Published in final edited form as:

Curr Epidemiol Rep. 2016 December ; 3(4): 297–305. doi:10.1007/s40471-016-0086-4.

Dependent Happenings: A Recent Methodological Review

M. Elizabeth Halloran^{1,2} and Michael G. Hudgens³

¹Center for Inference and Dynamics of Infectious Diseases, Vaccine and Infectious Disease Division, Fred Hutchinson Cancer Research Center

²Department of Biostatistics, School of Public Health, University of Washington

³Department of Biostatistics, Gillings School of Global Public Health, University of North Carolina at Chapel Hill

Abstract

One hundred years ago Sir Ronald Ross published his treatise on a general Theory of Happenings. Dependent happenings are those in which the frequency depends on the number already affected. When there is dependency of events, interventions can have different types of effects. Interventions such as vaccination can have direct protective effects for the person receiving the treatment, as well as indirect/spillover effects for others in the population. Causal inference is a framework for carefully defining the causal effect of a treatment, exposure, or policy, and then determining conditions under which such effects can be estimated from the observed data. We consider here scenarios in which the potential outcomes of an individual can depend on the treatment of other individuals in the population, known as causal inference with interference. Much of the research so far has assumed the population is divided into groups or clusters, and individuals can interfere with others within their clusters but not across clusters. Recent developments have assumed more general forms of interference. We review some of the different types of effects that have been defined for dependent happenings, particularly using the methods of causal inference with interference. Many of the methods are applicable across disciplines, such as infectious diseases, social sciences, and economics.

Keywords

causal inference; counterfactual; dependent happenings; experimental design; herd immunity; indirect effects; networks; peer influence effects; potential outcome; spillover effects; SUTVA

1 Introduction

One hundred years ago Sir Ronald Ross (1916) published his treatise on a general Theory of Happenings. The happening element corresponds in modern parlance to an occurrence or event. Ross differentiated independent from dependent happenings. Independent happenings were “those in which the frequency of the happening is *independent* of the number of individuals already affected” (Ross 1916), p. 211). Dependent happenings were those in which the frequency depends on the number already affected. To the first class belong such happenings as noninfectious diseases and accidents. To the second class belong “infectious diseases, membership of societies and sects with propagandas, trade-unions, political parties,

etc., due to propagation from within, that is, individual to individual” (Ross 1916), p. 211). Ross had been awarded the second Nobel prize in Medicine in 1902 for elucidating the role of mosquito vectors in the transmission cycle of malaria. He was also an amateur mathematician and developed the original malaria models to quantify the effects of interventions (Ross 1911, 1915). Ross had published an addendum in the second edition of his book on the prevention of malaria on a preliminary general Theory of Happenings in 1911.

A principled approach to evaluating the effects of interventions that has fostered clarity in thinking about both experimental and observational studies is causal inference (Neyman 1923, Rubin 1980, Holland 1986, Robins 1986). Causal inference is a framework for carefully defining the causal effects of a treatment, exposure, or policy, then determining conditions under which such effects can be estimated from the observed data. In causal inference, each individual is assumed to have a potential outcome that would occur under each of the different possible treatments. The outcome is potential in that it only occurs if, in fact, a particular treatment is given. Usually the treatment assignment of one person is assumed to not affect the potential outcomes of another person, called the assumption of *no interference* by Cox (1958). The assumption of no interference is a component of the Stable Unit Treatment Value Assumption (SUTVA) (Rubin 1978). The assumption of no interference is also called the individualistic treatment response (Manski 2013). Under dependent happenings, interference will likely be present. In the presence of interference, treatments can have direct effects on the person receiving the treatment as well as indirect effects on people not receiving treatment. An individual may have many different potential outcomes depending on the treatment status of the others.

Several branches of research have their roots in the work by Ross. One is the broad field of infectious disease transmission models, including modeling the effects of interventions in populations, which Ross called *a priori* pathometry. Another is the large and growing literature on estimating parameters for these transmission models. Another area is methods to evaluate interventions in populations based on data from field studies. In this review, our focus is on the latter, in particular, on drawing inference about effects of interventions in the presence of interference, that is, for dependent happenings.

Here we consider examples largely in the epidemiological context, in particular infectious diseases and the effects of vaccines. However, much of the research on causal inference with interference is motivated by applications in the social sciences and economics, which also come under the general heading of dependent happenings in the sense of Ross. Many of the methods are applicable across disciplines.

