paper is to ground the estimation of peer effects using network data in valid statistical estimation procedures with well-understood properties, and it is not clear how to quantify the bias that would be introduced by residual dependence across groups. However, this same kind of bias-variance tradeoff is an important and interesting direction for future research.

6. Simulations.

6.1. *Independent groups*. We ran simulations for three different sample sizes, $K = 200$, $K = 500$, and $K = 1000$ independent groups. Each group comprised an alter, an ego, and n_k mutual contacts. First, we generated K contact group sizes n_k by sampling from a Poisson distribution with mean $\lambda = 3$. Next, we assigned vaccination statuses to each individual in each group, including the alters and egos, with probability 0.4. We simulated the behavior of each group during a flu epidemic over 100 days. For the purposes of the simulation, we assumed that each member of a group had contact with all other members of the same group. Each day, an uninfected member of a group had a baseline probability of p_o of being infected from outside of the group, a baseline probability of p_u of being infected by any infectious, unvaccinated member of the same group and a baseline probability of p_v of being infected by any infectious, vaccinated member of the same group. If vaccinated, an individual's probability of being infected by any source was multiplied by $\delta \leq 1$. If infected on day t, an individual was infectious from

day $t+1$ through day $t+4$ and incapable of being infected or transmitting infection from day $t+5$ until the end of follow-up. This corresponds to an incubation period of $b=1$ and an infectious period of $f=3$, and it mimics the flu, for which the incubation period is between one and three days and the infectious period is between three and six days [Earn, Dushoff and Levin (2002)].

In all simulations, we fixed $p_o = 0.01$. We specified two different simulation settings for the parameters δ, p_v, and p_u, one setting corresponding to the null of no infectiousness or contagion effects ($\delta = 1$; $p_v = p_u = 0.5$) and one setting corresponding to the presence of protective contagion and infectiousness effects ($\delta = 0.2$; $p_v = 0.05$, $p_u = 0.5$). This choice of δ is consistent with some of the higher estimates of flu vaccine efficacy found in the literature [Osterholm et al. (2012)], and this choice of p_u is consistent with estimates of secondary attack rates for influenza [e.g., Yang et al. (2009)]. We are not aware of any empirical research that could inform our choice of p_v; simulations with $p_v = 0.1$ and $p_v = 0.15$ gave qualitatively similar results. We simulated 500 epidemics each under of the two scenarios, and for each simulation we estimated the infectiousness and contagion effects *Con* and *Inf* as follows: Among the subset of groups with $Y_a^T = 1$ and using a log-linear link function (assuming Poisson errors), we regressed Y_e^{T+s} on V_a and on the set of potential confounders comprised by the ego's vaccination status, the sum U_a^{T-b} of unvaccinated mutual contacts who were infectious at time $T-b$, and the sum L_a^{T-b} of vaccinated mutual contacts who were infectious at time $T-b$. We regressed Y_a^T on the same covariates using a logistic link function (assuming Bernoulli errors). The contagion and infectiousness effects are identified by the expressions given in (3.7) and (3.8) (with the list of covariates expanded to include U_a^{T-b} and L_a^{T-b}), evaluated at the sample mean values of the covariates. We bootstrapped the standard errors with 500 bootstrap replications. We also estimated the modified contagion and infectiousness effects *Con** and *Inf** by conditioning on U_e^{T+f} and L_e^{T+f} in addition to the covariates listed above.

The results are given in Table 1. For each simulation setting, that is, for each sample size (K) and for both the null hypothesis and the alternative hypothesis, we present the mean point estimates for the infectiousness and contagion effects on the ratio scale and the mean bootstrap standard error estimator. We also calculated a 95% confidence interval based on the 2.5th and 97.5th bootstrap quantiles for each simulation setting. For simulations under the null hypothesis, we report percent coverage of the null value 1, and for simulations under the alternative we report power, given by 100% minus the percent coverage. The point estimates are stable across sample sizes and the coverage of the basic bootstrap confidence interval is close to 95% under the null for all K. The power under the alternative is close 100% for both contagion effects, but for the infectiousness effects, and in particular for *Inf**, power is low when $K = 200$. This suggests that larger sample sizes may be needed to attain adequate power for tests of infectiousness effects. Under H_A, the estimated values of *Con** and *Inf** are incrementally closer to the null value

