Estimating Spatial Effects

Peter M. Aronow and Cyrus Samii*

March 9, 2017

Abstract

We consider causal inference in settings where treatments have effects that transmit over space, potentially in ways that overlap and in violation of the standard "no interference" assumption for many causal inference methods. Applications include public health interventions with dissipating herd immunity or environmental conservation interventions where deterrent effects degrade over distance from points of intervention. We define a spatial "marginalized individualistic response," which characterizes how, on average, units of observation that are a specified distance from an intervention point are affected by treatments at that point, averaging over effects emanating from other intervention points. We establish conditions for non-parametric identification and non-parametric consistency. We also develop a semi-parametric kriging estimator for situations where data are collected only at a subset of observation points. Inference is based on Fisher-type randomization methods.

Version 0.1
(Preliminary and incomplete)

*Peter M. Aronow is Assistant Professor, Departments of Political Science and Biostatistics, Yale University, 77 Prospect St., New Haven, CT 06520 (Email: peter.aronow@yale.edu). Cyrus Samii (contact author) is Assistant Professor, Department of Politics, New York University, 19 West 4th St., New York, NY 10012 (Email: cds2083@nyu.edu).

1 Introduction

Consider a public health intervention that inoculates all individuals at certain locations on a set of days against a contagious disease. "Herd immunity" effects result in protection against the disease for the non-inoculated being a function of the density of nearby neighbors who have been inoculated. Suppose people tend on a regular basis to travel outward from the inoculation locations. This would result in general immunity effects—that is, the combination of direct-inoculation immunity and herd immunity—being a function of proximity to the various places where the inoculations took place.

Or, consider a forest conservation intervention that stations forest rangers at certain locations on the forest edge. Rangers regularly observe forest conditions in the immediate vicinity of the stations, and then conduct occasional patrols outward into the forest. The deterrent effect against illegal logging is a function of the probability of being detected by the rangers. This contribution of any given station to this deterrent effect will depend on distance, and a given point in the forest is vulnerable to illegal logging on the basis of proximity to the various stations.

This paper considers estimating effects in such spatial settings. Treatments are allocated to intervention nodes in a space. The treatments generate effects that transmit outward from the intervention nodes. As such, a point in space may be subject to effects transmitting outward from multiple intervention nodes. The implication is that any given point in the space is, in principle, under the effect of the distribution of treatments over *all* intervention nodes. For the latter example, suppose we situate ourselves at geographical coordinate point $x = (x_1, x_2)$ in the forest. Generally speaking, it would be problematic to simply map point x to a nearby intervention node and assume that the treatment assignment at that intervention node completely characterizes what will happen at point x. We should probably account for the accumulated effects on x emanating from all intervention nodes. What kinds of effects are retrievable in such a setting? Under what conditions?

This paper addresses those questions by defining a spatial "marginalized individualistic re-

sponse," which characterizes how, on average, units of observation that are a specified distance from an intervention point are affected by treatments at that point, averaging over effects emanating from other intervention points. We establish conditions for non-parametric identification and non-parametric consistency for such effects. We also develop a semi-parametric kriging estimator for situations where data are collected only at a subset of observation points. Inference is based on Fisher-type randomization methods. We illustrate with an application to a field experiment on deworming treatments for children in rural Kenya.

The analysis here is most closely connected to the current literature on causal inference under "interference." Interference refers to situations where units are sensitive not only to treatments assigned to them or in their immediate vicinity, but to the array of treatments assigned to all units or over the entirety of the space in which the experiment is being conducted (Cox, 1958). A seminal contribution, and one that inspires our current investigation, is due to Hudgens and Halloran (2008), who characterize effects that are identifiable in situations where the structure of interference is unknown, but rather all that is known is the degree of saturation of treatments assigned at random within blocks. In such settings, one can identify effects that measure, on average, how no exposure differs from direct exposure to treatment, indirect exposure at a given rate of saturation, and the combination of the latter two. Like Hudgens and Halloran, we also consider effects defined in a manner that marginalizes over different profiles of indirect exposure. But a restriction in Hudgens and Halloran's set-up is that indirect exposure is assumed to be restricted to occur only with well-defined blocks, and not between them. This yields what Sobel (2006) refers to as a "partial interference" setting. The current paper considers less restrictive forms of interference that emanate over space, where the strength of exposure may dissipate in a more continuous fashion. Aronow and Samii (2017) also consider interference in settings more general than partial interference and they characterize marginalized effects (see their Prop. 8.1). However their attention is focused on cases where one can specify rather precisely the nature of the interference via an exposure mapping. Our use of marginalized effects avoids such specification. Given the complexities

of statistical inference in this setting, we rely in permutation inference, in a manner that applies insights developed by Aronow (2012) and Athey et al. (2016).

The marginalized individualistic response is an estimand that varies with the share of intervention nodes assigned to treatment in the overall experiment and the spatial distribution of intervention nodes, as these features of the experiment define background levels of "ambient exposure" to treatments. By establishing the marginalized individualistic response as our target of inference, we are, therefore, moving the goal posts away from trying to estimate effects that are more invariant to features of the experimental context. But this is inevitable when trying to identify effects in a manner that is highly agnostic about the nature of interference. Moreover, we show below how to estimate effects conditional on levels of local treatment saturation. These conditional effects may, in some cases, yield results that are more generalizable.