2. Partial Interference

Motivated by an interest in dependent happenings, Halloran and Struchiner (1991, 1995) defined direct, indirect, total and overall effects in the presence of interference (Figure 1). Consider a cluster (or group) of individuals under two scenarios. Under scenario A on the left, a certain portion of individuals in the cluster is vaccinated and the rest remains unvaccinated. Under scenario B on the right, no one in the cluster is vaccinated. The *direct*

effect of vaccination under scenario A is defined by comparing the average outcome when an individual is vaccinated with the average outcome when an individual is not vaccinated. The *indirect effect* is defined as a contrast between the average outcome when an individual is not vaccinated under scenario A compared with scenario B. The *total effect* is defined by comparing the average outcome when an individual is vaccinated under scenario A to the average outcome when an individual is not vaccinated in scenario B. In general the total effect can be decomposed into a function of the direct and indirect effects (Hudgens and Halloran 2008, VanderWeele and Tchetgen Tchetgen 2011a). The *overall effect* is defined by the contrast in the average outcome in the entire cluster under scenario A compared to the average outcome of the entire cluster under scenario B.

With vaccination, indirect effects are produced by an increase in herd immunity, the collective immunological status of the population. In other fields, indirect effects are sometimes called *spillover effects* or *peer influence effects* (Toulis and Kao 2013). In mediation analysis, the terms direct, indirect, and overall effects have different meanings than used here (VanderWeele et al 2012).

The potential outcome approach to causal inference typically assumes no interference between individuals. Assuming no interference, if there is one treatment and a control, then an individual has two potential outcomes. The individual causal effect, which in general cannot be observed, is the difference between the two potential outcomes under treatment and control. The average causal effect is the average of the difference between the potential outcomes if everyone received treatment and if everyone received control. Drawing inference about treatment effects generally requires knowledge or modeling of the mechanism by which individuals select or are assigned treatment. Under an independent assignment mechanism, such as randomization, one can construct an unbiased estimator of the average causal effect from the observed average outcomes in the treatment and control groups (Neyman 1923).

The situation is more complicated under dependent happenings in the presence of interference. Halloran and Struchiner (1995) proposed individual-level causal estimands in the presence of interference by letting the potential outcomes for any individual depend on the vector of treatment assignments to other individuals in the group (Rubin 1978, 1990). However, with a binary outcome and one treatment and one control, if there are N people in a population, there are 2^N possible treatment vectors and 2^N possible potential outcomes. Causal inference in the presence of so many potential outcomes would seem difficult, if not impossible, without making additional assumptions.

One possible additional assumption which simplifies the problem considerably is to assume individuals can be partitioned into groups (or clusters) wherein there is no interference between individuals in different groups. Sobel (2006) referred to this assumption as *partial interference*. As a motivating example, he considered interference in a housing mobility experiment in poor neighborhoods in five cities. Eligible residents were randomly assigned to receive one of two forms of relocation assistance or no assistance (control). It was plausible that many of the participants knew each other and might influence each other's behavior. Sobel assumed interference could occur within but not across sites, i.e., there was

partial interference. In the next sections, we consider approaches that assume partial interference.

2.1 Two-stage randomization

Suppose there are two different vaccination strategies (or policies or programs) under consideration. The first strategy might be to vaccinate 60% of the individuals in a population, the other strategy might be to vaccinate no one. A question of interest then is what are the direct, indirect, total and overall effects of the 60% vaccination strategy compared to no vaccination strategy? Assuming partial interference, Hudgens and Halloran (2008) defined group- and population-level causal estimands for direct, indirect, total, and overall causal effects of treatment under two different treatment allocations. To obtain unbiased estimators of the population-level causal estimands, Hudgens and Halloran (2008) proposed a two-stage randomization scheme, the first stage at the group level, the second stage at the individual level within groups. For example, suppose there are 20 groups of individuals. At the first stage, we could randomize 10 groups to the 60% vaccination strategy and the remaining groups to the no vaccination strategy. In the second stage within 10 groups we would randomly assign 60% of individuals to receive vaccine, and in the other 10 groups no one would receive vaccine. Unbiased estimates of the direct, indirect, and total effects can be obtained by contrasts in average outcomes among vaccinated and unvaccinated individuals under the different vaccination strategies. Likewise, contrasts in average outcomes of all individuals under the two vaccination strategies yield unbiased estimates of the overall effect.

To draw inference about the four effects, variance estimators and confidence intervals are needed. Hudgens and Halloran (2008) developed variance estimators under an additional assumption they call *stratified interference*, which assumes the indirect treatment effects may depend only on the proportion of other individuals in the group which receive treatment. Tchetgen Tchetgen and VanderWeele (2012), Liu and Hudgens (2014), and Rigdon and Hudgens (2015) develop exact and asymptotic confidence intervals.