TABLE 1
Simulation results for independent groups

Under H_0

Number of groups	*Inf* (SE)	Coverage	*Con* (SE)	Coverage	Inf*(SE)	Coverage	*Con** (SE)	Coverage
$K = 200$	1.00 (0.14)	94%	1.01 (0.21)	93%	1.00 (0.14)	95%	1.01 (0.22)	92%
$K = 500$	1.01 (0.09)	94%	1.00 (0.12)	95%	1.01 (0.09)	94%	1.00 (0.13)	95%
$K = 1000$	1.01 (0.06)	95%	1.01 (0.09)	95%	1.01 (0.06)	95%	1.01 (0.09)	95%

Under H_A

Number of groups	*Inf* (SE)	Power	*Con* (SE)	Power	*Inf** (SE)	Power	*Con** (SE)	Power
$K = 200$	0.50 (0.18)	69%	0.42 (0.11)	99%	0.53 (0.19)	60%	0.44 (0.11)	95%
$K = 500$	0.50 (0.11)	98%	0.43 (0.06)	100%	0.54 (0.12)	92%	0.45 (0.07)	100%
$K = 1000$	0.50 (0.08)	100%	0.43 (0.05)	100%	0.53 (0.08)	100%	0.46 (0.05)	100%

of 1 than *Con* and *Inf*. The incubation and infectiousness periods that we specified and that are consistent with the flu leave little time for infectiousness or contagion effects of V_a on Y_e^{T+s} to operate through the flu statuses of mutual friends; this may explain the negligible differences between *Con* and *Con** and between *Inf* and *Inf** under H_A.

6.2. *Social network data*. We ran simulations for three different network sizes: 12,000 nodes, 10,000 nodes, and 8000 nodes. We simulated a network of 10,000 nodes as follows: first, we simulated 2000 independent groups of 5 nodes, with each group being fully connected (i.e., there are ties between each pair of nodes in the group of 5). For each node, we then added a tie to each out-of-group node with probability 0.0001. Because ties are undirected (if node i is tied to node j, then by definition node j is tied node i), this results in approximately 2 expected out-of-group ties per node. To simulate networks of size 12,000 and 8000, we simulated 2400 and 1600 independent groups, respectively, and scaled the probability of an out-of-group tie to maintain an expected value of approximately 2 for each node. This network structure could represent a sample of families living in a city, where individuals are fully connected to the members of their family and occasionally connected to members of other families. After running step 1 of the procedure outlined in Section 5.2, we were left with $K = 707$ alter-ego pairs for the network of size 12,000, $K = 581$ for the network of size 10,000, and $K = 466$ for the network of size 8000.

The procedure proposed in Section 5.2 for hypothesis testing using social network data suffers from low power due to the extraction of conditionally independent pairs of nodes from the network. As these simulations illustrate, this results

in a dramatic reduction in the sample size used for analysis. Because infectious outcomes sampled from nodes in a network are dependent, the effective sample size for inference about such outcomes will always be smaller than the observed number of nodes, and how much more information about the parameters of interest is available depends on the specific setting. Important areas for future research include determining the effective sample size when observations are sampled from a network and are therefore dependent, and developing methods that make use of all available information.

On each of these three fixed networks, we simulated 200 epidemics under the null of no infectiousness or contagion effect and 200 epidemics under the alternative. For each simulation, we assigned vaccination statuses to each individual in the network with probability 0.5. We then simulated the behavior of each group during a flu epidemic over 100 days. An uninfected node had a probability of $p_o = 0.01$ of being infected from outside of the network on day 1 and there were no outside infections thereafter. Under the alternative, on each day an uninfected node had a baseline probability of $p_u = 0.5$ of being infected by any infectious, unvaccinated contact and group and a baseline probability of $p_v = 0.01$ of being infected by any infectious, vaccinated contact. If vaccinated, an individual's probability of being infected by any source was multiplied by $\delta = 0.2$. Under the null, on each day an uninfected node had a probability of $p_u = p_v = 0.5$ of being infected by any infectious contact (i.e., node with which it shared a tie). To ensure that the contagion effect was null, we specified that $\delta = 1$, that is, that vaccination had no protective effect against contracting the flu. In both settings, if infected on day t an individual was infectious from day $t + 1$ through day $t + 4$ and incapable of being infected or transmitting infection from day $t + 5$ until the end of follow-up.

