

Designing Experiments to Measure Spillover Effects*

Sarah Baird[†], J. Aislinn Bohren[‡], Craig McIntosh[§], Berk Özler[¶]

February 7, 2014

Abstract

This paper formalizes the design of experiments intended specifically to study spillover effects. By first randomizing the intensity of treatment within clusters and then randomly assigning individual treatment conditional on this cluster-level intensity, a novel set of treatment effects can be identified. We develop a formal framework for consistent estimation of these effects, and provide explicit expressions for power calculations. We show that the power to detect average treatment effects declines precisely with the quantity that identifies the novel treatment effects. A demonstration of the technique is provided using a cash transfer program in Malawi.

KEYWORDS: Experimental Design, Networks, Cash Transfers

JEL: C93, O22, I25

*We are grateful for the useful comments received from seminar participants at Caltech, Monash, Namur, Paris School of Economics, Stanford, University of British Columbia, UC Berkeley, University of Colorado, University of Melbourne, and Yale.

[†]George Washington University and University of Otago, sarah.baird@otago.ac.nz

[‡]University of Pennsylvania, abohren@sas.upenn.edu

[§]University of California, San Diego, ctmcintosh@ucsd.edu

[¶]World Bank and University of Otago, berk.ozler@otago.ac.nz

Designing Experiments to Measure Spillover Effects

Multiple economic disciplines have begun to explore the empirical issues raised by spillover effects from one individual to another. What [Charles Manski \(1993\)](#) refers to as endogenous effects are explored in different ways – by empirical studies permitting general equilibrium effects, the analysis of medical treatments that provide herd immunity, or by studies of network effects. An increasingly useful lens on this problem is experimental policy trials that explicitly consider interference between individuals. Once we permit interference in this context, the impact of a program only on its beneficiaries becomes an unsatisfying answer to the real policy impact. Thus, it becomes more important to understand spillovers and the overall effect on the entire population. What if a program creates benefits to some only by diverting them from others? How do individuals respond to the intensity of treatment within a population? Does the study even have an unpolluted counterfactual?

The possibility of interference between individuals has traditionally been seen as the Achilles heel of randomized experiments; standard experimental designs are unable to identify and measure spillovers.¹ Given these concerns, a new wave of empirical work has emerged in the past decade trying to relax the strong assumption of no interference, or that individuals are not affected by the treatment status of others. This literature includes studies that uncover network effects using experimental variation across treatment groups ([Matteo Bobba and Jeremie Gignoux 2013](#); [Edward Miguel and Michael Kremer 2004](#)), leave some members of a group untreated ([Manuela Angelucci and Giacomo De Giorgi 2009](#); [Felipe Barrera-Osorio, Marianne Bertrand, Leigh Linden and Francisco Perez-Calle 2011](#); [Gustavo J. Bobonis and Frederico Finan 2009](#); [Esther Duflo and Emmanuel Saez 2003](#); [Rafael Lalive and M. A. Cattaneo 2009](#)), exploit plausibly exogenous variation in within-network treatments ([Philip S. Babcock and John L. Hartman 2010](#); [Lori A. Beaman 2012](#); [Timothy G. Conley and Christopher R. Udry 2010](#); [Esther Duflo and Emmanuel Saez 2002](#); [Kaivan Munshi 2003](#)), or intersect an experiment with pre-existing networks ([Abhijit Banerjee, Arun G. Chandrasekhar, Esther Duflo and Matthew O. Jackson 2013](#); [Jiehua Chen, Macartan Humphries and Vijay Modi 2010](#); [Karen Macours and Renos Vakis 2008](#); [Emily Oster and Rebecca Thornton 2012](#)).

A *partial population* experiment ([Robert A. Moffitt 2001](#)), in which some clusters are assigned to control and a subset of individuals are offered treatment within clusters assigned to treatment, partially overcomes this challenge and yields valid estimates of treatment and spillover effects. But such experiments provide no exogenous variation in treatment

¹In the presence of spillovers, the blocked design produces biased estimates. The clustered design is not biased, but provides no information to estimate the extent of spillovers.

saturation to estimate the extent to which program effects are driven by the percentage of individuals treated in treatment clusters.² Consequently, the most recent empirical approach has been to conduct a two-level randomization in which the share of individuals assigned to treatment within treated clusters is directly varied.³

In this paper, we provide a formal presentation of these *randomized saturation* (RS) designs. We define the relevant set of treatment effects a researcher should consider in the presence of spillovers, and present a clear set of assumptions under which the RS design can consistently measure these effects. In a RS design, each cluster is randomly assigned a treatment saturation, and each individual within the cluster is randomly assigned a treatment status, given the assigned cluster saturation. This design allows for the consistent estimation of a rich set of treatment and spillover effects across the distribution of treatment saturations, which include the intention to treat effect (ITT), spillovers on the non-treated (SNT), and the total causal effect (TCE) as well as novel estimands such as the treatment effect on the uniquely treated (TUT) and the spillover effect on the treated (ST). In the process, a RS design also allows the researcher to discover the extent to which observed correlations in outcomes within clusters were caused by endogenous effects, thereby offering a solution to the reflection problem, albeit after the fact, and informing the design of future studies.

Next, we develop a technical framework to guide researchers through the various choices in RS designs. All (non-trivial) RS designs yield consistent estimates of treatment and spillover effects, but the power of each design varies with the saturation profile and the share of clusters assigned to each saturation. We impose a random effects variance structure and derive explicit expressions for the minimum detectable treatment and slope effects, which determine the statistical power of different RS designs. Power is a function of some standard quantities, such as the effect size and the intra-cluster correlation (ICC) of outcomes, as well as of some unique features of the RS design such as the share of individuals assigned to each treatment saturation and the variance of saturations. Power for the average ITT and SNT is decreasing in precisely the quantity that identifies the novel effects, namely, variation in

²Most extant partial population experiments feature cluster-level saturations that are either endogenous (Oportunidades) or fixed (Duflo and Saez (2003), where they are typically set at 50%). PROGRESA/Oportunidades (Mexico) is perhaps the most-studied example of a partial population experiment. This program features a treatment decision at the cluster (village) level and an objective poverty eligibility threshold at the household level, so both eligible and ineligible individuals in treatment villages can be compared to their counterparts in the pure control group. PROGRESA has been used to examine spillover effects in several contexts (Jennifer Alix-Garcia, Craig McIntosh, Katharine R. E. Sims and Jarrod R. Welch 2013; Angelucci and De Giorgi 2009; Bobonis and Finan 2009). Other partial population experiments include Duflo and Saez (2003) and Peter Kuhn, Peter Kooreman, Adriaan Soetevent and Arie Kapteyn (2011).

³Abhijit Banerjee, Raghavendra Chattopadhyay, Esther Duflo, Daniel Keniston and Nina Singh (2012); Bruno Crepon, Esther Duflo, Marc Gurgand, Roland Rathelot and Philippe Zamora (2013); Xavier Gine and Ghazala Mansuri (2012); Betsy Sinclair, Margaret McConnell and Donald P. Green (2012).

the intensity of treatment across clusters. RS designs therefore generate a tradeoff compared with the standard blocked, clustered, and partial-population designs: while they allow the researcher to identify novel effects on both the treated and the non-treated, this comes at the cost of reduced power to detect average effects. With input on the magnitude of ICCs and the relative importance of the various estimands, we use the power calculations to provide insight on design choices such as the optimal degree of variation in saturations and the size of the pure control. ⁴

We conclude the theoretical presentation with three additional uses of the RS design. First, we show that one can recover an estimate of the treatment on the compliers effect (TOC) by assuming that the observed spillover effects on those not offered treatment are a reasonable proxy for the spillovers experienced by non-compliers. This technique is critical because interference between units within clusters violates the exclusion restriction in the standard technique of instrumenting for treatment with randomized assignment to identify the TOC. Second, we consider experiments that use within-cluster controls to form the counterfactual. Imposing a functional form assumption on the saturation distribution allows the researcher to project the desired counterfactual outcome: untreated clusters with a saturation of zero. This value can be used to correct the naive estimate of the ITT, even in studies without a pure control. Finally, we show that an RS design implemented on a non-overlapping network also produces exogenous variation in the treatment saturation of overlapping networks (for example, social groups), variation that is generally superior to what would be obtained from *blocked* or *clustered* designs.

We close with an empirical application of these techniques using a cash transfer experiment in Malawi, wherein the fraction of eligible school-aged girls offered treatment was randomized across clusters. The study seeks to understand whether cash transfers could help adolescent girls improve schooling outcomes as well as delay marriage and pregnancy. In previous work, we have shown that, compared with a pure control group, conditional cash transfers (CCTs) significantly improved schooling outcomes while unconditional cash transfers (UCTs) caused substantial reductions in marriage and fertility rates among program beneficiaries (Sarah Baird, Craig McIntosh and Berk Özler 2011). In this paper, we exploit the sample of within-cluster controls and the RS design to investigate spillover effects on both program beneficiaries and eligible non-beneficiaries. Spillovers are a central concern for two distinct reasons. First, a large literature indicates that schooling cash transfer programs can alter the welfare of non-beneficiaries due to congestion effects in the classroom (Jere R Behrman, Piyali Sengupta and Petra Todd 2005), shifts in local norms around education

⁴In the Supplemental Appendix, we provide a Matlab program that allows a researcher to calculate the power of different potential RS designs.

(George A. Akerlof and Rachel E. Kranton 2002), income spillovers (Manuela Angelucci, Giacomo De Giorgi, Marcos A. Rangel and Imran Rasul 2010), or general equilibrium changes to prices (Jesse M. Cunha, Giacomo De Giorgi and Seema Jayachandran 2011) and production (Alix-Garcia et al. 2013). Second, and more specific to the Malawian context, cash transfers can decrease young women’s dependence on men for financial assistance (Winford Masanjala 2007) and/or the need for ‘transactional sex’ (Michelle J. Poulin 2007; Ann Swindler and Susan Watkins 2007), thereby reducing the incidence of teen pregnancies and early marriages among program beneficiaries, but with ambiguous spillovers to non-beneficiaries in the same communities.

We find that while average spillover effects are muted for all outcomes, they generally intensify with treatment saturation: positive treatment effects on beneficiaries are accompanied by positive spillovers on non-beneficiaries, which increase with treatment intensity. On the other hand, treatment effects among beneficiaries themselves decline with treatment saturation. More importantly, we find no evidence for higher rates of marriage or pregnancy among within-cluster controls, suggesting that diversionary effects do not counter the documented beneficial effects of UCTs on these outcomes. Finally, taking advantage of exogenous variation generated by the RS design in the number of treated friends of individuals, we confirm that spillover effects are similarly muted in social networks.

The remainder of the paper is structured as follows. Section 1 formally models a randomized saturation design, outlines the assumptions required to use this design, defines novel estimands related to spillovers, presents closed-form expressions for the power of these estimands, and discusses the critical design tradeoffs. Sections 2.1 and 2.2 discuss the use of randomized saturation designs in the absence of a pure control group, while Section 2.3 demonstrates the use of randomized saturation designs in a broader class of networks. Section 3 presents an application of the technique and Section 4 concludes. All proofs are in the Appendix.

1 A Randomized Saturation Design

One of the most basic design choices in any multi-level experiment is the question of allocating treatment to N individuals distributed across C clusters. The conventional wisdom focuses on the ‘design effect’, whereby a positive correlation between the outcomes of individuals in the same cluster, i.e. intra-cluster correlation (ICC), causes a power loss if treatment is assigned at the cluster level. It would be easy to conclude that a *blocked* design, in which half of individuals in each cluster is treated and the other half is used as the counterfactual, is preferable. Critically, however, individuals in the same cluster may behave similarly because

they are influenced by the behavior of others in the group (endogenous effects), their behavior reflects the exogenous characteristics of the group (contextual effects), or because they share similar characteristics or face similar institutional environments (correlated effects) (Manski 1993). The entire thrust of the ‘reflection problem’ introduced by Manski (1993) is the impossibility of separating these effects using observational data that is typically available to the researcher at baseline. If only contextual or correlated effects are responsible for the observed ICC, indeed the blocked design proves optimal. However, if endogenous effects are present, then a blocked design is the wrong choice because the counterfactual is contaminated by interference from treated individuals. A *clustered* design, in which some clusters are assigned to treatment while others to control, would produce unbiased treatment effects if there is no interference across clusters, but with the loss of statistical power arising from cross-cluster identification. Thus this most basic of design choices ends up on the horns of the reflection problem: because neither the blocked nor the clustered design actually reveals the extent of interference, researchers learn little from a given study as to the optimal design of subsequent studies. The RS design provides a solution to this conundrum.

A *randomized saturation* (RS) design is an experiment with two stages of randomization. Take as given a set of N individuals divided into C non-overlapping groups, or clusters.⁵ The first stage randomizes the treatment saturation of each cluster, and the second stage randomizes the treatment status of each individual in the cluster, according to the realized saturation of the cluster. Formally, in the first stage, each cluster $c = 1, \dots, C$ is assigned a treatment saturation $\pi_c \in \Pi \subset [0, 1]$ according to the distribution F , with mean $\mu = E[\pi]$ and variance $\eta^2 = Var(\pi)$. In the second stage, each individual $i = 1, \dots, n$ in cluster c is assigned a treatment status $T_{ic} \in \{0, 1\}$, where $T_{ic} = 1$ represents a treated individual.⁶ The realized treatment saturation of stage 1 specifies the distribution of the treatment status in stage 2 for each cluster, $P(T_{ic} = 1 | \pi_c = \pi) = \pi$. Let f be the probability mass function for distribution F .⁷ A RS design ω is completely characterized by the pair $\{\Pi, f\}$.

The saturation $\pi_c = 0$ represents a cluster with no treatment individuals, or a pure control cluster. A within-cluster control is defined as an untreated individual in a cluster with treated individuals: $S_{ic} = \mathbb{1}\{T_{ic} = 0, \pi_c > 0\}$. This results in the following distribution

⁵The RS design and the studies discussed here use a simple, spatially defined definition of ‘cluster’ that is mutually exclusive and exhaustive. This is distinct from the issue of randomizing saturations with overlapping social networks (Peter Aronow 2012), which typically require a more complex sequential randomization routine (Panos Toulis and Edward Kao 2013). However, an additional advantage of this design is that it will also create exogenous variation in the saturation of any network that is correlated with given cluster, even if this other network is overlapping. This is discussed in more depth in Section 2.3.

⁶This notation implicitly assumes each cluster is of equal size. This is for notational convenience; the results easily extend to unequally sized clusters.

⁷For expositional simplicity, we present the theoretical results in a discrete saturation support framework, although the analysis easily generalizes to continuous or mixed distributions.

over the three possible treatment statuses:

$$\begin{aligned}
\mathbf{Treatment\ Individual:} & \quad P(T_{ic} = 1) = \mu \\
\mathbf{Pure\ Control:} & \quad P(S_{ic} = 0, T_{ic} = 0) = \psi \\
\mathbf{Within-cluster\ Control:} & \quad P(S_{ic} = 1) = 1 - \mu - \psi := \mu_S
\end{aligned}$$

where $\psi := f(0)$. We say a randomized saturation design has a *pure control* if $\psi > 0$.

A RS design introduces correlation between the treatment status of two individuals in the same cluster. This correlation is proportional to the variance of the cluster level treatment saturations, $\rho_T = \eta^2 / (\mu(1 - \mu))$, where η^2 can be split into the variance in treatment saturation across treated clusters, $\eta_T^2 = Var(\pi | \pi > 0)$, and the variance from pure control clusters:

$$\eta^2 = (1 - \psi) \eta_T^2 + \left(\frac{\psi}{1 - \psi} \right) \mu^2$$

Section 1.3.1 shows that in the presence of intra-cluster correlation (ICC), η^2 affects the power of the design.

The RS design nests several common experimental designs, including the clustered, blocked and partial population designs.⁸ The blocked design is biased in the presence of spillovers, and it is not possible to measure spillovers with either design. Therefore, we must put some restrictions on the RS design in order to be able to identify treatment and spillover effects. We say a RS design is *non-trivial* if it has at least two saturations, at least one of which is strictly interior.

Definition 1. *A randomized saturation design is **non-trivial** if $|\Pi| \geq 2$ and $\exists \pi \in \Pi$ such that $\pi \in (0, 1)$.*

Multiple saturations guarantee a comparison group to determine whether effects vary with treatment saturation, and an interior saturation guarantees the existence of within-cluster controls to identify spillovers on the untreated ($\mu_S > 0$). Note that the blocked and clustered designs are trivial, while the partial population design is non-trivial.

Remark 1. *Before turning to our formal framework, it is important to clarify the population in which the researcher is measuring spillovers. The RS design defines the treatment saturation of a cluster as the share of the study sample that is offered treatment. If spillovers*

⁸Fixing the probability of treatment at P , the clustered design corresponds to $\Pi = \{0, 1\}$ and $f(1) = P$, the blocked design corresponds to $\Pi = \{P\}$ and $f(P) = 1$ and the partial population design corresponds to $\Pi = \{0, \pi\}$ and $f(\pi) = P/\pi$. In the clustered design, there is perfect correlation between the treatment status of two individuals in the same cluster and in the blocked design, there is no correlation. Note $\eta_T^2 = 0$ for all three.

occur within the study sample, then this is the appropriate saturation measure.⁹ Alternatively, if there is a ‘gateway to treatment’ and not all eligible individuals are sampled into the study, or spillovers occur on a larger population within the cluster, then it is necessary to distinguish between the true treatment saturation (the share of treated individuals in the spillover network) and the assigned treatment saturation (the share of treated individuals in the study population).¹⁰ If sampling rates or the share of the spillover population eligible for treatment are constant across clusters, the true saturation is the sampling rate times the assigned saturation. If the sampling rates are driven by cluster characteristics, then the true saturation is endogenous. In this case, the researcher can instrument for the true saturation with the assigned saturation. To streamline the remainder of the theoretical analysis, we assume that the assigned and true saturations coincide.

1.1 Defining Treatment and Spillover Effects

Let Y_{ic} represent the outcome for individual i in cluster c . In a general framework, outcomes can depend in an arbitrary way on an individual’s own treatment status, as well as the treatment status of all other individuals in the study:

$$Y_{ic} = g(T_{ic}, R_{ic}, \{T_{jd}, R_{jd}\}_{jd \neq ic}; X_{ic}, \varepsilon_{ic})$$

where $R_{ic} \in \{0, 1\}$ indicates whether an individual complies with treatment, X_{ic} is a vector of covariates and ε_{ic} is an error term.¹¹

To use the RS design for causal inference requires an assumption on how the treatment status of others impacts Y_{ic} . We relax the stable unit treatment value assumption (SUTVA) within clusters, but maintain it across clusters: spillovers may flow within a cluster, but do not flow between clusters. This ensures that pure control clusters provide a valid counterfactual for treated clusters and that cross cluster comparisons can identify how spillovers depend on the intensity of treatment saturation.

