Biometrika (2018), **105**, 4, *pp.* 849–858
Printed in Great Britain
doi: 10.1093/biomet/asy036
Advance Access publication 6 August 2018

Model-assisted design of experiments in the presence of network-correlated outcomes

By GUILLAUME W. BASSE

Department of Statistics, Harvard University, 1 Oxford Street, Cambridge, Massachusetts 02138, U.S.A.
gbasse@fas.harvard.edu

and EDOARDO M. AIROLDI

Department of Statistical Science, Fox School of Business, Temple University, 1810 Liacouras Walk, Philadelphia, Pennsylvania 19122, U.S.A.
airoldi@alumni.harvard.edu

Summary

In this paper we consider how to assign treatment in a randomized experiment in which the correlation among the outcomes is informed by a network available pre-intervention. Working within the potential outcome causal framework, we develop a class of models that posit such a correlation structure among the outcomes. We use these models to develop restricted randomization strategies for allocating treatment optimally, by minimizing the mean squared error of the estimated average treatment effect. Analytical decompositions of the mean squared error, due both to the model and to the randomization distribution, provide insights into aspects of the optimal designs. In particular, the analysis suggests new notions of balance based on specific network quantities, in addition to classical covariate balance. The resulting balanced optimal restricted randomization strategies are still design-unbiased when the model used to derive them does not hold. We illustrate how the proposed treatment allocation strategies improve on allocations that ignore the network structure.

Some key words: Causal inference; Degree distribution; Network balance; Network data; Optimal treatment allocation; Randomized experiment; Rerandomization.

1. Introduction

The past decade has witnessed a surge of interest in causal analyses in the context of social networks, social media platforms and online advertising (Christakis & Fowler, 2007; Aral et al., 2009; Bakshy et al., 2011, 2012; Bond et al., 2012; Gui et al., 2015; Kim et al., 2015; Phan & Airoldi, 2015; Cavusoglu et al., 2016). From a statistical perspective, the challenging aspect of these applications is how to account for the presence of connections, or network data, observed pre-intervention, possibly with uncertainty. While there is a well-developed literature on several aspects of the statistical analysis of network data (Wasserman & Faust, 1994; Bickel & Chen, 2009; Goldenberg et al., 2010; Kolaczyk & Csárdi, 2014), the literature on methods for experimentation and causal analyses that use observed connections is still nascent (Rosenbaum, 2007; Hudgens & Halloran, 2008; Toulis & Kao, 2013; Ogburn & VanderWeele, 2017).

The need to account for network connections in causal analyses has led scholars to focus on two specific problem settings: network interference (Toulis & Kao, 2013; Ugander et al., 2013;

Aronow & Samii, 2017; Eckles et al., 2017), where the potential outcomes of unit i are functions of the treatment assigned to unit i and of the treatments assigned to other units that are related to unit i through the network; and network-correlated outcomes (McPherson et al., 2001; Shalizi & Thomas, 2011; Manski, 2013), where the network informs the correlation among the potential outcomes because the potential outcomes of unit i depend on its covariates, and the covariates of units that are connected are more similar than the covariates of those that are not. In this paper, we focus on the second setting, which has received less attention.

Restricted randomization as a way to increase the precision of estimates has a long tradition (Yates, 1948; Youden, 1972; Simon, 1979; Bailey, 1983; Higham et al., 2015). The basic idea is that some assignments are considered problematic and should be excluded; stratified randomization, for instance, implicitly excludes assignments for which certain covariates of interest are unbalanced between the treatment arms. In networks, the challenge is to identify the features that must be balanced, which makes it difficult to know how to restrict the randomization. Drawing inspiration from the model-assisted survey sampling literature (e.g., Särndal et al., 2003), we propose a model-assisted approach to design. We posit a working model for the potential outcomes specified conditionally on a network observed pre-intervention, and then restrict the randomization to assignments for which the estimator of interest achieves a low mean squared error. The class of models we propose leads to analytical expressions for the mean squared error that suggest new notions of balance in terms of network statistics related to the degree distribution. We also develop theoretical results showing that our model-assisted restricted randomization approach maintains the design-unbiasedness of the difference-in-means estimator even when the model is misspecified, and reduces its expected variance when the model holds.

