Online supplementary material: Policy design in experiments with unknown interference

Davide Viviano[*] Jess Rudder[†]

October, 2023

This online Appendix contains additional extensions and results not included in the main Appendix.

Appendix E Additional Extensions

E.1 Staggered adoption

In this section, we sketch the experimental design in the presence of staggered adoption, i.e., when treatments are assigned only once to individuals and post-treatment outcomes are collected once. The algorithm works similarly to what was discussed in Section 4 with one small difference: every period, we only collect information from a given clusters' pair and update the policy for the subsequent pair and proceed in an iterative fashion.

[*]Department of Economics, Harvard University. Correspondence: dviviano@fas.harvard.edu.

[†]Department of Economics, University of Chicago. Email: jrudder@uchicago.edu

Theorem E.1 (In-sample regret)**.** *Let the conditions in Theorem 4.3 hold and let* $\beta \in \mathbb{R}$*, with* $\check{\beta}^t$ *estimated as in Algorithm E.1. Then* $P\Big(\frac{1}{T}\sum_{t=1}^T \Big[W(\beta^*) - W(\check{\beta}^t)\Big] \le \bar{C}\frac{p\log(T)}{T}\Big) \ge 1 - 1/n$ *for a finite constant* $\bar{C} < \infty$*.*

See Appendix F.2 for the proof. The disadvantage of the staggered adoption case is that we cannot control the in-sample regret *worst-case* over all clusters as in Section 4, but only the average regret across clusters.

Algorithm E.1 Adaptive Experiment with staggered adoption

Require: Starting value $\beta \in \mathbb{R}$, K clusters, $T+1$ periods of experimentation.

1: Create pairs of clusters $\{k, k+1\}, k \in \{1, 3, \cdots, K-1\}$;

2: $t = 0$:

a: For n units in each cluster observe the baseline outcome $Y_{i,0}^{(h)}, h \in \{1, \cdots, K\}, \check{\beta}^0 = \beta$.

b: Initalize a gradient estimate $\widehat{M}_t = 0$

3: **while** $1 \le t \le T$ **do**

a: Sample without replaceent one pair of clusters $\{k, k+1\}$ not observed in previous iterations;

b: Define $\check{\beta}^t = \check{\beta}^{t-1} + \alpha_t \widehat{M}_t$;

c: Assign treatments as

$$D_{i,t}^{(h)} \sim \pi(1, \beta_t), \quad \beta_t = \begin{cases} \check{\beta}^t + \eta_n \text{ if } h \text{ is even} \\ \check{\beta}^t - \eta_n \text{ if } h \text{ is odd} \end{cases}, \quad n^{-1/2} < \eta_n \le n^{-1/4}$$

d: For n units in each cluster $h \in \{1, \cdots, K\}$ observe $Y_{i,t}^{(h)}$.

4: **end while**

5: Return $\hat{\beta}^* = \check{\beta}^T$

E.2 Extensions for the network formation model

Consider the following equation:

$$(X_i^{(k)}, U_i^{(k)}) \sim_{i.i.d.} F_{U|X} F_X, \quad A_{i,j}^{(k)} = l\Big(X_i^{(k)}, X_j^{(k)}, U_i^{(k)}, U_j^{(k)}, \omega_{i,j}^{(k)}\Big) 1\{i_k \leftrightarrow j_k\} \tag{E.1}$$

where $\omega_{i,j}^{(k)} \Big| \Big\{\omega_{u,v}^{(k)}\Big\}_{(u,v)\neq(i,j),(u,v)\neq(j,i)}, X^{(k)}, U^{(k)}, \nu^{(k)} \sim_{i.i.d.} F_\omega$.

Intuitively, Equation (E.1) states that the connections form also based on unobservables $\omega_{i,j}$ which are drawn independently for each pair (i, j) (note we can have $\omega_{i,j} = \omega_{j,i}$). We can now state the following lemma.

Lemma E.2 (Outcomes)**.** *Under Assumption 2.1, with a network formation as in Equation (E.1), Assumption 2.2, under an assignment in Assumption 2.3 with exogenous (i.e., not data-dependent)* $\beta_{k,t}$*, Lemma B.4 (and so Lemma 2.1) hold.*

The proof is in Appendix F.3. Lemma E.2 implies that our results (and derivations) for inference and estimation directly extend with network formation as in Equation (E.1). One exception is the theorem in Section 5 which instead holds under some minor modification to the proof, which we omit for brevity.[1]

Finally, it is also possible to extend Lemma 2.1 to settings where the network also depends on others' unobservables. Specifically, consider the following extension:

$$(X_i^{(k)}, U_i^{(k)}) \sim_{i.i.d.} F_{U|X}F_X, \quad A_{i,j}^{(k)} = l\Big(X_i^{(k)}, X_j^{(k)}, U_i^{(k)}, U_j^{(k)}, U_{v:1\{i_k \sim v_k\}=1}, U_{v':1\{j_k \sim v'_k\}=1}\Big)1\{i_k \leftrightarrow j_k\}. \tag{E.2}$$

Equation (E.2) states that the connection between individual i and j can depend not only on unobservables of i and j, but also unobservables of all the possible connections of both individual i and individual j. Following a similar argument of Lemma B.4, we can show

$$Y_{i,t}^{(k)} = y\Big(X_i^{(k)}, \beta_{k,t}\Big) + \varepsilon_{i,t}^{(k)} + \alpha_t + \tau_k, \quad \mathbb{E}_{\beta_{k,t}}\Big[\varepsilon_{i,t}^{(k)}|X_i^{(k)}\Big] = 0, \tag{E.3}$$

for some function $y(\cdot)$ unknown to the researcher. Therefore, our results for estimation and inference hold also in this setting.

E.3 Selection of η_n: rule of thumb

In this subsection we provide a rule of thumb for selecting η_n. Following Theorem 3.1 and following Lemma B.3 which provides exact constants, with probability at least $1 - 1/n$

$$\Big|\widehat{M}_{(k,k+1)} - M(\beta)\Big| \le \sqrt{\frac{2\sigma^2\gamma_N \log(2\gamma_N n)}{n\eta_n^2}} + c\eta_n, \quad c = \Big|\Big|\frac{\partial^2 m(d,x,\beta)}{\partial\beta^2}\Big|\Big|_\infty,$$

where the l_∞ is taken with respect to each element of the Hessian, x, β.[2] we cannot directly minimize the upper bound since otherwise we would violate the condition that $\eta_n = o(n^{-1/4})$. Instead, we minimize $\min_{\eta_n} \sqrt{\frac{2\sigma^2\gamma_N\log(2\gamma_N n)}{n\eta_n^2}} + c\eta_n/s_n^2$, $s_n = o(1)$, where s_n penalizes the bias by an $o(1)$ component and chosen below. It follows that for given penalization s_n the minimizer of the expression is $\eta_n^2 = \sqrt{\frac{2s_n\sigma^2\gamma_N\log(2\gamma_N n)}{nc}}$. Let $s_n = \gamma_N/(n^{1/4}\log(n\gamma_N))$, assumed to be $o(1)$ for inference by assumption. we can then write the solution to the optimization problem as $\gamma_N\sqrt{\frac{2\sigma^2}{c}}n^{-5/16} \approx \gamma_N\sqrt{\frac{2\sigma^2}{c}}n^{-1/3}$. Here, we can replace σ^2 and c with some out-of-sample estimates of the outcomes' variance and curvature. Whenever the researcher does

[1] In particular, the difference is in one step of the proof where we need to show concentration of $\frac{1}{\gamma_N}\sum_{i=1}^N A_{i,j}$ around its expectation, here also taking into account also the component $\omega_{i,j}$. Under the assumption that $\omega_{i,j}$ are independent for all $j \in \{1, \cdots, N\}$ the argument in the proof, i.e., concentration of the edges conditional on (X_i, U_i) directly follows also for this step, since conditional on X_i, U_i, we still obtain independence of $A_{i,j}$, $j \in \{1, \cdots, N\}$, and we can then follow the same argument in the derivation.

[2] The constants for the upper bound for the variance follow from Lemma B.3, while the component $c\eta_n$ captures the bias obtained from a second-order Taylor expansion to $m(\cdot)$.

not have a good guess for γ_N, we recommend choosing $\eta_n = \sqrt{\frac{2\sigma^2}{c}} n^{-1/3}$ (without the term γ_N) which also leads to valid inference, but slightly smaller small-sample bias (and larger small-sample variance) than the optimal choice. Finally, since researcher may impose small sample upper bound on the bias, the suggested rule of thumb is

$$\eta_n = \begin{cases} \sqrt{\frac{2\sigma^2}{c}} n^{-1/3} & \text{if } \sqrt{\frac{2\sigma^2}{c}} n^{-1/3} \leq H \\ H \text{ otherwise} \end{cases}$$

where Hc denotes an upper bound on the bias of the estimator imposed by the researcher.