Baird et al (2014) considered a similar two-stage randomized design in the context of economic experiments to measure indirect/spillover effects. They refer to the level of (treatment) coverage in a cluster as the *saturation level* and their study design as the *randomized saturation design*. Their causal estimands, the intention to treat effect, spillover on the non-treated effect, and total causal effect, are analogous to the total, indirect, and overall effects, respectively, defined above. Baird et al also consider optimal design of two stage randomized trials, particularly choosing the optimal coverage level, to detect the different effects. Based on these considerations, they designed an experiment in Malawi to assess whether cash transfers help adolescent girls improve schooling outcomes and also delay marriage and pregnancy. Their study design accounted for possible spillover effects to non-beneficiaries as well as the participating adolescent girls in the same communities.

2.2 Partially randomized and observational studies

Often studies are not randomized at two stages, but only at the individual level, the cluster level, or neither. For example, Sur et al (2009) describe a cluster-randomized study in India

where 80 geographic clusters were randomized to receive either typhoid vaccine or control (hepatitis A) vaccine. A subset of individuals in these clusters chose to receive the study vaccine (typhoid or control). Similarly, Moulton et al (2001) designed a cluster-randomized study to estimate total, indirect and overall effectiveness of pneumococcal vaccine in infants. The design entailed randomizing clusters to receive the pneumococcal vaccine or a control vaccine, with individuals self-selecting to enroll or not in the trial. See Hayes and Moulton (2009) for a general overview of cluster-randomized trials, and Halloran (2012) for discussion of the related minicommunity design for assessing indirect effects of vaccination.

In studies which do not utilize two-stage randomization, the estimators described in the previous section would in general be biased or inconsistent due to potential confounding between treatment and the outcome. In the observational setting where the treatment assignment mechanism is not known and there is no interference, the propensity score, the probability an individual receives a treatment assignment based on a function of observed pretreatment covariates, is one method to adjust for confounding (Rosenbaum and Rubin 1983). Recently propensity score methods have been extended to the setting where interference may be present. For example, Hong and Raudenbush (2006) considered interference in the context of the effect on reading scores of children of being retained in kindergarten versus being promoted to the first grade. They classified schools by whether they retained a high proportion or a low proportion of kindergartners. Interference within a school was assumed to be summarized in the dichotomous school treatment of high or low retention rate. The study was observational at two levels: schools were not randomized to have high or low retention, and students at risk to be retained were not randomized to be retained. Assuming partial interference, Hong and Raudenbush used a multilevel propensity score stratification allowing for interference. The school propensity of adopting a high retention rate was estimated using pretreatment information. Then the child propensity of repetition in high-retention schools and the child propensity of repetition in low-retention schools were estimated also based on pretreatment information. Estimation of the causal effects used stratifications by the estimated school and child propensity scores.

Tchetgen Tchetgen and VanderWeele (2012) also used group-level propensity scores to develop inverse probability weighted (IPW) estimators of the direct, indirect, total and overall causal effects in observational studies in the absence of two-stage randomization. These IPW estimators can be viewed as a generalization of the usual IPW estimator of the causal effect of a treatment in the absence of interference. Lundin and Karlsson (2014) proposed similar IPW estimators of direct, indirect, and total effects under interference where treatment assignment is randomized only at the first stage, and in some groups, all individuals remain untreated (Figure 1). They were motivated by the Positive Parenting Program, an Australian parenting support program tested in pre-schools in Uppsala, Sweden. Perez-Heydrich et al (2014) and Liu et al (2016) considered the asymptotic properties of different IPW estimators in the presence of partial interference.

2.3 Example

To illustrate estimates of direct, indirect, total and overall effects of vaccination, we present results of an analysis by Perez-Heydrich et al (2014). They used the IPW estimators

proposed by Tchetgen Tchetgen and VanderWeele (2012) to estimate the effects of cholera vaccination in a study in Matlab, Bangladesh. All children and women were individually randomized to either of two cholera vaccines or a placebo, though not all participated in the vaccine trial. Unvaccinated individuals included eligible non-participants and placebo recipients. Vaccinated individuals included recipients of either vaccine. This was an individually randomized study, so neighborhoods (groups) were defined by using a clustering algorithm on geo-referenced data of the location of individuals' households. Effect estimates were calculated based on the difference of the IPW average outcomes in the relevant groups (Figure 2). Levels of vaccine coverage are denoted by α and α' .