⁹For example, [Banerjee et al. \(2012\)](#) study interventions to improve performance among constables in Rajasthan police stations. [Sinclair, McConnell and Green \(2012\)](#) study sending social-pressure mailings to registered voters in a congressional district.

¹⁰For example, [Gine and Mansuri \(2012\)](#) sample every fourth household in a neighborhood, and randomly offer treatment to 80 percent of these households. This causes the true treatment saturation to be 20 percent rather than the assigned 80 percent. Other examples include unemployed individuals on official unemployment registries form a small portion all unemployed individuals in an administrative region ([Crepon et al. 2013](#)); neighborhoods eligible for infrastructure investments comprise only 3 percent of all neighborhoods ([Craig McIntosh, Tito Alegria, Gerardo Ordonez and Rene Zenteno 2013](#)); and malaria prevention efforts target vulnerable individuals, who account for a small share of total cluster population ([GF Killeen, TA Smith, HM Ferguson, H Mshinda, S Abdulla et al. 2007](#)).

¹¹As is standard, R_{ic} is only observed for individuals with $T_{ic} = 1$.

Assumption 1. *There is no cross-cluster interference in outcomes: Y_{ic} is independent of $\{T_{jd}, R_{jd}\}$ for all $d \neq c$.*

Assumption 1 simplifies the framework so that outcomes only depend on the treatment of other individuals in the same cluster,

$$Y_{ic} = g(T_{ic}, R_{ic}, \{T_{jc}, R_{jc}\}_{j \neq i}; X_{ic}, \varepsilon_{ic}).$$

Given Assumption 1, we can formally define several treatment and spillover effect measures, both at specific saturations and pooled across multiple saturations. The **Intention to Treat** (ITT) effect is the difference between the expected outcome for individuals offered treatment in a cluster with saturation π and the expected outcome for pure control individuals,

$$ITT(\pi) := E(Y_{ic} | T_{ic} = 1, \pi_c = \pi) - E(Y_{ic} | T_{ic} = 0, \pi_c = 0).$$

The corresponding term for the **Spillover on the Non-Treated** (SNT) effect is the difference between the expected outcome for individuals not offered treatment in a cluster with saturation π and the expected outcome for pure control individuals,

$$SNT(\pi) := E(Y_{ic} | T_{ic} = 0, \pi_c = \pi) - E(Y_{ic} | T_{ic} = 0, \pi_c = 0).$$

The **Total Causal Effect** (TCE) measures the overall cluster-level difference between treated and pure control clusters,

$$TCE(\pi) := E(Y_{ic} | \pi_c = \pi) - E(Y_{ic} | \pi_c = 0) = \pi * ITT(\pi) + (1 - \pi) * SNT(\pi).$$

Individuals offered treatment will experience two types of treatment effects, a direct treatment effect from the program as well as a spillover effect that arises from the treatment of other individuals in their cluster. A natural way to formalize these two effects is to decompose the ITT into two components: the **Treatment on the Uniquely Treated** (TUT) measures the ITT on a sole individual offered treatment within a cluster.¹²

$$TUT := E(Y_{ic} | T_{ic} = 1, \pi_c = 0) - E(Y_{ic} | T_{ic} = 0, \pi_c = 0) = ITT(0),$$

The **Spillover on the Treated** (ST) measures the saturation-dependent spillover effect on

¹²The saturation of a cluster includes all treated individuals in the cluster. When the size of a cluster is finite, it is impossible to simultaneously have a treatment individual and a saturation of zero - technically, $ITT(1/n)$ captures the isolated impact of treatment. We use $TUT = ITT(0)$ for notational simplicity.

individuals offered treatment,

$$ST(\pi) := E(Y_{ic} | T_{ic} = 1, \pi_c = \pi) - E(Y_{ic} | T_{ic} = 1, \pi_c = 0).$$

The ITT is the sum of these two components, $ITT(\pi) = TUT + ST(\pi)$.

It is also possible to pool across saturations and estimate an average effect for the entire experiment. Given a RS design ω , define \overline{ITT}_ω as the difference between the expected outcome for individuals offered treatment in each saturation π , weighted by the share of treated clusters with saturation π , and the expected outcome for pure control individuals,

$$\begin{aligned} \overline{ITT}_\omega &:= \sum_{\Pi \setminus 0} E(Y_{ic} | T_{ic} = 1, \pi_c = \pi) \left(\frac{f(\pi)}{1 - \psi} \right) - E(Y_{ic} | T_{ic} = 0, \pi_c = 0) \\ &= \sum_{\Pi \setminus 0} ITT(\pi) \left(\frac{f(\pi)}{1 - \psi} \right). \end{aligned}$$

with analogous definitions for \overline{SNT}_ω , \overline{TCE}_ω and \overline{ST}_ω . This measure depends on the distribution and support of saturations, and will vary across RS designs.¹³

We can now formalize what we refer to as *spillover effects*. There are spillover effects on the untreated (treated) if there exists a π such that $SNT(\pi) \neq 0$ ($ST(\pi) \neq 0$). A sufficient condition to test for the presence of spillovers is $\overline{SNT} \neq 0$ or $\overline{ST} \neq 0$.

1.2 Consistent Estimates of Treatment and Spillover Effects

Next, we establish that a RS design yields consistent estimates of treatment and spillover effects, both at individual saturations and pooled across multiple saturations. Sufficient conditions for consistency are a design with a pure control and an interior saturation, and no interference between clusters.

Result 1. *Assume Assumption 1 and let ω be a non-trivial randomized saturation design with a pure control. Then ω generates unbiased, consistent estimators for $ITT(\pi)$, $SNT(\pi)$ and $TCE(\pi)$ at each $\pi \in \Pi$.*

In order to estimate the pooled effects described in Section 1.1, we must introduce weights. When data are pooled, this unintentionally places a disproportionate weight on treated individuals in high saturation clusters and untreated individuals in low saturation clusters. Saturation weights correct for this distortion.¹⁴

¹³We make this dependence explicit by indexing the pooled measure with ω ; this index is suppressed at times for expositional simplicity.

¹⁴One could define many different pooled effects, including the pooled effect that results from using

Definition 2. Saturation weights apply weight $s_\pi^T = 1/\pi$ to treated individuals and weight $s_\pi^U = 1/(1 - \pi)$ to untreated individuals in treated clusters.

For example, a cluster with $\pi = 2/3$ has twice as many treated individuals as a cluster with $\pi = 1/3$. Weighting the treated individuals by $s_{2/3}^T = 3/2$ and $s_{1/3}^T = 3$ allows one to calculate a pooled estimate that places equal weight on both clusters, rather than twice as much weight the $\pi = 2/3$ clusters.

Result 2. Assume Assumption 1 and let ω be a non-trivial randomized saturation design with a pure control. Then using saturation weights, ω generates unbiased, consistent estimators for \overline{ITT}_ω and \overline{SNT}_ω , and without saturation weights, \overline{TCE}_ω .

We need an additional condition on the RS design to obtain a consistent estimate of the TUT and ST. It is possible to estimate the TUT by either including clusters with very low saturations, or imposing a functional form on $ITT(\pi)$ and deriving $TUT = ITT(0)$ from estimates at other saturations.

Result 3. Assume Assumption 1 and let ω be a non-trivial randomized saturation design with a pure control. If $T\hat{U}T$ is unbiased and consistent, then $S\hat{T}(\pi) = I\hat{T}T(\pi) - T\hat{U}T$ and $\overline{S\hat{T}}_\omega = \overline{I\hat{T}T}_\omega - T\hat{U}T$ are unbiased, consistent estimators.

1.3 Calculating Variances: Stratified Interference and Random Effects

Estimating the variance of treatment and spillover effects requires an assumption on the nature of interference between units and the variance of the data generating process. Within a cluster, we observe a single realization of the many potential configurations of individual treatment assignment at a given saturation.¹⁵ We follow Eric J. Tchetgen and Tyler VanderWeele (2010) in using the ‘Stratified Interference’ assumption proposed by Michael Hudgens and Elizabeth Halloran (2008). This assumption says that the outcome of an individual is independent of the identity of the other individuals assigned to treatment.

Assumption 2. Fixing $\{\pi_c, T_{ic}, R_{ic}, X_{ic}, \varepsilon_{ic}\}$, $Y_{ic} = y$ for any permutation of the treatment status of individuals $j \neq i$.

unweighted data. The definition we propose has two advantages: (i) it is comparable across treatment and within-cluster controls, in that the pooled ITT and SNT give the same weight $f(\pi)/(1 - \psi)$ to each saturation-specific effect $ITT(\pi)$ or $SNT(\pi)$, and (ii) it facilitates an easy test for the shape of the effect (linearity, convexity, etc.) by comparing the pooled ITT to the ITT at the expected saturation.

¹⁵This is not an issue with non-interference, as each unit has only two potential outcomes.

This assumption significantly simplifies the analysis and allows inference without possessing information about the underlying network structure within a cluster.¹⁶

Second, we parameterize the nature of interference within clusters with a random effects error structure.

Assumption 3. *The data generating process has a random effects error structure, with $\varepsilon_{ic} = v_c + w_{ic}$, common cluster component $v_c \sim (0, \tau^2)$, individual component $w_{ic} \sim (0, \sigma^2)$ and (v_c, w_{ic}) orthogonal to $(\pi_c, T_{ic}, R_{ic}, X_{ic})$.*

A random effects framework combined with a RS design decomposes the clustering of outcomes into two components: (i) the extent to which outcomes are endogenously driven by treatment of others in the same cluster, and (ii) the statistical random effect in outcomes, which reduces the power of the clustered estimates but does not imply interference between units.

Remark 2. *This approach mirrors regression techniques typically used to analyze economic and medical experiments, and enables a direct comparison of the power of RS designs to the power of the canonical blocked and clustered designs, making explicit the impact that randomizing saturations has on power. It differs from the approach taken by the recent statistics literature (Hudgens and Halloran 2008), as well as in the paper most similar to ours (Sinclair, McConnell and Green 2012), both of which use randomization inference techniques (Ronald A. Fisher 1935).*

Given Assumptions 1, 2 and 3, we can express Y_{ic} as:

$$Y_{ic} = g(T_{ic}, R_{ic}, \pi_c; X_{ic}) + v_c + w_{ic}.$$

The random effects assumption provides the additional structure needed to characterize the relationship between the RS design, the data generating process and the *Minimum Detectable Effect* (MDE), the smallest treatment or spillover effect that it is possible to distinguish from zero (Howard S. Bloom 1995). Suppose that the true effect is nonzero for some treatment or spillover effect β . Given statistical significance level α , the null hypothesis that $\beta = 0$ is rejected with probability γ (the power) for values of β that exceed:

$$MDE = [t_{1-\gamma} + t_\alpha] * SE(\hat{\beta}).$$

¹⁶In the absence of this assumption, a researcher would need to observe the complete network structure in each cluster, understand the heterogeneity in networks across clusters, and use a model of network-driven spillovers to simulate the variance in outcomes that could be generated by these networks.

In the next two subsections, we characterize the MDE for the treatment and spillover effect measures defined in Section 1.1, show how the MDE depends on the structure of the RS design, establish properties of the optimal RS design to measure each effect (the design that yields the smallest MDE), and illustrate the trade-off between measuring pooled and slope effects.

1.3.1 The Minimum Detectable Pooled Effect

A simple regression-based estimator of the pooled effects is:

$$Y_{ic} = \beta_0 + \beta_1 T_{ic} + \beta_2 S_{ic} + \phi X_{ic} + \varepsilon_{ic} \quad (1)$$

For any non-trivial RS design with a pure control, this model identifies the pooled treatment effect and the pooled spillover effect on untreated individuals, but not the pooled spillover effect on treated individuals. The coefficients depend on the empirical distribution of saturations; given design ω , equation 1 with saturation weights returns $\overline{ITT}_\omega = \hat{\beta}_1$ and $\overline{SNT}_\omega = \hat{\beta}_2$ and equation 1 without saturation weights returns $\overline{TCE}_\omega = (\mu/(1-\psi))\hat{\beta}_1 + ((1-\mu-\psi)/(1-\psi))\hat{\beta}_2$.

The following theorem characterizes the MDE of the pooled ITT and SNT.

Theorem 1. *Assume Assumptions 1, 2 and 3 and let ω be a non-trivial randomized saturation design with a pure control. Then, given statistical significance level α and power γ , the MDE of \overline{ITT}_ω is:*

$$MDE_\omega^T = (t_{1-\gamma} + t_\alpha) \sqrt{\frac{1}{nC} \left\{ (n-1) \tau^2 \left(\frac{1}{(1-\psi)\psi} + \left(\frac{1-\psi}{\mu^2} \right) \eta_T^2 \right) + (\tau^2 + \sigma^2) \left(\frac{\psi + \mu}{\mu\psi} \right) \right\}}$$

The MDE of \overline{SNT}_ω (MDE_ω^S) is similar, substituting μ_S for μ .

The MDE depends on the size of the treatment and control group, and the within-cluster variation in treatment status, η_T^2 . This expression illustrates the relationship between the random effects structure and the RS design. The first term in the brackets captures the variation in $\hat{\beta}$ due to the common cluster component of the error term, and the second term captures the variation in $\hat{\beta}$ due to individual variation. Introducing randomization into the treatment saturation of clusters results in a power loss when there is a common cluster component to the error. Otherwise, if $\tau^2 = 0$, the standard error only depends on the size of the treatment and control groups, but is independent of how treatment is distributed across clusters.




FIGURE 1. Partial Population Design

Sufficient tests for the presence of treatment effects and spillover effects on the untreated are $\overline{ITT}_\omega \neq 0$ and $\overline{SNT}_\omega \neq 0$. The following set of Corollaries derive the optimal RS design to test for these effects. Consider the partial population design in which a cluster is treated with probability $1 - \psi$, and treated clusters all have the same treatment saturation P . This design minimizes the variation in treatment saturation, and therefore, the MDE for treatment and spillover effects.

Corollary 1. *Let Ω be the set of non-trivial RS designs with a pure control and suppose $\tau^2 > 0$. Then, fixing μ and ψ , the design with $\Pi = \{0, P = \mu/(1 - \psi)\}$ and $f = \{\psi, 1 - \psi\}$ (a partial population design) jointly minimizes MDE_ω^T and MDE_ω^S .*

The optimality of a partial population design stems from a positive ICC.

Choosing the optimal treatment saturation P involves a trade-off. The power of the pooled ITT increases with P , while the power of the pooled SNT decreases with P . The relative importance of detecting these two effects, as well as their expected magnitudes, will determine the optimal P .

Corollary 2. *Let Ω be the set of non-trivial RS designs with a pure control and suppose $\tau^2 > 0$. Then, fixing ψ , a partial population experiment with $P = 1/2$ minimizes $MDE_\omega^T + MDE_\omega^S$. In this design, $MDE_{PP}^T = MDE_{PP}^S$.*

The optimal size of the control group depends on the relative magnitude of the common cluster component of error to the individual component of error.

Corollary 3. *Let Ω be the set of non-trivial RS designs with a pure control. The size of the control group that minimizes $MDE_\omega^T + MDE_\omega^S$ depends on τ^2 , σ^2 and n :*

1. If $\tau^2 = 0$, then $\psi^* = \sqrt{2} - 1 \approx 0.41$

2. If $\sigma^2 = 0$, then $\psi^* = \sqrt{n(1+n)} - n$ which converges to 1/2.
3. If $\tau^2 > 0$ and $\sigma^2 > 0$, then $\psi^* \in \left(\sqrt{2} - 1, \sqrt{n(1+n)} - n\right)$

The optimal size of the control therefore lies in a relatively narrow range. Designating about 40% of individuals as pure controls yields the smallest sum of standard errors when there is no common cluster component to the error, while designating close to 50% is preferable when there is no individual component to error. It is always optimal to have the control be more than a third because it serves as the counterfactual for both treatment and spillover groups. As τ^2 increases, the optimal number of control clusters increases. This comparative static arises because the variance in $\hat{\beta}$ due to individual error is proportional to the total number of *individuals* in each treatment group, while the variance in $\hat{\beta}$ due to correlated error is proportional to the total number of *clusters* in each treatment group.

Moving away from the partial population design to a design with variation in the treatment saturation leads to a power loss in the ability to measure pooled effects. Corollary 4 characterizes the rate at which this power loss occurs.

Corollary 4. *Fix μ and ψ . Then $\text{Var}(\hat{\beta})$ increases linearly with respect to η_T^2 .*

Taken together, these corollaries provide important insights on experimental design. If the researcher is only interested in detecting treatment effects and spillover effects on the untreated, then a partial population experiment has the smallest MDE, and Corollary 3 specifies the optimal control group size. However, partial population designs have the drawback that they only measure effects at a single saturation. When researchers care about the effects at multiple saturations, they will need to introduce variation in the treatment saturation. Corollary 4 establishes the rate at which the power of the pooled effects declines from this increase in treatment saturation variance.

1.3.2 The Minimum Detectable Slope Effect

Now suppose that a researcher would like to determine how treatment and spillover effects vary with treatment intensity, or measuring spillover effects on the treated. This section presents two methods to estimate these measures: (1) a non-parametric model that estimates an individual treatment and spillover effect at each non-zero saturation; and (2) a linearized model that estimates the first order effect that changing the treatment saturation has on treatment and spillover effects. Identification of these models requires a RS design with multiple interior treatment saturations and a pure control.