Appendix F Additional proofs for the extensions

F.1 Proof of Theorem A.8

In this subsection, we derive the theorem for the gradient descent method under Assumption A.5. The derivation is split into the following lemmas.

Definition F.1 (Oracle gradient descent). We define for positive constants $\infty > \mu, \kappa > 0$, κ as defined in Lemma F.1, arbitrary $v \in (0,1)$, $\alpha_w = \frac{J}{\check{T}^{1/2-v/2}||M(\beta^*_{w-1})||}$, $J < 1$

$$\beta^*_w = \begin{cases} P_{\mathcal{B}_1, \mathcal{B}_2}\Big[\beta^*_{w-1} + \alpha_{w-1} M(\beta^*_{w-1})\Big] \text{ if } ||M(\beta^*_w)||_2 \geq \frac{\kappa}{\mu \check{T}^{1/2-v/2}} \\ \beta^*_{w-1} \text{ otherwise} \end{cases} \quad \beta^*_1 = \beta_0, \tag{F.1}$$

Lemma F.1 (Adaptive gradient descent for quasi-concave functions and locally strong concave). *Let $\mathcal{B}$ be compact. Define $G = \max\{\sup_{\beta \in \mathcal{B}} 2||\beta||^2, 1\}$. Let Assumption 3.1, 4.1, A.5 hold. Let κ be a positive finite constant, defined as in Equation (F.2). Then for any $v \in (0,1)$, for $\check{T} \geq ((G+1)/J)^{1/v}$, the following holds: $||\beta^*_{\check{T}} - \beta^*||^2 \leq \kappa \check{T}^{-1+v}$.*

Proof of Lemma F.1. To prove the statement, we use properties of gradient descent methods with gradient norm rescaling (Hazan et al., 2015), with modifications to the original arguments to explicitely obtain a rate T^{-1+v} for an arbitrary small v.

Preliminaries Clearly, if the algorithm terminates at w, under Assumption A.5 (B), this implies that $||\beta^*_w - \beta^*||_2^2 \leq \kappa \check{T}^{-1+v}$, proving the claim. Therefore, assume that the algorithm did not terminate at time w. This implies that for any $\check{w} \geq 1$, $||\beta^*_{\check{w}} - \beta^*||_2^2 > \kappa \check{T}^{-1+v}$. Define $\epsilon = \check{T}^{-1+v}$ and let ∇_w be the gradient evaluated at β^*_w. For every $\beta \in \mathcal{B}$, define $H(\beta)\Big|_{[\beta^*, \beta]}$ the Hessian evaluated at some point $\tilde{\beta} \in [\beta^*, \beta]$, such that

$$W(\beta) = W(\beta^*) + \frac{1}{2}(\beta - \beta^*)^\top H(\beta)\Big|_{[\beta^*,\beta]}(\beta - \beta^*),$$

which always exists by the mean-value theorem and differentiability of the objective function. Define

$$\frac{1}{2}(\beta - \beta^*)^\top H(\beta)\Big|_{[\beta^*, \beta]}(\beta - \beta^*) = f(\beta) \leq 0,$$

where the inequality follows by definition of β^* (note that $f(\beta)$ also depends on $\tilde{\beta}$, whose dependence we implicitely suppressed). Finally, note that $-|\lambda_{max}|||\beta - \beta^*||^2 \leq f(\beta) \leq -|\lambda_{\min}|||\beta - \beta^*||^2$ for constants $\lambda_{\max} > \lambda_{\min} > 0$. The lower bound follows directly by Assumption 3.1, while the upper bound follows directly from Assumption A.5 (C).

Cases Define

$$\kappa = \frac{|\lambda_{max}|}{|\lambda_{min}|} \geq 1. \tag{F.2}$$

Observe now that if $||\beta_w^* - \beta^*||^2 \leq \epsilon\kappa$, the claim trivially holds. Therefore, consider the case where $||\beta_w^* - \beta^*||^2 > \epsilon\kappa$.

Comparisons within the neighborhood Take $\tilde{\beta} = \beta^* - \sqrt{\epsilon}\frac{\nabla_w}{||\nabla_w||_2}$. Observe that

$$\begin{aligned} W(\tilde{\beta}) - W(\beta_w^*) &= \frac{1}{2}(\tilde{\beta} - \beta^*)^\top H(\tilde{\beta})\Big|_{[\beta^*, \tilde{\beta}]}(\tilde{\beta} - \beta^*) - \frac{1}{2}(\beta_w^* - \beta^*)^\top H(\beta_w^*)\Big|_{[\beta^*, \beta_w^*]}(\beta_w^* - \beta^*) \\ &\geq -|\lambda_{\max}|\epsilon + |\lambda_{\min}|\epsilon\kappa = 0. \end{aligned}$$

As a result, for all $\beta_w^* : ||\beta_w^* - \beta^*||^2 > \epsilon\kappa$, using quasi-concavity

$$\nabla_w^\top(\tilde{\beta} - \beta_w^*) \geq 0 \Rightarrow \nabla_w^\top(\beta^* - \beta_w^*) \geq \sqrt{\epsilon}||\nabla_w||_2 \tag{F.3}$$

Plugging in the above expression in the definition of β_w^* By construction of the algorithm, we write

$$||\beta^* - \beta_{w+1}^*||^2 \leq ||\beta^* - \beta_w^*||^2 - 2\alpha_w J \nabla_w^\top(\beta^* - \beta_w^*) + J^2\alpha_w^2||\nabla_w||^2.$$

By Equation (F.3), we can write $||\beta^* - \beta_{w+1}^*||^2 \leq ||\beta^* - \beta_w^*||^2 - 2J\alpha_w\sqrt{\epsilon}||\nabla_w||_2 + J^2\alpha_w^2||\nabla_w||^2$. Plugging in the expression for α_w, and using the fact that $J \leq 1$, we have $||\beta^* - \beta_{w+1}^*||^2 \leq ||\beta^* - \beta_w^*||^2 - J\epsilon$.

Recursive argument Recall that since the algorithm did not terminate, $||\beta^* - \beta_{\check{w}}^*||^2 > \epsilon\kappa$, for all $\check{w} \leq w$. Using this argument recursively, we obtain

$$||\beta^* - \beta_{\check{T}}^*||^2 \leq ||\beta^* - \beta_0||^2 - J\sum_{s=1}^{\check{T}}\epsilon = 2\max_{\beta \in \mathcal{B}}||\beta||^2 - J\check{T}^v \leq G + 1 - J\check{T}^v.$$

Whenever $\check{T} > (G/J + 1/J)^{1/v}$, we have a contradiction. The proof completes. □

Lemma F.2. *Let Assumption 2.1, 2.2, 4.1, A.5 hold. Assume that*

$$\epsilon_n \geq \sqrt{p}\Big[\bar{C}\sqrt{\gamma_N \frac{\log(\gamma_N \check{T} K/\delta)}{\eta_n^2 n}} + \eta_n\Big], \quad \frac{1}{4\mu \check{T}^{1/2-v/2}} - \epsilon_n \geq 0$$

for a finite constant $\bar{C} < 0$.

Then, with probability at least $1-\delta$, for any $\delta \in (0,1)$ for any $w \leq \check{T}$,

$$\text{either (i)}\Big|\Big|\check{\beta}_k^w - \beta_w^*\Big|\Big|_\infty = \mathcal{O}(P_w(\delta) + p\eta_n), \text{ or (ii)}\Big|\Big|\check{\beta}_k^w - \beta^*\Big|\Big|_2^2 \leq \frac{p}{\check{T}^{1-v}}$$

where $P_1(\delta) = \mathrm{err}(\delta)$ and $P_w(\delta) = \frac{2\sqrt{p}}{\nu_n} Bp \frac{1}{\check{T}^{1/2-v/2}} P_{w-1}(\delta) + P_{w-1}(\delta) + \frac{2\sqrt{p}}{\nu_n}\frac{1}{\check{T}^{1/2-v/2}}\mathrm{err}(\delta)$, for a finite constant $B < \infty$, and $\mathrm{err}(\delta) \leq c_0\Big(\sqrt{\gamma_N \frac{\log(\gamma_N p \check{T} K/\delta)}{\eta_n^2 n}} + p\eta_n\Big)$, with $\nu_n = \frac{1}{\mu \check{T}^{1/2-v/2}} - 2\epsilon_n$, and a finite constant $c_0 < \infty$.

Proof of Lemma F.2. First, by Lemma 4.1, the estimated coefficients are exogenous. Hence, by invoking Lemma B.7 and the union bound, we can write for every k and t, $\delta \in (0,1)$, $|\check{M}_{k,w}^{(j)} - M^{(j)}(\check{\beta}_{k+2}^w)| \leq c_0\Big(\sqrt{\gamma_N \frac{\log(\gamma_N K\check{T}/\delta)}{\eta_n^2 n}} + \eta_n\Big)$. We now proceed by induction. We first prove the statement, assuming that the constraint is always attained. We then discuss the case of the constrained solution. Define $B = \sup_\beta \Big|\Big|\frac{\partial^2 W(\beta)}{\partial \beta^2}\Big|\Big|_\infty$.