The direct effect estimates (Figure 2A) generally decrease with increasing α . The indirect effect estimates (Figure 2B) comparing incidence of cholera among unvaccinated individuals at coverage levels α and α' tend to increase as the difference in the coverage $\alpha - \alpha'$ increases. The total effect contours (Figure 2C) are roughly vertical, suggesting the estimated risk of cholera when vaccinated tends to be the same regardless of coverage. Of particular interest, the total effect estimate of being *vaccinated* in a neighborhood with $\alpha' = 60\%$ coverage compared to being *unvaccinated* in a neighborhood with $\alpha = 32\%$ coverage is 5.9 (95% confidence interval 3.0, 8.8), an order of magnitude greater than the estimated direct effect of 0.6 at 60% coverage (Figure 2A). Thus taking indirect effects into account leads to a much different conclusion about the vaccine. The estimates of overall effects (Figure 2D) exhibit a similar pattern to the indirect estimates. An R package "inference" is available for conducting this type of analysis (Saul 2015).

3. General Interference

Individuals do not necessarily congregate in identifiable groups wherein they interact only within those groups but not with individuals in other groups, as is assumed with partial interference. Recent methodological developments for casual inference in the presence of interference allow for more general forms of interference. In the general interference setting, each individual may be assumed to have a unique set of other individuals whose treatment might affect the outcome of the individual. This interference structure can be represented by a network.

Network theory, also called graph theory, provides a formal framework for describing interactions between individuals. A network is defined by a collection of N nodes or vertices, and the edges or ties connecting the nodes. In the context here, the nodes are individuals and the edges are connections to other individuals which may give rise to interference.

Interference and dependent happenings can be propagated through networks. Computational social science is a growing field, fueled by massive data sets on interactions that form social networks (Lazer et al 2009). Companies such as Google and Facebook collect huge data sets on interactions and perform randomized experiments within these contexts. Wireless sensor network technology (Salathé et al 2010) and geographical positions from cell phone data (Wesolowski et al 2015) make electronic monitoring of close contacts and movement

possible. Infectious diseases also propagate across networks on multiple scales (Balcan et al 2009).

Recent approaches to assessing treatment effects in the presence of general interference, some of which rely on network theory, include Rosenbaum (2007), Toulis and Kao (2013), Ugander et al (2013), Eckles et al (2014), van der Laan (2014), Sofrygin and van der Laan (2015), Aronow and Samii (2015), and Liu et al (2016) among others. These general interference methods typically assume one finite population of N individuals. For each individual, a set of other individuals is defined that may interfere with that individual. Different names have been used to describe these sets, including: *interference sets* (Liu et al 2016), *neighborhoods* (Toulis and Kao 2013, Ugander et al 2013), and *friends* (van der Laan 2014, Sofrygin and van der Laan 2015) of the individuals. From the network perspective, the interference sets can be represented by an adjacency matrix. Typically the interference sets for each individual in the population are assumed known and fixed. It is also usually assumed that the number of individuals in these interference sets is smaller than N and that any indirect or spillover effects on an individual, if present, emanate from his/her known interference set.

Aside from these commonalities, there is considerable variability in recent approaches to causal inference with interference regarding how treatment/exposure effects (i.e., causal estimands) are defined and the mode of inference adopted. A comprehensive comparison is beyond this review, both in scope and technical detail. Below we highlight a few approaches recently proposed.

3.1 Randomized Experiments

Aronow and Samii (2015) considered randomized experiments in the presence of arbitrary interference. Their approach depends on knowing the experimental design by which individuals are assigned treatment. They define an *exposure mapping* which relates the individual experimental treatment assignments to the exposures received by individuals depending on with whom they interact and other underlying features of the population. Their exposure mapping is equivalent to Manski's (2013) "effective treatments" function. A population is considered to have a number of possible exposure conditions and causal estimands are defined as contrasts between average individual potential outcomes for two different exposure levels. The average potential outcome at any exposure level can be estimated because, by design, the probability of the exposure conditions is known for each individual. Aronow and Samii (2015) proposed IPW estimators for inference in this setting.

Ugander et al (2013) and Eckles et al (2014) considered randomized experiments on networks where treatment of one individual could have indirect/spillover effects on neighboring individuals. Their causal estimand of interest is the average of the difference in the individual potential outcomes under two extreme assignments, one where every individual in a network receives treatment and one where no individuals receive treatment. If individuals were independently randomized to treatment and control, for individuals who share many edges with other individuals, the probability all of their neighbors would (i) all receive treatment or (ii) all receive control would be low. Thus Ugander et al considered a partition of the network into clusters of individuals and proposed randomizing some of the

clusters to receive treatment and the remaining clusters to receive control. This is called *graph cluster randomization*. Note these clusters are not the same as the interference sets described above. Ugander et al introduced the notion of *network exposure* wherein an individual is network exposed to treatment (control) if the individual's response under a particular assignment vector is the same as if everyone in the network had received the treatment (control). An unbiased estimator of the average treatment effect using inverse probability weighting was derived for any randomization design for which the network exposure probabilities can be explicitly computed.