The *Minimum Detectable Slope Effect* (MDSE) is the smallest rate of change δ in the effect, with respect to π , that it is possible to distinguish from zero. Suppose that the true

slope is nonzero. Given statistical significance level α , the null hypothesis that the effect is constant, $\delta = 0$, is rejected with probability γ for values of δ that exceed:

$$MDSE = [t_{1-\gamma} + t_\alpha] * SE(\hat{\delta}).$$

A Non-Parametric Model: A regression based estimator for the treatment and spillover effect at each saturation can be obtained through:

$$Y_{ic} = \beta_0 + \sum_{\Pi \setminus \{0\}} \beta_{1\pi} T_{ic} * \mathbb{1}\{\pi_c = \pi\} + \sum_{\Pi \setminus \{0\}} \beta_{2\pi} S_{ic} * \mathbb{1}\{\pi_c = \pi\} + \phi \cdot X_{ic} + \varepsilon_{ic}, \quad (2)$$

which returns $I\hat{T}T(\pi) = \hat{\beta}_{1\pi}$, $S\hat{N}T(\pi) = \hat{\beta}_{2\pi}$ and $T\hat{C}E(\pi) = \pi\hat{\beta}_{1\pi} + (1 - \pi)\hat{\beta}_{2\pi}$ for each $\pi \in \Pi \setminus \{0\}$.¹⁷ The support of the RS design determines which saturation specific estimates are identified, but unlike equation 1, the definition of the coefficients is independent of the empirical distribution of saturations $f(\pi)$. This model introduces the possibility to test for the presence of spillover effects on treated individuals. A hypothesis test of $\beta_{1\pi_j} = \beta_{1\pi_k}$ determines whether the ITT varies with the treatment saturation. By definition, $\beta_{1\pi_k} - \beta_{1\pi_j} = ST(\pi_k) - ST(\pi_j)$, so this hypothesis also tests for the presence of spillover effects on treated individuals. Similarly, $\beta_{2\pi_j} = \beta_{2\pi_k}$ tests whether the SNT varies with the treatment saturation.

We can also use equation 2 to estimate the change in spillover effects between saturations. Given saturations π_j and π_k , the rate of change of the spillover effect on treated individuals is $\delta_{jk}^T = (\beta_{1\pi_k} - \beta_{1\pi_j}) / (\pi_k - \pi_j)$, with an analogous definition for the within-cluster controls. If spillover effects are affine, then this is a measure of the slope of the spillover effect, $dITT(\pi)/d\pi$ or $dST(\pi)/d\pi$; in the case of a non-linear spillover effect, one can view δ_{jk} as a first order approximation of the slope.

Similar to Theorem 1, we can characterize the MDSE of the ITT and SNT between any pair of saturations $\pi_j, \pi_k \in \Pi$, which is proportional to $SE(\hat{\delta}_{jk}^T)$ or $SE(\hat{\delta}_{jk}^S)$.¹⁸

Theorem 2. *Assume Assumptions 1, 2 and 3 and let ω be a randomized saturation design with $\kappa \geq 2$ interior saturations. Then, given statistical significance level α and power γ , the*

¹⁷No saturation weights are necessary to estimate individual saturation effects.

¹⁸Recall the MDSE of the ITT and ST are equivalent, by definition. It is also possible to calculate the MDE of $ITT(\pi)$ and $SNT(\pi)$ for each saturation π ; this result is similar to the pooled MDE and is presented in the Appendix.

MDSE between saturations π_j and π_k for the treated group is:

$$MDSE_{\omega}^T(\pi_j, \pi_k) = \frac{(t_{1-\gamma} + t_{\alpha})}{\pi_k - \pi_j} \sqrt{\frac{1}{nC} * \left\{ (n-1) \tau^2 \left(\frac{1}{f(\pi_j)} + \frac{1}{f(\pi_k)} \right) + (\tau^2 + \sigma^2) \left(\frac{1}{\mu_j} + \frac{1}{\mu_k} \right) \right\}}$$

where $\mu_k := \pi_k f(\pi_k)$. An similar expression characterizes the MDSE for the within-cluster control group as $MDSE_{\omega}^S$, substituting $\mu_k^S := (1 - \pi_k) f(\pi_k)$ for μ_k .

As the distance between two saturations increases, it is possible to detect smaller slope effects. At the same time, increasing the spread of saturations has a countervailing effect by making the number of treatment (within-cluster control) individuals very small at low (high) saturations. The latter effect dominates at saturations close to zero or one. When the cluster component of error is large, the share of clusters assigned to each saturation, $f(\pi_j)$, plays a larger role in determining the MDSE - a more equal distribution leads to a smaller MDSE. When the individual component of error is large, the share of treated and control individuals assigned to each saturation, μ_j , is more important. Note that while a pure control is required to identify treatment and spillover effects at each saturation in equation 2, it is not required to identify the slope effects.

There are two steps to the design choice for the non-parametric model: selecting which saturations to use (the support of Π), and deciding how to allocate individuals into each saturation bin (the distribution $f(\pi)$). A researcher can either fix a hypothesized slope size and determine how far apart saturations must be to detect this slope, or fix the distance between two saturations and calculate the smallest detectable slope size. Although a partial population design with a saturation of $\pi = 1/2$ is optimal for detecting pooled effects, this design does not identify slope effects. Moving away from the partial population design to a design with two interior saturations, Corollary 5 determines how we should assign the saturations.

Corollary 5. *Let Ω be the set of RS designs with at least two interior saturations. Then, fixing $f(\pi_j) = f(\pi_k)$, the saturations (π_j^*, π_k^*) that minimize $MDSE_{\omega}^T(\pi_j, \pi_k) + MDSE_{\omega}^S(\pi_j, \pi_k)$ are symmetric about $1/2$. The optimal distance $\Delta^* = \pi_k^* - \pi_j^*$ depends on τ^2 , σ^2 and n :*

1. If $\tau^2 = 0$, then $\Delta^* = \sqrt{2}/2 \approx 0.71$.
2. If $\tau^2 > 0$, then $\Delta^* \in (\sqrt{2}/2, 1)$ and $\lim_{n \rightarrow \infty} \Delta^* = 1$.
3. Δ^* is increasing in τ^2 and n , and decreasing in σ^2 .

Therefore, $\pi_j^* = (1 - \Delta^*)/2$ and $\pi_k^* = (1 + \Delta^*)/2$.

Although Theorem 2 is generally too intractable to yield broader analytical insights about optimal design questions, it is possible to numerically calculate the MDSE for designs with more than two saturations. Given κ saturations, a researcher could use Theorem 2 to answer questions like (i) fixing equal sized bins $f(\pi_1) = \dots = f(\pi_\kappa)$, what is the optimal spacing of saturations; or (ii) fixing equally spaced saturations π_1, \dots, π_κ , what share of clusters should be assigned to each bin? This model also allows for hypothesis tests on the shape of the $ITT(\pi)$ and $SNT(\pi)$. For example, a test of concavity requires three interior saturations.

It is possible to use the expression for the $MDE(\pi)$ to calculate the optimal control group size numerically, given an estimate for τ^2 and σ^2 .¹⁹ Similar to the pooled model, the optimal size of the control group will be smaller in the presence of only individual error than in the presence of only cluster-level error, and will lie in between for intermediate error distributions. The optimal control will be smaller than the size of any treatment saturation $\psi^* < f(\pi)$, but will be larger than any treatment or within-cluster control group, $\psi^* > \max\{\pi f(\pi), (1 - \pi)f(\pi)\}$.

An Affine Model: It is also possible to measure slope effects by imposing a functional form on the shape of the $ITT(\pi)$ and $SNT(\pi)$. For example, we could use an affine model to estimate the first order slope effect:

$$Y_{ic} = \delta_0 + \delta_1 T_{ic} + \delta_2 S_{ic} + \delta_3 (T_{ic} * \pi_c) + \delta_4 (S_{ic} * \pi_c) + \phi \cdot X_{ic} + \varepsilon_{ic} \quad (3)$$

This regression identifies the TUT as the intercept of the treatment effect, $T\hat{U}T = \hat{\delta}_1$. The coefficients δ_3 and δ_4 are slope terms estimating how effects change with the saturation, $d\hat{S}T(\pi)/d\pi = \hat{\delta}_3$ and $d\hat{S}\hat{N}T(\pi)/d\pi = \hat{\delta}_4$. The intercept δ_2 estimates spillover effects at saturation zero. There should be no spillover effect on untreated individuals if the saturation of treatment is zero ($SNT(0) = 0$ by definition), so $\delta_2 = 0$ serves as a hypothesis test for the linearity of the spillover relationship. A test for $d\hat{S}T/d\pi = d\hat{S}\hat{N}T/d\pi$ is given by an F-test of the hypothesis that $\delta_3 = \delta_4$.

Similar to Theorem 2, identification of equation 3 requires a RS design with two interior saturations and a pure control. We present an analogous result to Theorem 2 in the Appendix, which characterizes the analytical expression for the MDSE, proportional to $SE(\hat{\delta}_3)$ and $SE(\hat{\delta}_4)$.

It is also possible to test for linearity, or identify non-linear relationships with a similar regression to equation 3. For example, including a squared term $T_{ic} * \pi_c^2$ would identify a quadratic relationship. In simulations, the affine $MDSE$ is smaller than the non-parametric $MDSE$ for detecting these higher moments. Another advantage of the affine model is that

¹⁹The expression for $MDE(\pi)$ is in the proof of Theorem 2 in the Appendix.

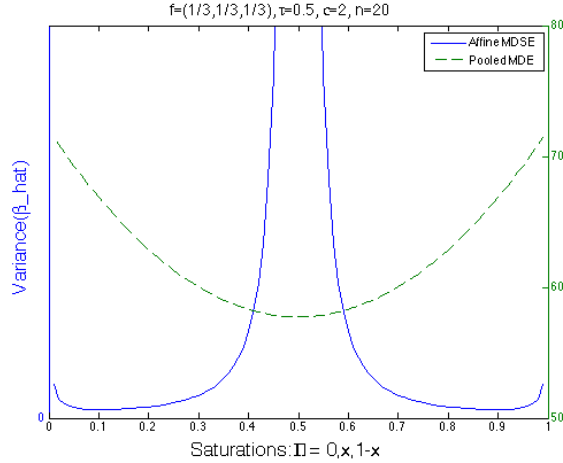


FIGURE 2. Trade-off between Pooled MDE and MDSE

it can be estimated with data from a RS design in which saturations are assigned from a continuum.²⁰

The optimal RS design for a pooled analysis stands in sharp contrast to that for a slope analysis, most obviously in the extent of variation in treatment saturation. A graphical representation of the tradeoff between detecting pooled and slope effects is presented in Figure 2. The optimal RS design to identify both slope and pooled effects will depend on the relative importance that the researcher places on each effect, as well as the expected size of each effect. To facilitate actual implementation of an RS experiment, we created a Matlab program to calculate the minimum detectable effects in the pooled, non-parametric and affine models for different designs. The researcher specifies the relative importance of measuring (i) pooled versus slope effects and (ii) treatment versus spillover on the untreated effects. The program then calculates the optimal support of the RS design, Π , and the optimal allocation of clusters to each saturation bin, $f(\pi)$.

1.4 Estimating the Treatment on the Compliers Effect

This section returns to the general framework of Section 1.2, and derives the **Treatment on the Compliers** (TOC) effect in a model with spillovers.²¹ The TOC is the difference between the expected outcome for individuals who comply with treatment and the expected

²⁰This design is necessary when using a Chow test to identify threshold effects.

²¹This is more commonly known as the Treatment on the Treated (TOT) effect. Throughout this paper, we use the term ‘treated’ to refer to the group offered treatment; therefore, to avoid confusion, we refer to the impact on those actually receiving treatment as the Treatment on the Compliers effect.

outcome for pure control individuals who would have complied with treatment,

$$TOC(\pi) = E(Y_{ic} | T_{ic} = 1, R_{ic} = 1, \pi_c = \pi) - E(Y_{ic} | T_{ic} = 0, R_{ic} = 1, \pi_c = 0).$$

A similar expression defines the pooled effect \overline{TOC} .

In a model with spillovers, the non-compliers in a treatment cluster may be affected by the treatment of compliers, and don't necessarily have the same expected outcome as non-compliers in a control cluster. Define the **Spillover on the Non-Compliers** (SNC) as:

$$SNC(\pi) = E(Y_{ic} | T_{ic} = 1, R_{ic} = 0, \pi_c = \pi) - E(Y_{ic} | T_{ic} = 0, R_{ic} = 0, \pi_c = 0),$$

a spillover term which is conceptually similar to the SNT. Combining these expressions, the $TOC(\pi)$ can be expressed as the difference between the $ITT(\pi)$ and $SNC(\pi)$, weighted by the compliance rate $r(\pi)$:²²

$$TOC(\pi) = \frac{ITT(\pi) - (1 - r(\pi))SNC(\pi)}{r(\pi)}.$$

These expressions have no empirical counterpart because compliance in the control is not observed, and interference between units invalidates the usual strategy of estimating the TOC from the ITT. With no interference, $SNC(\pi) = 0$, and the standard approach of instrumenting for compliance with being offered the treatment produces a valid estimate of the $TOT(\pi)$. With interference, we need an estimate of $SNC(\pi)$ to estimate the $TOT(\pi)$.

An alternative way forward is to assume that spillovers on within-cluster non-compliers are similar to spillovers on within-cluster controls, which are empirically identifiable.²³

Assumption 4. $SNC(\pi) = SNT(\pi)$

This assumption effectively replaces the IV estimator's assumption that $SNC(\pi) = 0$ with an estimate of the spillover effect on untreated individuals, and allows us to recover an estimate of the $TOC(\pi)$.²⁴

²²If compliance varies across the saturation distribution, then changes in $ITT(\pi)$ will be driven by this as well as changes in the underlying $TOC(\pi)$ and $SNC(\pi)$. Indeed, in some cases, such as adoption of a new technology, the most important saturation-driven heterogeneity may come from variation in uptake across the saturation distribution.

²³Unlike many extant partial population experiments in which the within-cluster controls are *ineligible* for the treatment, in a RS design the within-cluster controls come from the same population as the treatment sample, so this assumption may be more warranted.

²⁴Crepon et al. (2013) estimate the treatment on the treated effect by assuming that the externality on an untreated worker is independent of his treatment status, which is equivalent to Assumption 4.

Result 4. Assume Assumption 1 and 4. A non-trivial randomized saturation design with a pure control yields a consistent estimate of the TOC at saturation π ,

$$TO\hat{C}(\pi) = \frac{I\hat{T}T(\pi) - (1 - \hat{r}(\pi))S\hat{N}T(\pi)}{\hat{r}(\pi)}$$

where $\hat{r}(\pi) = \sum_{i,c} \mathbb{1}\{T_{ic} = 1, R_{ic} = 1, \pi_c = \pi\} / \sum_{i,c} \mathbb{1}\{T_{ic} = 1, \pi_c = \pi\}$ is a consistent estimate of the compliance rate at saturation π .

If the compliance rate is constant with respect to treatment saturation, then an analogous expression exists for \overline{TOC} as a function of \overline{ITT} and \overline{SNT} .²⁵

Similar to the ITT, we can break the TOC into two effects: a direct treatment effect from the program, the **Treatment on the Unique Complier** (TUC), and a spillover effect, the **Spillover on the Compliers** (SC). An analogous result to Result 3 identifies these effects.

Returning to the random effects model and maintaining Assumption 4, we can back out estimates of the TOC, SC and TUC. From equation 2, $T\hat{O}C(\pi) = (\hat{\beta}_{1\pi} - (1 - \hat{r}(\pi))\hat{\beta}_{2\pi})/\hat{r}(\pi)$. If we assume the compliance rate is constant with respect to π , equation 1 identifies $\overline{TOC} = ((\hat{\beta}_1 - (1 - \hat{r})\hat{\beta}_2))/\hat{r}$. Equation 3 identifies $d\hat{S}C(\pi)/d\pi = (\hat{\delta}_3 - (1 - r)\hat{\delta}_4)/\hat{r}$ and $T\hat{U}C = \hat{\delta}_1/\hat{r}$. Cross-equation hypothesis testing can be performed using either Seemingly Unrelated Regression or GMM.

In conclusion, the RS framework provides an empirical resolution of why units within a cluster behave similarly. A study that finds high ICCs but no spillover effects can attribute clustering to correlated or contextual effects, while a study with the same ICCs but large spillovers should attribute clustering to endogenous effects. In this way the randomization of saturations resolves the reflection problem (albeit after the fact), and informs optimal design of subsequent experiments in similar contexts.

2 Extensions of the RS Design

2.1 Using Within-cluster Controls as Counterfactuals

Suppose there is no evidence of spillovers on untreated individuals – the estimate of $SNT(\pi)$ is a precise zero for all π . Then the within-cluster controls are not subject to interference

²⁵Estimating the pooled TOC is tricky if the compliance rate varies with treatment saturation:

$$\overline{TOC} = \sum_{\Pi \setminus \{0\}} \left[\left(\frac{1}{r(\pi)} \right) ITT(\pi) + \left(\frac{1 - r(\pi)}{r(\pi)} \right) SNC(\pi) \right] \left(\frac{f(\pi)}{1 - \psi} \right).$$

It is not possible to express \overline{TOC} as a function of \overline{ITT} and \overline{SNT} . We must either estimate $ITT(\pi)$ and $SNC(\pi)$ for each π , or weight observations to take into account the varying compliance rate.

from the treatment and they can be used as counterfactuals.

Assumption 5. $SNT(\pi) = 0$ for all $\pi \in \Pi$.

This assumption is testable using any RS design that identifies a consistent estimate of the $S\hat{N}T(\pi)$.

When Assumption 5 holds, the researcher can pool within-cluster and pure controls, and estimate a simpler model to measure treatment effects:

$$Y_{ic} = \beta_0 + \beta_1 T_{ic} + \phi \cdot X_{ic} + \varepsilon_{ic} \quad (4)$$

Given RS design ω , this regression returns $\overline{ITT}_\omega = \hat{\beta}_1$.²⁶ Power is significantly improved by the larger counterfactual, particularly when the ICC is high.

Theorem 3 characterizes the pooled MDE when the within-cluster controls are included in the counterfactual.

Theorem 3. *Assume Assumptions 1, 2, 3 and 5 and let ω be a randomized saturation design. Then, given statistical significance level α and power γ , the MDE of \overline{ITT}_ω is:*

$$MDE_\omega^T = (t_{1-\gamma} + t_\alpha) \sqrt{\frac{1}{nC} \left\{ \left(\frac{(1 + \rho(n-1))}{\mu(1-\mu)} \right) \tau^2 + \left(\frac{1}{\mu(1-\mu)} \right) \sigma^2 \right\}}$$

where $\rho = \eta^2 / \mu(1-\mu)$ is the correlation in treatment status between two individuals in the same cluster.

Theorem 3 nests the familiar expressions for the MDE of the blocked and clustered designs, and provides context for two well-known results. Fixing the treatment probability μ , the expression for the MDE is decreasing in the variance of the treatment saturation η^2 , and minimized when this variation is zero, which corresponds to the blocked design. Second, fixing η^2 , the MDE is minimized when $\mu(1-\mu)$ is maximized, which occurs at $\mu = 1/2$. Therefore, in the absence of spillovers, the optimal design is a blocked study with equal size treatment and control groups.

