



A Recipe for inference: Start with Causal Inference. Add Interference. Mix Well with R.

Bradley C. Saul
UNC Chapel Hill

Michael G. Hudgens
UNC Chapel Hill

Abstract

In causal inference, interference occurs when the treatment of one subject affects the outcome of other subjects. Interference can distort research conclusions about causal effects when not accounted for properly. In the absence of interference, inverse probability weighted (IPW) estimators are commonly used to estimate causal effects from observational data. Recently, IPW estimators have been extended to handle interference. [Tchetgen Tchetgen and VanderWeele \(2012\)](#) proposed IPW methods to estimate direct and indirect (or spillover) effects that allow for interference between individuals within groups. In this paper, we present **inference**, an R package that computes these IPW causal effect estimates when interference may be present within groups. We illustrate use of the package with examples from political science and infectious disease.

Keywords: causal inference, interference, R, SUTVA.

1. Introduction

Interference occurs when the treatment (or exposure) of one subject affects the outcome of other subjects ([Cox 1958](#)). Without accounting for interference, measuring only a treatment's direct effect may be misleading. For example, a vaccine's direct effect on an individual in a group with a large proportion of vaccinated individuals can be small. However, the protective, indirect effect from other group members' vaccinations may be large. In this case the vaccine may be judged to be ineffective based on the direct effect despite possibly having great public health utility due to the indirect effect ([Clemens, Shin, and Ali 2011](#)). Other areas where interference may be present include criminology (e.g., [Sampson 2010](#); [Verbitsky-Savitz and Raudenbush 2012](#)), developmental psychology (e.g., [Duncan, Boisjoly, Kremer, Levy, and Eccles 2005](#); [Foster 2010](#)), econometrics (e.g., [Sobel 2006](#); [Manski 2013](#)), education (e.g., [Hong and Raudenbush 2006](#); [VanderWeele, Hong, Jones, and Brown 2013](#)), imaging (e.g., [Luo, Small, shan R. Li, and Rosenbaum 2012](#)), political science (e.g., [Sinclair, McConnell,](#)

and Green 2012; Bowers, Fredrickson, and Panagopoulos 2013), social media and network analysis (e.g., VanderWeele and An 2013; Toulis and Kao 2013; Eckles, Karrer, and Ugander 2017; Kramer, Guillory, and Hancock 2014), sociology (e.g., Aronow and Samii 2017), and spatial analyses (e.g., Zigler, Dominici, and Wang 2012; Graham, McCoy, and Stephens 2013).

Inverse probability weighted (IPW) methods are often used to estimate causal effects when interference is absent (Rosenbaum 1987; Robins, Hernan, and Brumback 2000; Lunceford and Davidian 2004; Cole and Hernan 2008). Recent developments have extended IPW estimators to estimate causal effects when interference may be present in either randomized or observational studies. Tchetgen Tchetgen and VanderWeele (2012) proposed estimators for observational studies assuming partial interference (Sobel 2006), i.e., individuals can be partitioned into groups where there may be interference between individuals in the same group but not between individuals in different groups. Partial interference could be reasonable, for example, in study of bovine disease where physical separation of herds precludes pathogen transmission between herds. On the other hand, if birds or farm workers could spread the pathogen between herds, then partial interference may be questionable. The Tchetgen Tchetgen and VanderWeele (2012) IPW estimators require a model for the group-level propensity score (i.e., the probability of a group’s observed treatment allocation). The large sample properties of these estimators were derived by Perez-Heydrich, Hudgens, Halloran, Clemens, Ali, and Emch (2014).

To date, software for analysis of causal effects in the presence of interference is limited. Without interference, the R (R Core Team 2017) package **ipw** provides tools to compute IPW estimators (Van der Wal and Geskus 2011). Existing interference-related R packages, including **interferenceCI** (Rigdon 2015) and **blockTools** (Moore 2016), were designed for analysis of randomized experiments. In this paper, we present the R package **inference** (Saul 2017) which computes the Tchetgen Tchetgen and VanderWeele (2012) IPW estimators and the large sample variance estimators developed by Perez-Heydrich *et al.* (2014). Package **inference** is available from the Comprehensive R Archive Network (CRAN) at <https://CRAN.R-project.org/package=inference>.

The outline for the remainder of this paper is as follows. The next section provides background on interference and an overview of the mathematical concepts and notation. Section 3 describes the package’s main features. Sections 4 and 5 demonstrate the software with examples from public health and political science. The example in Section 5 shows advanced features of the package. We discuss computational issues with IPW estimators when groups have large numbers of individuals in Section 6. We conclude with a brief discussion and future directions in Section 7.

2. Preliminaries

2.1. A brief history of interference

Much of causal inference assumes that the exposure of one individual does not affect the outcomes of other individuals, i.e., there is no interference between individuals. Rubin (1980) bundled no interference with the assumption that treatments for all units are comparable (no hidden forms of treatment) into the “stable unit treatment value assumption” (SUTVA). Despite sporadic efforts, the research community gave little attention to this assumption until

the early 2000s. One approach to relaxing the no interference assumption is to assume partial interference. Under this assumption, space, time, and/or social network groupings preclude interference between individuals in different groups, but interference may occur within a group.

Sobel (2006), Hudgens and Halloran (2008), and Tchetgen Tchetgen and VanderWeele (2012) developed methods based on the assumption of partial interference to estimate causal effects in the presence of interference. When an experimenter randomizes units – by design – at the group and individual levels, Hudgens and Halloran (2008) defined estimators that, under certain assumptions, are unbiased for a treatment’s direct and indirect (or spillover) effects.

Observational studies complicate estimation of interference effects. Tchetgen Tchetgen and VanderWeele (2012) proposed IPW estimators of causal effects based on group-level propensity scores for non-randomized treatment allocation. They showed these estimators to be unbiased when the propensity score is known. Perez-Heydrich *et al.* (2014) derived the large sample properties of these estimators when the propensity scores are unknown but correctly modeled. They applied these results to draw inference about the direct and indirect effects of cholera vaccination in Matlab, Bangladesh.

2.2. Basic partial interference setup

Consider N individuals partitioned into m groups, each with n_i individuals for $i = 1, \dots, m$. The triplet $(Y_{ij}, A_{ij}, \mathbf{X}_{ij})$ represents the observed outcome, treatment, and baseline covariate vector, respectively, for individual j in group i . We let capitalized letters denote random variables, and lowercase letters (e.g., $(y_{ij}, a_{ij}, \mathbf{x}_{ij})$) denote observed or realized values. Let \mathbf{X}_i and \mathbf{A}_i be the matrix of baseline covariates and vector of treatment allocations for members of group i . Let $\mathbf{A}_{i,-j} = (A_{i1}, \dots, A_{ij-1}, A_{ij+1}, \dots, A_{in_i})$ represent a group’s treatment allocation excluding the j th subject. Let $Y_{ij}(a_{ij}, \mathbf{a}_{i,-j}) = Y_{ij}(\mathbf{a}_i)$ be the potential outcome for individual j in group i if, possibly contrary to fact, group i received \mathbf{a}_i . By causal consistency, $Y_{ij} = Y_{ij}(\mathbf{A}_i)$ (Pearl 2010). Let \mathbf{Y}_i be the vector of potential outcomes for group i . By assuming no interference between groups, an individual’s potential outcome may depend only on the treatment allocation of its group. The set $\mathcal{A}(n_i)$ contains all of group i ’s possible treatment vectors. With a binary treatment, $A_{ij} \in \{a_1, a_2\}$, this set has 2^{n_i} elements.