2. Analytical insights for evaluating allocations

2·1. *Causal inference set-up*

We work within the potential outcomes framework (Rubin, 1974; Holland, 1986; Imbens & Rubin, 2015). We consider a population of N units, a binary treatment, denoted by $Z_i = 1$ if unit i is assigned to treatment, and real-valued outcomes, denoted by Y_i. The corresponding vectors are denoted by Y and Z. We make the stable unit-treatment-value assumption, so the outcome of unit i is a function only of the treatment assigned to it, $Y_i(Z) = Y_i(Z_i)$, thus excluding interference (Rubin, 1974). We consider a finite population setting, where the potential outcomes $Y(Z)$ are unknown constant quantities, given Z. The only source of variation is how treatment is allocated to units, which is done according to a distribution on the space of all binary vectors of length N, called the randomization distribution.

To illustrate model-assisted restricted randomization, we consider the average treatment effect as the inferential target of interest, $\tau^* = N^{-1}\sum_{i=1}^{N}\{Y_i(1) - Y_i(0)\}$, and the difference-in-means estimator of the average treatment effect,

$$\hat{\tau}(Y \mid Z) = \frac{\sum_{i=1}^{N} Z_i Y_i}{\sum_{i=1}^{N} Z_i} - \frac{\sum_{i=1}^{N}(1 - Z_i) Y_i}{\sum_{i=1}^{N}(1 - Z_i)}. \tag{1}$$

2·2. *The normal-sum model*

The model-assisted approach to experimental design requires a model, which is used to improve the inferential properties of the difference-in-means estimator when the model holds. We posit a model that depends on a network, which is available at the design stage.

Consider N units and an undirected network $\mathcal{G}$ among them or, equivalently, a binary adjacency matrix A of size $N \times N$ with the added constraint that $A_{ii} = 1$ for all i; we call A the extended adjacency matrix. The neighbourhood of a unit i is defined as the index set $\mathcal{N}_i = \{j : A_{ij} = 1 \text{ or } A_{ji} = 1\}$. Consider the model

$$X_j \sim N(\mu, \sigma^2), \tag{2}$$

$$Y_i(0) \mid X \sim N\big(\textstyle\sum_{j \in \mathcal{N}_i} X_j, \gamma^2\big), \tag{3}$$

$$Y_i(1) = Y_i(0) + \tau, \tag{4}$$

where $N(m, v)$ denotes the normal distribution with mean m and variance v. The network induces correlation among the control potential outcomes because the mean of each $Y_i(0)$ is the sum of the covariate values, X_j, of units j in a neighbourhood of unit i. The effect of treatment is additive. Equations (2)–(4) define the normal-sum model. The implied model for the observed outcomes, Y^{obs}, is given in the Supplementary Material. We generalize this model in § 2·5 but will otherwise focus on the normal-sum model for clarity in presenting the restricted randomization approach.

The normal-sum model provides a useful abstraction for exploring the problem of optimal design of experiments in the presence of network-correlated outcomes. An illustration will help to anchor the intuition. The normal-sum model arises naturally, for example, when considering the time users spend on a social media platform. Consider the binary treatment Z_i to be the exposure to a new feature of the website designed to increase engagement and time spent online, and let $Y_i(Z_i)$ be the time spent online by user i when assigned to treatment Z_i. The causal effect of interest τ is then the effect of the new feature on the time spent online. Let us assume a constant, additive treatment effect for simplicity. In the absence of network connections and in the absence of treatment, X_i is the expected value of $Y_i(0)$ conditional on X_i. So X_i can be thought of as the intrinsic propensity of user i to spend time on the website. The model then captures the fact that the time spent on the website by user i increases with the number of his or her neighbours, with their propensities to spend time on the website, and with the exposure to the new feature if the treatment has an effect. The rest of the paper explores the implications of the normal-sum model for designing optimal treatment allocation strategies.

2·3. *Interpretation of the mean squared error for a fixed treatment allocation vector*

We compute the mean squared error of the difference-in-means estimator according to the normal-sum model for Y^{obs}, defined as $\text{MSE}(\hat{\tau} \mid Z) = E\{(\hat{\tau} - \tau^*)^2 \mid Z\}$ for a fixed treatment allocation vector Z; we call this quantity the conditional mean squared error. We have

$$\text{MSE}(\hat{\tau} \mid Z) = \underbrace{\mu^2\{\delta_{\mathcal{N}}(Z)\}^2}_{\text{bias}^2} + \underbrace{\gamma^2\omega(Z)^{\mathrm{T}}\omega(Z) + \sigma^2\omega(Z)^{\mathrm{T}}A^{\mathrm{T}}A\omega(Z)}_{\text{variance}}.$$

We can identify desirable assignments by evaluating their conditional mean squared error. This idea is the basis for the model-assisted restricted randomization strategies in § 3·2.