An immediate result of Theorem 3 is that the power of the pooled treatment effect in any RS design lies between the power of the treatment effect in the blocked and clustered designs.

Corollary 6. *Let ω be a randomized saturation design with treatment probability μ . Then*

$$MDE_B^T < MDE_\omega^T < MDE_C^T,$$

²⁶Saturation weights are necessary if there are spillover effects on treated individuals, $ST(\pi) \neq 0$ for some $\pi \in \Pi$.

where MDE_B^T is the MDE in a blocked design with saturation μ and MDE_C^T is the MDE in a clustered design with share of treatment clusters μ .

2.2 Using the RS Design to Estimate the Pure Control Outcome

If a study has no pure control group, the counterfactual is at the mercy of within-cluster spillovers. In this context, the RS design has the distinct advantage of allowing a researcher to test for the presence of spillover effects and estimate the unperturbed counterfactual. If the spillover effect is continuous at zero, the researcher can use the variation in treatment saturation to project what would happen to untreated individuals as the saturation approaches zero.²⁷ With this unperturbed counterfactual in hand, we can then estimate $\overline{S\hat{N}T}$, and use this value to correct the estimate of the $\overline{I\hat{T}T}$.

Assumption 6 provides a simple way to estimate the pure control by assuming that the outcome variable is linear with respect to treatment saturation.

Assumption 6. $E(Y|T, \pi)$ is an affine (linear) function of π .

While it is possible to use a more flexible functional form and the specification can be tested, the linear case provides simple intuition for the technique.²⁸

Given Assumption 6, it is natural to estimate:

$$Y_{ic} = \delta_0 + \delta_1 T_{ic} + \delta_2 * \pi_c + \delta_3 (T_{ic} * \pi_c) + \phi \cdot X_{ic} + \varepsilon_{ic} \quad (5)$$

Given RS design ω with no pure control, estimating equation 4 with saturation weights and equation 5, the hypothesis test $\delta_2 = 0$ determines whether there is variation in the control outcome across saturations. If spillovers are present on untreated individuals, then the counterfactual needs to be corrected. The coefficient $\hat{\delta}_0$ is an estimate of the desired ‘pure’ control outcome, $E(Y_{ic} | T_{ic} = 0, \pi_c = 0)$, while $\hat{\beta}_0$ is an estimate of the within-cluster control outcome actually used as the counterfactual, $E(Y_{ic} | T_{ic} = 0, \pi_c > 0)$. The difference between $\hat{\beta}_0$ and $\hat{\delta}_0$ is the $\overline{S\hat{N}T}$, which can be used to derive an unbiased estimate of the $\overline{I\hat{T}T}$.

²⁷Although continuity is a reasonable assumption, it is not universally applicable. Consider signalling in a ground-hog colony. Individuals are ‘treated’ by being alerted to the presence of a nearby predator, and the possible individual-level outcomes are ‘aware’ and ‘not aware’. The animal immediately signals danger to the rest of the colony, and control outcomes will be universally ‘aware’ for any positive treatment saturation, but ‘unaware’ when the saturation is exactly zero.

²⁸In a panel difference in difference regression, the quantity giving the desired counterfactual would be the un-interacted ‘post-treatment’ dummy. This is the change the control group would have experienced at saturation zero

Result 5. Assume Assumption 1, 2, 3 and 6, and let ω be a randomized saturation design with no pure control and $\kappa \geq 2$ interior saturations. Then ω generates consistent estimators of $\overline{ITT}_\omega = \hat{\beta}_1 + \hat{\beta}_0 - \hat{\delta}_0$ and $\overline{SNT}_\omega = \hat{\beta}_0 - \hat{\delta}_0$, where $\hat{\beta}_0$, $\hat{\beta}_1$ and $\hat{\delta}_0$ are the estimates from equation 4 with saturation weights and equation 5.

Similar estimates for the *ITT* and *SNT* at a specific saturation are generated by estimating equation 4 on a single saturation.

The RS design opens up unique empirical possibilities even when there is no pure control group. This is particularly important for settings in which a pure control is not feasible due to regulatory requirements or other exogenous restrictions.²⁹

2.3 Spillover Effects in Overlapping Networks

The RS design we present must be implemented in a non-overlapping network (such as villages or schools), but many networks of interest do not satisfy this strong requirement (such as peer networks or extended families). However, an RS design implemented on a non-overlapping network also produces exogenous variation in the treatment saturation of overlapping networks, variation that is always superior to what would be obtained from a *blocked* design and generally superior to *clustered* designs. This variation depends on the structure of both networks – it increases as the correlation between the two networks increases. As implementing a RS design using non-overlapping clusters is much more straightforward than the sequential randomization required to conduct a RS design in overlapping networks (Toulis and Kao 2013), this provides an attractive way of generating random variation in treatment saturation even when the true network of interest is overlapping.

Figure 3 illustrates the treatment saturation distributions in an overlapping network that results from implementing either a blocked, clustered or RS design on a non-overlapping network.³⁰ Using an overlapping network with five links per individual, we plot the share of individuals at each treatment saturation in the non-overlapping network, where the treatment saturation captures the share of an individual’s links who receive treatment. We use the probability that a link in the overlapping network connects two individuals in the same cluster (the unit of the non-overlapping network) to measure the correlation between networks.³¹

²⁹For example, in McIntosh et al. (2013), a Mexican government rule required that each participating cluster (municipality) be guaranteed at least one treated sub-unit (neighborhood).

³⁰We use a blocked design in which 50% of individuals in each cluster are treated, a clustered design in which 50% of clusters are treated at either 100% or 0% saturation, and a RS design in which an equal share of clusters are treated at saturations 0%, 33%, 67% or 100%. Each assignment rule results in the same overall fraction (one half) of the sample being treated.

³¹The specific structure of the network is irrelevant. Any network with the same number of links and correlation measure will achieve the same saturation distribution.




FIGURE 3. Treatment Saturation of Alternate Network

As can be seen in Figure 3, the *blocked* design produces little overall variation in treatment saturations; the saturations are centered around 50%, independent of the correlation. The *clustered* design suffers from the opposite problem: because treatment has taken place at the cluster level, it is dominated by nodes that have either high or low treatment saturations when there is correlation between networks. Finally, the RS design produces a more even distribution of saturations when there is correlation between networks. In the limit, when there is no correlation between networks, the three designs produce the same saturation distributions (left panel of Figure 3).

3 Empirical Application

The Schooling, Income, and Health Risk (SIHR) is a randomized saturation study designed to understand the role that Conditional and Unconditional Cash Transfers (CCTs and UCTs) play in improving schooling outcomes and reducing early marriage and pregnancy among unmarried, school-age females. We now present an analysis of all of the estimands developed in this paper using the RS design to understand how these programs altered outcomes for the within-cluster controls as well as for the treated. The study took place in the Zomba district of Malawi. Before the start of the intervention, 176 EAs were selected from urban (Zomba city, 29 EAs) and rural (147 EAs) strata for inclusion in the study. ³²

In the 176 study EAs, each dwelling was visited to take a census of all never-married females aged 13-22 years. Within this eligible population we defined two cohorts: those enrolled in school at baseline (baseline schoolgirls), and those not enrolled in school at baseline (baseline dropouts). All baseline dropouts were selected for inclusion in the study

³²Each EA contains an average of 250 households spanning several villages.

		Control Enumeration Areas (N=88)	Treatment Enumeration Areas (N=88)						
			0% Saturation (N=15)	33% Saturation (N=24)	66% Saturation (N=25)	100% Saturation (N=24)			
Study Strata:	Baseline Schoolgirls	Pure Control 1,495	Within-village Control 200	15 CCT 87	9 UCT 68	16 CCT 143	9 UCT 87	15 CCT 276	9 UCT 128
	Baseline Dropouts	Pure Control	CCT	Within-village Control 173	135	Within-village Control 70	44	CCT	CCT

FIGURE 4. Research Design

Shaded cells indicate treatment and numbers give sample sizes at the individual level per cell. Household transfer amounts randomized at the EA level, monthly values of \$4, \$6, \$8, \$10. Participant transfer amounts randomized at the individual level, monthly values of \$1, \$2, \$3, \$4, \$5.

due to the small size of this cohort (approximately five per EA, accounting for about 15% of the target population), while we sampled within the larger cohort of baseline schoolgirls. The percentage of this cohort randomly selected for inclusion in the study was just above 60% and varied by geographical stratum and age group.³³ This sampling procedure yielded 3,796 individuals, who were enrolled in the study and completed baseline interviews at the end of 2007. Of these study participants, 889 were baseline dropouts and 2,907 were the baseline schoolgirls who we analyze here.

Out of the 176 EAs, 88 EAs were assigned to pure control and 88 to treatment. All baseline dropouts in treatment EAs were offered CCTs. The randomized saturation experiment as well as the UCT/CCT experiment was conducted only among baseline schoolgirls. 46 EAs had CCT saturations randomized, 27 EAs had the UCT saturations randomized, and 15 EAs saw only baseline dropouts treated.³⁴ In EAs assigned to CCT, 15 are treated at 33%, 16 are treated at 67%, and 15 are treated at 100%, while there were 9 UCT EAs in each saturation bin. The 15 EAs in which only baseline dropouts are treated provide a 0% CCT saturation, measuring the spillover from CCT treatment of baseline dropouts on baseline schoolgirls. Within each EA, we then selected the integer number of treatments that made the EA-level sample saturation as close as possible to that assigned. Figure 3 presents a schematic of the randomized saturation study design.

In the CCT arm, households were offered cash transfers of between \$5 and \$15 per month if the study participant attended school at least 80% of the days her school was in session

³³The sampling rate varied from 14% to 45% in urban EAs and 70% to 100% in rural ones.

³⁴Due to funding constraints for the transfers, the study included a larger pure control group than would have been ideal for power purposes alone.

during the past month. The UCT arm featured the same transfer system, but the cash transfers were offered unconditionally.³⁵ The cash transfer program started in early 2008, and continued for two years.

We performed the RS experiment only among baseline schoolgirls sampled into the study, meaning that the inclusion rules and sampling rates form a ‘gateway to treatment’ for the true saturation within both the eligible population and the overall population. This has two distinct implications. First, while the conduit for the spillover effects may be an ineligible group such as potential male partners, a gateway-to-treatment study can only hope to capture spillover effects that are both generated and experienced by eligibles. Second, a sampling rate of less than 100% pushes down treatment saturations within the entire eligible population relative to those assigned within the study sample. Because this study featured a high overall sampling rate of 68%, the true saturations are only slightly lower than the assigned. The correlation coefficient between the assigned and true saturations at the EA level is 0.86, while the assigned saturations are completely orthogonal to the sampling weights with a correlation coefficient of 0.03. The actual saturation experiment is thus dampened by one-third from the assigned experiment. To recover marginal effects in the correct units of the true saturation, we instrument for the true saturations with the assigned.

Our analysis utilizes data from three sources. First, the annual SIHR Household Survey provides three rounds of data (baseline, 12-month follow-up, and 24-month follow-up). This survey provides data on the core respondent’s marital status and fertility, as well as her network of friends. Second, we visited the schools of all study participants, who reported being enrolled in school during the 12-month follow-up interviews, and collected data on their enrollment and attendance directly from their schools.³⁶ Finally, to obtain an objective measure of learning, we administered independent tests for English, mathematics, and cognitive skills to study participants in their homes at the 24-month follow-up. The tests were developed by a team of experts at the Human Sciences Research Council according to the Malawian curricula for these subjects for Standards 5-8 and Forms 1-2.³⁷ The outcomes used in the empirical analysis, then, are enrollment, average test scores, and self-reported marriage and pregnancy.

Table 1 shows balance tests with the same specifications to be used in the analysis of spillover effects. All results are shown separately for CCTs and UCTs, providing cross-

³⁵See Baird, McIntosh and Özler (2011) for more details on intervention design.

³⁶While a school survey was also conducted at the 24-month follow-up, this was done only for a random sub-sample of study participants due to budget constraints. Hence, the outcome variable of number of terms enrolled goes from a minimum of zero to a maximum of three – an indicator of school attendance during the first year of the program.

³⁷Primary school in Malawi is from Standard 1 to 8, while secondary school is from Form 1 to 4.

sectional baseline comparisons at the individual level while clustering standard errors at the EA level to account for the design effect. The set of 10 variables for which we examine baseline balance between various treatment groups is the same set reported in Baird, McIntosh and Özler (2011). Panel A shows the simple balance tests; the spillover sample is generally similar to the pure control at baseline. In Panel B we include linear these slope terms, meaning that the top half of Panel B tests for the difference of the 0% saturation (observed in the CCT, extrapolated in the UCT) from the pure control. This provides falsification for the intercept and slope terms to be used in the saturation analysis. Overall the experiment appears well balanced.

3.1 Analysis of Treatment and Spillover Effects

We now present the treatment and spillover effects that can be identified using the randomized saturation design. Table 2 estimates equations (1) and (3), with two modifications. First, we allow the CCT and UCT arms to have separate treatment and spillover effects. Second, we instrument for the true saturation within the eligible population using the randomly assigned saturation in the sample so as to provide marginal effects in the units of the true saturation. Our analysis includes all baseline schoolgirls, controlling for a basic set of baseline covariates and clustering standard errors at the EA level. We present two sets of results for each outcome, first showing simple *ITT* and *SNT* effects by estimating equation (1) in the odd-numbered columns, and then proceeding to test for the presence of saturation slope effects using equation (3) in the even-numbered columns. The bottom two panels of Table 2 explicitly calculate the treatment effects that were developed in Section 1.³⁸

The regression coefficients on the treatment saturations give the linearized slope effects for each outcome. The pooled *ITT* is presented in columns 1 and 2 and the *SNT* in columns 3 and 4. The *TUT* is the intercept term, given by the first two rows in the even-numbered columns. We can divide the *TUT* by the respective compliance rates to calculate the *TUC*, the treatment effect on the unique complier; and calculate the *ToC*, the pooled treatment on the compliers effect, using Assumption 4. These two estimands allow us to calculate the pooled spillovers on the compliers: $SC = TOC - TUC$. Finally, we perform F-tests on each of these estimands, which are linear combinations of regression coefficients across equations. Estimation conducted using Seemingly Unrelated Regressions with OLS models or two-step GMM with IV models provide identical results for the significance levels in the bottom panels of Table 2.

The cluster-level pooled spillover on the non-treated effects (*SNT*) are given by the co-

³⁸The compliance rate was 77.4% for the CCT arm and 99% for the UCT arm.

efficients on the within-cluster control indicators in the odd-numbered columns. Despite the sizable pooled intention to treat effects (ITT), we find no average spillover effects on the non-treated. Furthermore, for each statistically significant treatment effect the average spillover effect on within-cluster controls has the same sign, indicating no evidence of detrimental spillover effects among untreated individuals in treated clusters. The total causal effect, which is a weighted average of the ITT and SNT , presented in the bottom two panels confirms this finding: the TCE closely tracks the ITT in statistical significance and typically appears close to the ITT multiplied times .65, the average treatment saturation in clusters with any treatment. Saturation effects presented in the even-numbered columns suggest that these spillovers on the non-treated increase with treatment saturation, although none of these slope estimates are significant at the 10% level.

When we move to examining spillover effects on the treated (ST), we find some evidence that the beneficial treatment effects decline with treatment intensity. For example, the treatment on the uniquely treated (TUT) effect on enrollment in the CCT group is 0.25 terms, while the pooled ITT estimate is 0.133. Effects on test scores are evident in estimates presented in columns 3 and 4. Similarly, for marriage and pregnancy, the TUT effects in the UCT arm are consistently higher (in absolute value) than the pooled ITT effects, suggesting that beneficial intention to treat effects wear off as more eligible individuals are treated within the cluster. Intriguingly, this indicates that $\frac{d(SNT(\pi))}{d\pi}$ and $\frac{d(ST(\pi))}{d\pi}$ have opposite signs for all four outcome variables, suggesting a welfare tradeoff between treated and untreated units that becomes more pronounced as the treatment saturation increases.

Underlying the estimation of the saturation slope terms in Table 2 are the discrete distributions of the treatment saturations assigned in our experiment: 33%, 67%, 100%, and a 0% CCT cell that estimates the spillover on schoolgirls from CCT baseline dropout treatment alone. We calculate the non-parametric cell-specific $ITT(\pi)$ for each treatment saturation and $SNT(\pi)$ for each saturation below 100%. Table 3 presents this fully granular analysis of impact, showing coefficient estimates for each combination of treatment arm and saturation separately, using the non-parametric regression model in equation 2. In the first column we provide the average true saturation rate within the eligible population for each assigned saturation bin. The impact estimates in columns 2-5 reinforce the findings from Table 2, which used the affine model to estimate saturation effects: spillover effects on the non-treated are generally strongest (and have the same sign as the pooled intention to treat effects) for the cells with the highest treatment saturation (see, e.g. column 2 row 9 or column 4 row 11). Furthermore, again consistent with the earlier findings from Table 2, intention to treat effects are highest in the cells with low saturation, becoming insignificant for the highest saturations (see, e.g. the $ITT(\pi)$ estimates for schooling outcomes in the CCT arm and

those for marriage and pregnancy in the UCT arm in the top panel of Table 3).

3.2 Analysis of Spillover Effects in Friends Network

The finding of weak spillovers within relatively large spatial units could mask the presence of stronger spillovers within social networks. Using data collected at baseline on the closest friends of each study participant, we show that spillover effects are equally muted within this more intimate social network.

At baseline, we asked each study participant to list their five closest friends and to provide some basic information about these friends. We matched the friends to our study sample to determine their treatment status. Restricting our sample to the set of individuals who (a) lived in study EAs, (b) were eligible for the treatment (i.e. never-married females aged 13-22), and (c) are either themselves in the study sample or were listed at baseline among the five closest friends gave us a sample of 8,981 individuals in 176 EAs. As described in Section II.C, because there is a positive correlation between the locations of the study participants and their friends, the RS design generates exogenous variation in treatment intensity within each individual’s social network.

In Table 4, we present program effects as a function of the number of treated friends. The covariates included in the analysis of social networks must reflect the fact that we failed to link friends in a non-negligible number of cases and that we only observe treatment status for friends linked to our sample. To account for this endogenous variation in our ability to link friends in the study sample, we include fully flexible controls for the distribution of the number of matched friends. The findings mirror those presented in Table 2: for beneficiaries and non-beneficiaries alike, none of the outcomes is responsive to the number of treated friends. The similarity in the pattern of spillovers in the spatial and social networks provides some support for stratified interference described in Assumption 2.