Estimands

Without interference, researchers often estimate an average treatment effect, which contrasts the average outcome for two treatment allocations: the entire population treated versus the entire population untreated. With interference, causal estimands may be defined in terms of the continuum of treatment allocation strategies between those extremes. In **interference**, we consider Bernoulli-type allocation strategies proposed by Tchetgen Tchetgen and VanderWeele (2012), where individuals independently receive treatment with probability α . For this allocation strategy, the probability of a group’s treatment vector is denoted as $\pi_i(\mathbf{A}_i; \alpha) = \prod_{j=1}^{n_i} \alpha^{A_{ij}} (1 - \alpha)^{1 - A_{ij}}$ and, excluding the j th subject, $\pi_i(\mathbf{A}_{i,-j}; \alpha) = \prod_{k=1, k \neq j}^{n_i} \alpha^{A_{ik}} (1 - \alpha)^{1 - A_{ik}}$. The analyst may compute the estimators described below over a range of α ’s to explore hypothetical underlying treatment allocations. In their analysis of a cholera vaccine study, Perez-Heydrich *et al.* (2014) examined strategies between $\alpha = 0.3$ and $\alpha = 0.6$ because 75% of the groups had observed vaccine coverages in that range.

Define an individual’s average potential outcome when assigned treatment a under strategy

α by

$$\bar{Y}_{ij}(a; \alpha) = \sum_{\mathbf{a}_{i,-j} \in \mathcal{A}(n_i-1)} Y_{ij}(a, \mathbf{a}_{i,-j}) \pi_i(\mathbf{a}_{i,-j}; \alpha).$$

In words, $\bar{Y}_{ij}(a; \alpha)$ is a weighted average of individual j 's potential outcomes under possible treatment vectors of the other $n_i - 1$ subjects in group i weighted by the probability of each treatment vector. Similarly, define the marginal individual average potential outcome by

$$\bar{Y}_{ij}(\alpha) = \sum_{\mathbf{a}_i \in \mathcal{A}(n_i)} Y_{ij}(\mathbf{a}_i) \pi_i(\mathbf{a}_i; \alpha).$$

Here, the weighted average of individual j 's potential outcomes is across all group treatment vectors in $\mathcal{A}(n_i)$.

A simple mean of individual average potential outcomes within a cluster defines group average potential outcomes. Then group-level estimands are averaged to make population-level estimands. For example, $\bar{Y}(a; \alpha) = \sum_{i=1}^m \{ \sum_{j=1}^{n_i} \bar{Y}_{ij}(a; \alpha) / n_i \} / m$ is the population-level average outcome when individuals receive treatment a and their group adopts allocation strategy α . Likewise, $\bar{Y}(\alpha) = \sum_{i=1}^m \{ \sum_{j=1}^{n_i} \bar{Y}_{ij}(\alpha) / n_i \} / m$ is the population-level average outcome when groups adopt allocation strategy α .

Contrasts of the population average potential outcomes define causal effects. [Hudgens and Halloran \(2008\)](#) describe four causal effects: direct, indirect, total, and overall (see also [Tchetgen Tchetgen and VanderWeele 2012](#)). The direct (or unit-level treatment) effect compares average potential outcomes within a single allocation strategy:

$$\overline{DE}(\alpha) = \bar{Y}(a_1; \alpha) - \bar{Y}(a_2; \alpha).$$

An indirect effect compares a treatment's average potential outcomes under different allocation strategies:

$$\overline{IE}(\alpha, \alpha') = \bar{Y}(a_1; \alpha) - \bar{Y}(a_1; \alpha').$$

For a binary treatment, there are two indirect effects for a fixed (α, α') pair: one for a_1 and one for a_2 . If interference is not present, then the indirect effect equals zero. The total effect accounts for both the direct and indirect effects:

$$\overline{TE}(\alpha, \alpha') = \bar{Y}(a_1; \alpha) - \bar{Y}(a_2; \alpha').$$

The overall effect contrasts the marginal average potential outcomes for two allocation strategies:

$$\overline{OE}(\alpha, \alpha') = \bar{Y}(\alpha) - \bar{Y}(\alpha').$$

See [Hudgens and Halloran \(2008\)](#) and [Tchetgen Tchetgen and VanderWeele \(2012\)](#) for further discussion of these causal estimands.

2.3. IPW estimation

[Tchetgen Tchetgen and VanderWeele \(2012\)](#) proposed IPW estimators of the causal estimands defined above assuming partial interference. Their estimator weights an individual's outcome by the inverse of the group-level propensity score, $P(\mathbf{A}_i | \mathbf{X}_i)$, the probability of a group's treatment allocation given the covariates of the group's individuals. [Tchetgen Tchetgen and VanderWeele \(2012\)](#) showed the IPW estimators to be unbiased when the group-level propensities are known, under the following assumptions:

1. Conditional independence: $P(\mathbf{A}_i = \mathbf{a}_i | \mathbf{X}_i, \mathbf{Y}_i) = P(\mathbf{A}_i = \mathbf{a}_i | \mathbf{X}_i)$.
2. Positivity: $P(\mathbf{A}_i = \mathbf{a}_i | \mathbf{X}_i) > 0 \quad \forall \mathbf{a}_i \in \mathcal{A}(n_i)$.

The true propensity scores are not generally known in observational studies and must be estimated. Tchetgen Tchetgen and VanderWeele (2012) suggested estimating $P(\mathbf{A}_i | \mathbf{X}_i)$ using a generalized mixed effects model. We denote these models as $f_{\mathbf{A}_i | \mathbf{X}_i}(\mathbf{A}_i | \mathbf{X}_i; \boldsymbol{\theta}_x, \theta_s)$, where $\boldsymbol{\theta}_x$ represents fixed effects parameters and θ_s a group random effect parameter. Model parameters may be estimated by maximum likelihood methods, which we denote $\hat{\boldsymbol{\theta}} = (\hat{\boldsymbol{\theta}}_x, \hat{\theta}_s)$. For a binary treatment, a model for the group's propensity score might be:

$$f_{\mathbf{A}_i | \mathbf{X}_i}(\mathbf{A}_i | \mathbf{X}_i; \boldsymbol{\theta}_x, \theta_s) = \int \prod_{j=1}^{n_i} h_{ij}(b_i; \boldsymbol{\theta}_x)^{A_{ij}} \{1 - h_{ij}(b_i; \boldsymbol{\theta}_x)\}^{1-A_{ij}} f_b(b_i; \theta_s) db_i, \quad (1)$$

where $h_{ij}(b_i; \boldsymbol{\theta}_x) = P(A_{ij} = 1 | \mathbf{X}_{ij}, b_i, \boldsymbol{\theta}_x) = \text{logit}^{-1}(\mathbf{X}_{ij} \boldsymbol{\theta}_x + b_i)$ and $f_b(\cdot; \theta_s)$ is the density of a normal random variable with mean 0 and variance θ_s . This is the default group-level propensity score model in **inference**; the examples below show how the user can modify the default group propensity score model. Validity of inferences drawn using the methods described below requires correct specification of the group propensity score model. Therefore, it is important in practice to conduct diagnostics to assess the fit of the model employed. For example, if (1) is assumed, then the Tchetgen Tchetgen and Coull (2006) diagnostic test can be used to assess whether the random effects are normally distributed.