Unconstrained case Consider $w = 1$. Then since all clusters start from the same starting point β_0 recall that ($\beta_1^* = \beta_0$), we can write with probability $1-\delta$, by the union bound over p (which hence enters in the $\log(p)$ component of err_n) and Lemma B.7 $\Big|\Big|\check{M}_{k,1} - M(\beta_1^*)\Big|\Big|_\infty \leq \mathrm{err}(\delta)$. Consider now the case where the algorithm stops. This implies that it must be that $||\check{M}_{k,1}||_2 \leq \frac{1}{\mu \check{T}^{1/2-v/2}} - \epsilon_n$. By Lemma B.7

$$||M(\beta_1^*)||_2 \leq ||\check{M}_{k,1}||_2 + \sqrt{p}\mathrm{err}(\delta) \leq \frac{1}{\mu \check{T}^{1/2-v/2}} - \epsilon_n + \sqrt{p}\mathrm{err}(\delta) \leq \frac{1}{\mu \check{T}^{1/2-v/2}}. \tag{F.4}$$

since $\epsilon_n \geq \sqrt{p}\mathrm{err}(\delta)$. As a result, also the oracle algorithm stops at β_1^* by construction of ϵ_n. Suppose the algorithm does not stop. Then it must be that $||\check{M}_{k,1}|| \geq \frac{1}{\mu \check{T}^{1/2-v/2}} - \epsilon_n$ and

$$||V_1(\beta_1^*)|| \geq \frac{1}{\mu \check{T}^{1/2-v/2}} - \epsilon_n - \sqrt{p}\mathrm{err}_1 \geq \frac{1}{\mu \check{T}^{1/2-v/2}} - 2\epsilon_n := \nu_n > 0.$$

Observe now that

$$\begin{aligned} \Big|\Big|\frac{\check{M}_{k,1}}{||\check{M}_{k,1}||_2} - \frac{M(\beta_1^*)}{||M(\beta_1^*)||_2}\Big|\Big|_\infty &\leq \Big|\Big|\frac{\check{M}_{k,1} - M(\beta_1^*)}{||M(\beta_1^*)||_2}\Big|\Big|_\infty + \Big|\Big|\frac{\check{M}_{k,1}(||\check{M}_{k,1}||_2 - ||M(\beta_1^*)||_2)}{||M(\beta_1^*)||_2 ||\check{M}_{k,1}||_2}\Big|\Big|_\infty \\ &\leq \Big|\Big|\frac{\check{M}_{k,1} - M(\beta_1^*)}{||M(\beta_1^*)||_2}\Big|\Big|_\infty + \sqrt{p}\Big|\Big|\frac{\check{M}_{k,1} - M(\beta_1^*)}{||M(\beta_1^*)||_2}\Big|\Big|_\infty. \end{aligned} \tag{F.5}$$

The last inequality follows from the reverse triangular inequalities and standard properties of the norms. Then with probability at least $1-\delta$, for any $\delta \in (0,1)$ (F.5) $\leq \frac{1}{\nu_n} \times 2\sqrt{p}\mathrm{err}(\delta)$.

completing the claim for $w = 1$. Consider now a general w. Define the error until time $w-1$ as P_{w-1}.Then for every $j \in \{1, \cdots, p\}$, by Assumption 3.1, we have with probability at least $1 - w\delta$ (using the union bound), and letting $\mathbf{1}_p = [1, \cdots, 1] \in \mathbb{R}^p$,

$$\begin{aligned}
&\check{M}_{k,w}^{(j)} = M^{(j)}(\check{\beta}_{k+2}^{w}) + \mathrm{err}(\delta) = M^{(j)}(\beta_w^* + \mathbf{1}_p P_w(\delta)) + \mathrm{err}(\delta) \\
&\Rightarrow \Big|\Big|\check{M}_{k,w} - M(\beta_w^*)\Big|\Big|_\infty \leq BpP_w(\delta) + \mathrm{err}(\delta),
\end{aligned}$$

where the above inequality follows by the mean-value theorem and Assumption 3.1. Suppose now that $||\check{M}_{k,w}||_2 \leq \frac{1}{\mu \check{T}^{1/2-v/2}} - \epsilon_n$. Then for the same argument as in Equation (F.4), we have $||M(\check{\beta}_k^w)||_2 \leq \frac{1}{\mu \check{T}^{1/2-v/2}}$. Under Assumption A.5 (B) this implies that $||\check{\beta}_k^w - \beta^*||_2^2 \leq \frac{1}{\check{T}^{1-v}}$, which proves the statement. Suppose instead that the algorithm does not stop. Then we can write by the induction argument

$$\Big|\Big|\check{\beta}_k^w + \frac{1}{\check{T}^{1/2-v/2}}\frac{\check{M}_{k,w}}{||\check{M}_{k,w}||_2} - \beta_w^* - \frac{1}{\check{T}^{1/2-v/2}}\frac{M(\beta_w^*)}{||M(\beta_w^*)||_2}\Big|\Big|_\infty \leq P_w(\delta) + \frac{1}{\check{T}^{1/2-v/2}}\underbrace{\Big|\Big|\frac{\check{M}_{k,w}}{||\check{M}_{k,w}||_2} - \frac{M(\beta_w^*)}{||M(\beta_w^*)||_2}\Big|\Big|_\infty}_{(B)}. \tag{F.6}$$

Using the same argument in Equation (F.5), we have with probability at least $1 - \delta$, $(B) \leq \frac{2\sqrt{p}}{\nu_n}\Big[\mathrm{err}(\delta) + BpP_w(\delta)\Big]$, which completes the proof for the unconstrained case. The $\check{T}$ component in the error expression follows from the union bound across all $\check{T}$ events.

Constrained case Since the statement is true for $w = 1$, we can assume that it is true for all $s \leq w-1$ and prove the statement by induction. Since $\mathcal{B}$ is a compact space, we can write

$$\begin{aligned}
&\Big|\Big|P_{\mathcal{B}_1,\mathcal{B}_2-\eta_n}\Big[\sum_{s=1}^{w}\alpha_{k,s}\check{M}_{k,s}\Big] - P_{\mathcal{B}_1,\mathcal{B}_2}\Big[\sum_{s=1}^{w}\alpha_s M(\beta_s^*)\Big]\Big|\Big|_\infty \\
&\leq \Big|\Big|P_{\mathcal{B}_1,\mathcal{B}_2-\eta_n}\Big[\sum_{s=1}^{w}\alpha_{k,s}\check{M}_{k,s}\Big] - P_{\mathcal{B}_1,\mathcal{B}_2-\eta_n}\Big[\sum_{s=1}^{w}\alpha_s M(\beta_s^*)\Big]\Big|\Big|_\infty + p\mathcal{O}(\eta_n) \\
&\leq 2\Big|\Big|\sum_{s=1}^{w}\alpha_{k,s}\check{M}_{k,s} - \sum_{s=1}^{w}\alpha_s M(\beta_s^*)\Big|\Big|_\infty + p\mathcal{O}(\eta_n).
\end{aligned}$$

For the first component in the last inequality, we follow the same argument as above. □

Lemma F.3. *Let the conditions in Lemma F.2 hold. Then with probability at least $1-\delta$, for any $k \in \{1, \cdots, K\}$, for any $v \in (0,1), \delta \in (0,1)$, $\check{T} \geq \zeta^{1/v}$,*

$$||\beta^* - \check{\beta}_k^{\check{T}}||_2^2 \leq \frac{\kappa}{\check{T}^{1-v}} + \check{T}e^{Bp\sqrt{p}\check{T}} \times c_0\Big(\gamma_N\frac{\log(p\gamma_N\check{T}K/\delta)}{\eta_n^2 n} + p^2\eta_n^2\Big),$$

with $0 < \zeta, \kappa, B < \infty$ being constants independent on $(n, \check{T})$ and ϵ_n as defined in Lemma F.2, and a finite constant $c_)$ $< \infty$.

Proof. We invoke Lemma F.2. Observe that we only have to check that the result holds for (i) in Lemma F.2, since otherwise the claim trivially holds. Using the triangular inequality, we can write $||\beta^* - \check{\beta}_k^{\check{T}}||_2^2 \leq ||\beta^* - \beta^*_{\check{T}}||_2^2 + ||\check{\beta}_k^{\check{T}} - \beta^*_{\check{T}}||_2^2$. The first component on the right-hand side is bounded by Lemma F.1, with $\check{T} \geq \zeta^{1/v}, \zeta$ being a constant defined in Lemma F.1.