Returning to clusters defined over spatial units, such as EAs, the purity of a cross-cluster counterfactual will be compromised if the regional intensity of treatment has an effect on outcomes. To test for this, we conclude by following Miguel and Kremer (2004) and Bobba and Gignoux (2013) in using GIS data on the locations of the EA centroids to count the number of treatment and control EAs within distance bands of <3km and 3-6km from each EA. Since the treated number of EAs is randomized conditional on the total number within each band, we can use this variation to look for cross-cluster spillovers that would violate Assumption 1. Table 5 demonstrates that this cash transfer experiment did not generate strong cross-cluster effects. Coefficient estimates for the number of treated EAs within the two distance bands are always small and statistically insignificant, implying that there are

no spillovers for enrollment, test scores, marriage, or fertility across clusters (columns 1, 3, 5, and 7).³⁹ Exploiting incidental randomization across clusters we confirm Assumption 1 and, as in the within-cluster analysis, find little evidence of spillover effects.

4 Conclusion and Discussion

In recent years, empirical researchers have become increasingly concerned with the problem of interference between subjects. Experiments designed to rigorously estimate spillovers opens up a fascinating set of research questions and provides policy-relevant information about program design. Research designs and RCTs that fail to account for spillovers can be biased; finding meaningful treatment effects but failing to observe deleterious spillovers can lead to misconstrued policy conclusions. This paper attempts to push the frontier of research designs by formalizing the analysis of randomized saturation experiments.

The benefit of randomizing treatment saturations is the ability to generate direct experimental evidence on the nature of spillover and threshold effects. The cost of doing so is statistical power. Having laid out the assumptions necessary to estimate both the mean and variance of spillover effects, we develop explicit, closed-form expressions for the power of RS experiments. We first provide a general expression for power when we seek to estimate treatment and spillover effects jointly. The power loss from randomizing saturations is directly related to the variation in treatment saturation, and so is an inherent feature of the design. Our explicit power calculation formulae provide concrete guidance for optimal research design depending on whether the researcher is primarily interested in measuring pooled treatment and spillover effects or slope effects (which necessitates more partially treated clusters). When spillover effects are found to be muted, this bolsters the credibility of causal inference from clustered designs.

Our empirical application provides little evidence of spillover effects within clusters, or indeed across clusters. This suggests that the significant decreases in marriage and fertility amongst schoolgirls in the unconditional cash transfer treatment group (Baird, McIntosh and Özler 2011) are causal in a larger sense, and are not arising because the treatment diverts such behavior to others girls in the study. For marriage, and pregnancy, the coefficient on treatment saturations for the within-cluster controls is in fact negative, indicating a slight protective effect of the program on nearby individuals who do not receive the treatment.

The framework presented here serves as an important guide to policy questions. For

³⁹In contrast to Bobba and Gignoux (2013), who find large spillover effects of PROGRESA in Mexico but only on treated individuals, we find no consistent evidence that program beneficiaries experience spillovers from adjacent clusters that are any different from untreated individuals (columns 1, 3, 5, and 7). In other words, the ST and the SNT measured cross-cluster are both zero.

example, if a researcher is implementing a program with fixed resources and can either treat 100% of five villages or 50% of ten villages, which treatment allocation will maximize the total benefit? In the Malawi cash transfer program, our results suggest that they would have the same total effect, and the *TCE* of the program is closely approximated by the *ITT* times the average saturation rate, independent of how individuals are assigned to treatment. Small policy trials conducted on a subset of the population can miss important scale or congestion effects that will accompany the full-scale implementation of a program. To the extent that varying the cluster-level saturation leads to differential impacts on prices, norms, and congestion effects, the randomized saturation design provides an experimental framework that can bolster both external and internal validity.

References

- Akerlof, George A., and Rachel E. Kranton.** 2002. "Identity and Schooling: Some Lessons for the Economics of Education." *Journal of Economic Literature*, 40(4): 1167–1201.
- Alix-Garcia, Jennifer, Craig McIntosh, Katharine R. E. Sims, and Jarrod R. Welch.** 2013. "The Ecological Footprint of Poverty Alleviation: Evidence from Mexico's Oportunidades Program." *The Review of Economics and Statistics*, 95(2): 417–435.
- Angelucci, Manuela, and Giacomo De Giorgi.** 2009. "Indirect Effects of an Aid Program: How Do Cash Transfers Affect Ineligibles' Consumption?" *American Economic Review*, 99(1): 486–508.
- Angelucci, Manuela, Giacomo De Giorgi, Marcos A. Rangel, and Imran Rasul.** 2010. "Family networks and school enrolment: Evidence from a randomized social experiment." *Journal of Public Economics*, 94(3-4): 197–221.
- Aronow, Peter.** 2012. "A General Method for Detecting Interference in Randomized Experiments." *Sociological Methods Research*, 41(1): 3–16.
- Babcock, Philip S., and John L. Hartman.** 2010. "Networks and Workouts: Treatment Size and Status Specific Peer Effects in a Randomized Field Experiment." National Bureau of Economic Research, Inc NBER Working Papers 16581. NBER Working Papers.
- Baird, Sarah, Craig McIntosh, and Berk Özler.** 2011. "Cash or Condition? Evidence from a Cash Transfer Experiment." *The Quarterly Journal of Economics*, 126(4): 1709–1753.
- Banerjee, Abhijit, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson.** 2013. "The Diffusion of Microfinance." *Science*, 341(6144).
- Banerjee, Abhijit, Raghavendra Chattopadhyay, Esther Duflo, Daniel Keniston, and Nina Singh.** 2012. "Can Institutions be Reformed from Within? Evidence from

a Randomized Experiment with the Rajasthan Police.” National Bureau of Economic Research, Inc NBER Working Papers 17912.

Barrera-Osorio, Felipe, Marianne Bertrand, Leigh Linden, and Francisco Perez-Calle. 2011. “Improving the Design of Conditional Cash Transfer Programs: Evidence from a Randomized Education Experiment in Colombia.” *American Economic Journal: Applied Economics*, 3(2): 167–195.

Beaman, Lori A. 2012. “Social Networks and the Dynamics of Labour Market Outcomes: Evidence from Refugees Resettled in the U.S.” *The Review of Economic Studies*, 79(1): 128–161.

Behrman, Jere R, Piyali Sengupta, and Petra Todd. 2005. “Progressing through PROGRESA: An Impact Assessment of a School Subsidy Experiment in Rural Mexico.” *Economic Development and Cultural Change*, 54(1): 237–75.

Bloom, Howard S. 1995. “Minimum Detectable Effects: A Simple Way to Report the Statistical Power of Experimental Designs.” *Evaluation Review*, 19(5): 547–556.

Bobba, Matteo, and Jeremie Gignoux. 2013. “Policy Evaluation in the Presence of Spatial Externalities: Reassessing the Progresa Program.” Working Paper.

Bobonis, Gustavo J., and Frederico Finan. 2009. “Neighborhood Peer Effects in Secondary School Enrollment Decisions.” *The Review of Economics and Statistics*, 91(4): 695–716.

Chen, Jiehua, Macartan Humphries, and Vijay Modi. 2010. “Technology Diffusion and Social Networks: Evidence from a Field Experiment in Uganda.” Working Paper.

Conley, Timothy G., and Christopher R. Udry. 2010. “Learning about a New Technology: Pineapple in Ghana.” *American Economic Review*, 100(1): 35–69.

Crepon, Bruno, Esther Duflo, Marc Gurgand, Roland Rathelot, and Philippe Zamora. 2013. “Do Labor Market Policies have Displacement Effects? Evidence from a Clustered Randomized Experiment.” *The Quarterly Journal of Economics*, 128(2): 531–580.

Cunha, Jesse M., Giacomo De Giorgi, and Seema Jayachandran. 2011. “The Price Effects of Cash Versus In-Kind Transfers.” National Bureau of Economic Research, Inc NBER Working Papers 17456. NBER Working Papers.

Duflo, Esther, and Emmanuel Saez. 2002. “Participation and investment decisions in a retirement plan: the influence of colleagues’ choices.” *Journal of Public Economics*, 85(1): 121–148.

Duflo, Esther, and Emmanuel Saez. 2003. “The Role Of Information And Social Interactions In Retirement Plan Decisions: Evidence From A Randomized Experiment.” *The Quarterly Journal of Economics*, 118(3): 815–842.

- Fisher, Ronald A.** 1935. *The Design of Experiments*. Oxford, England: Oliver & Boyd.
- Gine, Xavier, and Ghazala Mansuri.** 2012. “Together we will : experimental evidence on female voting behavior in Pakistan.” Working Paper.
- Hudgens, Michael, and Elizabeth Halloran.** 2008. “Towards Causal Inference with Interference.” *Journal of the American Statistical Association*, 103(482): 832–842.
- Killeen, GF, TA Smith, HM Ferguson, H Mshinda, S Abdulla, et al.** 2007. “Preventing childhood malaria in Africa by protecting adults from mosquitoes with insecticide-treated nets.” *PLoS Med*, 4(7): e229.
- Kuhn, Peter, Peter Kooreman, Adriaan Soetevent, and Arie Kapteyn.** 2011. “The Effects of Lottery Prizes on Winners and Their Neighbors: Evidence from the Dutch Postcode Lottery.” *American Economic Review*, 101(5): 2226–2247.
- Lalive, Rafael, and M. A. Cattaneo.** 2009. “Social Interactions and Schooling Decisions.” *The Review of Economics and Statistics*, 91(3): 457–477.
- Macours, Karen, and Renos Vakis.** 2008. “Changing Households’ Investments and Aspirations through Social Interactions: Evidence from a Randomized Transfer Program in a Low-Income Country.” World Bank Working Paper 5137.
- Manski, Charles.** 1993. “Identification of Endogenous Social Effects: The Reflection Problem.” *Review of Economic Studies*, 60(3): 531–542.
- Masanjala, Winford.** 2007. “The PovertyHIV/AIDS Nexus in Africa: A Livelihood Approach.” *Social Science and Medicine*, 64(5): 1032–1041.
- McIntosh, Craig, Tito Alegria, Gerardo Ordonez, and Rene Zenteno.** 2013. “Infrastructure Impacts and Budgeting Spillovers: The Case of Mexico’s Habitat Program.” Working Paper.
- Miguel, Edward, and Michael Kremer.** 2004. “Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities.” *Econometrica*, 72(1): 159–217.
- Moffitt, Robert A.** 2001. “Policy Interventions, Low-Level Equilibria And Social Interactions.” 45–82. MIT Press.
- Munshi, Kaivan.** 2003. “Networks in the Modern Economy: Mexican Migrants in the U.S. Labor Market.” *Quarterly Journal of Economics*, 118(2): 549–599.
- Oster, Emily, and Rebecca Thornton.** 2012. “Determinants of Technology Adoption: Peer Effects in Menstrual Cup Take-Up.” *Journal of the European Economic Association*, 10(6): 1263–1293.
- Poulin, Michelle J.** 2007. “Sex, money, and premarital relationships in southern Malawi.” *Social Science and Medicine*, 65(11): 2383–2393.

- Sinclair, Betsy, Margaret McConnell, and Donald P. Green.** 2012. “Detecting Spillover Effects: Design and Analysis of Multilevel Experiments.” *American Journal of Political Science*, 56(4): 1055–1069.
- Swindler, Ann, and Susan Watkins.** 2007. “Ties of Dependence: AIDS and Transactional Sex in Rural Malawi.” *Studies in Family Planning*, 38(3): 147–163.
- Tchetgen, Eric J., and Tyler VanderWeele.** 2010. “On Causal Inference in the Presence of Interference.” *Statistical Methods in Medical Research*, 21(1): 55–75.
- Toulis, Panos, and Edward Kao.** 2013. “Estimation of Causal Peer Influence Effects.” *Journal of Machine Learning Research*, 28. Proceedings of the 30th International Conference on Machine Learning Research.

A Mathematical Appendix

A.1 Proofs from Section 1.2, 1.4 and 2.2

Proof of Result 1: Let $\bar{y}_{1,\pi}$ and $\bar{y}_{0,\pi}$ be the sample averages for treated and untreated observations, respectively, in clusters with saturation π . Note for $\pi > 0$,

$$\bar{y}_{1,\pi} = \frac{1}{nC\pi f(\pi)} \sum_{c=1}^C \sum_{i=1}^n Y_{ic} \mathbb{1}_{T_{ic}=1, \pi_c=\pi}$$

with an analogous definition for $\bar{y}_{0,\pi}$ and $\bar{y}_{0,0}$. Then $I\hat{T}T(\pi) = \bar{y}_{1,\pi} - \bar{y}_{0,0}$ converges to $E(Y_{ic} | T_{ic} = 1, \pi_c = \pi) - E(Y_{ic} | T_{ic} = 0, \pi_c = 0) = ITT(\pi)$ by the strong law of large numbers. The results for $S\hat{N}T(\pi) = \bar{y}_{0,\pi} - \bar{y}_{0,0}$ and $T\hat{C}E(\pi) = \pi I\hat{T}T(\pi) + (1 - \pi)S\hat{N}T(\pi)$ are analogous. Q.E.D.

Proof of Result 2: Let \bar{y}_1 and $\bar{y}_{0,\pi>0}$ be the sample averages for treated and within-cluster control observations across all saturations, respectively, weighted with saturation weights. Note that

$$\bar{y}_1 = \frac{1}{nC \sum_{i,c} s_{\pi_c}^T \mathbb{1}_{T_{ic}=1, \pi_c>0}} \sum_{c=1}^C \sum_{i=1}^n Y_{ic} s_{\pi_c}^T \mathbb{1}_{T_{ic}=1, \pi_c>0} = \frac{1}{1 - \psi} \sum_{\Pi \setminus 0} f(\pi) \bar{y}_{1,\pi}$$

Therefore,

$$\bar{y}_1 - \bar{y}_{0,0} = \frac{1}{1 - \psi} \sum_{\Pi \setminus 0} f(\pi) (\bar{y}_{1,\pi} - \bar{y}_{0,0}) = \frac{1}{1 - \psi} \sum_{\Pi \setminus 0} f(\pi) I\hat{T}T(\pi)$$

and from Result 1, $I\hat{T}T(\pi)$ is a consistent, unbiased estimate of $ITT(\pi)$. Therefore, $\overline{I\hat{T}T} = \bar{y}_1 - \bar{y}_{0,0}$ converges to \overline{ITT} . Similarly, $\overline{S\hat{N}T} = \bar{y}_{0,\pi>0} - \bar{y}_{0,0}$ converges to \overline{SNT} .

One must estimate \overline{TCE} from a pooled estimate of the ITT and SNT without saturation weights, because the shifting composition of the sample is integral to the definition of the TCE.. Let $\bar{y}_{\pi>0}$ be the sample average for pooled treated and within-cluster control observations across all saturations.

$$\bar{y}_{\pi>0} = \frac{1}{nC \sum_{i,c} \mathbb{1}_{\pi_c>0}} \sum_{c=1}^C \sum_{i=1}^n Y_{ic} \mathbb{1}_{\pi_c>0} = \frac{1}{1 - \psi} \sum_{\Pi \setminus 0} \pi f(\pi) \bar{y}_{1,\pi} + (1 - \pi) f(\pi) \bar{y}_{0,\pi}$$

Therefore,

$$\bar{y}_{\pi>0} - \bar{y}_{0,0} = \frac{1}{1 - \psi} \sum_{\Pi \setminus 0} f(\pi) \pi I\hat{T}T(\pi) + (1 - \pi) f(\pi) S\hat{N}T(\pi) = \frac{1}{1 - \psi} \sum_{\Pi \setminus 0} f(\pi) \hat{T}CE(\pi)$$

and from Result 1, $\hat{T}CE(\pi)$ is a consistent, unbiased estimate of $TCE(\pi)$. Therefore, $\overline{\hat{T}CE} = \bar{y}_{\pi>0} - \bar{y}_{0,0}$ converges to \overline{TCE} . Note that while it was possible to directly esti-

mate $T\hat{C}E(\pi)$ from $I\hat{T}T(\pi)$ and $S\hat{N}T(\pi)$ in Result 1, it is not possible to directly estimate $\overline{T\hat{C}E}$ from $\overline{I\hat{T}T}$ and $\overline{S\hat{N}T}$. Q.E.D.

$$\overline{T\hat{C}E} = (\mu/(1-\psi))\overline{I\hat{T}T}_{unweighted} + ((1-\mu-\psi)/(1-\psi))\overline{S\hat{N}T}_{unweighted}$$

Proof of Result 3: Analogous to Result 1. Q.E.D.

Proof of Result 4: Given Assumption 4, there is a consistent estimate of the SNC. The rest of the proof is analogous to Result 1. Q.E.D.

Proof of Result 5: Given Assumption 6, we can identify the slope of the ITT and SNT. The rest of the proof is analogous to Result 1. Q.E.D.

A.2 Preliminary Calculations

This section provides background material used to derive the MDE and MDSE in Theorems 1, 2 and 3.

A.2.1 Form of the MDE

The MDE depends on the standard error of $\hat{\beta}$:

$$MDE = [t_{1-\gamma} + t_{\alpha}] * SE(\hat{\beta})$$

To compute the MDE, we need to determine $SE(\hat{\beta})$. This depends on the data generating process and the randomization structure. Consider a model with a random effects error structure:

$$y_{ic} = \mathbf{x}'_{ic}\beta + v_c + w_{ic}$$

for a vector of treatment status covariates \mathbf{x}_{ic} , where v_c is the common cluster component of error and w_{ic} is the individual error. Let $X'_c X_c = \sum_{i=1}^n x_{ic} x'_{ic}$ and $u'_c = [u_{1c} \dots u_{nc}]$, where $u_{ic} = v_c + w_{ic}$. Then the standard error of $\hat{\beta}$ is:

$$SE(\hat{\beta}) = \sqrt{\frac{1}{nC} * A^{-1} B A^{-1}}$$

where

$$A := \text{prob} \lim \frac{1}{nC} \sum_{c=1}^C X'_c X_c \quad \text{and} \quad B := \text{prob} \lim \frac{1}{nC} \sum_{c=1}^C X'_c u_c u'_c X_c$$

Given that all clusters are identical ex-ante, $\frac{1}{N} \sum_{c=1}^C E[X'_c X_c] = \frac{1}{C} \sum_{c=1}^C \frac{1}{n} E[X'_c X_c] = \frac{1}{n} E[X'_c X_c]$. Also note that A and B are independent of whether one takes $n \rightarrow \infty$ or

$C \rightarrow \infty$. Therefore, using the formulas for matrices A and B yields:

$$A = \frac{1}{n}E[X'_c X_c] \quad \text{and} \quad B = \frac{1}{n}E[X'_c u_c u'_c X_c]$$

The matrix B can be decomposed into two matrices, $B = (n-1)\tau^2 D + (\tau^2 + \sigma^2)A$, which leads to the expression:

$$SE(\hat{\beta}) = \sqrt{\left(\frac{n-1}{nC}\right)\tau^2 A^{-1} D A^{-1} + \left(\frac{1}{nC}\right)(\tau^2 + \sigma^2)A^{-1}}$$

where A and D depend on the RS design. We will utilize this expression to calculate $SE(\hat{\beta})$ for different effects and RS designs.