The IPW estimator for the group-level average potential outcomes is a straightforward weighted sum,

$$\hat{Y}_i^{ipw}(a, \alpha) = \frac{\sum_{j=1}^{n_i} \pi_i(\mathbf{A}_{i,-j}; \alpha) I(A_{ij} = a) Y_{ij}}{n_i f_{\mathbf{A}_i | \mathbf{X}_i}(\mathbf{A}_i | \mathbf{X}_i; \hat{\boldsymbol{\theta}})}, \quad (2)$$

as is the estimator for group-level marginal potential outcomes,

$$\hat{Y}_i^{ipw}(\alpha) = \frac{\sum_{j=1}^{n_i} \pi_i(\mathbf{A}_i; \alpha) Y_{ij}}{n_i f_{\mathbf{A}_i | \mathbf{X}_i}(\mathbf{A}_i | \mathbf{X}_i; \hat{\boldsymbol{\theta}})}. \quad (3)$$

From (2) and (3), one constructs population-level average potential outcome and marginal population-level average potential outcome estimators by $\hat{Y}^{ipw}(a; \alpha) = \sum_{i=1}^m \hat{Y}_i^{ipw}(a; \alpha) / m$ and $\hat{Y}^{ipw}(\alpha) = \sum_{i=1}^m \hat{Y}_i^{ipw}(\alpha) / m$. Estimators for direct, indirect, total, and overall effects simply contrast population-level estimators,

$$\begin{aligned} \widehat{DE}(\alpha) &= \hat{Y}^{ipw}(a_1; \alpha) - \hat{Y}^{ipw}(a_2; \alpha) \\ \widehat{IE}(\alpha, \alpha') &= \hat{Y}^{ipw}(a_1; \alpha) - \hat{Y}^{ipw}(a_1; \alpha') \\ \widehat{TE}(\alpha, \alpha') &= \hat{Y}^{ipw}(a_1; \alpha) - \hat{Y}^{ipw}(a_2; \alpha') \\ \widehat{OE}(\alpha, \alpha') &= \hat{Y}^{ipw}(\alpha) - \hat{Y}^{ipw}(\alpha'). \end{aligned}$$

IPW variance estimation

Perez-Heydrich *et al.* (2014) derived asymptotic distributions of the IPW estimators using standard estimating equation theory. Briefly, the IPW estimators above are consistent and asymptotically normal as the number of groups m tends to infinity. When the group-level

propensity scores are known (as in the case of simulation or randomized studies), a large sample estimator of the variance of $\widehat{DE}(\alpha)$ is

$$\frac{1}{m^2} \sum_{i=1}^m \left\{ \widehat{DE}_i(\alpha) - \widehat{DE}(\alpha) \right\}^2 .$$

Results for \widehat{IE} , \widehat{TE} , and \widehat{OE} are analogous. As explained below, when `variance_estimation = "naive"` in the `inference` function, this formula is used to compute standard errors and Wald-type confidence intervals.

When the propensity scores are unknown and instead estimated using a parametric model, computing variance estimators is more complicated and involves derivatives of the group propensity with respect to each parameter and derivatives of the propensity model's log-likelihood. The supplementary materials in [Perez-Heydrich *et al.* \(2014\)](#) contain the mathematical details, and this method is available with the `variance_estimation = "robust"` option. The "robust" option computes consistent variance estimates which account for the estimation of the weights, whereas the "naive" option computes variance estimates described in the preceding paragraph which ignore estimation of the weights and are conservative (i.e., tend to be too large).

3. Using inference

3.1. User's guide

To start, install the package from CRAN using `install.packages("inference")`. The list below details the arguments for `inference`, the primary function in `inference`. Special attention should be given to the `propensity_integrand` and `formula` arguments.

- **formula**: formula used to define the causal model. `formula` has a minimum of 4 parts, separated by `|` and `~` in a specific structure: `outcome | exposure ~ covariates | group`. The order matters, and the pipes (`|`) split the data frame into corresponding pieces ([Zeileis and Croissant 2010](#)). The `exposure ~ covariates` piece is passed as a single formula to the chosen `model_method` (defined below) used to estimate or fix propensity parameters.
 - The following includes a random effect for the group: `outcome | exposure ~ covariates + (1 | group) | group`. In this instance, the group variable appears twice.
 - If the study design includes a “participation” variable (as in both examples below), this is easily added to the formula: `outcome | exposure | participation ~ covariates | group`.
- **propensity_integrand**: a function, which may be created by the user, used to compute the IP weights. This defaults to the function `logit_integrand()`, which calculates the product of inverse logits for individuals in a group: $\prod_{j=1}^{n_i} \{r h_{ij}(b_i)\}^{A_{ij}} \{1 - r h_{ij}(b_i)\}^{1-A_{ij}} f_b(b_i; \theta_s)$, where $h_{ij}(b_i) = \text{logit}^{-1}(\mathbf{X}_{ij} \boldsymbol{\theta}_x + b_i)$, b_i is a group-level random effect, $f_b(\cdot)$ is a $N(0, \theta_s)$ density, and r is a known constant. In an observational

study typically $r = 1$. The examples below include individual randomized experiments in which case r denotes the randomization probability among trial participants. `logit_integrand()` is the integrand of (1) where $h_{ij}(b_i)$ is scaled by a constant r term. If no random effect is included in the formula, `logit_integrand()` ignores the random effect. IP weights are computed by numerically integrating `propensity_integrand` over the random effect distribution using `stats::integrate()` to which arguments may be passed via `...` (see below). The default `logit_integrand()` also takes the following argument that can be passed via the `...` argument in `interference()`:

- **randomization**: a scalar. This is the r in the formula just above. It defaults to 1 in the case that a `participation` vector is not included. The vaccine study example in Section 4 demonstrates use of this argument.
- **loglikelihood_integrand**: a function, which may be created by the user, that defines the log-likelihood of the propensity score model. This should generally be the same function as `propensity_integrand`, which is the default.
- **allocations**: a vector of values in $[0, 1]$. These are the α 's defined in Section 2.2. Increasing the number of elements of the allocation vector increases computation time; however, a larger number of allocations will make plotted effect estimates smoother. A minimum of two allocations is required.
- **data**: the analysis data frame. This must include all the variables defined in the formula.
- **model_method**: the method used to estimate or set the propensity model parameters. Must be one of "glm", "glmer", or "oracle". For a fixed effects only model use "glm", or to include random effects use `lme4`'s "glmer" (Bates, Mächler, Bolker, and Walker 2015). `logit_integrand` only supports a single random effect for the grouping variable, corresponding to b_i . When the propensity parameters are known (as in simulations) or if estimating parameters for the propensity model outside of `interference`, use the "oracle" option. See `model_options` for details on how to pass the oracle parameters. Defaults to "glmer".
- **model_options**: a list of options passed to the function in `model_method`. Defaults to `list(family = binomial(link = "logit"))`. When `model_method = "oracle"`, the list must have two elements, `fixed.effects` and `random.effects`. If the model does not include random effects, set `random.effects = NULL`.
- **causal_estimation_method**: currently only supports and defaults to "ipw".
- **causal_estimation_options**: a list with a single item `variance_estimation`, which is either "naive" or "robust". See Section 2.3 for details. Defaults to "robust".
- **conf.level**: level for confidence intervals. Defaults to 0.95.
- **rescale.factor**: a scalar multiplication factor by which to rescale outcomes and effects. Defaults to 1.
- **integrate_allocation**: indicator of whether the integrand function uses the allocation parameter. Defaults to TRUE.
- **runSilent**: If FALSE, prints status of computations. Defaults to TRUE.
- **...**: used to pass additional arguments to internal functions such as `numDeriv::grad()` or `stats::integrate()`. Arguments can also be passed to the `propensity_integrand` and `loglikelihood_integrand` functions.

3.2. The ‘interference’ object

An `interference()` call results in an S3 object of class ‘interference’ which contains:

- `estimates`: a data frame of causal effect estimates;
- `models$propensity_model`: the object returned by `glm` or `glmer`;
- `summary`: a list of objects summarizing the causal model such as the number of groups, number of allocations, and the formula used in the `interference` call;
- `weights`: $(\# \text{ of groups}) \times (\# \text{ of allocations})$ matrix of group-level weights:

$$w_{i,k} = \frac{\pi_i(\mathbf{A}_i; \alpha_k)}{f_{\mathbf{A}_i|\mathbf{X}_i}(\mathbf{A}_i|\mathbf{X}_i; \hat{\boldsymbol{\theta}})}.$$

If `variance_estimation = "robust"`, then the object also includes:

- `weightd`: $(\# \text{ of groups}) \times (\# \text{ of allocations}) \times (\# \text{ of parameters})$ array of weights computed using derivatives of the propensity function with respect to each parameter;
- `scores`: $(\# \text{ of groups}) \times (\# \text{ of parameters})$ matrix of derivatives of the log-likelihood.

3.3. Utility functions

The package includes tools to extract effect estimates of interest from the S3 object. The functions `direct_effect`, `indirect_effect` (or `ie`), `total_effect` (or `te`), and `overall_effect` (or `oe`) select appropriate records from the `estimates` data frame in the ‘interference’ object. Section 4 shows an example.

4. Example: Vaccine study

This section illustrates the use of **inference** with an example drawn from vaccine research. The package includes a single dataset based on the same set of parameters used in the simulation study by [Perez-Heydrich *et al.* \(2014\)](#). The `vaccinesim` dataset consists of 3000 units in 250 groups and contains two covariates (`X1` = age in decades and `X2` = distance to river), a vaccination indicator (`A`), a participation indicator (`B`), a binary outcome (`Y`) indicating cholera infection (1 yes, 0 no), and the unit’s `group`.

```
R> library("inference")
R> head(vaccinesim)
```

	Y	X1	X2	A	B	group
1	1	5.3607405	1.715527	0	0	1
2	0	0.1964597	1.730802	0	1	1
3	0	0.4846243	1.769546	1	1	1
4	0	0.8012977	1.715527	0	1	1
5	0	2.1426629	1.772158	1	1	1
6	0	1.2861017	1.715527	0	1	1

Like the original study ([Ali *et al.* 2005](#)) that inspired the simulation, individuals were randomized to vaccine with a known probability of 2/3, but subjects could opt to not participate

in the trial. In essence, there are both experimental and observational aspects to the data. The `interference` function handles this design when `logit_integrand`'s `randomization` argument is used and a participation variable is included in the formula.

```
R> example1 <- interference(formula =
+   Y | A | B ~ X1 + X2 + (1 | group) | group,
+   allocations = c(0.3, 0.45, 0.6), data = vaccinesim,
+   randomization = 2/3, method = "simple")
```

The only arguments required for `interference` to run are `formula`, `allocations`, and `data`. When using the "robust" method (the default) to compute the variance, the internal workings call `numDeriv::grad` (Gilbert and Varadhan 2016) and `stats::integrate` frequently. The option `method = "simple"` greatly speeds up the `numDeriv::grad` function. For more accurate derivatives, leave out this option. See `?numDeriv::grad` for more options.

The `print.interference` function provides an overview of the causal effect estimates, estimated standard errors, and Wald-type confidence intervals. In the output, `alpha1` and `alpha2` refer to α and α' , while `trt1` and `trt2` refer to a_1 and a_2 , respectively.

```
R> example1
```

```
-----
                                Model Summary
-----
Formula: Y | A | B ~ X1 + X2 + (1 | group) | group
Number of groups: 250
3 allocations were used from 0.3 (min) to 0.6 (max)
-----

                                Causal Effect Summary
                                Confidence level: 0.95
                                Variance method: robust
-----

Direct Effects
alpha1 trt1 alpha2 trt2 estimate std.error conf.low conf.high
  0.30   0   0.30   1   0.1605   0.02474  0.11202   0.2090
  0.60   0   0.60   1   0.1086   0.01857  0.07219   0.1450
  0.45   0   0.45   1   0.1339   0.01778  0.09904   0.1687

Indirect Effects
alpha1 trt1 alpha2 trt2 estimate std.error conf.low conf.high
  0.30   0   0.60   0  0.15907   0.02617  0.10777   0.2104
  0.30   0   0.45   0  0.08657   0.01732  0.05263   0.1205
  0.45   0   0.60   0  0.07250   0.01408  0.04490   0.1001

Total Effects
alpha1 trt1 alpha2 trt2 estimate std.error conf.low conf.high
  0.30   0   0.60   1   0.2677   0.02435  0.2199   0.3154
```

0.30	0	0.45	1	0.2205	0.02469	0.1721	0.2688
0.45	0	0.60	1	0.1811	0.01841	0.1450	0.2172

Overall Effects

alpha1	trt1	alpha2	trt2	estimate	std.error	conf.low	conf.high
0.30	NA	0.60	NA	0.17607	0.019247	0.13835	0.2138
0.30	NA	0.45	NA	0.09867	0.014207	0.07083	0.1265
0.45	NA	0.60	NA	0.07740	0.008981	0.05980	0.0950

The utility functions return selected effect estimates.

```
R> direct_effect(example1, 0.3)
```

	alpha1	trt1	alpha2	trt2	estimate	std.error	conf.low	conf.high
1	0.3	0	0.3	1	0.1605036	0.02473782	0.1120184	0.2089888

```
R> ie(example1, 0.3)
```

	alpha1	trt1	alpha2	trt2	estimate	std.error	conf.low	conf.high
1	0.3	0	0.30	0	0.00000000	0.00000000	0.00000000	0.00000000
2	0.3	0	0.45	0	0.08656992	0.01731878	0.05262574	0.1205141
3	0.3	0	0.60	0	0.15906887	0.02617398	0.10776882	0.2103689

4.1. Plotting effect estimates

Plots of effect estimates over a range of α levels may be helpful in summarizing results. [Perez-Heydrich et al. \(2014\)](#) present several such graphical displays. Here we demonstrate how to generate similar plots of effect estimates using **inference**.