In the absence of specific constraints on the number of treated units, different treatment allocation vectors will generally have a different number of treated and untreated units, defined as $N_1 = \sum_i Z_i$ and $N_0 = \sum_i (1 - Z_i)$, respectively, both functions of Z. Then the bias term,

$$\mu\delta_{\mathcal{N}} = \mu\big(N_1^{-1} \textstyle\sum_{\{i:Z_i=1\}} |\mathcal{N}_i| - N_0^{-1} \sum_{\{i:Z_i=0\}} |\mathcal{N}_i|\big), \tag{5}$$

is proportional to the difference in the average neighbourhood sizes of treated and untreated units, and measures a lack of balance between the two groups in terms of the average degree. A larger

value of the mean μ amplifies the contribution of this imbalance to the mean squared error. Since the designer does not control μ, desirable treatment assignments minimize bias by balancing the average neighbourhood size between treated and untreated units. The first variance term is

$$\gamma^2 \omega^{\mathrm{T}} \omega = \gamma^2 \big(N_1^{-1} + N_0^{-1}\big), \tag{6}$$

which is minimized when $N_1 = N_0$. This term penalizes the difference between the number of treated and untreated units. A larger value of γ amplifies the contribution of this imbalance to the mean squared error. This result is consistent with classical results on the optimality of balanced randomization for estimating the average treatment effect in the absence of network-correlated outcomes (e.g., Imbens & Rubin, 2015, Ch. 6). The second variance term involves features of the network; it is

$$\sigma^2 \omega^{\mathrm{T}} A^{\mathrm{T}} A \omega = (\sigma^2/N_1^2) \sum_{\{i,j:\, Z_i = Z_j = 1\}} |\mathcal{N}_i \cap \mathcal{N}_j| \tag{7}$$

$$+ (\sigma^2/N_0^2) \sum_{\{i,j:\, Z_i = Z_j = 0\}} |\mathcal{N}_i \cap \mathcal{N}_j| \tag{8}$$

$$- \{2\sigma^2/(N_1 N_0)\} \sum_{\{i,j:\, Z_i = 1, Z_j = 0\}} |\mathcal{N}_i \cap \mathcal{N}_j|. \tag{9}$$

The term on the right-hand side of (7) is proportional to the average number of shared neighbours among pairs of units both assigned to the treatment group. The term (8) is proportional to the average number of shared neighbours among pairs of units both assigned to the control group. The term (9) is proportional to the average number of shared neighbours among pairs of units where one is assigned to treatment and the other to control. Considering the signs of these three factors, the second variance term may be minimized by assigning units with shared neighbours to different groups, and by avoiding the assignment of entire clusters of units that are densely connected to either treatment or control.

2·4. *Interpretation of the mean squared error averaged over allocation vectors*

Next, we compute the mean squared error of the difference-in-means estimator according to the normal-sum model and the distribution on the allocation vectors implied by a complete randomization strategy, which assigns equal probability to all of the treatment allocation vectors Z for which the numbers of units in treatment and control are fixed at (N_0, N_1). We refer to this quantity, $\mathrm{MSE}(\hat{\tau}) = E[E\{(\hat{\tau} - \tau^*)^2 \mid Z\}]$, as the marginal mean squared error. It is

$$\mathrm{MSE}(\hat{\tau}) = \big(N_1^{-1} + N_0^{-1}\big)\big(\gamma^2 + \sigma^2\big) \tag{10}$$

$$+ \big(N_1^{-1} + N_0^{-1}\big)\Bigg\{\underbrace{\sigma^2(|\bar{\mathcal{N}}| - 1)}_{C_1} - \underbrace{\frac{2\sigma^2}{N(N-1)} \sum_{i<j} |\mathcal{N}_i \cap \mathcal{N}_j|}_{C_2} + \underbrace{\frac{\mu^2 N}{N-1}(|\bar{\mathcal{N}}|^2 - |\bar{\mathcal{N}}|^{\langle 2 \rangle})}_{C_3}\Bigg\}.$$

The right-hand side of (10) is the mean squared error of the difference-in-means estimator due to a complete randomization strategy in the absence of a network, since $(\gamma^2 + \sigma^2)$ is the total variance implied by the network-sum model. The three terms C_1, C_2 and C_3 can be seen as contributions to the variance due to the presence of network-correlated outcomes. The term C_1

is proportional to the average degree of the nodes; thus networks with higher average degrees will tend to yield higher mean squared errors, ceteris paribus. The term C_2 is proportional to the average number of shared neighbours among all pairs of nodes; thus networks that are locally denser will tend to have lower mean squared error, ceteris paribus. The term C_3 is proportional to the variance of observed degrees; thus low variability in the degree of the nodes will lead to lower mean squared error, ceteris paribus. This contribution need not be positive, because of term C_2, which summarizes average local density.