Motivated by a randomized experiment by Facebook, Toulis and Kao (2013) propose causal estimands for peer influence (indirect) effects describing interference in a social network. For each individual, a neighborhood is defined by the other individuals with whom he/she shares an edge. If at least one neighbor receives treatment, then the individual is considered *exposed to peer influence effects*. The potential outcome for each individual can depend on his treatment and that of his neighbors. Toulis and Kao defined two main causal estimands. The causal estimand for the *primary effect* is the average over the whole population of the difference in the individual outcomes if an individual received treatment versus received control when everyone else in the neighborhood received control. The main causal estimands for peer influence effects are defined by fixing the specific number of neighbors who receive treatment. For example, if k neighbors receive treatment, the k -level causal estimand for peer-influence effects is averaged over individuals with at least k neighbors. Two estimation procedures are proposed: a frequentist model-based estimator assuming a certain sequential randomization design and known network, and a Bayesian approach which accounts for uncertainty in the network topology.

We note there is a large, related literature on peer effects in social sciences dating back to Manski's seminal paper (1993) on the "reflection" problem. Other recent developments of statistical methods for drawing inference about peer effects are described in An (2011) and VanderWeele and An (2013). O'Malley and Onnela (2014) give an overview, comparing approaches in social network analysis and network science.

3.2 Observational Studies

van der Laan (2014) and Sofrygin and van der Laan (2015) considered statistical inference about causal effects in the presence of general interference in the observational setting. They defined the population of interest to be a set of (possibly) dependent individuals and assumed only a single draw from the true data generating distribution is observed. That is, unlike traditional statistical inference, multiple independent and identically distributed (iid) replicates were not assumed. With partial interference one might assume the groups are iid, permitting application of existing statistical theory (e.g., see Perez-Heydrich et al 2014). However, with general interference observing iid replicates is generally not possible, such that standard large sample frequentist approaches do not apply. van der Laan (2014) and Sofrygin and van der Laan (2015) derived the asymptotic properties of targeted maximum likelihood estimators in this setting, providing a method for valid statistical inference in the presence of general interference. See also Ogburn and VanderWeele (2014b) for related work allowing for general interference in the observational setting.

4 Infectiousness and Contagion Effects

It is often of interest in causal inference to understand the mechanism or pathway through which an exposure or treatment has an effect on the outcome of interest. Below we describe two possible causal mechanisms of interference: the infectiousness and contagion effects.

The vaccine effect on infectiousness is the reduction in transmissibility from a vaccinated versus unvaccinated infected person during a contact with a susceptible person. In a study in Niakhar, Senegal, for example, Préziosi and Halloran (2003) estimated the relative reduction in infectiousness of a vaccinated pertussis case compared to an unvaccinated case to be 67% (95% confidence interval 29, 86). Even if vaccination is randomized, the infectiousness effect is measured only in people who become infected, a post-randomization variable, so the estimate could be subject to selection bias (Hudgens and Halloran 2006). Combining causal inference with interference with principal stratification (Frangakis and Rubin 2002), VanderWeele and Tchetgen Tchetgen (2011b) and Halloran and Hudgens (2012a) proposed causal estimands of the infectiousness effect in households of size two, assuming interference occurred only within households, that is partial interference. Though this causal effect is not identifiable without further assumptions, bounds can be identified (VanderWeele and Tchetgen Tchetgen, 2011b; Halloran and Hudgens 2012a, 2012b) and sensitivity analyses performed (VanderWeele et al 2014).

Widespread vaccination in a population can produce indirect effects either by preventing individuals from becoming infected or by reducing the ability of infected vaccinated individuals to transmit to others. VanderWeele et al (2012) differentiated these two components, calling the first the *contagion effect* and the second the *infectiousness effect* (as discussed in the preceding paragraph). Using ideas analogous to mediation analysis, they showed how the indirect effect can be decomposed into these two components in households of size two, assuming partial interference. Ogburn and VanderWeele (2014b) extended these methods to groups of arbitrary size and one large network without assuming partial interference.

5 Other Approaches to Causal Inference with Interference

Most of the approaches described above utilize the potential outcomes framework for causal inference to assess treatment or exposure effects in the presence of interference. Randomization-based inference or large sample frequentist methods are typically employed for drawing statistical inferences. Below brief descriptions are provided of some recently proposed alternative approaches to causal inference in the presence of interference.