This paper ties the literature on causal inference and interference to the literature on semi-parametric spatial estimation. In developing more efficient estimators of spatial effects, we apply current methods of non- and semi-parametric spatial regression, including kriging based on Gaussian process regression (Cressie, 1993). Our integration of these methods is related to current work on the use of machine learning methods to address other causal inference problems (Van der Laan and Rose, 2011).

We begin with a formalization of the inferential setting, followed by identification and consistency results, a brief discussion of inference, and then our application. The conclusion suggests avenues for further exploration.

2 Setting

Our setting is a sample of points in the $\mathbb{R}^2$ coordinate space, where this sample is denoted as $\mathscr{X} \subseteq \mathbb{R}^2$, and each point in $\mathscr{X}$ is a 2-vector $x = (x_1, x_2)$. Denote x-specific potential outcomes as $Y_x(\mathbf{z})$, where $\mathbf{z}$ is the vector of treatment indicators, with each element $z_i \in \{0,1\}$, assigned to a

set, $\mathscr{N}$, of intervention nodes indexed by $i = 1, ..., N$ and sampled from some large population of potential intervention nodes. We use the notation $(z, \mathbf{z}_{-i})$ to indicate that we are fixing the ith value of the vector $\mathbf{z}$ to z and the other $N-1$ values to the vector $\mathbf{z}_{-i}$. Following Hudgens and Halloran (2008) and Manski (2012), we characterize the treatment regime in terms of a parameter α for the joint distribution of treatment assignments, $\mathbf{Z} = (Z_1, ..., Z_N)$, where each Z_i is a binary random variable on $\{0, 1\}$. Define $N_1 = \sum_{i=1}^{N} Z_i$ as the sum of intervention nodes assigned to treatment. Then $\Pr(\mathbf{Z} = \mathbf{z}|\alpha)$ is the probability that treatment assignment vector $\mathbf{z}$ is realized from among the set of possible assignment vectors, $\mathscr{Z}^0$.

The outcome observed at point x is given by

$$\begin{aligned} Y_x &= \sum_{\mathbf{z} \in \mathscr{Z}} Y_x(\mathbf{z}) I(\mathbf{Z} = \mathbf{z}) \\ &= Z_i \sum_{\mathbf{z}_{-i} \in \mathscr{Z}_{-i}} Y_x(1, \mathbf{z}_{-i}) I(\mathbf{Z}_{-i} = \mathbf{z}_{-i}) + (1 - Z_i) \sum_{\mathbf{z}_{-i} \in \mathscr{Z}_{-i}} Y_x(0, \mathbf{z}_{-i}) I(\mathbf{Z}_{-i} = \mathbf{z}_{-i}), \end{aligned} \tag{1}$$

where $\mathscr{Z}_{-i}$ is the set of potential assignments for units other than i. The decomposition in (1) allows us to interpret an outcome in terms of treatment values for node i. We suppose that we observe Y_x for all $x \in \mathscr{X}$.

We can express an "individualistic" average of potential outcomes for point x, marginalizing over the treatment status of other units

$$Y_{ix}(z; \alpha) \equiv \mathrm{E}_{\mathscr{Z}_0}[Y_x(z, \mathbf{Z}_{-i})] = \mathrm{E}_{\mathscr{Z}}[Y_x(z, \mathbf{Z}_{-i})] = \sum_{\mathbf{z}_{-i} \in \mathscr{Z}_{-i}} Y_x(z, \mathbf{z}_{-i}) \Pr(\mathbf{Z}_{-i} = \mathbf{z}_{-i}|Z_i = z, \alpha). \tag{2}$$

In exactly the same manner as the types of average effects proposed by Hudgens and Halloran (2008), $Y_{ix}(z; \alpha)$ depends on the design. That is, it does not incorporate potential outcome values corresponding to $\mathbf{z}_{-i}$ values that are impossible by design and thus for which $\Pr(\mathbf{Z}_{-i} = \mathbf{z}_{-i}|Z_i = z, \alpha) = 0$. It is also sensitive to the way the experimental design, characterized by α, determines $\Pr(\mathbf{Z}_{-i} = \mathbf{z}_{-i}|Z_i = z, \alpha)$ over the $\mathbf{z}_{-i}$ values. One could construct an

individualistic average that is less sensitive to the design by using weights to counterbalance variation in $\Pr(\mathbf{Z}_{-i} = \mathbf{z}_{-i}|Z_i = z, \alpha)$ over the $\mathbf{z}_{-i}$ values, although the set of potential outcomes that one can incorporate may be intrinsic to the design. These points are clear in a simple case where there are only 2 treatment nodes. Then, the full schedule of potential outcomes is $(Y_x(0,0), Y_x(0,1), Y_x(1,0), Y_x(1,1))$. Now consider Bernoulli assignment, equivalent to using a separate coin flip (perhaps weighted) to determine treatment assignment for each intervention node. Under Bernoulli assignment, N_1 is random over the support $\{0,1,2\}$, and there would be the potential to observe each of $(Y_x(0,0), Y_x(0,1), Y_x(1,0), Y_x(1,1))$ and $Y_{ix}(z;\alpha) = (Y_x(z,z) + Y_x(z,1-z))/2$. The situation with complete random assignment and $N_1 = 1$ fixed is different, since units are assigned to treatment without replacement. With complete random assignment we would only ever observe $Y_x(0,1)$ and $Y_x(1,0)$, and so we simply have $Y_{ix}(z;\alpha) = Y_x(z,1-z)$.