2·5. *More general models of network-correlated potential outcomes*

The normal-sum model introduced in § 2·2 is a special case of a more general model that replaces (3) with the more general formulation

$$Y_i(0) \mid X \sim N\big[g\{(X_j)_{j\in\mathcal{N}_i}\}, \gamma^2\big],$$

with regularity conditions on the function g which essentially ensure that for any subset of nodes $\mathcal{S} \subset \mathcal{N}_i$, the conditional expectation $E[g\{(X_j)_{j\in\mathcal{N}_i}\} \mid (X_j)_{j\in\mathcal{S}}]$ is well behaved. We detail the positivity, symmetry and monotonicity properties as well as the general form of the mean squared error for this model in the Supplementary Material, and we show that the general form of the mean squared error suggests that good designs seek to decrease the number of neighbours shared within treatment groups and increase the number of units shared between treatment groups, while balancing the sizes of the groups and the distribution of neighbourhood sizes. These derivations indicate that the network balance criteria the proposed restricted randomizations are based upon extend well beyond the normal-sum model. Moreover, model-assisted strategies come with theoretical guarantees that hold regardless of the validity of the model, as we show next.

3. Methodology and theory

3·1. *Classical randomization and restricted randomization strategies*

Randomization strategies are probability distributions on the set of binary vectors $\mathcal{Z}$. Restricted randomization strategies are probability distributions implied by discarding allocation vectors $Z \in \mathcal{Z}$ according to a set of rules. According to a Bernoulli randomization strategy with parameter $p \in (0, 1)$, each treatment allocation vector $Z \in \mathcal{Z}$ has individual treatments Z_i drawn as independent Bernoulli random variables with probability of success p. A completely randomized design with parameters (N_0, N_1), where $N_0 + N_1 = N$, considers only treatment allocation vectors $Z \in \mathcal{Z}$ such that $\sum_{i=1}^{N} Z_i = N_1$, and assigns equal probability to them. If $N_0 = N_1 = N/2$, we refer to this as a balanced completely randomized design.

Restricted randomization strategies stem from the observation that when designing an experiment, it is often clear how to evaluate whether a treatment allocation vector is undesirable. For instance, when an allocation vector Z leads to statistical imbalance for one or more key covariates, it leaves the door open to confounding even in the presence of randomization (Gosset, 1938). Indeed, the most common form of restricted randomization is to discard treatment allocations that lead to covariate imbalances (Lock-Morgan & Rubin, 2012).

3·2. *Model-assisted restricted randomization strategies*

We introduce four model-assisted designs, which differ in the degree of reliance on the model. First, we consider balanced restricted randomization strategies, which discard treatment allocation vectors where the number of treated units N_1 differs from the number of untreated units N_0, or

differs by more than 1 when N is odd. This strategy aims at minimizing the contribution of the total variance to the conditional mean squared error, according to (6).

Second, we introduce unbiased restricted randomization strategies, which discard treatment allocation vectors where the average number of neighbours for treated units differs from the average number of neighbours for untreated units. This strategy aims at minimizing the contribution of the bias to the conditional mean squared error, as suggested by the discussion of (5).

Third, we introduce optimal restricted randomization strategies, which favour treatment allocation vectors that minimize the average number of shared neighbours among pairs of treated units, according to (7), minimize the average number of shared neighbours among pairs of untreated units, according to (8), and maximize the average number of shared neighbours among pairs of units one of which is treated and the other untreated, according to (9).

Let $\mathcal{Z} = \{0, 1\}^N$ be the set of all possible treatment allocation vectors on N units. Formally, we can define sets of allocations corresponding to the restricted randomizations defined above:

$$\mathcal{Z}^{\mathrm{b}} = \{Z \in \mathcal{Z} : N_1 - N_0 = 0\}, \tag{11}$$

$$\mathcal{Z}^{\mathrm{u}} = \{Z \in \mathcal{Z} : N_1^{-1} \textstyle\sum_{\{i:Z_i=1\}} |\mathcal{N}_i| - N_0^{-1} \sum_{\{i:Z_i=0\}} |\mathcal{N}_i| = 0\}, \tag{12}$$

$$\mathcal{Z}^{\mathrm{o}} = \{Z \in \mathcal{Z} : \mathrm{MSE}(\hat{\tau} \mid Z) \leqslant q_\alpha^{\mathrm{MSE}}\}, \tag{13}$$

where q_α^{MSE} is the α quantile of the distribution of the conditional mean squared error. These subsets of assignments depend on network statistics that the normal-sum model suggests as relevant for computing the conditional mean squared error, discussed in § 2·3.