Carnegie et al (2016) considered an epidemic (mathematical) modeling approach for estimating the overall treatment effect in the presence of interference both within and across clusters in cluster-randomized trials of infectious disease.

Ogburn and VanderWeele (2014a) extended causal diagrams (or graphs) to the interference setting. They considered how causal graphs can help distinguish among causal mechanisms that give rise to interference, including the infectiousness and contagion effects described above.

Manski (2016) considered decisions faced by a health planner concerned with vaccination policy, in particular whether to mandate vaccination or allow individuals to self-select whether to be vaccinated. An econometric perspective was adopted which includes consideration of the indirect effect of vaccination on persons not vaccinated. The health policy decision was framed as an optimization problem where a certain cost function is to be minimized which accounts for both the social cost of illness and the social cost of vaccination. Manski then considered the optimal decision in settings where the indirect effect is assumed monotone with increasing vaccine coverage but otherwise unknown. Seasonal influenza vaccination was used as an example. This econometric perspective differs from the potential outcomes based approaches described above in that Manski did not seek to estimate or draw inference about the indirect effects of vaccination. Instead the goal was to determine the optimal vaccination policy given only partial knowledge of the indirect effect.

Laber et al (2016) developed a Bayesian on-line estimation method to determine optimal treatment allocation strategies for control of white-nose syndrome in bats, a usually fatal disease caused by a fungus. The possible interference structure is described by a network where the nodes are counties in the United States. Interference and spillover effects are governed by a model describing the probability of spatial spread of white-nose syndrome between counties.

6 Conclusions

One hundred years after publication of Ross' paper on dependent happenings, research on estimating treatment effects in the presence of interference is burgeoning. One current trend is to develop methods in the presence of general interference. These methods often involve using concepts from network theory. Applications motivating this work includes social media, economics and infectious diseases.

Future studies to evaluate causal effects under interference may continue to rely on cluster-randomization with the accompanying assumption of partial interference. But increasingly information on how people travel between clusters will be collected and analyzed to better understand contamination across clusters. Methods are needed that incorporate such information into both the causal estimands of interest and methods for inference. In the infectious disease context, genomic data is increasingly being gathered on the pathogen sequences. Phylogenetic information could also contribute to understanding the network structure of contacts and transmission. Much research remains to be done in developing methods for causal inference in the presence of interference.

Acknowledgments

This work was partially funded by National Institutes of Health grants R01 AI085073 and R37 AI032042.

References

An, W. The Sage Handbook of Social Network Analysis. London: Sage; 2011. Models and methods to identify peer effects; p. 515-532.

- Aronow PM, Samii C. Estimating average causal effects under interference between units. arXiv: 1305.6156v2. 2015:1–35. Developed a general approach for estimating effects of interventions in the presence of general forms of interference.
- Baird S, Bohren A, McIntosh C, Ozler B. Designing experiments to measure spillover effects. PIER Working Paper 14-032. 2014 Considered design of randomized experiments to draw inference about causal indirect (spillover), total, and overall effects in the context of economic studies with an empirical example.
- Balcan D, Colizza V, Goncalves B, Hud H, Ramasco JJ, Vespignani A. Multiscale mobility networks and the spatial spreading of infectious diseases. *Proceedings of the National Academy of Sciences*. 2009; 106(51):21484–21489.
- Carnegie NB, Wang R, De Gruttola V. Estimation of the overall treatment effect in the presence of interference in cluster-randomized trials of infectious disease prevention. *Epidemiologic Methods*. 2016 in press.
- Cox, DR. *Planning of Experiments*. New York: John Wiley and Sons, Inc; 1958.
- Eckles D, Karrer B, Ugander J. Design and analysis of experiments in networks: Reducing bias from interference. arXiv:1404.7530. 2015
- Frangakis CE, Rubin DB. Principal stratification in causal inference. *Biometrics*. 2002; 58:21–29. [PubMed: 11890317]
- Halloran ME. The minicommunity for assessing indirect effects of vaccination. *Epidemiologic Methods*. 2012; 1(1):83–105. [PubMed: 23599908]
- Halloran ME, Hudgens MG. Causal vaccine effects for infectiousness. *International Journal of Biostatistics*. 2012a; 8(2) Article 6.
- Halloran ME, Hudgens MG. Comparing bounds on vaccine effects on infectiousness. *Epidemiology*. 2012b; 23:931–932. [PubMed: 23038122]
- Halloran ME, Struchiner CJ. Study designs for dependent happenings. *Epidemiology*. 1991; 2:331–338. [PubMed: 1742381]
- Halloran ME, Struchiner CJ. Causal inference for infectious diseases. *Epidemiology*. 1995; 6:142–151. [PubMed: 7742400]
- Hayes, RJ., Moulton, LH. *Cluster Randomised Trials: A Practical Approach*. Boca Raton, FL: Chapman and Hall/CRC; 2009.
- Holland PW. Statistics and causal inference. *J Am Stat Assoc*. 1986; 81:945–960.
- Hong S, Raudenbush G. Evaluating kindergarten retention policy: A case study of causal inference for multilevel observational data. *Journal of the American Statistical Association*. 2006; 101:901–910.
- Hudgens MG, Halloran ME. Causal vaccine effects on binary post-infection outcomes. *J Am Stat Assoc*. 2006; 101:51–64. [PubMed: 19096723]
- Hudgens MG, Halloran ME. Towards causal inference with interference. *J Am Stat Assoc*. 2008; 103:832–842. [PubMed: 19081744]
- Laber EB, Meyer NJ, Reich BJ, Pacifici K, Collazo JM, Drake JM. On-line estimation of an optimal treatment allocation strategy for the control of white-nose syndrome in bats. 2016 Unpublished manuscript.
- Lazer D, Pentland A, Adamic L, et al. Computational social science. *Science*. 2009; 323:721–723. [PubMed: 19197046]
- Liu L, Hudgens MG. Large sample randomization inference of causal effects in the presence of interference. *Journal of the American Statistical Association*. 2014; 109:288–301. [PubMed: 24659836]
- Liu L, Hudgens MG, Becker-Dreps S. On inverse probability weighted estimators in the presence of interference. 2016 Unpublished manuscript.
- Lundin M, Karlsson M. Estimation of causal effects in observational studies with interference between units. *Stat Methods Appl*. 2014; 23:417–433.
- Manski CF. Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies*. 1993; 60(3):531–542.
- Manski CF. Identification of treatment response with social interactions. *The Econometrics Journal*. 2013; 16:S1–S23.