A.2.2 Form of Matrices

$$u_c u'_c = \begin{bmatrix} u_{1c}^2 & u_{1c}u_{2c} & u_{1c}u_{nc} \\ u_{1c}u_{2c} & \dots & \dots \\ \dots & \dots & \dots \\ u_{1c}u_{nc} & \dots & u_{nc}^2 \end{bmatrix}$$

$$E[u_c u'_c] = \begin{bmatrix} \tau^2 + \sigma^2 & \tau^2 & \tau^2 \\ \tau^2 & \tau^2 + \sigma^2 & \dots \\ \dots & \dots & \dots \\ \tau^2 & \dots & \tau^2 + \sigma^2 \end{bmatrix} = \tau^2 + \begin{bmatrix} \sigma^2 & 0 & 0 \\ 0 & \sigma^2 & \dots \\ \dots & \dots & \dots \\ 0 & \dots & \sigma^2 \end{bmatrix}$$

Given $x'_{ic} = [1 \ T_{ic} \ S_{ic} \ \dots]$, (a $1 \times k$ vector), we can write:

$$X_c = \begin{bmatrix} x'_{1c} \\ x'_{2c} \\ \dots \\ x'_{nc} \end{bmatrix} = \begin{bmatrix} 1 & T_{1c} & S_{1c} & \dots \\ 1 & T_{2c} & S_{2c} & \dots \\ \dots & \dots & \dots & \dots \\ 1 & T_{nc} & S_{nc} & \dots \end{bmatrix} \text{ an } n \times k \text{ matrix}$$

$$X'_c X_c = \sum_{i=1}^n x_{ic} x'_{ic} = \sum_{i=1}^n \begin{bmatrix} 1 & T_{ic} & S_{ic} & \dots \\ T_{ic} & T_{ic}^2 & T_{ic} S_{ic} & \dots \\ S_{ic} & T_{ic} S_{ic} & S_{ic}^2 & \dots \\ \dots & \dots & \dots & \dots \end{bmatrix} \text{ a } k \times k \text{ matrix}$$

A.2.3 Relevant Expectations

Distribution of treatment status: ⁴⁰

- $E[T_{ic}] = P(T_{ic} = 1) = \mu$

⁴⁰We implicitly assume that realized saturation is equal to assigned saturation, i.e. $\pi_c = \frac{1}{n} \sum_{i=1}^n T_{ic}$, so given assigned saturation, there is no variation in realized saturation.

- $E[S_{ic}] = 1 - \mu - \psi = \mu_S$
- $E[C_{ic}] = \psi$
- $E[T_{ic}^x] = E[T_{ic}] = \mu$
- $E[S_{ic}^x] = E[S_{ic}] = \mu_S$
- $E[T_{ic}S_{ic}] = 0$

Variance of treatment status:

- $Var[T_{ic}] = E[T_{ic}^2] - E[T_{ic}]^2 = \mu(1 - \mu)$
- $Var[S_{ic}] = (1 - \mu - \psi)(\mu + \psi)$
- $Var[C_{ic}] = \psi(1 - \psi)$

Within cluster treatment status:

- $E[T_{ic}T_{jc}] = P(T_{ic} = 1, T_{jc} = 1) = \sum_{\Pi} P(T_{ic} = 1, T_{jc} = 1, \pi_c = \pi) = \sum_{\Pi} \pi^2 f(\pi) = E[\pi^2]$

where the second equality follows from the chain rule of probability and the third equality follows from the fact that randomization at the individual level is independent within a cluster i.e. T_{ic} is independent of T_{jc} , conditional on π_c .

- $E[S_{ic}S_{jc}] = 1 - 2\mu + E[\pi^2] - \psi$

$$\begin{aligned}
E[S_{ic}S_{jc}] &= P(S_{ic} = 1, S_{jc} = 1) \\
&= \sum_{\Pi} P(S_{ic} = 1, S_{jc} = 1, T_c = \pi) \\
&= \sum_{\Pi \setminus \{0\}} (1 - \pi)^2 f(\pi) \\
&= E[(1 - \pi)^2] - \psi \\
&= 1 - 2\mu + E[\pi^2] - \psi
\end{aligned}$$

- $E[C_{ic}C_{jc}] = \psi$
- $E[T_{ic}S_{jc}] = \mu - E[\pi^2]$

$$\begin{aligned}
E[T_{ic}S_{jc}] &= P(T_{ic} = 1, S_{jc} = 1) \\
&= \sum_{\Pi} P(T_{ic} = 1, S_{jc} = 1, T_c = \pi) \\
&= \sum_{\Pi \setminus \{0\}} \pi(1 - \pi) f(\pi) \\
&= E[\pi(1 - \pi)] - 0 * \psi \\
&= \mu - E[\pi^2]
\end{aligned}$$

Across cluster treatment status:

- $E [T_{ic}T_{jd}] = \mu^2$
since $E [T_{ic}T_{jd}] = E [T_{ic}] E [T_{id}]$ by independence
- $E [S_{ic}S_{jd}] = (1 - \mu - \psi)^2$
- $E [C_{ic}C_{jd}] = \psi^2$

Correlation with saturation π_c is

- $E [T_{ic}\pi_c^x] = \sum_{\Pi} \pi^x P(T_{ic} = 1, \pi_c = \pi) = \sum_{\Pi} \pi^{x+1} f(\pi) = E[\pi^{x+1}]$
- $E [T_{ic}^x \pi_c] = E [T_{ic} \pi_c]$
- $E [T_{ic}T_{jc}\pi_c^x] = \sum_{\Pi} \pi^x P(T_{ic} = 1, T_{jc} = 1, \pi_c = \pi) = \sum_{\Pi} \pi^{x+2} f(\pi) = E[\pi^{x+2}]$
- $E [S_{ic}\pi_c^x] = \sum_{\Pi} \pi^x P(S_{ic} = 1, \pi_c = \pi) = \sum_{\Pi \setminus \{0\}} \pi^x (1 - \pi) f(\pi) = E [(1 - \pi)\pi^x]$
- $E [S_{ic}^x \pi_c] = E [S_{ic} \pi_c]$
- $E [S_{ic}S_{jc}\pi_c^x] = \sum_{\Pi} \pi^x P(S_{ic} = 1, S_{jc} = 1, \pi_c = \pi) = \sum_{\Pi} \pi^x (1 - \pi)^2 f(\pi) = E[\pi^x (1 - \pi)^2]$
- $E [T_{ic}S_{jc}\pi_c^x] = \sum_{\Pi} \pi^x P(T_{ic} = 1, S_{jc} = 1, \pi_c = \pi) = \sum_{\Pi} \pi^x \pi (1 - \pi) f(\pi) = E[\pi^{x+1} (1 - \pi)]$

Correlation of treatment status between two girls in the same cluster:

$$\begin{aligned} \rho_T &= \frac{E [T_{ic}T_{jc}] - E [T_{ic}] E [T_{ic}]}{Var [T_{ic}]} = \frac{\eta^2}{\mu(1 - \mu)} \\ \rho_S &= \frac{E [S_{ic}S_{jc}] - E [S_{ic}] E [S_{ic}]}{Var [S_{ic}]} \\ &= \frac{(1 - 2\mu + \eta^2 + \mu^2 - \psi) - (1 - \mu - \psi)^2}{(1 - \mu - \psi)(\mu + \psi)} \\ &= \frac{\eta^2 + \psi(1 - 2\mu - \psi)}{(1 - \mu - \psi)(\mu + \psi)} \\ \rho_C &= 1 \end{aligned}$$

Distribution of u_{ic} :

- $E [u_{ic}^2] = \tau^2 + \sigma^2$
- $E [u_{ic}u_{jc}] = \tau^2$ if $i \neq j$ which is $Cov(u_{ic}u_{jc})$
- $E [u_{ic}u_{jd}] = 0$ if $c \neq d$
- T_{ic} or S_{ic} is independent of $u_{ic} \Rightarrow E [f(u_{ic})g(T_{ic})] = E [f(u_{ic})] * E [g(T_{ic})]$

Sum of the error terms and treatment status within each cluster:

$$\begin{aligned} \frac{1}{n} E \left[\left(\sum_{i=1}^n u_{ic} \right)^2 \right] &= E [u_{ic}^2] + (n-1) E [u_{jc} u_{ic}] = (n\tau^2 + \sigma^2) \\ \frac{1}{n} E \left[\left(\sum_{i=1}^n u_{ic} \right) \left(\sum_{i=1}^n T_{ic} u_{ic} \right) \right] &= E [u_{ic}^2 T_{ic}] + (n-1) E [u_{jc} u_{ic} T_{ic}] = (n\tau^2 + \sigma^2) \mu \\ \frac{1}{n} E \left[\left(\sum_{i=1}^n u_{ic} \right) \left(\sum_{i=1}^n S_{ic} u_{ic} \right) \right] &= (n\tau^2 + \sigma^2) (1 - \mu - \psi) \end{aligned}$$

since T_{ic} and S_{ic} are independent of u_{ic} .

$$\begin{aligned} \frac{1}{n} E \left[\left(\sum_{i=1}^n T_{ic} u_{ic} \right)^2 \right] &= E [u_{ic}^2 T_{ic}^2] + (n-1) E [u_{jc} u_{ic} T_{ic} T_{jc}] \\ &= (\tau^2 + \sigma^2) \mu + (n-1) \tau^2 E[\pi^2] \\ \frac{1}{n} E \left[\left(\sum_{i=1}^n S_{ic} u_{ic} \right)^2 \right] &= (\tau^2 + \sigma^2) (1 - \mu - \psi) + (n-1) \tau^2 (1 - 2\mu + E[\pi^2] - \psi) \\ \frac{1}{n} E \left[\left(\sum_{i=1}^n T_{ic} u_{ic} \right) \left(\sum_{i=1}^n S_{ic} u_{ic} \right) \right] &= (n-1) \tau^2 (\mu - E[\pi^2]) \end{aligned}$$

since $T_{ic} S_{ic} = 0$ for all i, c .

A.2.4 Variance of Treatment Saturation

The marginal distribution of saturations across treatment clusters (removing control clusters) is:

$$g(\pi) = \frac{f(\pi)}{1 - \psi}$$

with support $\Pi \setminus \{0\}$.

$$\begin{aligned} E_g[\pi^2] &= \frac{1}{1 - \psi} \sum_{\Pi \setminus \{0\}} \pi^2 f(\pi) \\ E_g[\pi] &= \frac{1}{1 - \psi} \sum_{\Pi \setminus \{0\}} \pi f(\pi) \end{aligned}$$

The component of total variation in cluster saturation due to variation in the saturation of treated clusters is:

$$\begin{aligned}
\eta_T^2 &: = \text{Var}(\pi|\pi > 0) = E_g[\pi^2] - E_g[\pi]^2 \\
&= \left(\frac{1}{1-\psi}\right) \left[\sum_{\Pi \setminus \{0\}} \pi^2 f(\pi) - \left(\frac{1}{1-\psi}\right) \left(\sum_{\Pi \setminus \{0\}} \pi f(\pi) \right)^2 \right] \\
&= \left(\frac{1}{1-\psi}\right) \eta^2 - \left(\frac{\psi}{(1-\psi)^2}\right) \mu^2
\end{aligned}$$

Then the total variation can be expressed as the sum of the variation in the saturation of treated clusters and the variation between treated and control clusters, weighted by the size of the control group:

$$\eta^2 = (1-\psi) \eta_T^2 + \left(\frac{\psi}{1-\psi}\right) \mu^2$$

A.3 Proof of Theorem 1

We want to compute matrices A and B for the model with $x'_{ic} = [1 \ T_{ic} \ S_{ic}]$. Using the calculations from Section A.2, we can calculate:

$$\begin{aligned}
A &= \frac{1}{n} \sum_{i=1}^n E \begin{bmatrix} n & T_{ic} & S_{ic} \\ T_{ic} & T_{ic}^2 & T_{ic}S_{ic} \\ S_{ic} & T_{ic}S_{ic} & S_{ic}^2 \end{bmatrix} = \begin{bmatrix} 1 & \mu & \mu_S \\ \mu & \mu & 0 \\ \mu_S & 0 & \mu_S \end{bmatrix} \\
B &= \frac{1}{n} E \begin{bmatrix} (\sum_{i=1}^n u_{ic})^2 & (\sum_{i=1}^n u_{ic}) (\sum_{i=1}^n T_{ic} u_{ic}) & (\sum_{i=1}^n u_{ic}) (\sum_{i=1}^n S_{ic} u_{ic}) \\ (\sum_{i=1}^n u_{ic}) (\sum_{i=1}^n T_{ic} u_{ic}) & (\sum_{i=1}^n T_{ic} u_{ic})^2 & (\sum_{i=1}^n T_{ic} u_{ic}) (\sum_{i=1}^n S_{ic} u_{ic}) \\ (\sum_{i=1}^n u_{ic}) (\sum_{i=1}^n S_{ic} u_{ic}) & (\sum_{i=1}^n T_{ic} u_{ic}) (\sum_{i=1}^n S_{ic} u_{ic}) & (\sum_{i=1}^n S_{ic} u_{ic})^2 \end{bmatrix} \\
&= (n-1)\tau^2 \begin{bmatrix} 1 & \mu & \mu_S \\ \mu & \eta^2 + \mu^2 & \mu - \mu^2 - \eta^2 \\ \mu_S & \mu - \mu^2 - \eta^2 & \mu_S - \mu + \eta^2 + \mu^2 \end{bmatrix} + (\tau^2 + \sigma^2) A
\end{aligned}$$

Using mathematica to compute $SE(\hat{\beta}) = \sqrt{\frac{1}{nC} * A^{-1}BA^{-1}}$, taking the diagonal entries and plugging in the expression relating η^2 and η_T^2 yields the result. Q.E.D.

Proof of Corollary 1: Fixing μ and ψ , $SE(\hat{\beta}_1)$ and $SE(\hat{\beta}_2)$ are both minimized at $\eta_T^2 = 0$. This corresponds to a partial population experiment with a control group of size ψ and a treatment saturation of $P = \mu/(1-\psi)$. Q.E.D.

Proof of Corollary 2: Fixing ψ , a partial population design has the smallest sum of standard errors, for any treatment size μ . Therefore, we can restrict attention to the set of partial population designs, and the expression for $Var(\hat{\beta})$ simplifies to:

$$MDE_{\omega}^T = [t_{1-\gamma} + t_{\alpha}] \sqrt{\frac{1}{nC} * \left\{ (n-1) \tau^2 \left(\frac{1}{(1-\psi)\psi} \right) + (\tau^2 + \sigma^2) \left(\frac{\psi + \mu}{\mu\psi} \right) \right\}}$$

$$MDE_{\omega}^S = [t_{1-\gamma} + t_{\alpha}] \sqrt{\frac{1}{nC} * \left\{ (n-1) \tau^2 \left(\frac{1}{(1-\psi)\psi} \right) + (\tau^2 + \sigma^2) \left(\frac{1-\mu}{(1-\mu-\psi)\psi} \right) \right\}}$$

The sum of these expressions is minimized at $\mu = \mu_S = (1 - \psi)/2$, which corresponds to a partial population experiment with $P = 1/2$. Q.E.D.

Proof of Corollary 3: In a partial population design with $\mu = (1 - \psi)/2$,

$$MDE_{\omega}^T = MDE_{\omega}^S = [t_{1-\gamma} + t_{\alpha}] \sqrt{\frac{1}{nC} * \left\{ (n-1) \tau^2 \left(\frac{1}{(1-\psi)\psi} \right) + (\tau^2 + \sigma^2) \left(\frac{\psi + 1}{(1-\psi)\psi} \right) \right\}}$$

When there is no inter-cluster correlation, $\tau^2 = 0$, the expression simplifies to:

$$[t_{1-\gamma} + t_{\alpha}] \sqrt{\frac{1}{nC} * \sigma^2 \left(\frac{\psi + 1}{(1-\psi)\psi} \right)}$$

which is minimized at $\psi^* = \sqrt{2} - 1$. When there is no individual error, $\sigma^2 = 0$, the expression simplifies to:

$$[t_{1-\gamma} + t_{\alpha}] \sqrt{\frac{1}{nC} * \left\{ \tau^2 \left(\frac{n + \psi}{(1-\psi)\psi} \right) \right\}}$$

which is minimized at $\psi^* = \sqrt{n(1+n)} - n$. Note $\lim_{n \rightarrow \infty} \sqrt{n(1+n)} - n = 1/2$. Given that $(\psi + 1)/((1 - \psi)\psi)$ and $(\psi + n)/((1 - \psi)\psi)$ are both convex with unique minimums, any weighted sum of these functions is minimized at a value ψ^* that lies between the minimum of each function. Therefore, when $\tau^2 > 0$ and $\sigma^2 > 0$, $\psi^* \in (\sqrt{2} - 1, \sqrt{n(1+n)} - n)$. Q.E.D.

Proof of Corollary 4: Follows directly from Theorem 1. Q.E.D.