Using Lemma F.2, we bound with probability at least $1-\delta$, the second component as follows $||\check{\beta}_k^{\check{T}} - \beta^*_{\check{T}}||_2^2 \leq p||\check{\beta}_k^{\check{T}} - \beta^*_{\check{T}}||_\infty^2 = p \times \mathcal{O}(P_{\check{T}}^2(\delta))$. We conclude the proof by explicitly defining recursively, for all $1 < w \leq \check{T}$,

$$P_w = (1 + \frac{2Bp\sqrt{p}}{\nu_n \check{T}^{1/2-v/2}})P_{w-1} + \frac{1}{\check{T}^{1/2-v/2}}\tilde{\text{err}}_n(\delta), \quad P_1 = \tilde{\text{err}}_n(\delta).$$

where $\tilde{\text{err}}_n(\delta) = \frac{2\sqrt{p}}{\nu_n} c_0(\sqrt{\gamma_N \frac{\log(pTK/\delta)}{\eta_n^2 n}} + p\eta_n)$, and $B < \infty$ denotes a finite constant. Using a recursive argument, we obtain $P_w = \tilde{\text{err}}_n(\delta) \sum_{s=1}^w \alpha_s \prod_{j=s}^w (\frac{2Bp\sqrt{p}}{\nu_n \check{T}^{1/2-v/2}} + 1)$. Recall now that $\nu_n \geq \frac{1}{2\mu \check{T}^{1/2-v/2}}$, for ε_n as in Lemma F.2. As a result we can bound the above expression as

$$\sum_{s=1}^w \alpha_s \prod_{j=s}^w (\frac{2Bp\sqrt{p}}{\nu_n \check{T}^{1/2-v/2}} + 1) \leq \sum_{s=1}^w \alpha_s \prod_{j=s}^w (\frac{4\mu \check{T}^{1/2-v/2} Bp\sqrt{p}}{\check{T}^{1/2-v/2}} + 1) \leq \sum_{s=1}^w \alpha_s \exp\Big(\sum_{j=s}^w 4\mu Bp\sqrt{p}\Big).$$

Now we have $\exp\Big(\sum_{j=s}^w 4\mu Bp\sqrt{p}\Big) \leq \exp\Big(4\mu \check{T} Bp\sqrt{p}\Big)$, since $w \leq \check{T}$. We now write

$$P_w(\delta) \leq \tilde{\text{err}}_n(\delta) \sum_{s=1}^w \alpha_s \exp\Big(4\mu \check{T} Bp\sqrt{p}\Big) \leq \tilde{\text{err}}_n(\delta) \check{T}^{1/2+v} \exp\Big(8\mu^2 \check{T} Bp\sqrt{p}\Big).$$

□

Corollary 1. *Theorem A.8 holds.*

Proof. Consider Lemma F.2 where we choose $\delta = 1/n$. Observe that we choose $\epsilon_n \leq \frac{1}{4\mu \check{T}^{1/2-v/2}}$, which is attained by the conditions in Lemma F.2 as long as n is small enough such that

$$\sqrt{p}\Big[\bar{C}\sqrt{\log(n)\gamma_N \frac{\log(p\gamma_N \check{T}K)}{\eta_n^2 n}} + \eta_n\Big] \leq \frac{1}{4\mu \check{T}^{1/2-v/2}}$$

attained under the assumptions stated in Lemma F.2. As a result, we have $\nu_n = \frac{1}{4\mu \check{T}^{1/2-v/2}}$. By Lemma F.3 for all k, with probability at least $1-1/n$, $||\check{\beta}_k^{\check{T}} - \beta^*||^2 \lesssim \frac{p}{\check{T}^{1-v}}$. Also, we have $||\beta^* - \frac{1}{K}\sum_k \check{\beta}_k^{\check{T}}||_2^2 \leq \frac{1}{K}\sum_k ||\check{\beta}_k^{\check{T}} - \beta^*||^2$. The proof concludes by Theorem F.3 and Assumption 3.1, after observing that $W(\beta^*) - W(\hat{\beta}^*) \lesssim ||\beta^* - \hat{\beta}^*||_2^2$. □

F.2 Proof of Theorem E.1

The proof mimics the proof of Theorem 4.3.

Consider Lemma B.10 where we choose $\delta = 1/n$. Note that we can directly apply Lemma B.10 also to the gradient estimated with Algorithm E.1, since, by the circular-cross fitting argument, each parameter $\check{\beta}_k^w$ is estimated using sequentially pairs of different clusters as in Algorithm E.1. The rest of the proof follows verbatim from the one of Theorem 4.3.

F.3 Proof of Lemma E.2

The proof follows similarly to the proof of Lemma B.4, here taking into account also the component $\omega_{i,j}$. Under Assumption 2.2, we can write for some function g,

$$r\Big(D_{i,t}^{(k)}, D_{j:A_{i,j}^{(k)}>0,t}^{(k)}, X_i^{(k)}, X_{j:A_{i,j}^{(k)}>0}^{(k)}, A_i^{(k)}, U_i^{(k)}, U_{j:A_{i,j}^{(k)}>0}^{(k)}, \nu_{i,t}^{(k)}\Big) = g(Z_{i,t}^{(k)}).$$

Here, $Z_{i,t}^{(k)}$ depends on $A_i^{(k)}$, i.e., the edges of individual i, and on unobservables and observables of all those individuals such that $A_{i,j}^{(k)} > 0$, namely,

$$Z_{i,t}^{(k)} = \Big[D_{i,t}^{(k)}, X_i^{(k)}, U_i^{(k)}, \nu_{i,t}^{(k)}, A_i^{(k)} \otimes \Big(X^{(k)}, U^{(k)}, D_t^{(k)}\Big), \Big\{\Big[X_j^{(k)}, U_j^{(k)}, \omega_{i,j}^{(k)}\Big], j : 1\{i_k \leftrightarrow j_k\} = 1\Big\}\Big].$$

Importantly, under Equation (E.1), $A_i^{(k)}$ is a function of $\Big\{\Big[X_j^{(k)}, U_j^{(k)}, \omega_{i,j}^{(k)}\Big], j : 1\{i_k \leftrightarrow j_k\} = 1\Big\}$, only, and each entry depends on $(X_j, U_j, X_i, U_i, \omega_{i,j})$ through the same function l for each individual. What is important, is that $\sum_j 1\{i_k \leftrightarrow j_k\} = \gamma_N^{1/2}$ for each unit i. Therefore, for some function $\tilde{g}$ (which depends on l in Assumption 2.1), we can equivalently write

$$Z_{i,t}^{(k)} = \tilde{g}(D_{i,t}^{(k)}, \nu_{i,t}^{(k)}, X_i^{(k)}, U_i^{(k)}, \tilde{Z}_{i,t}^{(k)}), \quad \tilde{Z}_{i,t}^{(k)} = \Big\{\Big[X_j^{(k)}, U_j^{(k)}, D_{j,t}^{(k)}, \omega_{i,j}^{(k)}\Big], j : 1\{i_k \leftrightarrow j_k\} = 1\Big\},$$

where $\tilde{Z}_{i,t}^{(k)}$ is the vector of $\Big[X_j^{(k)}, U_j^{(k)}, D_{j,t}^{(k)}, \omega_{i,j}\Big]$ of all individuals j with $1\{i_k \leftrightarrow j_k\} = 1$.

Now, observe that since $(U_i^{(k)}, X_i^{(k)}) \sim_{i.i.d.} F_{X|U}F_U$, $\omega_{i,j} \sim F_\omega$, $\{\nu_{i,t}\}$ are $i.i.d.$ conditionally on $U^{(k)}, X^{(k)}, \omega^{(k)}$ and treatments are randomized as in Assumption 2.3, we have

$$G_{i,j}^{(k)} = \Big[X_j^{(k)}, U_j^{(k)}, \nu_{j,t}^{(k)}, D_{j,t}^{(k)}, \omega_{i,j}^{(k)}\Big]\Big|\beta_{k,t} \sim \mathcal{D}(\beta_{k,t})$$

are distributed with some distribution $\mathcal{D}(\beta_{k,t})$ which only depends on the exogenous coefficient $\beta_{k,t}$ governing the distribution of $D_{i,t}^{(k)}$ under Definition 2.3. Also, $G_{i,j}^{(k)}$ are independent across j (but not i) by the independence assumption of $\omega_{i,j}$. As a result for $\beta_{k,t}$ being exogenous, Lemma 2.1 holds since $\sum_j 1\{i_k \leftrightarrow j_k\} = \gamma_N^{1/2}$ for all i, hence $\tilde{Z}_{i,t}$ are identically (but not independently) distributed across units i, since $\tilde{Z}_{i,t}$ is a vector of γ_N $i.i.d$ random variables, each having the same marginal distribution which does not depend on i (therefore $\tilde{Z}_{i,t}$ has the same joint distribution across i).