In addition to data on the treatment vector, $\mathbf{Z}$, and outcomes, Y_x for $x \in \mathscr{X}$, we also observe covariate data. We define two specific types of covariates that we will use below. First, let d_{ix} be a measure of the distance between treatment node i and x. In the application below we will use Euclidean distance although other distance metrics are possible. Second, define

$$N_x(d) = \sum_{i=1}^{N} I(d_{ix} \leq d),$$

that is, the number of intervention nodes that are within distance d of point x. The covariate $N_x(d)$ is useful for characterizing the potential level of treatment saturation to which a point x could be exposed from intervention nodes within distance d. A point in the coordinate space can be characterized more generally in terms of a covariate vector, W_x, of which the d_{ix} and $N_x(d)$ terms are elements.

We focus on the case of a randomized experiment such that for the assignment vector $\mathbf{Z}$ we have

$$\mathbf{Z} \perp\!\!\!\perp (Y_x(\mathbf{z}), Y_x(\mathbf{z}), W_x), \tag{3}$$

and

$$0 < \Pr(\mathbf{Z} = \mathbf{z}|\alpha) < 1 \tag{4}$$

for all $\mathbf{z} \in \mathscr{Z} \subseteq \mathscr{Z}^0$, where $\mathscr{Z}$ is non-empty. Conditions (3) and (4) are satisfied by complete or Bernouilli random assignment. Condition (3) would violated if, for example, intervention nodes in areas that are more sensitive to having treatment nearby were either more or less likely to receive treatment than nodes in areas where sensitivity is lower. Generalizations of this setting could consider sets of feasible assignments, $\mathscr{Z}$, that are further restricted to yield randomizations that induce certain patterns in the distribution of exposure to treatment assignments over $\mathscr{X}$. For example, there may be interest in inducing certain forms of variation in the saturation of treatments over $\mathscr{X}$.

3 Causal Effects and Identification

A point-specific causal effect at x of switching treatment at node i from 0 to 1 is given by,

$$\tau_{ix}(\alpha) = Y_{ix}(1;\alpha) - Y_{ix}(0;\alpha).$$

Of course, this is an unobservable quantity, as treatment at node i must either be 0 or 1. As usual, then, our target of inference will be an average of such causal effects. Specifically, we define an average causal effect that marginalizes over points x at distance d from all intervention nodes as follows:

$$\tau(d;\alpha) = \mathrm{E}_{\mathscr{N}}\left[\mathrm{E}_{\mathscr{X}}\left[\tau_{ix}(\alpha)|d_{ix} = d\right]\right]. \tag{5}$$

The inner expectation averages over over x points about intervention node i at radius d, and then the outer expectation averages over intervention nodes. Following Manski (2012) we call $\tau(d;\alpha)$ the "marginalized individualistic response" (MIR) for locations in the space at distance d from an intervention node.

We can also define a conditional marginalized individualistic response (CMIR) as,

$$\tau(d;\alpha;w) = \mathrm{E}_{\mathscr{N}}[\mathrm{E}_{\mathscr{X}}[\tau_{ix}(\alpha)|d_{ix} = d, W_x = w]],$$

where again the inner expectation averages over x points about intervention node i at radius d and for which $W_x = w$, and the outer expectation averages over the intervention nodes.

We begin with a basic identification result for the CMIR under random assignment and positivity assumptions. This result also covers identification of the MIR as a special case.

Proposition 1. *Define the following difference in integrated conditional means,*

$$\begin{aligned}\tau(d;\alpha;w;\mathbf{Z}) =& \mathrm{E}_{\mathscr{N}|Z_i=1}\left[\mathrm{E}_{\mathscr{X}}[Y_x|d_{ix} = d, Z_i = 1, W_x = w]\right] \\ &- \mathrm{E}_{\mathscr{N}|Z_i=0}\left[\mathrm{E}_{\mathscr{X}}[Y_x|d_{ix} = d, Z_i = 0, W_x = w]\right]. \end{aligned} \tag{6}$$

Then, under (3) *and* (4)*, the difference in integrated conditional means is unbiased for* $\tau(d;\alpha;w)$*, as*

$$\tau(d;\alpha;w) = \mathrm{E}_{\mathscr{Z}}[\tau(d;\alpha;w;\mathbf{Z})], \tag{7}$$

where the expectations marginalizes over treatment assignments.

Proof. Note that by (3),

$$Z_i \perp\!\!\!\perp (Y_x(1;\alpha), Y_x(0;\alpha), W_x) \tag{8}$$

and

$$\mathbf{Z}_{-i} \perp\!\!\!\perp (Y_x(1,\mathbf{z}_{-i}), Y_x(0,\mathbf{z}_{-i}), W_x) \text{ for all } \mathbf{z}_{-i} \in \underline{\mathscr{Z}}_{-i}. \tag{9}$$