For each simulation, we estimated the infectiousness and contagion effects Con^* and Inf^* following the procedure described in Section 5.2. We evaluated these effects at the sample mean value of the covariates U_a^{T-b}, L_a^{T-b}, U_e^{T+f} and L_e^{T+f}. We bootstrapped the standard errors with 500 bootstrap replications. The results are given in Table 2. For each simulation setting, that is, for each network size and for both the null hypothesis and the alternative hypothesis, we present the mean point estimates for the infectiousness and contagion effects on the ratio scale, the mean bootstrap standard error estimator, and the percent coverage (of the null value 1) of the 95% confidence interval based on the 2.5th and 97.5th bootstrap quantiles. For simulations under the alternative hypothesis we calculated the power, given by 100% minus the percent coverage. For the 8000- and 10,000-node networks, there were 6 and 1 simulations, respectively, out of 200, for which the GLMs used to estimate the parameters involved in the contagion and infectiousness effects did not converge due to empty strata of the predictors. We omit these simulations from the results in Table 2, but note that in extreme cases convergence could be an issue in addition to power.

TABLE 2
Simulation results for network data

Under H_0

Network size	*Inf** (SE)	Coverage	*Con** (SE)	Coverage
8000 nodes	0.996 (0.001)	100%	1.205 (1.657)	96%
10,000 nodes	1.000 (0.001)	100%	1.183 (1.183)	94%
12,000 nodes	1.001 (0.001)	100%	1.166 (0.842)	94%

Under H_A

Network size	*Inf** (SE)	Power	*Con** (SE)	Power
8000 nodes	0.650 (0.259)	45%	0.168 (0.017)	99%
10,000 nodes	0.616 (0.072)	53%	0.164 (0.013)	100%
12,000 nodes	0.609 (0.054)	63%	0.164 (0.010)	100%

The point estimates are stable across network sizes and the coverage of the basic bootstrap confidence interval is close to or above 95% under the null for all network sizes. The estimates of contagion are biased under the null, with decreasing bias as network size increases. Given that coverage is close to the nominal level for all three network sizes, this may simply be a reflection of the moderate effective sample size: 707 alter-ego pairs are used in the analysis of the 12,000-node network, and at this sample size mild bias is perhaps not surprising. The power under the alternative is close to 100% for the contagion effect for all network sizes, but for the infectiousness effect power is low: 45% for the network of size 8000, increasing to 63% for the network of size 12,000. The low power combined combined with 100% coverage under the null suggests that our testing procedure for infectiousness effects may be conservative in some settings.

One concern that has been raised about previous uses of statistical models like GLMs and GEEs for network data is the possibility that the models lack any power to reject the null hypothesis when the alternative is true [Shalizi (2012)]. This is a concern because the models are inherently misspecified under the alternative hypothesis due to dependence among observations, even if they are correctly specified under the null hypothesis. Because the methods we propose here can be correctly specified under both the null and the alternative hypotheses, they can be powered to reject the null hypothesis when the infectiousness or contagion effect is present. We simulated a contagious process spreading through groups or networks of individuals rather than simulating from the GLMs that we use to analyze the data. This data-generating process may not be compatible with the parametric form of the GLMs that we fit. However, this is an issue of parametric misspecification; similar parametric mismatches are likely to plague many applications of GLMs to real data.