- Manski, CF. Mandating vaccination with unknown indirect effects; Working paper; 2016.
- Moulton LH, O'Brien KL, Kohberger R, Chang I, Reid R, Weatherholtz R, Hackell JG, Siber GR, Santosham M. Design of a group-randomised *Streptococcus pneumoniae* vaccine trial. *Contr Clin Trials*. 2001; 22:438–452.
- Neyman JS. On the application of probability theory to agricultural experiments. Essay on principles. Section 9, translated by D.M. Dabrowska and T.P. Speed. *Stat Sci*. 1923; 5:465–472. [reprinted 1990].
- Ogburn EL, VanderWeele TJ. Causal diagrams for interference. *Statistical Science*. 2014a; 29(4):559–578. Extended causal graphs to allow for interference.
- Ogburn EL, VanderWeele TJ. Vaccines, contagion, and social networks. *arXiv:1403.1241*. 2014b
- O'Malley AJ, Onnela JP. Topics in social network analysis and network science. *arXiv:1404.0067*. 2014
- Perez-Heydrich C, Hudgens MG, Halloran ME, Clemens JD, Ali M, Emch ME. Assessing effects of cholera vaccination in the presence of interference. *Biometrics*. 2014; 70(3):731–741. [PubMed: 24845800]
- Préziosi M-P, Halloran ME. Effects of pertussis vaccination on transmission: vaccine efficacy for infectiousness. *Vaccine*. 2003; 21:1853–1861. [PubMed: 12706669]
- Rigdon J, Hudgens MG. Exact confidence intervals in the presence of interference. *Statistics and Probability Letters*. 2015; 105:130–135. [PubMed: 26190877]
- Robins JM. A new approach to causal inference in mortality studies with sustained exposure period - Application to control of the healthy worker survivor effect. *Math Model*. 1986; 7:1393–1512.
- Rosenbaum PR. Interference between units in randomized experiments. *Journal of the American Statistical Association*. 2007; 1002:191–200.
- Rosenbaum PR, Rubin DB. The central role of the propensity score in observational studies. *Biometrika*. 1983; 70:41–55.
- Ross, R. The Prevention of Malaria. 2nd. London: John Murray; 1911.
- Ross R. Some *a priori* pathometric equations. *Brit Med J*. 1915; 1:546–547. [PubMed: 20767555]
- Ross R. An application of the theory of probabilities to the study of *a priori* pathometry, Part 1. *Proc R Soc Series A*. 1916; 92:204–230.
- Rubin DB. Bayesian inference for causal effects: The role of randomization. *Ann Stat*. 1978; 7:34–58.
- Rubin DB. Discussion of “Randomization analysis of experimental data in the Fisher randomization test” by Basu. *J Am Stat Assoc*. 1980; 75:591–593.
- Rubin DB. Comment: Neyman (1923) and causal inference in experiments and observational studies. *Stat Sci*. 1990; 5:472–480.
- Salathé M, Kazandjieva M, Lee JW, Levis P, Feldman MW, Jones JH. A high-resolution human contact network for infectious disease transmission. *Proceedings of the National Academy of Sciences*. 2010; 107(51):22020–22025.
- Saul, B. inference: Methods for Causal Inference with Interference. R package version 0.4.62, <http://CRAN.R-project.org/package=inference>. 2015.
- Sobel M. What do randomized studies of housing mobility demonstrate? Causal inference in the face of interference. *J Am Stat Assoc*. 2006; 101:1398–1407.
- Sofrygin, O., van der Laan, M. Semi-parametric estimation and inference for the mean outcome of the single time-point intervention in a causally connected population; U.C. Berkeley Division of Biostatistics Working Paper Series. Working Paper 344; 2015. <http://biostats.bepress.com/ucbbiostat/paper344>
- Sur D, Ochiai R, Bhattacharya SK, Ganguly NK, Ali M, Manna B, Dutta S, Donner A, Kanungo S, Park JK, et al. A cluster-randomized effectiveness trial of Vi typhoid vaccine in India. *New England Journal of Medicine*. 2009; 361:335–344. [PubMed: 19625715]
- Tchetgen Tchetgen EJ, VanderWeele TJ. On causal inference in the presence of interference. *Stat Methods Med Res*. 2012; 21(1):55–75. [PubMed: 21068053] Reviewed literature on causal inference in the presence of interference. Also developed new methods, including an inverse probability weighted estimator for observational data with partial interference.