First, we estimate the effects over a dense sequence of allocations so that lines will be smooth.

```
R> example2 <- interference(
+   formula = Y | A | B ~ X1 + X2 + (1 | group) | group,
+   allocations = seq(0.2, 0.8, by = 0.01), data = vaccinesim,
+   randomization = 2/3, method = "simple")
```

In [Figure 1](#), we present the direct and indirect effect estimates over this range of allocations. For direct effects, a simple scatterplot showing the point-wise confidence intervals suffices. One approach with indirect effects fixes α and plots estimates over a range of α' , whereas a contour plot displays all pairwise (α, α') comparisons over a range of allocation strategies.

```
R> deff <- direct_effect(example2)
R> x <- deff$alpha1
R> y <- as.numeric(deff$estimate)
R> u <- as.numeric(deff$conf.high)
```




Figure 1: Plots of the estimates from `example2`. The top plot shows the direct effect estimates. The bottom two plots demonstrate different ways of viewing the indirect effect estimates. The shaded regions show the point-wise 95% confidence intervals.

```
R> l <- as.numeric(deff$conf.low)
R> plot(c(min(x), max(x)), c(-0.15, 0.25), type = "n", bty = "l",
+       xlab = expression(alpha), ylab = "" )
R> title(ylab = expression(widehat(DE) * "(" * alpha * ")"), line = 2)
R> polygon(c(x, rev(x)), c(u, rev(l)), col = "skyblue", border = NA)
R> lines(x, y, cex = 2)
R> ieff.4 <- ie(example2, allocation1 = 0.4)
R> x <- ieff.4$alpha2
R> y <- as.numeric(ieff.4$estimate)
R> u <- as.numeric(ieff.4$conf.high)
R> l <- as.numeric(ieff.4$conf.low)
R> plot(c(min(x), max(x)), c(-0.15, 0.25), type = "n", bty = "l",
+       xlab = expression(alpha * "'"), ylab = "")
R> title(ylab =
+       expression(widehat(IE) * "(" * 0.4 * "," * alpha * "'" * ")"),
+       line = 2)
R> polygon(c(x, rev(x)), c(u, rev(l)), col = "skyblue", border = NA)
R> lines(x, y, cex = 2)
```




Figure 2: A histogram of group-level weights, $w_{i,k}$, for $\alpha_k = 0.5$.

```
R> ieff <- subset(example2[["estimates"]], effect == "indirect")
R> x <- sort(unique(ieff$alpha1))
R> y <- sort(unique(ieff$alpha2))
R> z <- xtabs(estimate ~ alpha1 + alpha2, data = ieff)
R> contour(x, y, z, xlab = expression(alpha),
+   ylab = expression(alpha * "'"), bty = "l")
```

4.2. Diagnostics

IPW estimators are known to be unstable if the weights range greatly. The package includes a basic utility to check the performance of the group-level weights, $w_{i,k}$, for multiple allocations. The function `diagnose_weights` plots histograms of weights for chosen allocation levels. If the `allocations` argument is left `NULL`, the function plots histograms for five allocation levels used in the `interference` call. Figure 2 shows the resulting histogram for a single allocation. The analyst should examine groups with extreme weights, which may unduly influence population-level estimates.

```
R> diagnose_weights(example2, allocations = 0.5, breaks = 30)
```

5. Example: Voting experiment

The preceding example used the default `logit_integrand` function to define the group-level propensities. The following example demonstrates how to customize the propensity score function.

Nickerson (2008) reported an experiment on voter behavior to examine peer-to-peer indirect effects on voting participation. The experiment randomized households with only two registered voters in Denver and Minneapolis to receive one of three assignments: voting

encouragement, recycling encouragement, or nothing. Canvassers knocked on doors of households randomized to the voting or recycling groups a week before the 2002 primary. If a registered voter answered the door, the canvassers delivered a scripted message about voting (treatment) or recycling (control). The researchers used voter turnout records to determine if each member of the household then voted in the election. Nickerson was interested in the potential spillover effect of the voting encouragement to the untreated individual via the treated individual. From analysis of the observed data, he concluded there was a “secondary effect” where the household members not contacted by the canvassers voted more often in the treatment groups compared to the control groups.

The dataset `voters` contains information for 3861 households, 2549 in Denver and 1312 in Minneapolis, including covariates such as age, gender, previous voting history, and party affiliation. Our estimand of interest involves average voting outcomes when households receive voting encouragement compared to when households receive the recycling message, hence we exclude households not contacted by canvassers. We also exclude the single household where a canvasser appears to have contacted both registered voters.

```
R> voters <- within(voters, {
+   treated <- (treatment == 1 & reached == 1) * 1
+   c_age <- (age - mean(age)) / 10
+ })
R> reach_cnt <- tapply(voters$reached, voters$family, sum)
R> voters <- voters[!(voters$family %in% names(reach_cnt[reach_cnt > 1])), ]
R> voters <- voters[voters$hsecontact == 1, ]
```

5.1. Household-level propensity

Unlike the vaccine study example, in this data set randomization occurred at the group level but individual level treatment was not randomized. With only two subjects, $\mathbf{A}_i = (A_{i1}, A_{i2})$ is the treatment allocation for group i and $\mathbf{X}_i = (\mathbf{X}_{i1}, \mathbf{X}_{i2})$ is the matrix of individuals’ covariate matrices for group i . Let $\mathbf{B}_i = (B_{i1}, B_{i2})$ be indicators of being reached by a canvasser in group i . Since we only consider households where someone answered the door, $\mathbf{B}_i \in \{(1, 0), (0, 1)\}$ and $P(B_{i1} = 1|\mathbf{X}_i) + P(B_{i2} = 1|\mathbf{X}_i) = 1$. Let $h_{ij} = P(B_{ij} = 1|\mathbf{X}_i; \boldsymbol{\theta}_x) = \text{logit}^{-1}(\mathbf{X}_i \boldsymbol{\theta}_x)$. Let $G_i \in \{0, 1\}$ be the indicator that group i is randomized to treatment (1) or control (0). By design, $P(G_i = 1) = 0.5$. Since $P(A_{i1}|\mathbf{X}_i; \boldsymbol{\theta}_x) = 1 - P(A_{i2}|\mathbf{X}_i; \boldsymbol{\theta}_x)$, $P(\mathbf{A}_i|\mathbf{X}_i; \boldsymbol{\theta}_x)$ can arbitrarily be defined in terms of either household member. By convention we use the first subject (subject one) of each group found in the dataset. Among treated groups, the probability of subject one being treated is the probability that a canvasser reached subject one. That is, $P(A_{i1}|G_i = 1, \mathbf{X}_i; \boldsymbol{\theta}_x) = h_{i1}^{A_{i1}}(1 - h_{i1})^{1-A_{i1}}$. Thus, the group-level propensity can be expressed:

$$\begin{aligned} P(\mathbf{A}_i|\mathbf{X}_i; \boldsymbol{\theta}_x) &= P(\mathbf{A}_i|G_i = 1, \mathbf{X}_i; \boldsymbol{\theta}_x)P(G_i = 1) + P(\mathbf{A}_i|G_i = 0, \mathbf{X}_i; \boldsymbol{\theta}_x)P(G_i = 0) \\ &= 0.5\{P(\mathbf{A}_i|G_i = 1, \mathbf{X}_i; \boldsymbol{\theta}_x) + P(\mathbf{A}_i|G_i = 0, \mathbf{X}_i; \boldsymbol{\theta}_x)\} \\ &= \begin{cases} 0.5 & \text{if } A_{i1} = 0, A_{i2} = 0 \text{ and } G_i = 0, \\ 0.5h_{i1}^{A_{i1}}(1 - h_{i1})^{1-A_{i1}} & \text{if } (A_{i1} = 0, A_{i2} = 1 \text{ or } A_{i1} = 1, A_{i2} = 0) \text{ and } G_i = 1, \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

Thus, h_{i1} is sufficient to determine the group-level propensity. If we know whether or not the first subject was reached by a canvasser, then we know if the second was. Therefore, we can estimate parameters for h_{i1} with a dataset that includes only subject one from each group. To do this, we must estimate the parameters outside of **inference** and use `model_method = "oracle"`. We include centered age (in decades) in the propensity model for demonstration purposes.

```
R> voters1 <- do.call(rbind, by(voters, voters[, "family"],
+   function(x) x[1, ]))
R> coef.voters <- coef(glm(reached ~ c_age, data = voters1,
+   family = binomial(link = "logit")))
```

5.2. Coding the propensity function

Custom `propensity_integrand` and `loglikelihood_integrand` functions must have at least one argument:

- **b**: the first argument is the variable for which the `integrate` function computes the integral. As in this example, the function can be written so that the integral evaluates to 1 and has no effect.

For example, the following function will fix the group-level propensity to 0.5 for all groups when `variance_estimation = "naive"`:

```
R> fixed_propensity <- function(b) return(0.5 * dnorm(b))
```

For more realistic models, additional arguments may be passed to the custom function:

- **X**: the covariate matrix (determined by the `formula`) for the i th group;
- **A**: the vector of treatment indicators for the i th group;
- **parameters**: vector of estimated parameters from the `model_method`;
- **allocation**: the allocation level for which the propensity is currently being estimated;
- **...**: other arguments can be passed via the ellipsis in `inference`.

Now we have the pieces to write the propensity function for the voting example.

```
R> household_propensity <- function(b, X, A, parameters,
+   group.randomization = 0.5) {
+   if (!is.matrix(X)) {
+     X <- as.matrix(X)
+   }
+   if (sum(A) == 0) {
+     pr <- group.randomization
+   } else {
+     X.1 <- X[1, ]
+     A.1 <- A[1]
+     h <- plogis(X.1 %*% parameters)
```

```

+   pr <- group.randomization * dbinom(A.1, 1, h)
+ }
+ out <- pr * dnorm(b)
+ out
+ }

```

5.3. Evidence of a peer influence effect

The influence of the door opener on the non-door opener's voting behavior corresponds to an indirect effect. Though the Bernoulli-type parametrization of the estimands allows us to look at a range of allocations, $\widehat{IE}(0.5, 0)$ makes the sensible comparison between a world where individuals receive a voting message with probability 0.5 to a world where individuals have zero probability of receiving the voting message.

```

R> example3 <- interference(formula =
+   voted02p | treated | reached ~ c_age | family,
+   propensity_integrand = "household_propensity", data = voters,
+   model_method = "oracle",
+   model_options = list(fixed.effects = coef.voters, random.effects = NULL),
+   allocations = c(0, 0.5), integrate_allocation = FALSE,
+   causal_estimation_options = list(variance_estimation = "robust"),
+   conf.level = 0.9)
R> ie(example3, 0.5, 0)[, c("estimate", "conf.low", "conf.high")]

      estimate   conf.low  conf.high
1 0.03151501 0.002290586 0.06073944

```

The point estimate suggests an individual receiving the voting encouragement increases the voting likelihood of the other household member by 3.2%. The 90% confidence interval excludes zero, indicating a significant indirect effect corroborating the analysis in [Nickerson \(2008\)](#).

For comparison, suppose that a flip of a fair coin determined which registered voter opened the door. We exclude age as a covariate and instead set $h_{i1} = 0.5$. Here we assume to know the propensity score, so we use `variance_estimation = "naive"`.

```

R> example4 <- interference(formula =
+   voted02p | treated | reached ~ 1 | family,
+   propensity_integrand = "household_propensity", data = voters,
+   model_method = "oracle",
+   model_options = list(fixed.effects = 0, random.effects = NULL),
+   allocations = c(0, 0.5), integrate_allocation = FALSE,
+   causal_estimation_options = list(variance_estimation = "naive"),
+   conf.level = 0.9)
R> ie(example4, 0.5, 0)[, c("estimate", "conf.low", "conf.high")]

      estimate   conf.low  conf.high
1 0.03144654 0.002209019 0.06068406

```

Examining the group-level weights may help diagnose coding errors in the propensity score function. In the case of a fixed probability as in `example4`, the propensity weights are easily computed by hand. For example, for $\alpha = 0.5$,

$$w_i = \frac{\pi(\mathbf{A}_i; 0.5)}{f_{\mathbf{A}_i|\mathbf{X}_i}(\mathbf{A}_i|\mathbf{X}_i; \theta_x = 0)} = \begin{cases} \frac{0.5^2}{.5} = 0.5 & \text{if } G_i = 0 \\ \frac{0.5^2}{.5^2} = 1 & \text{if } G_i = 1 \end{cases},$$

which we can confirm that the software computed this.

```
R> G <- tapply(voters[1:12, "treated"], voters[1:12, "family"], sum)
R> W <- head(example4[["weights"]])[, 2]
R> cbind(G, W)
```

```
   G  W
2  1 1.0
4  0 0.5
5  0 0.5
6  0 0.5
9  1 1.0
10 0 0.5
```