The rest of the paper focuses on the first three model-assisted strategies: balanced restricted randomization, which assigns equal probability to all $Z \in \mathcal{Z}^{\mathrm{b}}$; balanced unbiased restricted randomization, which assigns equal probability to all $Z \in \mathcal{Z}^{\mathrm{b}} \cap \mathcal{Z}^{\mathrm{u}}$; and balanced unbiased optimal restricted randomization, which assigns equal probability to all $Z \in \mathcal{Z}^{\mathrm{b}} \cap \mathcal{Z}^{\mathrm{u}} \cap \mathcal{Z}^{\mathrm{o}}$.

The fourth model-assisted strategy, which we refer to as unconstrained optimal restricted randomization, aims to trade off small increases in bias for significant reductions in variance. It assigns equal probability to all $Z \in \mathcal{Z}^{\min}$, defined as

$$\mathcal{Z}^{\min} = \{Z \in \mathcal{Z} : \arg\min \mathrm{MSE}(\hat{\tau} \mid Z)\}. \tag{14}$$

The set $\mathcal{Z}^{\min}$ is usually either a singleton or a set of small cardinality. This makes it challenging to perform randomization-based inference using this design; in particular, the approach proposed in § 3·5 is often unfeasible in practice. This design is largely of theoretical interest.

3·3. *Model-based optimal treatment allocation strategies*

The model-assisted strategies in § 3·2 use a model for the outcomes to select allocations that improve properties of the difference-in-means estimator. The natural next step is to use the model to derive a better estimator for the average treatment effect, replacing for instance the difference-in-means estimator with the maximum likelihood estimator of τ under the normal-sum model. The estimator $\hat{\tau}_{\mathrm{MLE}}$ and its conditional mean squared error are derived in the Supplementary Material. The optimal maximum likelihood design is then the model-based restricted randomization strategy that assigns equal probability to all $Z \in \mathcal{Z}^{\mathrm{MLE}}$, defined as

$$\mathcal{Z}^{\mathrm{MLE}} = \{Z \in \mathcal{Z} : \arg\min \mathrm{MSE}(\hat{\tau}_{\mathrm{MLE}} \mid Z)\}. \tag{15}$$

In the Supplementary Material, we show that the maximum likelihood estimator for τ is not robust with respect to misspecification of the model or the network, unlike the model-assisted

restricted randomization designs. When evaluating the performance of model-based strategies, we fix parameters μ, σ and γ at their true values and treat τ as the only unknown parameter.

3·4. *Restricted randomizations via rerandomization*

A general approach to sampling from arbitrary restricted randomization designs, referred to as rerandomization, has recently been formalized by Lock-Morgan & Rubin (2012). Let ϕ be a binary function such that assignment Z belongs to the restricted randomization set if and only if $\phi(Z) = 1$. A simple way to sample from the restricted randomization design is via rejection sampling: draw an assignment Z from the original design, and then keep the assignment if $\phi(Z) = 1$, or reject it if $\phi(Z) = 0$. In our setting, the restricted sets in (11)–(13) can be defined in terms of different functions ϕ. Denote the indicator function by $I(\cdot)$; then

$$\phi^{\mathrm{b}}(Z) = I\Big\{\textstyle\sum_{i=1}^{N} Z_i = \sum_{i=1}^{N}(1 - Z_i)\Big\},$$
$$\phi^{\mathrm{u}}(Z) = I\big\{\mu\delta_{\mathcal{N}}(Z) = 0\big\},$$
$$\phi^{\mathrm{o}}(Z) = I\Big\{\mathrm{MSE}(\hat{\tau} \mid Z) \leqslant q_{\alpha}^{\mathrm{MSE}}\Big\}.$$

Thus rerandomization can be used to sample from the restricted randomization designs we proposed. It is particularly useful when performing exact tests and computing confidence intervals, as we show next.