7. Discussion. We proposed methods for consistently estimating contagion and infectiousness effects in independent groups of arbitrary size; these methods are easy to implement and perform well in simulations. We extended our methodology to groups sampled from social network data, providing a theoretically justified method for using GLMs to analyze network data. Note that the principles we applied to GLMs can be applied to GEEs as well, resulting in correctly specified GEEs for network data. The principles that justify our use of GLMs to estimate the contagion and infectiousness effects are easily extended to any estimand for which GLMs would be a desirable modeling tool. However, our network data methods require a large amount of data and are not appropriate for small or dense networks. On the one hand, this highlights the fact that dependence among observations in networks reduces effective sample size and necessitates larger samples; on the other hand methods should be developed that can harness more information from the data and increase the power to detect contagion, infectiousness, and other causal effects.

Acknowledgments. The authors would like to thank Shira Mitchell and two anonymous reviewers for their helpful comments, which resulted in many significant improvements to this work.

SUPPLEMENTARY MATERIAL

Application of methods from "Vaccines, contagion, and social networks" to Harvard flu data (DOI: 10.1214/17-AOAS1023SUPP; .pdf). In the supplementary material, we describe an analysis of the Harvard flu network data using our methods.

REFERENCES

ALI, M. M. and DWYER, D. S. (2009). Estimating peer effects in adolescent smoking behavior: A longitudinal analysis. *J. Adolesc. Health* **45** 402–408.

ANDERSON, R. M. and MAY, R. M. (1985). Vaccination and herd immunity to infectious diseases. *Nature* **318** 323–329.

ARONOW, P. M. and SAMII, C. (2013). Estimating average causal effects under general interference. Technical report, arXiv:1305.6156.

BOWERS, J., FREDRICKSON, M. M. and PANAGOPOULOS, C. (2013). Reasoning about interference between units: A general framework. *Polit. Anal.* **21** 97–124.

BRESLOW, N. E. (1996). Generalized linear models: Checking assumptions and strengthening conclusions. *Stat. Appl.* **8** 23–41.

CACIOPPO, J. T., FOWLER, J. H. and CHRISTAKIS, N. A. (2009). Alone in the crowd: The structure and spread of loneliness in a large social network. *J. Pers. Soc. Psychol.* **97** 977.

CHOI, D. S. (2014). Estimation of monotone treatment effects in network experiments. ArXiv Preprint, arXiv:1408.4102.

CHRISTAKIS, N. A. and FOWLER, J. H. (2007). The spread of obesity in a large social network over 32 years. *N. Engl. J. Med.* **357** 370–379.

CHRISTAKIS, N. A. and FOWLER, J. H. (2008). The collective dynamics of smoking in a large social network. *N. Engl. J. Med.* **358** 2249–2258.

CHRISTAKIS, N. A. and FOWLER, J. H. (2010). Social network sensors for early detection of contagious outbreaks. *PLoS ONE* **5** e12948.

CHRISTAKIS, N. A. and FOWLER, J. H. (2013). Social contagion theory: Examining dynamic social networks and human behavior. *Stat. Med.* **32** 556–577. MR3042499

COHEN-COLE, E. and FLETCHER, J. M. (2008). Is obesity contagious? Social networks vs. environmental factors in the obesity epidemic. *J. Health Econ.* **27** 1382–1387.

EAMES, K. T. D. and KEELING, M. J. (2002). Modeling dynamic and network heterogeneities in the spread of sexually transmitted diseases. *Proc. Natl. Acad. Sci. USA* **99** 13330–13335.

EAMES, K. T. D. and KEELING, M. J. (2004). Monogamous networks and the spread of sexually transmitted diseases. *Math. Biosci.* **189** 115–130. MR2065412

EARN, D. J. D., DUSHOFF, J. and LEVIN, S. A. (2002). Ecology and evolution of the flu. *Trends Ecol. Evol.* **17** 334–340.

ECKLES, D., KARRER, B. and UGANDER, J. (2014). Design and analysis of experiments in networks: Reducing bias from interference. ArXiv Preprint, arXiv:1404.7530.

EUBANK, S., GUCLU, H., KUMAR, V. S. A., MARATHE, M. V., SRINIVASAN, A., TOROCZKAI, Z. and WANG, N. (2004). Modelling disease outbreaks in realistic urban social networks. *Nature* **429** 180–184.