To show the local dependence result, note that $G_{i,j}^{(k)}$ is mutually independent of $G_{u,v}^{(k)}$ for all $\{(v,u) : (v,u) \notin \{(i,j),(j,i)\}, u \neq j\}$ because $\omega_{i,j}$ are independent for different entries

(i,j). In addition, $\varepsilon_{i,t}^{(k)}|\beta_{k,t}$ is a measurable function of a vector[3]

$$\Big[X_j^{(k)}, U_j^{(k)}, \nu_{j,t}^{(k)}, D_{j,t}^{(k)}, \omega_{i,j}^{(k)}\Big]_{j:1\{i_k \leftrightarrow j_k\}=1}.$$

As a result, under mutual independence of $G_{i,j}$ with $G_{u,v}$ for all $\{(v,u):(v,u)\notin\{(i,j),(j,i)\}, u\neq j\}$, conditional on $\beta_{k,t}$, $\varepsilon_{i,t}^{(k)}$ is mutally independent with $\varepsilon_{v,t}^{(k)}$ for all v such that they are not connected and do not share a common element $\Big[X_j^{(k)}, U_j^{(k)}, \nu_{j,t}^{(k)}, D_{j,t}^{(k)}, \omega_{v,j}^{(k)}\Big]$, that is, such that $\max_j 1\{i_k \leftrightarrow j_k\}1\{v_k \leftrightarrow j_k\} = 0$ (here $1\{v_k \leftrightarrow v_k\} = 1$ for notational convenience). This holds since the edge between i and v is zero almost surely if $1\{v_k \leftrightarrow i_k\} = 0$.

There are at most $\gamma_N^{1/2} + \gamma_N$ many of $\varepsilon_{v,t}^{(k)}$ which can share a common neighbor with $\varepsilon_{i,t}^{(k)}$ ($\gamma_N^{1/2}$ many neighhbors and γ_N many neighbors of the neighbors), which concludes the proof.

Appendix G Numerical studies: additional results

G.1 One-wave experiment

In Figure G.2 we report the power plot for $\rho = 6$. In Figure G.3 we report the welfare gain from increasing β by 5% upon rejection of H_0 for $\rho = 6$. In Figure G.6, we report comparisons for different values of η_n.

G.2 Multiple-wave experiment

In Table G.1, we provide comparison with competitors for $\rho = 6$. Results are robust as in the main text.

In Figure G.4 we report a comparison among different learning rates, which are the one which rescales by $1/t$, the one that rescales by $1/\sqrt{T}$ and the one that rescales by $1/\sqrt{t}$.

In Figure G.9 we study the adaptive experiment as the starting value is the optimum minus 5% and show that the out-of-sample regret is small and close to zero.

G.3 Calibrated experiment with covariates for cash transfers

In this subsection, we turn to a calibrated experiment where we also control for covariates, as discussed in Section 2.3. we use data from Alatas et al. (2012, 2016). we estimate a function heterogenous in the distance of the household's village from the district's center. we use information from approximately four hundred observations, whose eighty percent or

[3]Here for notational convenience convenience only, we are letting $1\{i_k \leftrightarrow i_k\} = 1$.

more neighbors are observed. We let $X_i \in \{0,1\}, X_i = 1$ if the household is far from the district's center than the median household, and estimate

$$Y_i|X_i = x \quad = \quad \phi_0 + \tilde{X}_i\tau + D_i\phi_{1,x} + \frac{\sum_{j\neq i} A_{j,i}D_j}{\max\{\sum_{j\neq i} A_{j,i}, 1\}}\phi_{2,x} + \Big(\frac{\sum_{j\neq i} A_{j,i}D_j}{\max\{\sum_{j\neq i} A_{j,i}, 1\}}\Big)^2\phi_{3,x} + \eta_i, \tag{G.1}$$

where η_i are unobservables centered on zero conditional on $X_i = x$, and $\tilde{X}_i$ denotes controls which also include X_i.[4] Using the estimated parameter, we can then calibrate the simulations as follows.

We let $\eta_{i,t}, \sim \mathcal{N}(0, \sigma^2)$, where σ^2 is the residual variance from the regression. We then generate the network and the covariate as follows:

$$A_{i,j} = 1\Big\{||U_i - U_j||_1 \leq 2\rho/\sqrt{N}\Big\}, \quad U_i \sim_{i.i.d.} \mathcal{N}(0, I_2), \quad X_i = 1\{U_i^{(1)} > 0\}.$$

Here, $U_i^{(1)}$ is continuous and captures a measure of distances. Individuals are more likely to be friends if they have similar distances from the center, and X_i is equal to one if an individual is far from the district's center from the median household. We fix $\rho = 1.5$ to guarantee that the objective's function optimum is approximately equal to the optimum observed from the data (in calibration, the optimum is $\beta \approx 0.26$, while $\beta^* \approx 0.29$ on the data). We then generate data

$$Y_{i,t}|X_i = x \quad = \quad D_i\hat{\phi}_{1,x} + \frac{\sum_{j\neq i} A_{j,i}D_j}{\max\{\sum_{j\neq i} A_{j,i}, 1\}}\hat{\phi}_{2,x} + \Big(\frac{\sum_{j\neq i} A_{j,i}D_j}{\max\{\sum_{j\neq i} A_{j,i}, 1\}}\Big)^2\hat{\phi}_{3,x} + \eta_{i,t}. \tag{G.2}$$

where we removed covariates that did not interact with the treatment rule (i.e., do not affect welfare computations). The policy function is $\pi(x;\beta) = x\beta + (1-x)(1-\beta)$ where β is the probability of treatment for individuals farer from the center. Here, we implicitely imposed a budget constraint $\beta P(X_i = 1) + (1-\beta)P(X_i = 0) = 1/2$, where, by construction $P(X_i = 1) = 1/2$.

We collect results for the one-wave experiment in Figure G.5, G.7 (left-panel), where we report power and the relative improvement from improving by 5% the treatment probability for people in remote areas as discussed in the main text. Welfare improvements (and power) are increasing in the cluster size and the number of clusters. However, such improvements are negligible as we increase clusters from twenty to forty, suggesting that twenty clusters

[4] We also control for the education level, village-level treatments, i.e., how individuals have been targeted in a village (i.e., via a proxy variable for income, a community-based method, or a hybrid), the size of the village, the consumption level, the ranking of the individual poverty level, the gender, marital status, household size, the quality of the roof and top (which are indicators of poverty).

are sufficient to achieve the largest welfare effects.[5] In the right-hand side panel of Figure G.7 we report the out-of-sample regret. The regret is generally decreasing in the number of iterations, especially as the regret is further away from zero. As the regret gets almost zero (0.06%), the regret oscillates around zero as the number of iterations increases due to sampling variation. This behavior is suggestive that for some applications, few iterations (in this case, ten) are sufficient to reach the optimum, up to a small error. In Table G.2, we observe perfect coverage for $n = 600$, and under-coverage by no more than five percentage points in the remaining cases.

G.4 Calibrated simulations to M-Turk experiment

In this subsection, we study the performance of the methods to increase vaccination against COVID-19 through information diffusion.

Specifically, at the beginning of March 2021 (before the vaccination campaign was extensively implemented), we ran an M-Turk experiment where each individual was assigned either of two arms.[6] A control arm, which consisted of a survey asking basic questions on characteristics of the participants, and a treatment arm. The treatment arm was first assigned simple survey questions. Then, individuals under treatment were assigned three questions about COVID, whose correct answer was rewarded a small economic incentive. The goal of these questions is to increase awareness of the severity of COVID.[7] Each correct answer rewarded a small bonus and was displayed right after the participant submitted her answer to the three questions (before the end of the survey). In addition, at the end of the survey, participants were asked again one of the three questions, whose correct answer rewards a bonus. Participants were made aware of the bonuses and the survey's structure. The scope of the treatment was to increase awareness of the severity of the disease by asking questions and showing the correct answers to facilitate information transmission. At the end of the survey, both controls and treated units were asked when they would have done the vaccine. The outcome of interest is binary and equals whether individuals would have done the vaccine either as soon as possible or during the spring. We estimate the model with 1035

[5]The order of magnitude of the welfare gain is smaller compared to simulations with the unconditional probability since, here, we always treat exactly half of the population. As a result, welfare oscillates between 0.24 and 0.29 only (as opposed to zero to one as in the unconditional case), as shown in Figure 2.

[6]The experiment was certified an IRB exempt by UCSD, Human Research Protections Program.

[7]These were multiple answers questions. The first question asked (i) Which of these events caused more deaths of Covid in the US? (more answers allowed), giving four options (World War I and II, 50 times more than 9/11, US Civil war); What is the percentage of people in the US who had Covid within the last year? (approximately); The number of people infected from Covid in the last year is comparable to ... (giving three options).