3·5. *Inference via inversion of a sequence of exact Fisher tests*

There are traditionally three types of confidence interval in randomization-based inference: Neyman intervals, bootstrap intervals, and Fisher intervals. Neyman intervals are usually obtained using an asymptotic normal approximation to the distribution of the difference-in-means estimator (Imbens & Rubin, 2015, Ch. 6). This approach works well for simple designs for which the asymptotic variance can be estimated, but is challenging with more complicated designs. Li et al. (2017) proposed an asymptotic theory of rerandomization. Unfortunately, the asymptotic regime considered in that paper is not compatible with our setting: it requires the number of covariates to be fixed in the asymptotic regime, whereas in our case the quantities that are analogous to covariates include the number of neighbours shared by each pair of units, which grows with the number of units in the asymptotic regimes of interest; it also requires the constraints to be a function only of the vector of differences in means between treated and control units for the observed covariates, and of the covariance matrix of that vector, which does not hold in our case. Bootstrap intervals are difficult to implement since the correlation structure of the outcomes may be complex.

Instead, we propose using Fisher intervals, which are obtained by inverting a sequence of Fisher exact tests (e.g., Rosenbaum, 2002). This can be accomplished by rerandomization (Lock-Morgan & Rubin, 2012, § 2.2) but with the proposed restricted randomization distributions as the permutation distributions.

We illustrate by simulation the potential gains from Fisher intervals based on restricted randomization. For a fixed network of 500 nodes, we generated 200 realizations of the potential outcomes according to the normal-sum model, as well as 200 observed assignments. For each realization, we computed Fisher confidence intervals based on balanced optimal restricted randomization, balanced unbiased restricted randomization, and balanced complete randomization, with a nominal test size of $\alpha = 5\%$. The intervals based on balanced optimal restricted randomization have a median length of 5·5 with 90% interquantile range [5·1, 5·7]; those based on balanced unbiased restricted randomization have a median length of 5·7 with 90% interquantile

range [5·4, 5·9]. In contrast, the intervals based on balanced complete randomization have median length 6·5 with 90% interquantile range [6·0, 6·6]. The coverage for all three methods is 95%, as expected. These results suggest that restricted randomization inference reduces the length of the intervals while maintaining nominal coverage. More details are given in the Supplementary Material.

3·6. *Theory*

Model-assisted designs have desirable inferential properties even when the model they rely on for evaluating treatment allocations is wrong. We show that the difference-in-means estimator is design-unbiased (Särndal et al., 2003) for the restricted randomization strategies developed in § 3·2.

Definition 1 (Design unbiasedness). *An estimator $\hat{\tau}$ is unbiased with respect to a distribution on $\mathcal{Z}$, typically referred to as a design on $\mathcal{Z}$, if $E_{\mathcal{Z}}(\hat{\tau} - \tau) = 0$.*

The main result is the following.

Theorem 1. *The difference-in-means estimator $\hat{\tau}$ defined in* (1) *is an unbiased estimator of the average treatment effect with respect to the following distributions:*

(i) *the uniform distribution on $\mathcal{Z}^{\mathrm{b}}$, which defines the balanced design;*
(ii) *the uniform distribution on $\mathcal{Z}^{\mathrm{b}} \cap \mathcal{Z}^{\mathrm{u}}$, which defines the balanced unbiased design;*
(iii) *the uniform distribution on $\mathcal{Z}^{\mathrm{b}} \cap \mathcal{Z}^{\mathrm{o}}$, which defines the balanced optimal design;*
(iv) *the uniform distribution on $\mathcal{Z}^{\mathrm{b}} \cap \mathcal{Z}^{\mathrm{u}} \cap \mathcal{Z}^{\mathrm{o}}$, which defines the balanced unbiased optimal design.*

As a consequence of design-unbiasedness and of the increasingly nested supports, we can compare variances of $\hat{\tau}$ implied by the designs in Theorem 1, in expectation.

Corollary 1. *The estimator $\hat{\tau}$ defined in* (1) *satisfies*

$$E\{\mathrm{var}_{\mathcal{Z}^{\mathrm{b}} \cap \mathcal{Z}^{\mathrm{o}}}(\hat{\tau} \mid Y)\} \leqslant E\{\mathrm{var}_{\mathcal{Z}^{\mathrm{b}}}(\hat{\tau} \mid Y)\}.$$

Similar inequalities can be derived for any pair of nested designs in Theorem 1. These results are based on symmetry arguments, which is why $\mathcal{Z}^{\mathrm{b}}$ is always part of the support of designs that make the difference-in-means estimator unbiased. This notion of symmetry is made precise in the following lemma.

Lemma 1. *For Z in $\mathcal{Z}^{\mathrm{b}}$, $\hat{\tau}(1 - Z) = 2\tau - \hat{\tau}(Z)$.*

As a consequence, if we required the unconstrained optimal design to be balanced by restricting its support to $\mathcal{Z}^{\mathrm{b}} \cap \mathcal{Z}^{\mathrm{min}}$, we would recover design-unbiasedness for the difference-in-means estimator. However, we do not consider balanced unconstrained optimal designs.