Moreover, by (3) and (4), for $z = 0, 1$ and all $\mathbf{z}_{-i} \in \underline{\mathscr{Z}}_{-i}$,

$$\Pr(\mathbf{Z}_{-i} = \mathbf{z}_{-i}|\alpha; Z_i = z) > 0. \tag{10}$$

Then, by (8), (9), and (10),

$$
\begin{aligned}
&\mathrm{E}_{\mathscr{Z}}\left[\mathrm{E}_{\mathscr{N}|Z_i=1}\left[\mathrm{E}_{\mathscr{X}}[Y_x|d_{ix}=d,Z_i=1,W_x=w]\right]\right] \\
&\qquad=\mathrm{E}_{\mathscr{Z}}\left[\mathrm{E}_{\mathscr{N}|Z_i=1}\left[\mathrm{E}_{\mathscr{X}}[\sum_{\mathbf{z}_{-i}\in\mathscr{Z}_{-i}} Y_x(1,\mathbf{z}_{-i})I(\mathbf{Z}_{-i}=\mathbf{z}_{-i})\Big|d_{ix}=d,Z_i=1,W_x=w]\right]\right] \\
&\qquad=\mathrm{E}_{\mathscr{Z}}\left[\mathrm{E}_{\mathscr{N}|Z_i=1}\left[\mathrm{E}_{\mathscr{X}}[\sum_{\mathbf{z}_{-i}\in\mathscr{Z}_{-i}} Y_x(1,\mathbf{z}_{-i})\Pr(\mathbf{Z}_{-i}=\mathbf{z}_{-i}|\alpha)\Big|d_{ix}=d,Z_i=1,W_x=w]\right]\right] \\
&\qquad=\mathrm{E}_{\mathscr{Z}}\left[\mathrm{E}_{\mathscr{N}|Z_i=1}\left[\mathrm{E}_{\mathscr{X}}[Y_{ix}(1;\alpha)|d_{ix}=d,Z_i=1,W_x=w]\right]\right] \\
&\qquad=\mathrm{E}_{\mathscr{Z}}\left[\mathrm{E}_{\mathscr{N}}\left[\mathrm{E}_{\mathscr{X}}[Y_{ix}(1;\alpha)|d_{ix}=d,W_x=w]\right]\right] \\
&\qquad=\mathrm{E}_{\mathscr{N}}\left[\mathrm{E}_{\mathscr{X}}[Y_{ix}(1;\alpha)|d_{ix}=d,W_x=w]\right].
\end{aligned}
$$

We can perform similar for $Z_i=0$. Putting the two together and taking expectations over intervention nodes yields the expression for $\tau(d;\alpha;w)$. □

Corollary 1. *Given conditions* (3) *and* (4)*,*

$$\tau(d;\alpha)=\mathrm{E}\left[\mathrm{E}[Y_x|d_{ix}=d,Z_i=1]\right]-\mathrm{E}\left[\mathrm{E}[Y_x|d_{ix}=d,Z_i=0]\right]. \tag{11}$$

4 Local Effects

The marginalized individualistic treatment effect is conditional on α, which defines an overall level of "ambient exposure" to treatment in the experiment. Proposition 1 demonstrates that the difference in integrated means, $\tau(d;\alpha;w;\mathbf{Z})$, identifies effects that are "local" in the sense of conditioning on covariate values, w. Other types of local effects are often of interest, in particular, effects that condition on the level of local treatment saturation. For example, we might ask: what is the effect of adding putting an additional intervention node under treatment given that we already have some number, s, of nearby intervention nodes under treatment? Effects of this variety are different

than what is defined in Proposition 1 insofar as saturation-specific effects require conditioning on values of $\mathbf{Z}$ rather than W_x values.

The conventional approach to estimating the effects of local treatment saturation, as in Hudgens and Halloran (2008), is to operate under an assumption of partial interference and then to make comparisons across levels of treatment saturation. Here we consider an alternative approach based on the marginalized individualistic response.

Define a local treatment saturation variable,

$$S_x(d) = \sum_{i=1}^{N} Z_i I(d_{ix} \leq d).$$

By definition $S_x(d) \leq N_x(d)$, where the latter is the number of intervention nodes within distance d from point x, as defined above. Random assignment of $\mathbf{Z}$ does not ensure that probabilities of $S_x(d)$ values are uniform over $\mathscr{X}$, since probability of a given realization of $S_x(d)$ depends on $N_x(d)$. Let $\pi_x(s,d) = \Pr(S_x(d) = s)$. For example, given complete random assignment,

$$\pi_x(s,d) = \binom{N_x(d)}{s}\binom{\max(N_1 - N_x(d), 0)}{\max(N_1 - s, 0)} \Big/ \binom{N}{N_1}.$$

Aronow and Samii (2017) show conditions under which we can identify effects of variations in $S_x(d)$ by accounting for variation in the $\pi_x(s,d)$ values. Their analysis focuses on the cases where one can assume that the potential outcome mapping is relatively simple in terms of $S_x(d)$ values. Here, we drop that assumption and relate the marginalized individualistic response to levels of local treatment saturation.