FINE, P. E. M. (1993). Herd immunity: History, theory, practice. *Epidemiol. Rev.* **15** 265–302.

FOWLER, J. H. and CHRISTAKIS, N. A. (2008). Estimating peer effects on health in social networks: A response to Cohen-Cole and Fletcher; and Trogdon, Nonnemaker, and Pais. *J. Health Econ.* **27** 1400–1405.

GILL, J. (2001). *Generalized Linear Models*: *A Unified Approach*. *Quantitative Applications in the Social Sciences* **134**. Sage Publications, Thousand Oaks, CA.

HALLORAN, M. E. and HUDGENS, M. G. (2012). Causal inference for vaccine effects on infectiousness. *Int. J. Biostat.* **8** Art. 6. MR2925328

HALLORAN, M. E. and STRUCHINER, C. J. (1991). Study designs for dependent happenings. *Epidemiology* **2** 331–338.

HALLORAN, M. E. and STRUCHINER, C. J. (1992). Modeling transmission dynamics of stage-specific malaria vaccines. *Parasitol. Today* (*Regul. Ed.*) **8** 77–85.

HALLORAN, M. E. and STRUCHINER, C. J. (1995). Causal inference in infectious diseases. *Epidemiology* **6** 142–151.

IMAI, K., KEELE, L. and TINGLEY, D. (2010). A general approach to causal mediation analysis. *Psychol. Methods* **15** 309–334.

JOHN, T. J. and SAMUEL, R. (2000). Herd immunity and herd effect: New insights and definitions. *Eur. J. Epidemiol.* **16** 601–606.

KEELING, M. J. and EAMES, K. T. (2005). Networks and epidemic models. *J. R. Soc. Interface* **2** 295–307.

KELLER, M. A. and STIEHM, E. R. (2000). Passive immunity in prevention and treatment of infectious diseases. *Clin. Microbiol. Rev.* **13** 602–614.

KLOVDAHL, A. S. (1985). Social networks and the spread of infectious diseases: The AIDS example. *Soc. Sci. Med.* **21** 1203–1216.

KLOVDAHL, A. S., POTTERAT, J. J., WOODHOUSE, D. E., MUTH, J. B., MUTH, S. Q. and DARROW, W. W. (1994). Social networks and infectious disease: The Colorado Springs study. *Soc. Sci. Med.* **38** 79–88.

LATORA, V., NYAMBA, A., SIMPORE, J., SYLVETTE, B., DIANE, S., SYLVERE, B. and MUSUMECI, S. (2006). Network of sexual contacts and sexually transmitted HIV infection in Burkina Faso. *J. Med. Virol.* **78** 724–729.

LAZER, D., RUBINEAU, B., CHETKOVICH, C., KATZ, N. and NEBLO, M. (2010). The coevolution of networks and political attitudes. *Polit. Commun.* **27** 248–274.

LYONS, R. (2011). The spread of evidence-poor medicine via flawed social-network analysis. *Statistics*, *Politics*, *and Policy* **2**.

NOEL, H. and NYHAN, B. (2011). The unfriending problem: The consequences of homophily in friendship retention for causal estimates of social influence. *Soc. Netw.* **33** 211–218.

O'BRIEN, K. L. and DAGAN, R. (2003). The potential indirect effect of conjugate pneumococcal vaccines. *Vaccine* **21** 1815–1825.

OGBURN, E. L. and VANDERWEELE, T. J. (2014). Causal diagrams for interference. *Statist. Sci.* **29** 559–578. MR3300359

OGBURN, E. L., VANDERWEELE, T. J. and CHRISTAKIS, N. A. (2017). Supplement to "Vaccines, contagion, and social networks." DOI:10.1214/17-AOAS1023SUPP.

OSTERHOLM, M. T., KELLEY, N. S., SOMMER, A. and BELONGIA, E. A. (2012). Efficacy and effectiveness of influenza vaccines: A systematic review and meta-analysis. *Lancet, Infect. Dis.* **12** 36–44.