4. Discussion

The idea behind model-assisted design is fairly general, two key elements being the estimator and the model. The theoretical guarantees in § 3·6 are limited to estimators satisfying the symmetry condition of Lemma 1, and to the model family introduced in § 2·5. Extending the theory to a larger class of estimators and models is conceptually feasible, although it would often lead to complex expressions for the mean squared error and hard-to-interpret balance criteria.

The designs presented in (14) and (15) can be seen as extreme versions of the model-assisted approach, perhaps closer in spirit to the optimal model-based design literature (Kiefer, 1959). However, the mean squared errors minimized in the model-based and model-assisted design of experiments are associated with different estimators. Randomization inference based on the restricted distributions that the designs in (14) and (15) imply is often impractical, since the sets $\mathcal{Z}^{\text{MLE}}$ and $\mathcal{Z}^{\min}$ are generally too small. Moreover, even when feasible, inference based on these designs relies heavily on model-induced constraints, by requiring stringent balance of terms appearing in the conditional mean squared error, and is in general not robust with respect to model misspecification.

In practice, there are often additional issues to consider, which we have ignored for simplicity of exposition. Covariates will have to be taken into account, and the parameters μ, σ^2 and γ will need to be specified or estimated. One option is to specify point priors (Box & Lucas, 1959); another option is to specify full priors and work with the integrated mean squared error. In both situations, historical data and pilot studies could be used to calibrate these priors, and are recommended for optimal design in practice (Kim et al., 2015). Our theory for a more general model of network-correlated outcomes, as well as simulation studies, both detailed in the Supplementary Material, show that the efficiency gains one can expect to achieve with model-assisted design of experiments are robust with respect to misspecification.

This paper has introduced model-assisted design of network experiments and illustrated its use in a simple setting. Although the model in § 2·5 is fairly general, it focuses on estimating the average treatment effect in the presence of network-correlated outcomes and homophily, ignoring other phenomena of interest, including network interference, peer influence and contagion. Developing model-assisted design strategies for estimating other causal effects in more complicated settings, such as the presence of network interference and confounding due to homophily, is one of the directions we are currently pursuing.

Acknowledgement

The authors thank Iavor Bojinov for helpful comments. We are grateful to the editor, associate editor, and two reviewers for constructive feedback. This work was partially supported by the U.S. National Science Foundation, the Office of Naval Research, a Shutzer Fellowship and a Sloan Research Fellowship. The work was performed while the authors were at Harvard University.

Supplementary material

Supplementary material available at *Biometrika* online includes proofs of all the theoretical results, simulations illustrating the robustness of the proposed strategies with respect to misspecification, and additional technical details.

References

Aral, S., Muchnik, L. & Sundararajan, A. (2009). Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proc. Nat. Acad. Sci.* **106**, 21544–9.

Aronow, P. M. & Samii, C. (2017). Estimating average causal effects under general interference, with application to a social network experiment. *Ann. Appl. Statist.* **11**, 1912–47.

Bailey, R. A. (1983). Restricted randomization. *Biometrika* **70**, 183–98.

Bakshy, E., Hofman, J. M., Mason, W. A. & Watts, D. J. (2011). Everyone's an influencer: Quantifying influence on Twitter. In *Proc. 4th ACM Int. Conf. Web Search and Data Mining*. New York: Association for Computing Machinery, pp. 65–74.