Now, define a local treatment saturation variable that defines ambient exposure with respect to an intervention node, i,

$$S_{x,i}(d) = \sum_{j \neq i} Z_j I(d_{jx} \leq d).$$

Proposition 2. *Define the following conditional difference in integrated means for $d \leq d'$:*

$$\begin{aligned}\tau(d;\alpha;n;\mathbf{Z},s) =& \mathrm{E}_{\mathscr{N}|Z_i=1}\left[\mathrm{E}_{\mathscr{X}}\left[Y_x|d_{ix}=d,Z_i=1,N_x(d')=n,S_{x,i}(d')=s\right]\right] \\ &- \mathrm{E}_{\mathscr{N}|Z_i=0}\left[\mathrm{E}_{\mathscr{X}}\left[Y_x|d_{ix}=d,Z_i=0,N_x(d')=n,S_{x,i}(d')=s\right]\right]. \qquad (12)\end{aligned}$$

Suppose the conditions of Proposition 1 hold and that $\mathbf{Z}$ *is distributed such that the following also hold:*

- $\pi_x(s,d) = \pi_{x'}(s,d)$ *for all* x, x' *such that* $N_x(d) = N_{x'}(d)$*, and*
- *for all* s *and* $(\mathbf{z},\mathbf{z}')$ *such that* $\mathbf{z} \neq \mathbf{z}'$ *and* $S_x(d) = s$, $\Pr(\mathbf{Z}=\mathbf{z}|S_x(d)=s) = \Pr(\mathbf{Z}=\mathbf{z}'|S_x(d)=s)$.

Then the conditional difference in integrated means is unbiased for the conditional local effect, as,

$$\mathrm{E}_{\mathscr{X}|N_x(d')=n}[\tau(d;\alpha)|N_x(d')=n,S_{x,i}(d')=s] = \mathrm{E}_{\mathscr{X}}[\tau(d;\alpha;n;\mathbf{Z},s)].$$

Proof. (Sketch.) The result follows directly from the independence of treatment assignment and uniform assignment probabilities. □

Complete random assignment and Bernoulli random assignment are examples of treatment assignment processes satisfying the conditions of Proposition 2. In addition, the proposition establishes the basis for designs that seek to generate optimal variation in local saturation along the lines of the designs discussed by Baird et al. (2016) in the partial interference setting.

5 Non-parametric estimation

For the remainder of the presentation, we will suppress the conditioning notation. All results can be restated in a conditional sense by reintroducing the conditioning into the respective expectations, covariance, and other such operators below.

Suppose $\mathscr{X}$ consists of a finite lattice of points. Then for an intervention node i, we have $n_i(d)$ points that are at distance d and

$$E_{\mathscr{X}}[Y_x|d_{ix}=d] = \frac{1}{n_i(d)} \sum_{x:d_{ix}=d} Y_x \equiv \bar{Y}_i(d),$$

and for potential outcomes,

$$E_{\mathscr{X}}[Y_{ix}(z;\alpha)|d_{ix}=d] = \frac{1}{n_i(d)} \sum_{x:d_{ix}=d} Y_{ix}(z;\alpha) \equiv \bar{Y}_i(z;d).$$

Let $p = N^{-1}\sum_{i=1}^{N} Z_i$, the proportion of intervention nodes that are treated. An analogue estimator for the right hand side of (12) is given by

$$\hat{\tau}(d) = \frac{1}{Np}\sum_{i=1}^{N} Z_i\bar{Y}_i(d) - \frac{1}{N(1-p)}\sum_{i=1}^{N}(1-Z_i)\bar{Y}_i(d). \tag{13}$$

The estimator amounts to an average difference of averages—that is the difference in the average, over nodes, of the averages, over points, of outcomes at a given radius about the treatment and control nodes.

Proposition 3. *Suppose the conditions of Proposition 1 hold and*

$$\sum_{i=1}^{N}\sum_{j\neq i} \mathrm{Cov}\left[Z_i\bar{Y}_i(d), Z_j\bar{Y}_j(d)\right] = o(N^2), \tag{14}$$

$$\sum_{i=1}^{N}\sum_{j\neq i} \mathrm{Cov}\left[\bar{Y}_i(d), \bar{Y}_j(d)\right] = o(N^2), \tag{15}$$

and

$$\mathrm{E}_{\mathscr{Z}}[Z_i] = \bar{p} \text{ for all } i = 1,...,N. \tag{16}$$

Then as $N \to \infty$, $\hat{\tau}(d) \to \tau(d;\alpha)$.

Proof. As $N \to \infty$, $p \to \bar{p}$. By Slutsky's theorem, $\hat{\tau}(d)$ has the same limit as,

$$\begin{aligned}\hat{\tau}_{\bar{p}}(d) &= \frac{1}{N\bar{p}} \sum_{i=1}^{N} Z_i \bar{Y}_i(d) - \frac{1}{N(1-\bar{p})} \sum_{i=1}^{N} (1-Z_i)\bar{Y}_i(d) \\ &= \frac{1}{N\bar{p}} \sum_{i=1}^{N} Z_i \bar{Y}_i(1;d) - \frac{1}{N(1-\bar{p})} \sum_{i=1}^{N} (1-Z_i)\bar{Y}_i(0;d),\end{aligned}$$

where by (16) and (3)

$$\mathrm{E}_{\mathscr{Z}}[\hat{\tau}_{\bar{p}}(d)] = \frac{1}{N\bar{p}} \sum_{i=1}^{N} \mathrm{E}\left[\mathrm{E}\left[Z_i|\bar{Y}_i(1;d)\right]\bar{Y}_i(1;d)\right] - \frac{1}{N(1-\bar{p})} \sum_{i=1}^{N} \mathrm{E}\left[\mathrm{E}\left[(1-Z_i)|\bar{Y}_i(0;d)\right]\bar{Y}_i(0;d)\right] = \tau(d;\alpha),$$

implying $\mathrm{E}_{\mathscr{Z}}[\hat{\tau}(d) - \tau(d;\alpha)] \to 0$ as $N \to \infty$. Rearranging terms,

$$\hat{\tau}_{\bar{p}}(d) = \left(\frac{1}{N\bar{p}} + \frac{1}{N(1-\bar{p})}\right) \sum_{i=1}^{N} Z_i \bar{Y}_i(d) - \frac{1}{N(1-\bar{p})} \sum_{i=1}^{N} \bar{Y}_i(d).$$