6. Computational issues with IPW estimators

We show in this section how computation of the group-level weights may affect estimation as the number of individuals in groups grows. To illustrate, consider the IPW estimator of the overall effect, which weights individual outcomes in group i with:

$$w_{1i} = \frac{\pi(\mathbf{A}_i; \alpha)}{f_{\mathbf{A}_i|\mathbf{X}_i}(\mathbf{A}_i|\mathbf{X}_i; \hat{\theta}_x)} = \frac{\prod_{j=1}^{n_i} \alpha^{A_{ij}} (1 - \alpha)^{1 - A_{ij}}}{\int \prod_{j=1}^{n_i} h_{ij}^{A_{ij}} (1 - h_{ij})^{1 - A_{ij}} f_b(b_i; \hat{\theta}_s) db_i}$$

or equivalently,

$$w_{2i} = \left\{ \int \prod_{j=1}^{n_i} \left(\frac{h_{ij}}{\alpha} \right)^{A_{ij}} \left(\frac{1 - h_{ij}}{1 - \alpha} \right)^{1 - A_{ij}} f_b(b_i; \hat{\theta}_s) db_i \right\}^{-1}$$

or,

$$w_{3i} = \left\{ \int \exp \left[\sum_{j=1}^{n_i} \left\{ A_{ij} \log \left(\frac{h_{ij}}{\alpha} \right) + (1 - A_{ij}) \log \left(\frac{1 - h_{ij}}{1 - \alpha} \right) \right\} \right] f_b(b_i; \hat{\theta}_s) db_i \right\}^{-1}.$$

While mathematically equivalent, these weights may be computationally dissimilar. In the case of w_{1i} , the product term within the integral entails multiplying probabilities and thus will tend to 0 as n_i increases, causing the denominator of w_{1i} to get arbitrarily large. In contrast, the product term in w_{2i} entails multiplying values which may be less than or greater than 1 and thus tends to be less susceptible to numerical instability. Summing $\log(h_{ij}/\alpha)$

or $\log(1 - h_{ij}) / (1 - \alpha)$) and then exponentiating the result may provide greater numerical stability. Internally, **inference** uses w_{3i} .

When group sizes are small, the differences between these weights tend to be infinitesimal, but as group sizes grow the differences become important. To be specific, consider the code below comparing w_{1i} , w_{2i} , and w_{3i} for increasing group sizes where $\alpha = 0.5$, all units are treated, there is no random effect, and h_{ij} is fixed at 0.5.

```
R> compare_weights <- function(n, alpha = 0.5, h = 0.5) {
+   pi <- rep(alpha, n)
+   PrA <- rep(h, n)
+   c(w1 = prod(pi) / prod(PrA), w2 = 1 / prod(PrA / pi),
+     w3 = 1 / exp(sum(log(PrA / pi))))
+ }
R> n <- c(50, 100, 500, 1074, 1075, 10000)
R> cbind(n, t(sapply(n, FUN = compare_weights)))
```

	n	w1	w2	w3
[1,]	50	1	1	1
[2,]	100	1	1	1
[3,]	500	1	1	1
[4,]	1074	1	1	1
[5,]	1075	NaN	1	1
[6,]	10000	NaN	1	1

For group sizes up to 1074 there is no difference, but when n reaches 1075, w_{1i} returns NaN while w_{2i} and w_{3i} correctly return 1.

Perez-Heydrich *et al.* (2014) used w_{1i} to calculate weights, but 15 groups in their analysis had over 1000 subjects. These groups had missing values for weights for all the values of α considered and were excluded from computing the average IPW estimate. Rather than computing the average IPW across 700 groups, they inadvertently took the average across 685 groups. Correcting the estimates by using w_{2i} or w_{3i} did not alter the conclusions in this case, but analysts should be aware of this issue when dealing with large groups.

7. Discussion

The R package **inference** computes inverse probability weighted estimators of causal effects in the presence of interference. The package currently supports the IPW methods of Tchetgen Tchetgen and VanderWeele (2012) and Perez-Heydrich *et al.* (2014). These methods require a model for the group-level propensity scores. The package provides useful defaults for the propensity models but allows for non-standard models.

Development and application of statistical methods for inferring causal effects in the presence of interference is an active area of research. Future versions of **inference** may incorporate other estimation methods, such as doubly robust methods and stratification. Also, additional methods for estimating variances and effect bounds may be incorporated into the software.

Acknowledgments

We would like to thank Carolina Perez-Heydrich (Meredith College) for sharing code from her analysis, which made debugging and validation much smoother. David Nickerson (University of Notre Dame) graciously agreed to share and distribute the voting experiment data in the **inference** package. Steve Cole (UNC), Brian Barkley (UNC), Betz Halloran (University of Washington), Wen Wei Loh (UNC), and Mary Kirk Wilkinson provided helpful comments. This work was partially supported by NIH grant R01 AI085073. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

References

- Ali M, Emch M, von Seidlein L, Yunus M, Sack DA, Rao M, Holmgren J, Clemens JD (2005). “Herd Immunity Conferred by Killed Oral Cholera Vaccines in Bangladesh: A Reanalysis.” *The Lancet*, **366**(9479), 44–49. doi:10.1016/s0140-6736(05)66550-6.
- Aronow PM, Samii C (2017). “Estimating Average Causal Effects under General Interference, with Application to a Social Network Experiment.” *The Annals of Applied Statistics*. Forthcoming.
- Bates D, Mächler M, Bolker B, Walker S (2015). “Fitting Linear Mixed-Effects Models Using **lme4**.” *Journal of Statistical Software*, **67**(1), 1–48. doi:10.18637/jss.v067.i01.
- Bowers J, Fredrickson MM, Panagopoulos C (2013). “Reasoning About Interference Between Units: A General Framework.” *Political Analysis*, **21**(1), 97–124. doi:10.1093/pan/mps038.
- Clemens J, Shin S, Ali M (2011). “New Approaches to the Assessment of Vaccine Herd Protection in Clinical Trials.” *The Lancet Infectious Diseases*, **11**(6), 482–487. doi:10.1016/s1473-3099(10)70318-2.
- Cole SR, Hernan MA (2008). “Constructing Inverse Probability Weights for Marginal Structural Models.” *American Journal of Epidemiology*, **168**(6), 656–664. doi:10.1093/aje/kwn164.
- Cox DR (1958). *Planning of Experiments*. John Wiley & Sons, New York.
- Duncan GJ, Boisjoly J, Kremer M, Levy DM, Eccles J (2005). “Peer Effects in Drug Use and Sex Among College Students.” *Journal of Abnormal Child Psychology*, **33**(3), 375–385. doi:10.1007/s10802-005-3576-2.
- Eckles D, Karrer B, Ugander J (2017). “Design and Analysis of Experiments in Networks: Reducing Bias from Interference.” *Journal of Causal Inference*, **5**(1). doi:10.1515/jci-2015-0021.
- Foster EM (2010). “Causal Inference and Developmental Psychology.” *Developmental Psychology*, **46**(6), 1454–1480. doi:10.1037/a0020204.