A.4 Proof of Theorem 2

We want to compute matrices A and B for the model with $x'_{ic} = [1 \ T_{1ic} \ S_{1ic} \ T_{2ic} \ S_{2ic}]$ where $T_{1ic} = \mathbb{1}(T_{ic} = 1, \pi_c = \pi_1)$, $S_{1ic} = \mathbb{1}(T_{ic} = 0, \pi_c = \pi_1)$, $T_{2ic} = \mathbb{1}(T_{ic} = 1, \pi_c = \pi_{2ic})$ and so forth. Using the calculations from Section A.2 and defining $\mu_k := \pi_k f(\pi_k)$, $p_k :=$

$(1 - \pi_k) f(\pi_k)$, $\eta_k := \pi_k^2 f(\pi_k)$ and $\rho_k := (1 - \pi_k)^2 f(\pi_k) = p_k - \mu_k + \eta_k$, we can calculate:

$$\begin{aligned}
A &= \frac{1}{n} \sum_{i=1}^n E \begin{bmatrix} n & T_{1ic} & S_{1ic} & T_{2ic} & S_{2ic} \\ T_{1ic} & T_{1ic}^2 & 0 & 0 & 0 \\ S_{1ic} & 0 & S_{1ic}^2 & 0 & 0 \\ T_{2ic} & 0 & 0 & T_{2ic}^2 & 0 \\ S_{2ic} & 0 & 0 & 0 & S_{2ic}^2 \end{bmatrix} = \begin{bmatrix} 1 & \mu_1 & p_1 & \mu_2 & p_2 \\ \mu_1 & \mu_1 & 0 & 0 & 0 \\ p_1 & 0 & p_1 & 0 & 0 \\ \mu_2 & 0 & 0 & \mu_2 & 0 \\ p_2 & 0 & 0 & 0 & p_2 \end{bmatrix} \\
B &= \frac{1}{n} E \begin{bmatrix} (\sum_{i=1}^n u_{ic}) \\ (\sum_{i=1}^n T_{1ic} u_{ic}) \\ (\sum_{i=1}^n S_{1ic} u_{ic}) \\ (\sum_{i=1}^n T_{2ic} u_{ic}) \\ (\sum_{i=1}^n S_{2ic} u_{ic}) \end{bmatrix} * \begin{bmatrix} (\sum_{i=1}^n u_{ic}) \\ (\sum_{i=1}^n T_{1ic} u_{ic}) \\ (\sum_{i=1}^n S_{1ic} u_{ic}) \\ (\sum_{i=1}^n T_{2ic} u_{ic}) \\ (\sum_{i=1}^n S_{2ic} u_{ic}) \end{bmatrix}^\top \\
&= (n-1)\tau^2 \begin{bmatrix} 1 & \mu_1 & p_1 & \mu_2 & p_2 \\ \mu_1 & \eta_1 & \mu_1 - \eta_1 & 0 & 0 \\ p_1 & \mu_1 - \eta_1 & \rho_1 & 0 & 0 \\ \mu_2 & 0 & 0 & \eta_2 & \mu_2 - \eta_2 \\ p_2 & 0 & 0 & \mu_2 - \eta_2 & \rho_2 \end{bmatrix} + (\tau^2 + \sigma^2) A
\end{aligned}$$

Using mathematica to compute $SE(\hat{\beta}) = \sqrt{\frac{1}{nC} * A^{-1} B A^{-1}}$ and taking the diagonal entries yields the MDE^T for each saturation π_j :

$$\begin{aligned}
MDE_\omega^T(\pi_j) &= (t_{1-\gamma} + t_\alpha) \sqrt{\frac{1}{nC} * \left\{ (n-1)\tau^2 \left(\frac{\eta_j}{\mu_j^2} + \frac{1}{\psi} \right) + (\tau^2 + \sigma^2) \left(\frac{\psi + \mu_j}{\psi \mu_j} \right) \right\}} \\
&= (t_{1-\gamma} + t_\alpha) \sqrt{\frac{1}{nC} * \left\{ (n-1)\tau^2 \left(\frac{1}{f(\pi_j)} + \frac{1}{\psi} \right) + (\tau^2 + \sigma^2) \left(\frac{1}{\mu_j} + \frac{1}{\psi} \right) \right\}}
\end{aligned}$$

Next, we can compute $MDSE_\omega^T(\pi_j, \pi_k)$ from

$$SE(\delta_{jk}^T) = SE[(\beta_{1\pi_k} - \beta_{1\pi_j}) / (\pi_k - \pi_j)] = SE(\beta_{1\pi_k} - \beta_{1\pi_j}) / (\pi_k - \pi_j)$$

as

$$\begin{aligned}
Cov(\beta_{1\pi_k}, \beta_{1\pi_j}) &= \frac{1}{nC} (n\tau^2 + \sigma^2) \frac{1}{\psi} \\
Var(\beta_{1\pi_k} - \beta_{1\pi_j}) &= Var(\beta_{1\pi_j}) + Var(\beta_{1\pi_k}) - 2Cov(\beta_{1\pi_k}, \beta_{1\pi_j}) \\
&= \frac{1}{nC} * \left\{ (n-1)\tau^2 \left(\frac{1}{f(\pi_j)} + \frac{1}{f(\pi_k)} \right) + (\tau^2 + \sigma^2) \left(\frac{1}{\mu_j} + \frac{1}{\mu_k} \right) \right\}
\end{aligned}$$

Plugging this into the expression for $MDSE^T$ yields the result. Calculating the MDE^S and $MDSE^S$ is analogous. Q.E.D.

Proof of Corollary 5: Fixing the size of each saturation bin $f(\pi_j) = f_j$ and $f(\pi_k) = f_k$ and the distance between two saturations $\pi_k - \pi_j = \Delta$, minimizing $MDSE_\omega^T(\pi_j, \pi_k) + MDSE_\omega^S(\pi_j, \pi_k)$ is equivalent to solving:

$$\min_{\pi_j} \left(\frac{1}{f_j \pi_j} + \frac{1}{f_k(\pi_j + \Delta)} + \frac{1}{f_j(1 - \pi_j)} + \frac{1}{f_k(1 - \Delta - \pi_j)} \right)$$

The minimum occurs at the π_j^* that solves $\pi_j^*(1 - \pi_j^*)f_j = (\pi_j^* + \Delta)(1 - \Delta - \pi_j^*)f_k$. When $f_j = f_k$, $\pi_j^* = (1 - \Delta)/2$ and $\pi_k^* = \pi_j^* + \Delta = (1 + \Delta)/2$, which is symmetric about $1/2$.

Fixing $f_j = f_k$, the Δ that minimizes $MDSE_\omega^T(\pi_j, \pi_k) + MDSE_\omega^S(\pi_j, \pi_k)$ is equivalent to solving:

$$\min_{\Delta} \frac{1}{\Delta^2} \left(\frac{(n-1)\tau^2}{n} + \frac{(\tau^2 + \sigma^2)}{n} \left(\frac{2}{(1-\Delta)(1+\Delta)} \right) \right)$$

The optimal Δ^* solves:

$$\frac{(n-1)\tau^2}{2(\tau^2 + \sigma^2)} = \frac{2(\Delta^*)^2 - 1}{(1 - (\Delta^*)^2)^2}$$

If $\tau^2 = 0$, then $2(\Delta^*)^2 - 1 = 0$, yielding $\Delta^* = \sqrt{2}/2$. Note that $(2\Delta^2 - 1)/((1 - \Delta^2)^2)$ is monotonically increasing for $\Delta \in [0, 1)$, and strictly positive for $\Delta > \sqrt{2}/2$. When $\tau > 0$, $((n-1)\tau^2)(2(\tau^2 + \sigma^2))$ is also strictly positive, increasing in τ^2 and decreasing in σ^2 . Therefore, $\Delta^* \in (\sqrt{2}/2, 1)$ for $\tau^2 > 0$ and finite n , Δ^* is increasing in τ^2 and n , and decreasing in σ^2 . If $\tau^2 > 0$, then the left hand side converges to ∞ as $n \rightarrow \infty$, which requires $\Delta^* \rightarrow 1$. Q.E.D.

A.5 Proof of Theorem 3

We want to compute matrices A and B for the model with $x'_{ic} = [1 \ T_{ic}]$. Using the calculations from Section A.2, we can calculate:

$$A = \frac{1}{n} E \left[\begin{array}{cc} n & \sum_{i=1}^n T_{ic} \\ \sum_{i=1}^n T_{ic} & \sum_{i=1}^n T_{ic}^2 \end{array} \right] = \left[\begin{array}{cc} 1 & \mu \\ \mu & \mu \end{array} \right]$$

and

$$\begin{aligned} B &= \frac{1}{n} E \left[\begin{array}{cc} (\sum_{i=1}^n u_{ic})^2 & (\sum_{i=1}^n u_{ic})(\sum_{i=1}^n T_{ic} u_{ic}) \\ (\sum_{i=1}^n u_{ic})(\sum_{i=1}^n T_{ic} u_{ic}) & (\sum_{i=1}^n T_{ic} u_{ic})^2 \end{array} \right] \\ &= \tau^2(n-1) \left[\begin{array}{cc} 1 & \mu \\ \mu & \eta^2 + \mu^2 \end{array} \right] + (\tau^2 + \sigma^2) A \end{aligned}$$

This can be used to compute

$$SE(\hat{\beta}_1) = \sqrt{\frac{1}{nC} * \left[\left(\frac{1}{\mu(1-\mu)} + \frac{(n-1)\eta^2}{\mu^2(1-\mu)^2} \right) \tau^2 + \left(\frac{1}{\mu(1-\mu)} \right) \sigma^2 \right]}$$

Using $\eta^2 = \rho\mu(1 - \mu)$, we can express $SE(\widehat{\beta}_1)$ in terms of μ and ρ .

$$SE(\widehat{\beta}_1) = \sqrt{\frac{1}{nC} * \left[\left(\frac{(1 + \rho(n - 1))}{\mu(1 - \mu)} \right) \tau^2 + \left(\frac{1}{\mu(1 - \mu)} \right) \sigma^2 \right]}$$

Fixing μ , this expression is minimized at $\eta^2 = 0$ or $\rho = 0$. Q.E.D.

Proof of Corollary 6: Follows directly from Theorem 3, noting that the blocked design corresponds to $\rho = 0$ and the clustered design corresponds to $\rho = 1$. Q.E.D.

A.6 Affine Model

Theorem 4. *Assume Assumptions 1, 2 and 3 and let ω be a randomized saturation design with $\kappa \geq 2$ interior saturations and a pure control. Then, given statistical significance level α and power γ , the MDSE for the treated group is:*

$$MDSE_{\omega}^T = (t_{1-\gamma} + t_{\alpha}) \sqrt{\frac{1}{nC} * \{(n - 1) \tau^2 h_1 + (\tau^2 + \sigma^2) h_2\}}$$

where

$$h_1 = \left(\frac{(\eta^2 + \mu^2)^2 - 2\mu(\eta^2 + \mu^2)E[\pi^3] + \mu^2 E[\pi^4]}{(\eta^2 + \mu^2)^2 - \mu E[\pi^3]^2} \right) \text{ and } h_2 = \left(\frac{\eta^2 + \mu^2}{(\eta^2 + \mu^2)^2 - \mu E[\pi^3]} \right)$$

A similar expression characterizes the MDSE for the within-cluster control group as $MDSE_{\omega}^S$.

Proof of Theorem 4: We want to compute matrices A and B for the model with

$$x'_{ic} = [1 \quad T_{ic} \quad T_{ic}\pi_c \quad S_{ic} \quad S_{ic}\pi_c]$$

Using the calculations from Section A.2, we can calculate:

$$\begin{aligned} A &= \frac{1}{n} \sum_{i=1}^n E \begin{bmatrix} 1 & T_{ic} & T_{ic}\pi_c & S_{ic} & S_{ic}\pi_c \\ T_{ic} & T_{ic}^2 & T_{ic}^2\pi_c & T_{ic}S_{ic} & T_{ic}S_{ic}\pi_c \\ T_{ic}\pi_c & T_{ic}^2\pi_c & T_{ic}^2\pi_c^2 & T_{ic}S_{ic}\pi_c & T_{ic}S_{ic}\pi_c^2 \\ S_{ic} & T_{ic}S_{ic} & T_{ic}S_{ic}\pi_c & S_{ic}^2 & S_{ic}^2\pi_c \\ S_{ic}\pi_c & T_{ic}S_{ic}\pi_c & T_{ic}S_{ic}\pi_c^2 & S_{ic}^2\pi_c & S_{ic}^2\pi_c^2 \end{bmatrix} \\ &= \begin{bmatrix} 1 & \mu & \eta^2 + \mu^2 & 1 - \mu - \psi & \mu - \eta^2 + \mu^2 \\ \mu & \mu & \eta^2 + \mu^2 & 0 & 0 \\ \eta^2 + \mu^2 & \eta^2 + \mu^2 & E[\pi^3] & 0 & 0 \\ 1 - \mu - \psi & 0 & 0 & 1 - \mu - \psi & \mu - \eta^2 + \mu^2 \\ \mu - \eta^2 + \mu^2 & 0 & 0 & \mu - \eta^2 + \mu^2 & \eta^2 + \mu^2 - E[\pi^3] \end{bmatrix} \end{aligned}$$

$$\begin{aligned}
B &= \frac{1}{n} E \begin{bmatrix} (\sum_{i=1}^n u_{ic}) \\ (\sum_{i=1}^n T_{ic} u_{ic}) \\ (\sum_{i=1}^n T_{ic} \pi_c u_{ic}) \\ (\sum_{i=1}^n S_{ic} u_{ic}) \\ (\sum_{i=1}^n S_{ic} \pi_c u_{ic}) \end{bmatrix} * \begin{bmatrix} (\sum_{i=1}^n u_{ic}) \\ (\sum_{i=1}^n T_{ic} u_{ic}) \\ (\sum_{i=1}^n T_{ic} \pi_c u_{ic}) \\ (\sum_{i=1}^n S_{ic} u_{ic}) \\ (\sum_{i=1}^n S_{ic} \pi_c u_{ic}) \end{bmatrix}^T \\
&= (n-1)\tau^2 D + (\tau^2 + \sigma^2) A
\end{aligned}$$

where

$$D = \begin{bmatrix} 1 & \mu & E[\pi^2] & 1 - \mu - \psi & \mu - E[\pi^2] \\ \mu & E[\pi^2] & E[\pi^3] & \mu - E[\pi^2] & E[\pi^2] - E[\pi^3] \\ E[\pi^2] & E[\pi^3] & E[\pi^4] & E[\pi^2] - E[\pi^3] & E[\pi^3] - E[\pi^4] \\ 1 - \mu - \psi & \mu - E[\pi^2] & E[\pi^2] - E[\pi^3] & 1 - 2\mu + E[\pi^2] - \psi & \mu - 2E[\pi^2] + E[\pi^3] \\ \mu - E[\pi^2] & E[\pi^2] - E[\pi^3] & E[\pi^3] - E[\pi^4] & \mu - 2E[\pi^2] + E[\pi^3] & E[\pi^2] - 2E[\pi^3] + E[\pi^4] \end{bmatrix}$$

Using mathematica to compute $SE(\hat{\delta}) = \sqrt{\frac{1}{nC} * A^{-1} B A^{-1}}$ and taking the diagonal entries yields the result. The $MDSE^T$ is a function of $SE(\hat{\delta}_3)$, while the $MDSE^S$ is a function of $SE(\hat{\delta}_4)$.

TABLE 1. Balance Tests

		Dependent Variable:									
		Asset Index	Female Headed Household	Mobile Phone Ownership	Age	Highest Grade at Baseline	Mother Alive	Father Alive	Never had Sex	Ever Pregnant	
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
PANEL A: Pooled Tests.											
	Household Size	0.027	0.330	-0.0889**	-0.034	-0.251**	-0.218	-0.0422*	-0.002	-0.005	0.004
		(0.195)	(0.298)	(0.040)	(0.051)	(0.113)	(0.169)	(0.025)	(0.039)	(0.027)	(0.009)
	UCT	0.219	0.464**	-0.0773*	-0.032	0.138	0.355**	-0.001	0.051	-0.017	0.006
		(0.162)	(0.221)	(0.042)	(0.064)	(0.136)	(0.154)	(0.026)	(0.038)	(0.038)	(0.009)
	Within CCT EA Control	-0.155	0.078	-0.0923**	0.014	0.115	-0.058	0.011	0.056	0.011	0.007
		(0.152)	(0.326)	(0.036)	(0.048)	(0.130)	(0.192)	(0.027)	(0.034)	(0.030)	(0.011)
	Within UCT EA Control	-0.150	0.392*	-0.074	0.055	0.012	0.136	0.015	0.003	0.019	0.005
		(0.236)	(0.234)	(0.049)	(0.059)	(0.170)	(0.171)	(0.036)	(0.057)	(0.025)	(0.021)
	Mean in Pure Control:	6.432	0.581	0.343	0.616	15.252	7.479	0.842	0.705	0.797	0.023
	Observations	2,651	2,651	2,651	2,651	2,653	2,652	2,653	2,648	2,653	2,652
PANEL B: Linear Slope Term.											
	Household Size	-0.691	-1.242**	-0.027	-0.273***	-0.176	-0.408	-0.023	0.105	-0.057	0.032
		(0.489)	(0.498)	(0.087)	(0.091)	(0.277)	(0.357)	(0.072)	(0.083)	(0.054)	(0.029)
	UCT	-0.025	0.374	-0.164*	-0.021	0.377	0.613*	-0.031	0.058	-0.063	0.019
		(0.377)	(0.506)	(0.087)	(0.201)	(0.372)	(0.326)	(0.068)	(0.097)	(0.097)	(0.025)
	Within CCT EA Control	-0.176	0.002	-0.104***	-0.006	0.065	0.108	-0.037	0.038	0.001	0.007
		(0.195)	(0.523)	(0.035)	(0.060)	(0.175)	(0.243)	(0.032)	(0.036)	(0.044)	(0.018)
	Within UCT EA Control	0.333	-0.566	-0.054	-0.107	0.329	0.659*	0.082	-0.110	0.0965*	-0.062
		(0.631)	(0.482)	(0.127)	(0.152)	(0.472)	(0.385)	(0.103)	(0.157)	(0.050)	(0.054)
	EA saturation,	0.873	2.068***	-0.082	0.314***	-0.098	0.248	-0.025	-0.141	0.068	-0.037
	CCT Treatment	(0.556)	(0.623)	(0.117)	(0.113)	(0.306)	(0.385)	(0.074)	(0.109)	(0.068)	(0.032)
	EA saturation,	0.348	0.123	0.124	-0.016	-0.340	-0.370	0.044	-0.010	0.066	-0.018
	UCT treatment	(0.425)	(0.524)	(0.121)	(0.219)	(0.437)	(0.451)	(0.081)	(0.119)	(0.126)	(0.029)
	EA saturation,	0.074	0.263	0.039	0.068	0.169	-0.571	0.167**	0.065	0.033	-0.002
	Within-CCT Control	(0.404)	(1.086)	(0.089)	(0.146)	(0.361)	(0.566)	(0.071)	(0.090)	(0.101)	(0.041)
	EA saturation,	-1.061	2.099**	-0.045	0.354	-0.696	-1.148	-0.147	0.249	-0.170	0.147
	Within-UCT Control	(1.130)	(0.809)	(0.325)	(0.311)	(0.979)	(0.781)	(0.181)	(0.250)	(0.132)	(0.155)
	Mean in Pure Control:	6.432	0.581	0.343	0.616	15.252	7.479	0.842	0.705	0.797	0.023
	Observations	2,651	2,651	2,651	2,651	2,653	2,652	2,653	2,648	2,653	2,652

Regressions are OLS using Round 1 data with robust standard errors clustered at the EA level. All regressions are weighted with both sampling and saturation weights to make the results representative of the target population in the study Eas.