Conditions (14) and (15) imply $\mathrm{Var}[N\hat{\tau}_{\bar{p}}(d)] \to 0$, by condition (14) for the sum on the left and condition (15) for the sum on the right and the difference in sums. Chebychev's inequality completes the proof. □

Effectively conditions (14) and (15) require "local dependence" in radius-specific outcome means. In a spatial context, this would be satisfied under (i) an asymptotic growth process that adds intervention nodes moving outward in space beyond the radial distances for which we are estimating effects and (ii) for which we have a sufficiently high rate of strong mixing of outcomes over distance.

Condition (16) could be loosened to allow for unequal probabilities of assignment to treatment, in which case consistent estimation would require weighting units by the inverse of the probability of assignment to their treatment conditions (Aronow and Samii, 2015).

6 Semi-parametric estimation

Often times we obtain measurements at only a subset of points in the spatial field. For example we may have data only on a sample of units arrayed over the spatial field (e.g., a sample of households that reside at specific points in the field, but for which there are other households at other points, not included in the sample). In such cases, non-parametric estimation may have very high variance, since each radius-specific mean may incorporate only a few observations. We could improve upon such estimates in terms of variance if we had a reasonable method for interpolating outcome values between points in the spatial field, albeit with the potential to introduce some bias.

We consider a semi-parametric, Gaussian process kriging estimator for interpolating outcome values at spatial points for which we have no direct observation data (Cressie, 1993, Ch. 3). The estimator is based on a model for spatial points where

$$Y_x = P(x) + Z(x) + e_x,$$

where $P(x)$ is a polynomial expansion of the coordinates, x, $Z(.)$ is a mean zero Gaussian field with a covariance function of arbitrary decay, and $e(.)$ is independent mean-zero error with finite variance.

Denote the interpolated values from the kriging estimator as $\hat{Y}_x$ for $x \in \mathscr{X}$. Operating again on $\mathscr{X}$ as a finite lattice of, denote as well the means of interpolated outcomes at radius d from node i as $\bar{\hat{Y}}_i(d)$. Using these interpolated values, we can construct the estimator,

$$\hat{\tau}_k(d) = \frac{1}{Np}\sum_{i=1}^{N} Z_i \bar{\hat{Y}}_i(d) - \frac{1}{N(1-p)}\sum_{i=1}^{N}(1-Z_i)\bar{\hat{Y}}_i(d), \tag{17}$$

where the $\bar{\hat{Y}}_i(d)$ are the d-specific radial means of the interpolated $\hat{Y}_x$ values about treated and control intervention nodes.

Proposition 4. *Suppose conditions of Propositions 1 and 3 hold. Denote the size of the sample of*

evaluation points as $|\mathscr{X}|$*, and suppose that as* $|\mathscr{X}| \to \infty$*,* $\bar{\hat{Y}}_i(d) \to \bar{Y}_i(d)$ *for all* i *and* d*. Then, as both* $N \to \infty$ *and* $|\mathscr{X}| \to \infty$*,* $\hat{\tau}_k(d) \to \tau(d;\alpha)$*.*

Proof. The results follows from application of Slutsky's theorem. □

Proposition 4 shows, simply, that consistency in $|\mathscr{X}|$ for the estimated radial means translates into consistency of $\hat{\tau}_k(d)$ as both the number of intervention nodes, N, and evaluation points, $|\mathscr{X}|$, grow. Efficiency gains follow from the fact that the kriging estimator interpolates information on evaluation points that would otherwise be ignored completely by the non-parametric estimator.

7 Randomization Inference

Our agnosticism as to the nature of interference admits dependencies in potential outcomes that may be highly complex. Following Aronow (2012), Athey et al. (2016), and Rosenbaum (2007), we side-step the need to characterize such dependencies explicitly by relying on Fisher-type randomization inference, which relies on permuting treatments under the sharp null and assessing effect estimates relative to the sharp null permutation distribution (Fisher, 1966; Imbens and Rubin, 2015, Ch. 5). That is, we set out the sharp null hypothesis,

$$H_0 : Y_x(\mathbf{z}) = Y_x(\mathbf{z}') \text{ for all } \mathbf{z}, \mathbf{z}' \in \mathscr{Z}.$$

If we have an estimator under assignment $\mathbf{z}$ given by $\hat{\tau}_{\mathbf{z}}(d)$, we can define values of this estimator given a permutation of $\mathbf{z}$, call it $\mathbf{z}'$, and applying H_0 as $\hat{\tau}^0_{\mathbf{z}'}(d)$. The value of $\hat{\tau}^0_{\mathbf{z}'}(d)$ can be estimated for all $\mathbf{z}' \in \mathscr{Z}$, as all it requires is taking the $\mathbf{z}'$ values and applying them against the observed Y_x values. Under uniform assignment and the sharp null, each $\mathbf{z}$ is equally likely and the permutation distribution function for $\hat{\tau}^0_{\mathbf{Z}}(d)$ induced by the treatment variable, $\mathbf{Z}$, is simply

$$\Pr(\hat{\tau}^0_{\mathbf{Z}}(d) \leq t | H_0) = \frac{1}{|\mathscr{Z}|} \sum_{\mathbf{z} \in \mathscr{Z}} I(\hat{\tau}^0_{\mathbf{z}}(d) \leq t).$$

For other $\mathbf{Z}$ with other distributions, the permutation distribution function, $\Pr(\hat{\tau}^0_{\mathbf{Z}}(d) \leq t|H_0)$, can be approximated to arbitrary precision by sampling independent draws of $\mathbf{Z}$ based on the distribution of $\mathbf{Z}$ and then using the empirical distribution of the $\tau^0_{\mathbf{Z}}(d)$ values computed from each of the treatment draws.