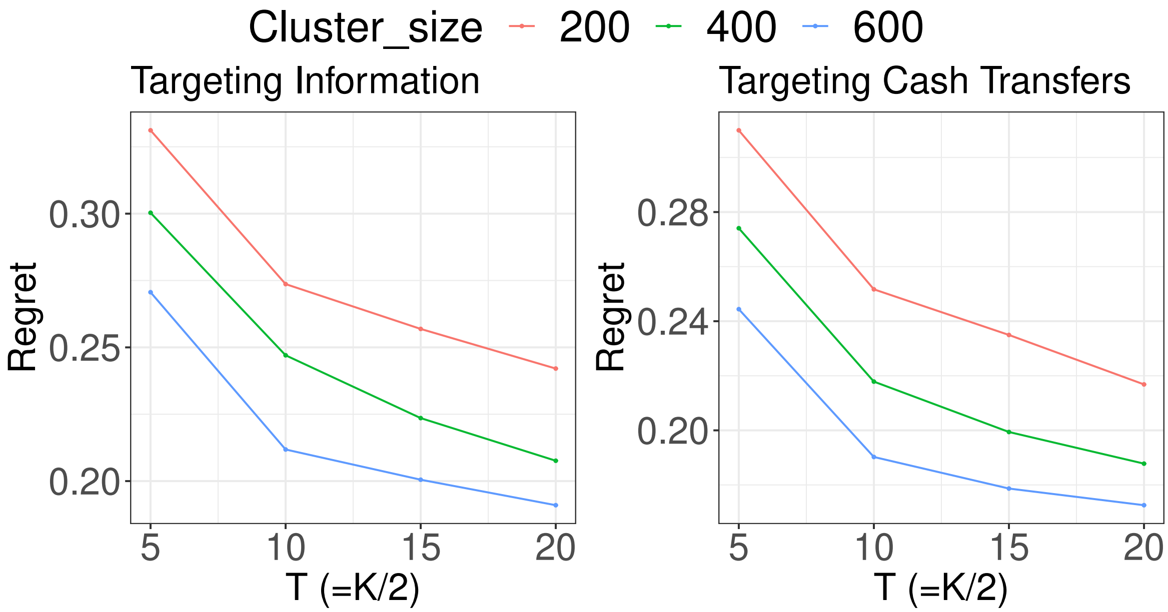

Figure G.1: Multi-wave experiment in Section 7. 200 replications. In-sample regret, average across clusters, for $\rho = 2$.

participants.[8] We estimate treatment effects by running a simple linear regression, where the treatment dummy interacts with the dummy, indicating whether the individual classifies herself as liberal, conservative, or "prefer not to say". We find substantial heterogeneity, with positive effects on early vaccination on liberals only. We consider a model and policy function as in Section G.3, where X, in this case, denotes whether an individual is liberal or conservative (drawn with the same DGP as in Section G.3 for simplicity), and $\rho = 2$ as in the main text. We calibrate $\hat{\phi}_{1,x}$ to the estimated direct treatment effect and fix the percentage of treatment units to fifty percent. A challenge here is that we do not know spillovers ϕ_{2x}, ϕ_{3x}. Therefore, we choose $\phi_{3,x} = r\phi_{2,x}$, where r is estimated from data on information diffusion from Cai et al. (2015) for simplicity. We choose $\phi_{2,x} = \max\{\alpha\phi_{1,x}, 0\}$, i.e., total spillovers equal direct effect $\phi_{1,x}$ times a constant $\alpha \in \{0.1, 0.2, 0.3, 0.4\}$ if these are positive, and zero otherwise.

We run our one wave experiment with these calibrations and collect results in Figure G.8. In the figure, we report the relative welfare improvement of increasing treatment probabilities for liberals of ten percent, upon rejection of the null hypothesis H_0 in Equation (19).

[8] We collected information from 2411 participants. We removed 158 observations that had already received the vaccine in March 2020 and 203 observations that took less than thirty seconds and more than five minutes to take the survey. We also removed all those observations which were not living in the US.

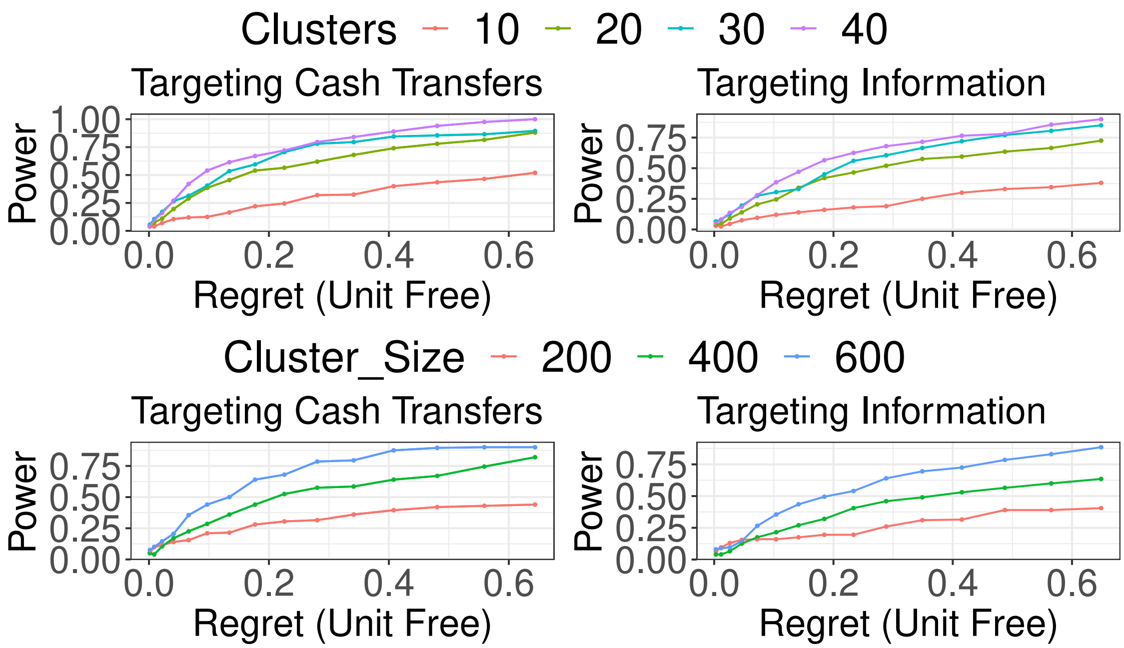

Figure G.2: One-wave experiment in Section 7. Power plot for $\rho = 6$. The panels at the top fix $n = 400$ and varies K. The panels at the bottom fix $K = 20$ and vary n.

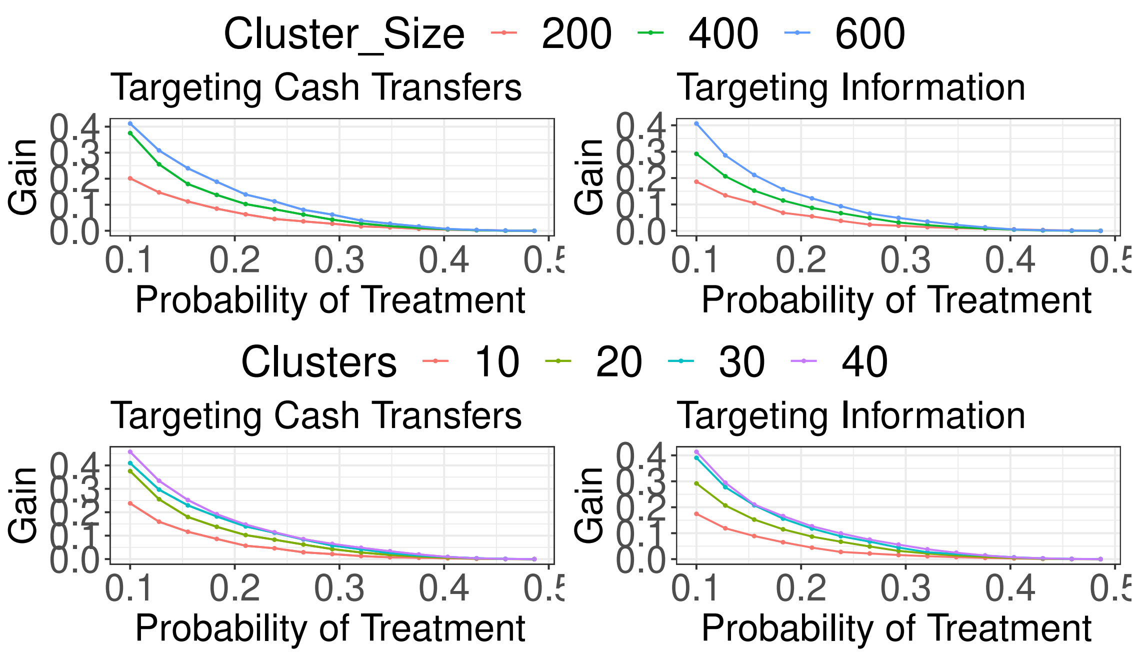

Figure G.3: One-wave experiment in Section 7. $\rho = 6$. Expected percentage increase in welfare from increasing the probability of treatment β by 5% upon rejection of H_0. Here, the x-axis reports $\beta \in [0.1, \cdots, \beta^* - 0.05]$. The panels at the top fix $n = 400$ and varies the number of clusters. The panels at the bottom fix $K = 20$ and vary n.