PEARL, J. (2001). Direct and indirect effects. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence* 411–420.

ROBINS, J. M. and GREENLAND, S. (1992). Identifiability and exchangeability for direct and indirect effects. *Epidemiology* **3** 143–155.

ROBINS, J. M. and RICHARDSON, T. S. (2010). Alternative graphical causal models and the identification of direct effects. In *Causality and Psychopathology*: *Finding the Determinants of Disorders and Their Cures* (P. Shrout, ed.). Oxford Univ. Press.

ROSENBAUM, P. R. (2007). Interference between units in randomized experiments. *J. Amer. Statist. Assoc.* **102** 191–200. MR2345537

ROSENQUIST, J. N., MURABITO, J., FOWLER, J. H. and CHRISTAKIS, N. A. (2010). The spread of alcohol consumption behavior in a large social network. *Ann. Intern. Med.* **152** 426–433.

SHALIZI, C. R. (2012). Comment on "Why and when 'flawed' social network analyses still yield valid tests of no contagion." *Statistics, Politics, and Policy* **3**.

SHALIZI, C. R. and THOMAS, A. C. (2011). Homophily and contagion are generically confounded in observational social network studies. *Sociol. Methods Res.* **40** 211–239. MR2767833

TOULIS, P. and KAO, E. (2013). Estimation of causal peer influence effects. In *Proceedings of the* 30*th International Conference on Machine Learning* 1489–1497.

UGANDER, J., KARRER, B., BACKSTROM, L. and KLEINBERG, J. (2013). Graph cluster randomization: Network exposure to multiple universes. In *Proceedings of the* 19*th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 329–337. ACM.

VALERI, L. and VANDERWEELE, T. J. (2013). Mediation analysis allowing for exposure-mediator interactions and causal interpretation: Theoretical assumptions and implementation with SAS and SPSS macros. *Psychol. Methods* **18** 137–150.

VANDERWEELE, T. J. (2011). Sensitivity analysis for contagion effects in social networks. *Sociol. Methods Res.* **40** 240–255. MR2767834

VANDERWEELE, T. J., OGBURN, E. L. and TCHETGEN TCHETGEN, E. J. (2012). Why and when "flawed" social network analyses still yield valid tests of no contagion. *Statistics, Politics, and Policy* **3** 1–11.

VANDERWEELE, T. J. and TCHETGEN TCHETGEN, E. J. (2011a). Bounding the infectiousness effect in vaccine trials. *Epidemiology* **22** 686–693.

VANDERWEELE, T. J. and TCHETGEN TCHETGEN, E. J. (2011b). Effect partitioning under interference in two-stage randomized vaccine trials. *Statist. Probab. Lett.* **81** 861–869. MR2793754

VANDERWEELE, T. J., TCHETGEN TCHETGEN, E. J. and HALLORAN, M. E. (2012). Components of the indirect effect in vaccine trials: Identification of contagion and infectiousness effects. *Epidemiology* **23** 751–761.

VAN DER LAAN, M. J. (2012). Causal inference for networks. U.C. Berkeley Division of Biostatistics Working Paper Series, Working Paper 300.

YANG, Y., SUGIMOTO, J. D., HALLORAN, M. E., BASTA, N. E., CHAO, D. L., MATRAJT, L., POTTER, G., KENAH, E. and LONGINI, I. M. (2009). The transmissibility and control of pandemic influenza A (H1N1) virus. *Science* **326** 729–733.

DEPARTMENT OF BIOSTATISTICS
JOHNS HOPKINS SCHOOL OF PUBLIC HEALTH
615 N. WOLFE ST.
BALTIMORE, MARYLAND 21205
USA
E-MAIL: eogburn@jhsph.edu

DEPARTMENT OF EPIDEMIOLOGY
HARVARD SCHOOL OF PUBLIC HEALTH
677 HUNTINGTON AVE
BOSTON, MASSACHUSETTS 02115
USA
E-MAIL: tvanderw@hsph.harvard.edu