Given an estimate $\hat{\tau}(d)$, we can compute an exact Fisher p-value for a test at level α of the sharp null as

$$p = \Pr(\hat{\tau}^0_{\mathbf{Z}}(d) \leq \hat{\tau}(d)|H_0),$$

where one rejects the null against a one-sided alternative that potential outcomes vary over treatment assignments in a manner such that $\tau(d,\alpha) > 0$ if $p > 1-\alpha$, one rejects the null against a one-sided alternative that $\tau(d,\alpha) < 0$ if $p < \alpha$, and one rejects the null against a two-sided alternative if p is outside the interval $[\alpha/2, 1-\alpha/2]$. In our application below, we also use graphical approach that plots the $\hat{\tau}(d)$ estimate against the distribution of $\hat{\tau}^0_{\mathbf{Z}}(d)$.

8 Application

Miguel and Kremer (2004) study the effects of treatments to inoculate children at schools in rural Kenya against intestinal worms. The deworming treatments are considered to exhibit herd immunity effects, given that the inoculation of some children in a community reduces the likelihood that non-inoculated children in that community come into contact with infection-causing intestinal worms. Given that children travel from communities at various distances to each school, the expectation is that there herd immunity will dissipate in distance to the intervention points (that is, the schools). The original Miguel and Kremer (2004) study has come under scrutiny for the assumptions behind their analysis of these types of spatial spillover effects. See Evans (2015) for a full account. In particular, Miguel and Kremer analyze spatial effects by defining strata based on a coarsening of the measure of distance from a school. Their strata were based on whether schools where within 3 km of a school and then between 3 and 6 km from a school. A question among crit-

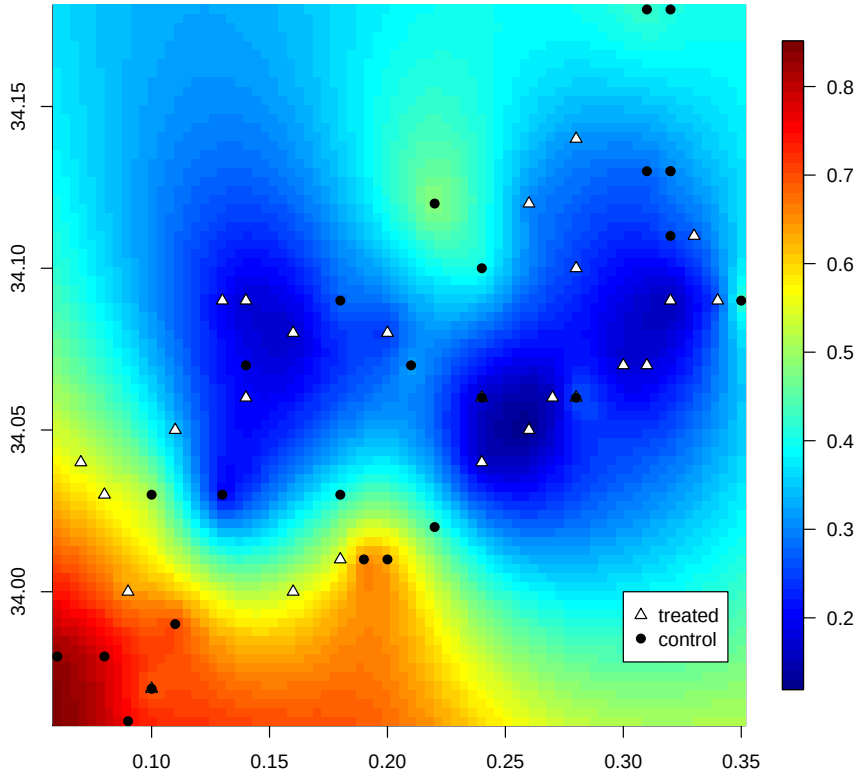

Figure 1: Kriging fit to Miguel and Kremer (2004) data. Outcome is worms infection rate at followup. Treated and control intervention points (schools) are plotted.

ics was whether the findings with respect to spillover effects are robust to other characterizations of the spillover. The methods we have defined here allow us to avoid having to make arbitrary specification assumptions.