- Gilbert P, Varadhan R (2016). **numDeriv**: *Accurate Numerical Derivatives*. R package version 2016.8-1, URL <https://CRAN.R-project.org/package=numDeriv>.
- Graham DJ, McCoy EJ, Stephens DA (2013). “Quantifying the Effect of Area Deprivation on Child Pedestrian Casualties by Using Longitudinal Mixed Models to Adjust for Confounding, Interference and Spatial Dependence.” *Journal of the Royal Statistical Society A*, **176**(4), 931–950. doi:10.1111/j.1467-985x.2012.01071.x.
- Hong G, Raudenbush SW (2006). “Evaluating Kindergarten Retention Policy: A Case Study of Causal Inference for Multilevel Observation Data.” *Journal of the American Statistical Association*, **101**(475), 901–910. doi:10.1198/016214506000000447.
- Hudgens MG, Halloran ME (2008). “Toward Causal Inference With Interference.” *Journal of the American Statistical Association*, **103**(482), 832–842. doi:10.1198/016214508000000292.
- Kramer ADI, Guillory JE, Hancock JT (2014). “Experimental Evidence of Massive-Scale Emotional Contagion through Social Networks.” *Proceedings of the National Academy of Sciences*, **111**(24), 8788–8790. doi:10.1073/pnas.1320040111.
- Lunceford JK, Davidian M (2004). “Stratification and Weighting via the Propensity Score in Estimation of Causal Treatment Effects: A Comparative Study.” *Statistics in Medicine*, **23**(19), 2937–2960. doi:10.1002/sim.1903.
- Luo X, Small DS, shan R Li C, Rosenbaum PR (2012). “Inference with Interference Between Units in an fMRI Experiment of Motor Inhibition.” *Journal of the American Statistical Association*, **107**(498), 530–541. doi:10.1080/01621459.2012.655954.
- Manski CF (2013). “Identification of Treatment Response with Social Interactions.” *The Econometrics Journal*, **16**(1), S1–S23. doi:10.1111/j.1368-423x.2012.00368.x.
- Moore RT (2016). **blockTools**: *Blocking, Assignment, and Diagnosing Interference in Randomized Experiments*. R package version 0.6-3, URL <https://CRAN.R-project.org/package=blockTools>.
- Nickerson DW (2008). “Is Voting Contagious? Evidence from Two Field Experiments.” *American Political Science Review*, **102**(1), 49–57. doi:10.1017/s0003055408080039.
- Pearl J (2010). “On the Consistency Rule in Causal Inference: Axiom, Definition, Assumption, or Theorem?” *Epidemiology*, **21**(6), 872–875. doi:10.1097/ede.0b013e3181f5d3fd.
- Perez-Heydrich C, Hudgens MG, Halloran ME, Clemens JD, Ali M, Emch ME (2014). “Assessing Effects of Cholera Vaccination in the Presence of Interference.” *Biometrics*, **70**(3), 731–741. doi:10.1111/biom.12184.
- R Core Team (2017). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Rigdon J (2015). **interferenceCI**: *Exact Confidence Intervals in the Presence of Interference*. R package version 1.1, URL <https://CRAN.R-project.org/package=interferenceCI>.

- Robins JM, Hernan MA, Brumback B (2000). “Marginal Structural Models and Causal Inference in Epidemiology.” *Epidemiology*, **11**(5), 550–560. doi:10.1097/00001648-200009000-00011.
- Rosenbaum PR (1987). “Model-Based Direct Adjustment.” *Journal of the American Statistical Association*, **82**(398), 387–394. doi:10.2307/2289440.
- Rubin DB (1980). “Randomization Analysis of Experimental Data: The Fisher Randomization Test Comment.” *Journal of the American Statistical Association*, **75**(371), 591–593. doi:10.2307/2287653.
- Sampson RJ (2010). “Gold Standard Myths: Observations on the Experimental Turn in Quantitative Criminology.” *Journal of Quantitative Criminology*, **26**(4), 489–500. doi:10.1007/s10940-010-9117-3.
- Saul BC (2017). *inference: Methods for Causal Inference with Interference*. R package version 1.0.0, URL <https://CRAN.R-project.org/package=inference>.
- Sinclair B, McConnell M, Green DP (2012). “Detecting Spillover Effects: Design and Analysis of Multilevel Experiments.” *American Journal of Political Science*, **56**(4), 1055–1069. doi:10.1111/j.1540-5907.2012.00592.x.
- Sobel ME (2006). “What Do Randomized Studies of Housing Mobility Demonstrate? Causal Inference in the Face of Interference.” *Journal of the American Statistical Association*, **101**(476), 1398–1407. doi:10.1198/016214506000000636.
- Tchetgen Tchetgen EJ, Coull BA (2006). “A Diagnostic Test for the Mixing Distribution in a Generalised Linear Mixed Model.” *Biometrika*, **93**(4), 1003–1010. doi:10.1093/biomet/93.4.1003.
- Tchetgen Tchetgen EJ, VanderWeele TJ (2012). “On Causal Inference in the Presence of Interference.” *Statistical Methods in Medical Research*, **21**(1), 55–75. doi:10.1177/0962280210386779.
- Toulis P, Kao E (2013). “Estimation of Causal Peer Influence Effects.” In *Proceedings of The 30th International Conference on Machine Learning*, pp. 1489–1497.
- Van der Wal WM, Geskus RB (2011). “**ipw**: An R Package for Inverse Probability Weighting.” *Journal of Statistical Software*, **43**(13), 1–23. doi:10.18637/jss.v043.i13.
- VanderWeele TJ, An W (2013). “Social Networks and Causal Inference.” In SL Morgan (ed.), *Handbook of Causal Analysis for Social Research*, Handbooks of Sociology and Social Research, pp. 353–374. Springer-Verlag.
- VanderWeele TJ, Hong G, Jones SM, Brown JL (2013). “Mediation and Spillover Effects in Group-Randomized Trials: A Case Study of the 4Rs Educational Intervention.” *Journal of the American Statistical Association*, **108**(502), 469–482. doi:10.1080/01621459.2013.779832.
- Verbitsky-Savitz N, Raudenbush SW (2012). “Causal Inference under Interference in Spatial Settings: A Case Study Evaluating Community Policing Program in Chicago.” *Epidemiologic Methods*, **1**(1), 107–130. doi:10.1515/2161-962x.1020.

Zeileis A, Croissant Y (2010). “Extended Model Formulas in R: Multiple Parts and Multiple Responses.” *Journal of Statistical Software*, **34**(1), 1–13. doi:10.18637/jss.v034.i01.

Zigler CM, Dominici F, Wang Y (2012). “Estimating Causal Effects of Air Quality Regulations Using Principal Stratification for Spatially Correlated Multivariate Intermediate Outcomes.” *Biostatistics*, **13**(2), 289–302. doi:10.1093/biostatistics/kxr052.

Affiliation:

Bradley C. Saul
UNC Chapel Hill
Department of Biostatistics
135 Dauer Drive 3101 McGavran-Greenberg
Hall, CB #7420 Chapel Hill, NC 27599-7420
E-mail: saulb@live.unc.edu
URL: <http://www.github.com/bsaul>