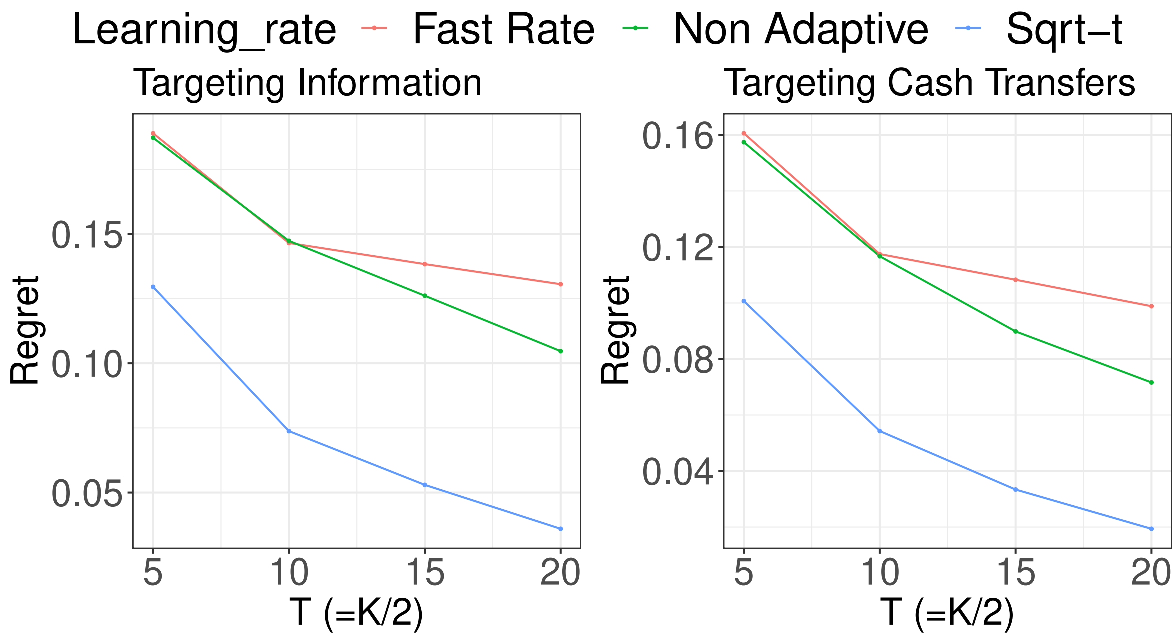

Figure G.4: Comparisons among different learning rates with experiment as in Section 7. 200 replications, $\rho = 2, n = 600, K = 2T$. Fast rate denotes a rescaling of order $1/t$; non-adaptive depends on a rescaling of order $1/\sqrt{T}$; the last one (Sqrt-t) depends on a rescaling of order $1/\sqrt{t}$.

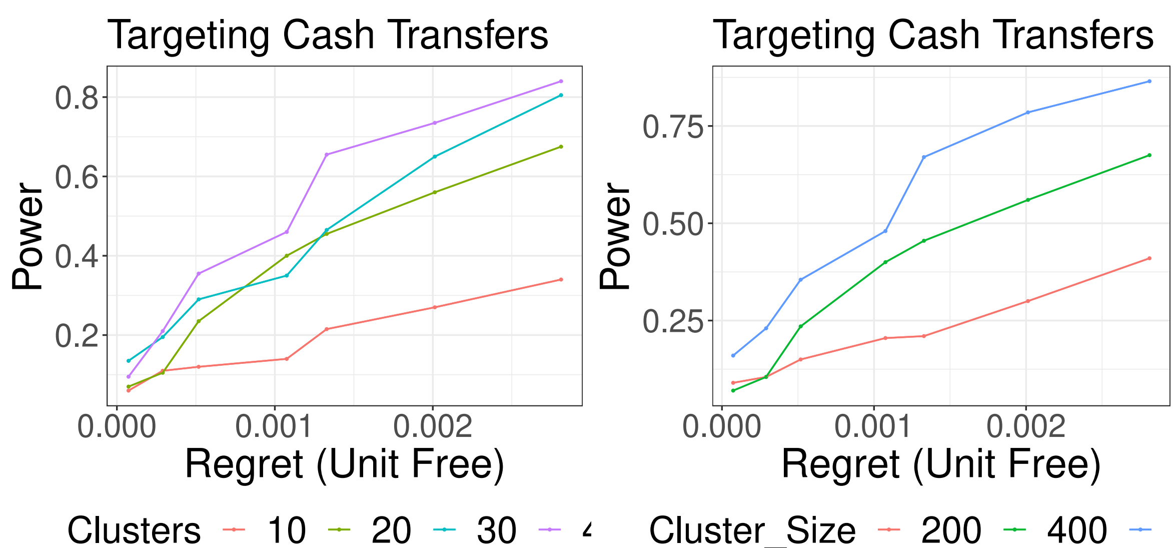

Figure G.5: Single-wave experiment in Section G.3. Power, 200 replications.

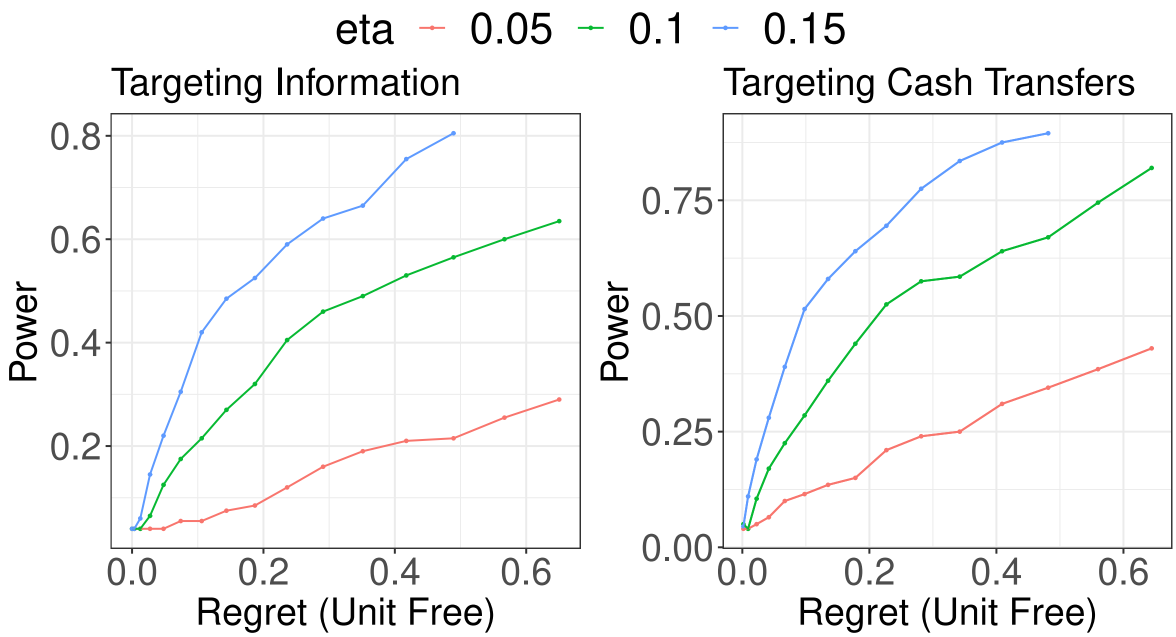

Figure G.6: One wave experiment calibrated to Alatas et al. (2012) and Cai et al. (2015). The plot reports power for different values of η_n varies, with $K = 200, n = 400$, with 200 replications.