TABLE 2. Linear Spillover Analysis

	Dependent Variable:							
	Terms Enrolled		Average Test Score		Ever Married		Ever Pregnant	
	OLS (1)	IV (2)	OLS (3)	IV (4)	OLS (5)	IV (6)	OLS (7)	IV (8)
CCT	0.133 (0.0454)***	0.250 (0.113)**	0.026 (0.00819)***	0.054 (0.0174)***	-0.004 (0.021)	0.001 (0.074)	0.034 (0.028)	0.088 (0.071)
UCT	0.071 (0.049)	0.267 (0.0933)***	0.008 (0.014)	0.009 (0.043)	-0.069 (0.0247)***	-0.144 (0.0482)***	-0.056 (0.0229)**	-0.125 (0.0563)**
Within CCT EA Control	0.018 (0.048)	-0.003 (0.082)	0.022 (0.014)	-0.008 (0.015)	0.009 (0.022)	0.007 (0.034)	0.008 (0.025)	-0.006 (0.032)
Within UCT EA Control	-0.088 (0.077)	0.011 (0.169)	-0.015 (0.024)	-0.068 (0.049)	-0.004 (0.031)	0.036 (0.070)	-0.018 (0.028)	0.078 (0.072)
True CCT treatment saturation, instrumented w/ assigned		-0.238 (0.227)		-0.061 (0.0352)*		-0.011 (0.128)		-0.113 (0.141)
True UCT treatment saturation, instrumented w/ assigned		-0.386 (0.167)**		-0.002 (0.076)		0.154 (0.115)		0.144 (0.110)
True CCT control saturation, instrumented w/ assigned		0.138 (0.324)		0.176 (0.108)		0.009 (0.124)		0.087 (0.166)
True UCT control saturation, instrumented w/ assigned		-0.294 (0.623)		0.162 (0.129)		-0.123 (0.180)		-0.288 (0.191)
Mean in Pure Control:		2.639		0.456		0.176		0.247
Observations	2,579	2,579	2,612	2,612	2,649	2,649	2,650	2,650
Estimates for CCT: (average R3 compliance rate of 77.4%, compliance defined as attending school regularly)								
Intention to Treat (ITT)	0.133 ***		0.026 ***		-0.004		0.034	
Treatment on Uniquely Treated (TUT)	0.250 **		0.054 ***		0.001		0.088	
Spillovers on the Treated (ST)	-0.117		-0.028 *		-0.005		-0.054	
Treatment on Compliers (ToC)	0.167 ***		0.027 **		-0.008		0.042	
Treatment on Unique Complier (TUC)	0.323 **		0.070 ***		0.001		0.113	
Spillover on Compliers (SC)	-0.156		-0.043 **		-0.009		-0.072	
Spillover on Non-Treated (SNT)	0.018		0.022		0.009		0.008	
Total Causal Effect (TCE)	0.082 **		0.021 ***		-0.002		0.016	
Estimates for UCT: (average R3 compliance rate of 99%, compliance defined as receiving transfer)								
Intention to Treat (ITT)	0.071		0.008		-0.069 ***		-0.056 **	
Treatment on Uniquely Treated (TUT)	0.267 ***		0.009		-0.144 ***		-0.125 **	
Spillovers on the Treated (ST)	-0.196 **		-0.001		0.075		0.069	
Treatment on Compliers (ToC)	0.072		0.008		-0.069 ***		-0.056 **	
Treatment on Unique Complier (TUC)	0.270 ***		0.009		-0.145 ***		-0.126 **	
Spillover on Compliers (SC)	-0.197 **		-0.001		0.076		0.070	
Spillover on Non-Treated (SNT)	-0.088		-0.015		-0.004		-0.018	
Total Causal Effect (TCE)	-0.002		-0.003		-0.039 **		-0.036 *	

Test scores are standardized to mean zero and standard deviation one in the control group. Odd-numbered columns are OLS regressions, even-numbered columns are IV, using Round 3 data with robust standard errors clustered at the EA level. All regressions except for the TCE are weighted with both sampling and saturation weights to make the results representative of the target population in the study EAs; TCE regressions use sampling weights only. TCE estimated through separate regression of outcomes on a dummy for treatment at the EA level. Baseline values of the following variables are included as controls: age dummies, strata dummies, household asset index, highest grade attended, and an indicator for ever had sex. Significance levels for cross-equation F-tests calculated using multiple equation two-step GMM estimation. Parameter estimates statistically different than zero at 99 percent (***), 95 percent (**), and 90 percent (*) confidence.

TABLE 3. Granular Spillover Analysis

	Dependent Variable:				
	True Saturation in Cell	Terms Enrolled	Average Test Score	Ever Married	Ever Pregnant
	(1)	(2)	(3)	(4)	(5)
CCT 33%	0.244 (0.0167)***	0.179 (0.0846)**	0.050 (0.0119)***	0.021 (0.056)	0.086 (0.0493)*
CCT 66%	0.372 (0.0487)***	0.185 (0.0548)***	0.019 (0.00881)**	-0.041 (0.0244)*	0.007 (0.043)
CCT 100%	0.657 (0.101)***	0.084 (0.063)	0.019 (0.012)	0.004 (0.021)	0.025 (0.040)
UCT 33%	0.246 (0.0208)***	0.195 (0.0585)***	0.007 (0.029)	-0.105 (0.0305)***	-0.082 (0.0380)**
UCT 66%	0.507 (0.0452)***	-0.003 (0.085)	0.012 (0.017)	-0.076 (0.0348)**	-0.089 (0.0386)**
UCT 100%	0.691 (0.0674)***	0.022 (0.070)	0.007 (0.019)	-0.038 (0.035)	-0.022 (0.026)
Spillover CCT 0%	-0.011 (0.030)	0.000 (0.085)	0.001 (0.014)	0.011 (0.035)	0.010 (0.032)
Spillover CCT 33%	0.247 (0.0163)***	0.027 (0.056)	0.010 (0.008)	-0.004 (0.038)	-0.038 (0.033)
Spillover CCT 66%	0.381 (0.0490)***	0.036 (0.073)	0.066 (0.0376)*	0.017 (0.029)	0.045 (0.053)
Spillover UCT 33%	0.244 (0.0207)***	-0.058 (0.063)	-0.032 (0.029)	0.009 (0.038)	0.013 (0.035)
Spillover UCT 66%	0.510 (0.0447)***	-0.138 (0.163)	0.015 (0.039)	-0.027 (0.046)	-0.071 (0.0389)*
Observations	2,650	2,579	2,612	2,649	2,650
R-Squared	0.820	0.099	0.418	0.144	0.200

Test scores have been standardized to have a mean of zero and a standard deviation of one in the control group. Regressions are OLS models using Round 3 data with robust standard errors clustered at the EA level. All regressions are weighted with both sampling and saturation weights to make the results representative of the target population in the study EAs. Baseline values of the following variables are included as controls in the regression analyses: age dummies, strata dummies, household asset index, highest grade attended, and an indicator for ever had sex. Parameter estimates statistically different than zero at 99 percent (***), 95 percent (**), and 90 percent (*) confidence.

TABLE 4. Spillover Analysis in Social Networks

	Terms Enrolled	Average Test Scores	Ever Married	Ever Pregnant
	(1)	(2)	(4)	(5)
CCT	0.164 (0.060)***	0.034 (0.024)	-0.017 (0.030)	0.012 (0.037)
UCT	0.122 (0.071)*	0.036 (0.025)	-0.075 (0.024)***	-0.065 (0.026)**
Within-Village Control CCT	0.01 (0.063)	0.016 (0.024)	0.015 (0.030)	-0.003 (0.028)
Within-Village Control UCT	-0.12 (0.114)	-0.003 (0.019)	0.005 (0.039)	-0.019 (0.048)
Number of Treated Friends for CCT Treatment Girls	-0.011 (0.055)	-0.021 (0.016)	0.006 (0.023)	0.015 (0.028)
Number of Treated Friends for UCT Treatment Girls	-0.031 (0.092)	-0.029 (0.019)	-0.006 (0.025)	-0.01 (0.024)
Number of Treated Friends for CCT Untreated Girls	0.109 (0.106)	0.008 (0.017)	-0.025 (0.047)	0.044 (0.058)
Number of Treated Friends for UCT Untreated Girls	0.196 (0.099)*	-0.002 (0.028)	-0.065 (0.042)	-0.052 (0.058)
Number of Treated Friends for Pure Control Girls	0.021 (0.145)	-0.032 (0.023)	0.036 (0.052)	0.084 (0.076)
Number of friends who are dropouts	-0.15 (0.047)***	-0.029 (0.008)***	0.085 (0.022)***	0.125 (0.022)***
Number of friends in same cluster	0.004 (0.015)	-0.02 (0.004)***	0.004 (0.006)	-0.006 (0.006)
1 Matched Friend	-0.008 (0.046)	-0.025 (0.009)***	0.037 (0.019)**	0.027 (0.022)
2 Matched Friends	0.001 (0.073)	0.003 (0.014)	0.055 (0.032)*	0.049 (0.032)
3 Matched Friends	-0.015 (0.098)	-0.003 (0.017)	0.087 (0.038)**	0.065 (0.041)
4 Matched Friends	-0.123 (0.196)	0.011 (0.030)	0.08 (0.074)	0.08 (0.089)
5 Matched Friends	0.227 (0.133)*	-0.016 (0.042)	-0.148 (0.057)**	-0.253 (0.066)***
Constant	2.639 (0.049)***	0.502 (0.012)***	0.133 (0.019)***	0.219 (0.019)***
Observations	2660	2620	2652	2653
R-squared	0.013	0.062	0.023	0.023

Analysis performed within social networks, as defined by the 'five closest friends' listed by core respondents at baseline. Regressions are OLS models with robust standard errors clustered at the EA level. All regressions are weighted with both sampling and saturation weights to make the results representative of the target population in the study EAs. Parameter estimates statistically different than zero at 99 percent (***), 95 percent (**), and 90 percent (*) confidence.

TABLE 5. Robustness check using cross-EA variation in treatment intensity

	Dependent Variable:											
	Terms Enrolled	Average Test Score	Ever Married	Ever Pregnant	(1)	(2)	(3)	(4)	(7)	(8)	(9)	(10)
CCT	0.119 (0.0431)***	0.126 (0.085)	0.022 (0.00896)**	0.008 (0.015)	0.000 (0.023)	-0.023 (0.044)	0.040 (0.026)	-0.011 (0.041)				
UCT	0.059 (0.050)	0.052 (0.111)	0.005 (0.013)	-0.019 (0.018)	-0.064 (0.0269)**	-0.090 (0.0484)*	-0.057 (0.0240)**	-0.114 (0.0508)**				
Within CCT EA Control	0.013 (0.047)	0.016 (0.047)	0.021 (0.014)	0.023 (0.0134)*	0.010 (0.023)	0.011 (0.023)	0.008 (0.026)	0.008 (0.025)				
Within UCT EA Control	-0.100 (0.074)	-0.095 (0.077)	-0.020 (0.023)	-0.015 (0.023)	0.000 (0.034)	0.002 (0.036)	-0.021 (0.029)	-0.020 (0.031)				
# of treated EAs within 3 km	-0.021 (0.018)	-0.020 (0.020)	-0.005 (0.005)	-0.002 (0.006)	0.005 (0.009)	0.005 (0.012)	0.004 (0.009)	0.003 (0.011)				
# of treated EAs between 3 & 6 km	0.010 (0.013)	0.019 (0.016)	0.001 (0.003)	0.006 (0.004)	-0.004 (0.006)	-0.002 (0.007)	-0.005 (0.006)	-0.003 (0.008)				
# of total EAs within 3 km	0.012 (0.012)	0.011 (0.013)	0.006 (0.00281)**	0.004 (0.004)	-0.003 (0.006)	-0.004 (0.007)	0.001 (0.006)	0.002 (0.007)				
# of total EAs between 3 & 6 km	-0.004 (0.007)	-0.008 (0.008)	-0.002 (0.002)	-0.004 (0.00216)*	0.000 (0.003)	-0.001 (0.004)	0.004 (0.003)	0.002 (0.004)				
Treated individual * # of treated EAs within 3 kilometers		0.003 (0.021)		0.005 (0.004)		-0.004 (0.011)		-0.007 (0.012)				
Treated individual * # of treated EAs between 3 and 6 kilometers		0.001 (0.040)		-0.006 (0.008)		0.013 (0.022)		0.009 (0.023)				
Treated individual * # of total EAs within 3 kilometers		-0.029 (0.026)		-0.014 (0.00467)***		-0.007 (0.015)		-0.004 (0.015)				
Treated individual * # of total EAs between 3 and 6 kilometers		0.012 (0.014)		0.007 (0.00276)**		0.004 (0.008)		0.007 (0.007)				
Observations	2,579	2,579	2,612	2,612	2,649	2,649	2,650	2,650				
R-squared	0.098	0.098	0.418	0.42	0.144	0.144	0.199	0.2				

Regressions are OLS models using Round 3 data with robust standard errors clustered at the EA level. All regressions are weighted with both sampling and saturation weights to make the results representative of the target population in the study EAs. Baseline values of the following variables are included as controls in the regression analyses: age dummies, strata dummies, household asset index, highest grade attended, and an indicator for ever had sex. Parameter estimates statistically different than zero at 99 percent (***), 95 percent (**), and 90 percent (*) confidence.

```

%PROGRAM TO ESTIMATE POWER OF A RANDOMIZED SATURATION STUDY
%AUTHOR: Aislinn Bohren
%SUPPLEMENTAL MATERIAL TO "DESIGNING EXPERIMENTS TO MEASURE SPILLOVER
%EFFECTS" By S. Baird, A. Bohren, C. McIntosh, B. Ozler

```

```

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

```

```

%USER INPUT
clear;clc;

```

```

n=20; %cluster size
C=5; %number of clusters
tau=0.5; %variance of cluster error
sigma=2; %variance of individual error
alpha=0.05 ; %significance
gamma=.9; %power

```

```

pi=[0,1/3,2/3]; %saturation bins
f=[1/3,1/3,1/3]; %distribution over bins

```

```

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

```

```

%CALCULATIONS
varN=tau+sigma;
varCo=(n-1)*tau;
varC=n*tau+sigma;

```

```

%Note: eta=E[pi^2] here; so eta in paper is eta-mu^2

```

```

x=1;
powerT=zeros(1,length(x));
powerS=zeros(1,length(x));
powerTonly=zeros(1,length(x));
powerSonly=zeros(1,length(x));
powerT1=zeros(1,length(x));
powerT2=zeros(1,length(x));
powerS1=zeros(1,length(x));
powerS2=zeros(1,length(x));
MDSE_T=zeros(1,length(x));
MDSE_S=zeros(1,length(x));
MDSE_TAffine=zeros(1,length(x));
MDSE_SAffine=zeros(1,length(x));
MU=zeros(1,length(x));
ETA=zeros(1,length(x));
etaT=zeros(1,length(x));

```

```

for j=1:length(x);
%Calculate distribution statistics
mu_ind=zeros(1,length(pi));
p=zeros(1,length(pi));
eta_ind=zeros(1,length(pi));
rho=zeros(1,length(pi));
c_ind=zeros(1,length(pi));
d_ind=zeros(1,length(pi));

```

```

for i=1:length(pi);
mu_ind(i)=pi(i) *f(i);
eta_ind(i)=pi(i)^2 *f(i);
c_ind(i)=pi(i)^3 *f(i);
d_ind(i)=pi(i)^4 *f(i);
p(i)=(1-pi(i))*f(i);
rho(i)=(1-pi(i))^2 * f(i);
end;

```

```

mu=sum(mu_ind);
eta=sum(eta_ind);
c=sum(c_ind);
d=sum(d_ind);
psi=0;if pi(1)==0; psi=f(1);end;
MU(j)=mu;
ETA(j)=eta;

```

```

%PROGRAM TO ESTIMATE POWER OF A RANDOMIZED SATURATION STUDY
%AUTHOR: Aislinn Bohren
%SUPPLEMENTAL MATERIAL TO "DESIGNING EXPERIMENTS TO MEASURE SPILLOVER
%EFFECTS" By S. Baird, A. Bohren, C. McIntosh, B. Ozler

```

```

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

```

```

%USER INPUT
clear;clc;

```

```

n=20; %cluster size
C=5; %number of clusters
tau=0.5; %variance of cluster error
sigma=2; %variance of individual error
alpha=0.05 ; %significance
gamma=.9; %power

```

```

pi=[0,1/3,2/3]; %saturation bins
f=[1/3,1/3,1/3]; %distribution over bins

```

```

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

```

```

%CALCULATIONS
varN=tau+sigma;
varCo=(n-1)*tau;
varC=n*tau+sigma;

```

```

%Note: eta=E[pi^2] here; so eta in paper is eta-mu^2

```

```

x=1;
powerT=zeros(1,length(x));
powerS=zeros(1,length(x));
powerTonly=zeros(1,length(x));
powerSonly=zeros(1,length(x));
powerT1=zeros(1,length(x));
powerT2=zeros(1,length(x));
powerS1=zeros(1,length(x));
powerS2=zeros(1,length(x));
MDSE_T=zeros(1,length(x));
MDSE_S=zeros(1,length(x));
MDSE_TAffine=zeros(1,length(x));
MDSE_SAffine=zeros(1,length(x));
MU=zeros(1,length(x));
ETA=zeros(1,length(x));
etaT=zeros(1,length(x));

```

```

for j=1:length(x);
%Calculate distribution statistics
mu_ind=zeros(1,length(pi));
p=zeros(1,length(pi));
eta_ind=zeros(1,length(pi));
rho=zeros(1,length(pi));
c_ind=zeros(1,length(pi));
d_ind=zeros(1,length(pi));

```

```

for i=1:length(pi);
mu_ind(i)=pi(i) *f(i);
eta_ind(i)=pi(i)^2 *f(i);
c_ind(i)=pi(i)^3 *f(i);
d_ind(i)=pi(i)^4 *f(i);
p(i)=(1-pi(i))*f(i);
rho(i)=(1-pi(i))^2 * f(i);
end;

```

```

mu=sum(mu_ind);
eta=sum(eta_ind);
c=sum(c_ind);
d=sum(d_ind);
psi=0;if pi(1)==0; psi=f(1);end;
MU(j)=mu;
ETA(j)=eta;

```