Bakshy, E., Rosenn, I., Marlow, C. & Adamic, L. (2012). The role of social networks in information diffusion. In *Proc. 21st Int. Conf. World Wide Web*. New York: Association for Computing Machinery, pp. 519–28.
Bickel, P. J. & Chen, A. (2009). A nonparametric view of network models and Newman–Girvan and other modularities. *Proc. Nat. Acad. Sci.* **106**, 21068–73.
Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D., Marlow, C., Settle, J. E. & Fowler, J. H. (2012). A 61-million-person experiment in social influence and political mobilization. *Nature* **489**, 295–8.
Box, G. E. P. & Lucas, H. (1959). Design of experiments in non-linear situations. *Biometrika* **46**, 77–90.
Cavusoglu, H., Phan, T. Q., Cavusoglu, H. & Airoldi, E. M. (2016). Assessing the impact of granular privacy controls on content sharing and disclosure on Facebook. *Info. Syst. Res.* **27**, 848–79.
Christakis, N. A. & Fowler, J. H. (2007). The spread of obesity in a large social network over 32 years. *New Engl. J. Med.* **357**, 370–9.
Eckles, D., Karrer, B. & Ugander, J. (2017). Design and analysis of experiments in networks: Reducing bias from interference. *J. Causal Infer.* **5**, DOI: 10.1515/jci-2015-0021.
Goldenberg, A., Zheng, A. X., Fienberg, S. E. & Airoldi, E. M. (2010). A survey of statistical network models. *Foundat. Trends Mach. Learn.* **2**, 129–233.
Gosset, W. J. (1938). Comparison between balanced and random arrangements of field plots. *Biometrika* **29**, 363–78.
Gui, H., Xu, Y., Bhasin, A. & Han, J. (2015). Network A/B testing: From sampling to estimation. In *Proc. 24th Int. Conf. World Wide Web*. New York: Association for Computing Machinery, 399–409.
Higham, R., Tharmanathan, P. & Birks, Y. (2015). Use and reporting of restricted randomization: A review. *J. Eval. Clin. Practice* **21**, 1205–11.
Holland, P. W. (1986). Statistics and causal inference. *J. Am. Statist. Assoc.* **81**, 945–60.
Hudgens, M. G. & Halloran, M. E. (2008). Toward causal inference with interference. *J. Am. Statist. Assoc.* **103**, 832–42.
Imbens, G. W. & Rubin, D. B. (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge: Cambridge University Press.
Kiefer, J. (1959). Optimum experimental designs. *J. R. Statist. Soc.* B **21**, 272–319.
Kim, D. A., Hwong, A. R., Stafford, D., Hughes, D. A., O'Malley, A. J., Fowler, J. H. & Christakis, N. A. (2015). Social network targeting to maximise population behaviour change: A cluster randomised controlled trial. *Lancet* **386**, 145–53.
Kolaczyk, E. D. & Csárdi, G. (2014). *Statistical Analysis of Network Data with R*. New York: Springer.
Li, X., Ding, P. & Rubin, D. B. (2017). Asymptotic theory of rerandomization in treatment-control experiments. *arXiv:* 1604.00698v4.
Lock-Morgan, K. & Rubin, D. B. (2012). Rerandomization to improve covariate balance in experiments. *Ann. Statist.* **40**, 1263–82.
Manski, C. F. (2013). Identification of treatment response with social interactions. *Economet. J.* **16**, S1–23.
McPherson, M., Smith-Lovin, L. & Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Ann. Rev. Sociol.* **27**, 415–44.
Ogburn, E. L. & VanderWeele, T. J. (2017). Vaccines, contagion, and social networks. *Ann. Appl. Statist.* **11**, 919–48.
Phan, T. Q. & Airoldi, E. M. (2015). A natural experiment of social network formation and dynamics. *Proc. Nat. Acad. Sci.* **112**, 6595–600.
Rosenbaum, P. R. (2002). Covariance adjustment in randomized experiments and observational studies. *Statist. Sci.* **17**, 286–327.
Rosenbaum, P. R. (2007). Interference between units in randomized experiments. *J. Am. Statist. Assoc.* **102**, 191–200.
Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *J. Educ. Psychol.* **66**, 688–701.
Särndal, C.-E., Swensson, B. & Wretman, J. (2003). *Model Assisted Survey Sampling*. New York: Springer.
Shalizi, C. R. & Thomas, A. C. (2011). Homophily and contagion are generically confounded in observational social network studies. *Sociol. Meth. Res.* **40**, 211–39.
Simon, R. (1979). Restricted randomization designs in clinical trials. *Biometrics* **35**, 503–12.
Toulis, P. & Kao, E. (2013). Estimation of causal peer influence effects. *J. Mach. Learn. Res.* **28**, 1489–97.
Ugander, J., Karrer, B., Backstrom, L. & Kleinberg, J. (2013). Graph cluster randomization: Network exposure to multiple universes. In *Proc. 19th ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*. New York: Association for Computing Machinery, pp. 329–37.
Wasserman, S. & Faust, K. (1994). *Social Network Analysis: Methods and Applications*. Cambridge: Cambridge University Press.
Yates, F. (1948). Contribution to the discussion of "The validity of comparative experiments" by F. J. Anscombe. *J. R. Statist. Soc.* A **111**, 204–5.
Youden, W. J. (1972). Randomization and experimentation. *Technometrics* **14**, 13–22.

[*Received on* 2 *July* 2015. *Editorial decision on* 20 *August* 2017]