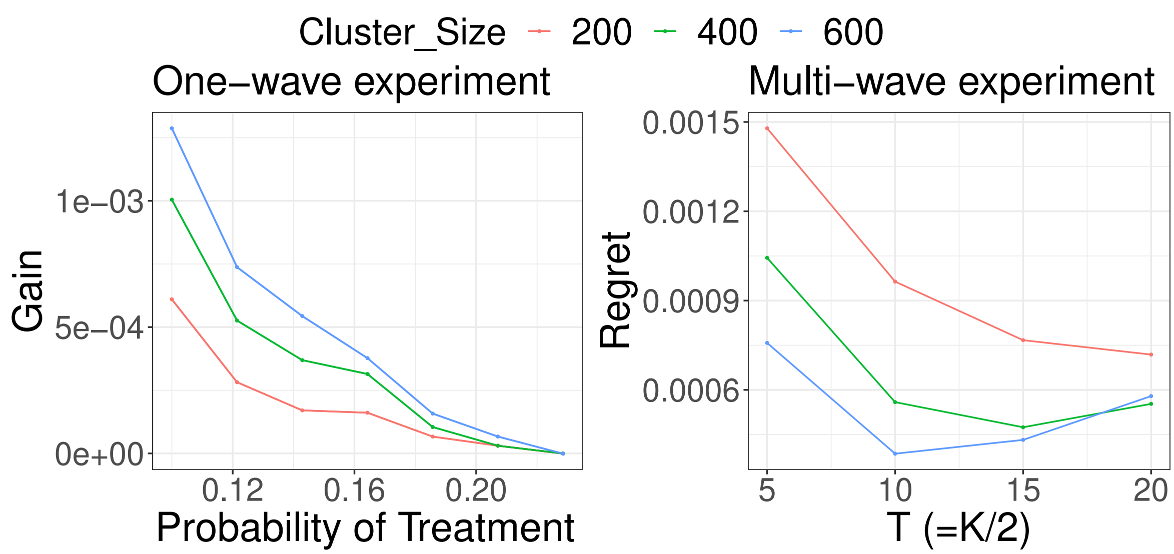

Figure G.7: Experiment in Section G.3. Left-hand side panel reports the expected percentage increase in welfare from increasing the probability of treatment β by 5% to individuals in remote areas upon rejection of H_0. Here, the x-axis reports $\beta \in [0.1, \cdots, \beta^* - 0.05]$. The right-hand side panel reports the in-sample regret. 400 replications.

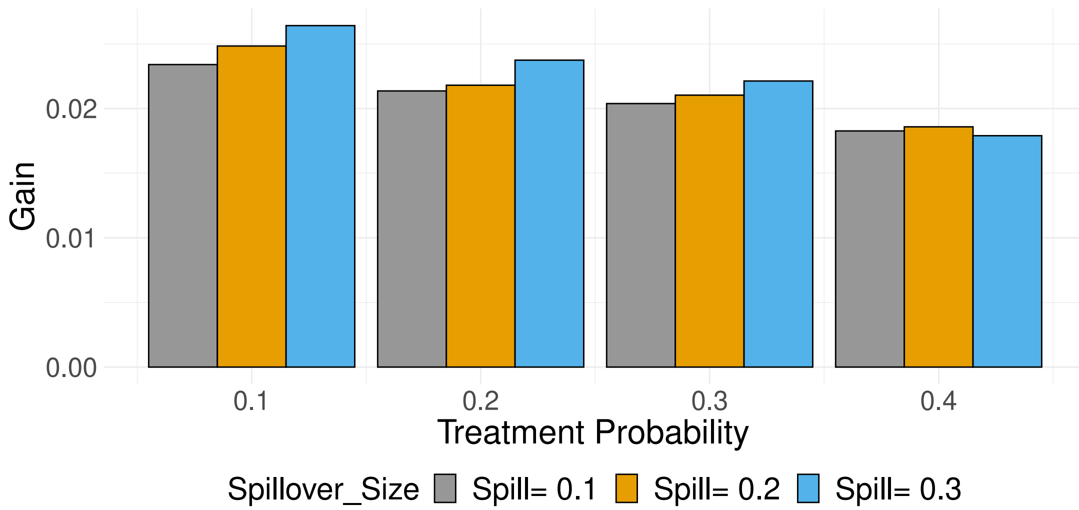

Figure G.8: One wave experiment in Section G.4. Relative welfare improvement for increasing the probability of treatment by ten percent, conditional on rejecting the null hypothesis. 200 replications. Different columns correspond to different levels of spillover effects (captured by α). Here $K = 20$.

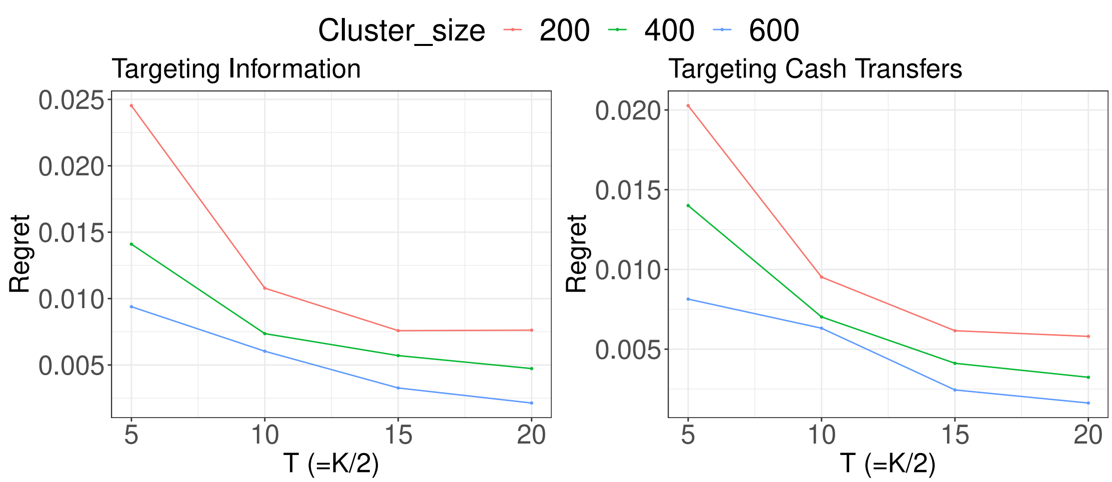

Figure G.9: Multi-wave experiment wave experiment in Section 7, as β is initialized at the optimum value minus 5%. Reported in the figure is the out-of-sample welfare. 200 replications.

Table G.1: Multiple-wave experiment in Section 7. Relative improvement in welfare with respect to best competitor for $\rho = 6$. The panel at the top reports the out-of-sample regret and the one at the bottom the worst case in-sample regret across clusters.

	Information				Cash Transfer			
$T =$	5	10	15	20	5	10	15	20
$n = 200$	0.03	0.105	0.243	0.156	0.233	0.243	0.264	0.287
$n = 400$	0.135	0.130	0.244	0.258	0.243	0.274	0.321	0.335
$n = 600$	0.217	0.214	0.281	0.344	0.261	0.313	0.343	0.360
$n = 200$	0.587	0.695	0.670	0.627	0.247	0.279	0.300	0.320
$n = 400$	0.551	0.667	0.830	0.869	0.266	0.306	0.343	0.352
$n = 600$	0.589	0.771	0.897	0.955	0.294	0.360	0.387	0.387

Table G.2: Single-wave experiment in Section G.3, 200 replications. Coverage for tests with size 5%.

$K =$	10	20	30	40
$n = 200$	0.955	0.935	0.900	0.905
$n = 400$	0.965	0.945	0.900	0.950
$n = 600$	0.935	0.965	0.920	0.965

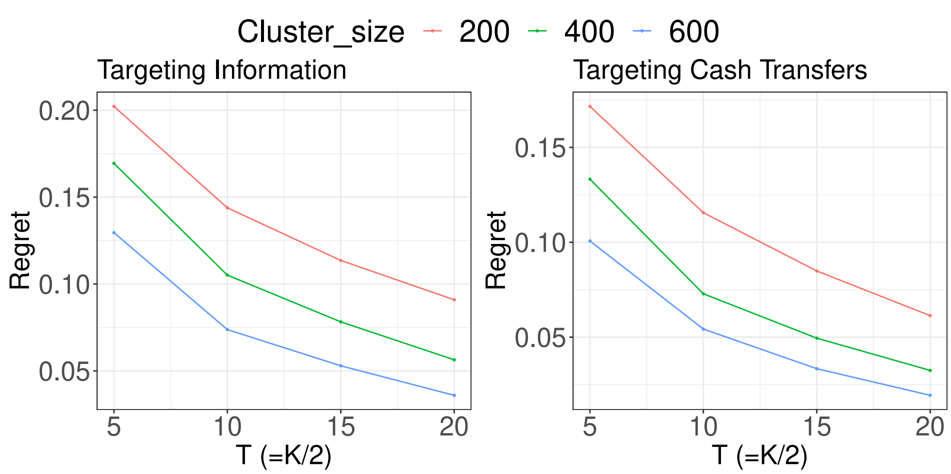

Figure G.10: Adaptive experiment $\rho = 2$. 200 replications. The panel reports the out-of-sample regret of the method as a function of the number of iterations.

References

Alatas, V., A. Banerjee, A. G. Chandrasekhar, R. Hanna, and B. A. Olken (2016). Network structure and the aggregation of information: Theory and evidence from indonesia.

American Economic Review 106(7), 1663–1704.

Alatas, V., A. Banerjee, R. Hanna, B. A. Olken, and J. Tobias (2012). Targeting the poor: evidence from a field experiment in indonesia. *American Economic Review 102*(4), 1206–40.

Cai, J., A. De Janvry, and E. Sadoulet (2015). Social networks and the decision to insure. *American Economic Journal: Applied Economics 7*(2), 81–108.

Hazan, E., K. Levy, and S. Shalev-Shwartz (2015). Beyond convexity: Stochastic quasi-convex optimization. In *Advances in Neural Information Processing Systems*, pp. 1594–1602.