Figure 1 shows the kriging fit to data on worms infections at follow-up, and Figure 2 shows our estimates of the MIR effects over distance from the school intervention points. When evaluated against the sharp null distribution, we find evidence for statistically significant effects at distances of less than about 5 km from schools, but not beyond that. These results are basically consistent with what Miguel and Kremer (2004) find. Our results make one more confident that these

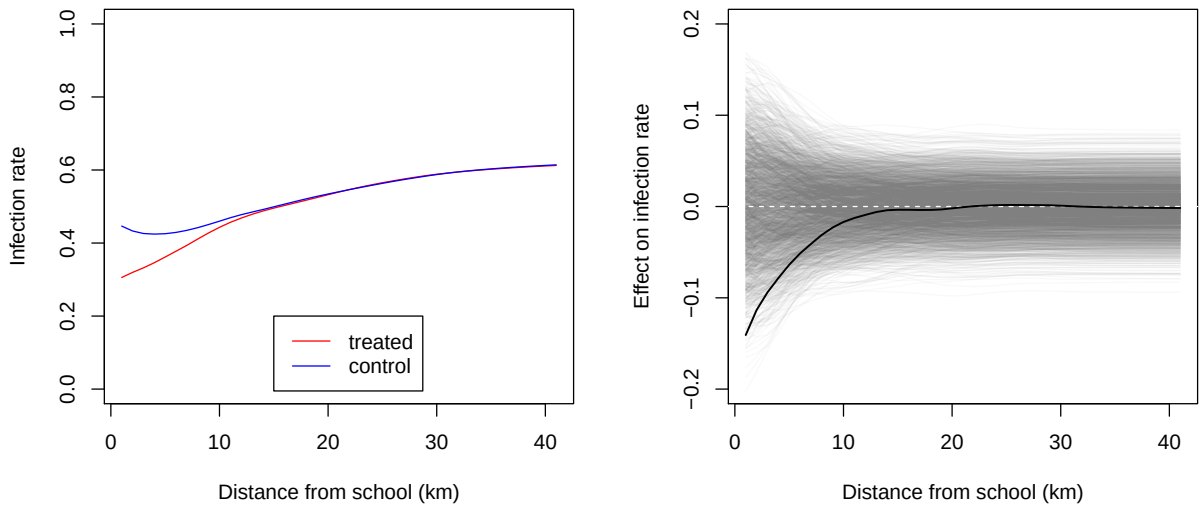

Figure 2: (left) Integrated treated and control means over distance from schools and (right) estimated marginalized individualistic response effects (black) with sharp null estimate distribution (gray).

conclusions are not based on arbitrary specification assumptions.

9 Conclusion

We have defined the "marginalized individualistic response" (MIR) as a type of effect that is identified in randomized experiments where outcomes for any unit of observation are affected by the entire set of treatment assignments. We develop results for the identification and estimation of a spatial variant of this effect. This characterizes how, on average, switching an intervention node from control to treatment would affect outcomes over space, averaging over ambient interference emanating from other nodes. The results admit various types of conditional effects. A spatial kriging estimator enhances efficiency in cases where spatial outcomes are observed at only a sample of points in the space. Fisher-type randomization inference provides a basis for principled inference.

The MIR effect has a clear policy implications: they characterize effect associated with the

marginal treated locality. As such, they address the question, what would be the expected value of adding one more treated node given a specific number of treated nodes already in place? We also examined a conditional version of the MIR effect, conditioning on local levels of treatment saturation. If the actual effects tend, in fact, to be relatively local, these conditional effects allow us to characterize incremental effects under various local treatment density settings. An experiment could be optimized in varying such local treatment densities. This would be a generalization over the types of experiments considered by Hudgens and Halloran (2008) and Baird et al. (2016), who operate under assumptions of strictly partial interference. The results here allow one to think in terms of local interference in a manner that resembles smooth mixing processes, rather than in terms of rigid blocks within which interference is contained.

References

Aronow, P. M. (2012). A general method for detecting interference between units in randomized experiments. *Sociological Methods and Research 41*(1), 3–16.

Aronow, P. M. and C. Samii (2015). Estimating average causal effects under general interference. Manuscript, Yale University and New York University.

Aronow, P. M. and C. Samii (2017). Estimating average causal effects under general interference, with application to a social network experiment. *Annals of Applied Statistics* (forthcoming).

Athey, S., D. Eckles, and G. W. Imbens (2016). Exact p-values of network inference. *Journal of the American Statistical Association* (forthcoming).

Baird, S., J. Bohren, C. McIntosh, and B. Ozler (2016). Optimal design of experiments in the presence of interference. Typescript, George Washington University, University of Pennsylvania, University of California-San Diego, and the World Bank.

Cox, D. R. (1958). *Planning of Experiments*. Wiley.

Cressie, N. A. C. (1993). *Statistics for Spatial Data*. New York, NY: Wiley.

Evans, D. (2015, August). Worm wars: The anthology. World Bank Development Impact Blog, http://blogs.worldbank.org/impactevaluations/worm-wars-anthology.

Fisher, R. A. (1966). *The Design of Experiments*. New York: Hafner.

Hudgens, M. G. and M. E. Halloran (2008). Toward causal inference with interference. *Journal of the American Statistical Association 103*(482), 832–842.

Imbens, G. W. and D. B. Rubin (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge: Cambridge University Press Press.

Manski, C. F. (2012). Identification of treatment response with social interactions. *The Econometrics Journal* (In press).

Miguel, E. and M. Kremer (2004). Worms: Identifying impacts on education and health in the presence of treatment externalities. *Econometrica 72*(1), 159–217.

Rosenbaum, P. R. (2007). Interference between units in randomized experiments. *Journal of the American Statistical Association 102*(477), 191–200.

Sobel, M. E. (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.

Van der Laan, M. and S. Rose (2011). *Targeted Learning: Causal Inference for Obervational and Experimental Data*. New York, NY: Springer.