- Toulis P, Kao E. Estimation of causal peer influence effects. *Proceedings of the 30th International Conference on Machine Learning. JMLR W&CP*. 2013; 28
- Ugander J, Karrer B, Backstrom L, Kleinberg J. Graph cluster randomization: Network exposure to multiple universes. *arXiv:1305.6979*. 2013 Considered design of randomized experiments on one large network.
- van der Laan MJ. Causal inference for a population of causally connected units. *Journal of Causal Inference*. 2014; 2:13–74. [PubMed: 26180755] Developed targeted maximum likelihood estimation for causal effects in the presence of general interference.
- VanderWeele TJ, An, W. *Handbook of Causal Analysis for Social Research*. Netherlands: Springer; 2013. Social networks and causal inference; p. 353-374.
- VanderWeele TJ, Tchetgen Tchetgen E, Halloran ME. Components of the indirect effect in vaccine trials: identification of contagion and infectiousness effects. *Epidemiology*. 2012; 23(5):751–761. [PubMed: 22828661]
- VanderWeele TJ, Tchetgen Tchetgen E, Halloran ME. Interference and sensitivity analysis. *Statistical Science*. 2014; 29(4):687–706. [PubMed: 25620841]
- VanderWeele TJ, Tchetgen Tchetgen EJ. Effect partitioning under interference in two-stage randomized vaccine trials. *Statistics and Probability Letters*. 2011a; 81:861–869. [PubMed: 21532912]
- VanderWeele TJ, Tchetgen Tchetgen EJ. Bounding the infectiousness effects in vaccine trials. *Epidemiology*. 2011b; 22(5):686–693. [PubMed: 21753730]
- Wesolowski A, Metcalf CJE, Eagle N, et al. Quantifying seasonal population fluxes driving rubella transmission dynamics using mobile phone data. *Proceedings of the National Academy of Sciences*. 2015; 112(35):11114–11119.

Figure 1.

Study designs for dependent happenings. A cluster is considered under two different scenarios. In the scenario A on the left, a certain portion of individuals in the cluster receive treatment, $Z = 1$, and the other portion of individuals receive control or nothing, $Z = 0$. In the scenario B on the right, everyone receives control or nothing. The direct, indirect, total, and overall effects of intervention are defined by the indicated contrasts (adapted from Halloran and Struchiner 1991, 1995)

Figure 2.

IPW estimates of (A) direct $DE(\alpha)$, (B) indirect $IE(\alpha, \alpha')$, (C) total $TE(\alpha, \alpha')$, and (D) overall $OE(\alpha, \alpha')$ effects based on the cholera vaccine trial data. In (A) the gray region represents approximate pointwise 95% confidence intervals. The histogram below depicts the distribution of observed neighborhood vaccine coverage. Units of the estimates are cases of cholera per 1000 individuals per year. (Reprinted with permission from Perez-Heydrich et al (2014))