Reducing gunshot victimization in high-risk social networks through direct effects and spillover effects

George Wood* and Andrew Papachristos†

Abstract

More than 60,000 people are victimized by gun violence each year in the United States. A large share of victims cluster in bounded and identifiable social networks. Despite a growing number of violence reduction programs that implicitly or explicitly leverage social networks to broaden programmatic effects, there is little evidence to suggest that reductions in gun violence are achieved through spillover effects on the peers of program participants. This study estimates the direct and spillover effects of a large gun violence field-intervention in Chicago, IL. Using a quasi-experimental design, we test whether assigning 2,349 seed individuals to a desistance-based program reduced the incidence of gunshot victimization. The study uses co-arrest network data to further test spillover effects on 6,132 unassigned individuals who did not participate in the program. Direct effects were associated with a 3.2 percentage point reduction in gun violence victimization among seeds over two years, while potential spillover was associated with a 1.5 percentage point reduction in victimization among peers. Findings suggest that peer-influence and the structure of networks might be leveraged to amplify gun violence reduction efforts.

1. Introduction

In 2016, more than 14,000 people in the United States were shot and killed by another person, while another 70,000 were wounded by guns in assaults (1, 2). Gun violence and its effects are not evenly distributed across the population, however. The trauma and consequences associated with gun violence disproportionately affect young minority men in socially and economically disadvantaged neighborhoods (1, 3, 4, 5). Recent research demonstrates that gunshot victimization further concentrates in small circumscribed social networks within high-risk populations (6, 7, 8). For example, 70% of all victims of non-fatal gunshot injuries in Chicago could be located in a co-arrest network comprised of less than 5% of the city's population (7). A study of a high-crime Boston community found that nearly 85% of all victims of fatal and non-fatal gunshot injuries could be located in a network comprised of 763 connected individuals (8). Importantly, the structure of such social networks and individuals' placement within them can severely elevate the risk of victimization and contribute to the diffusion of violence. Cascades of gun violence within high-risk networks—starting with one victim whose associate is subsequently victimized, and so on—accounted for 63% of victimizations over an eight-year period in Chicago (9).

The concentration of gun violence in social networks has implications for gun violence reduction efforts to the extent that identifying and leveraging such networks might enhance violence prevention and reduction practices. Many violence reduction efforts already implicitly or explicitly rely on this sort of networked-logic by concentrating efforts on high-risk individuals and groups within small geographic areas. The central premise of such programs is that gun violence might be reduced by dissuading specific individuals or groups from involvement in violence by providing information, resources, mediation, or other services directly to those at

*Department of Internal Medicine, Yale University
†Department of Sociology, Northwestern University

greatest levels of involvement in gun crime and violence (10, 11). Furthermore, some programs implicitly seek to engender spillover effects on unassigned individuals by placing outreach workers and mediators into the networks of individuals involved in violent disputes (12, 13) or else by encouraging individuals who are part of interventions to propagate a program's message within their personal networks (14).

While mounting evidence suggests that gun violence reduction programs focusing on small geographic areas or a small number of groups are effective in reducing aggregate levels of gunshot violence (15, 16), we know little about the magnitude of program effects on treated individuals and even less about whether victimization can be reduced through network spillover effects. To date, most evaluations have analyzed aggregate levels of gun violence either in geographic regions or among specific groups or gangs in which a program has been implemented, comparing the frequency of victimization before and after implementation (17). Such aggregated approaches are limited in estimating the magnitude of the effect on treated individuals because they capture victimization outcomes among individuals unaffected by the program while being affected by shifts in the incidence of gunshot violence caused by exogenous, non-program related factors (18, 19). In addition, it has been difficult to test whether programs reduce victimization among untreated but potentially-affected individuals through spillover effects, largely because it has been infeasible to identify potentially-affected individuals and difficult to establish a causal framework for estimating spillover (20). As such, a central question of precisely how violence reduction programs reduce the incidence of gunshot victimization remains open.

The intervention we report attempted to reduce gunshot victimization among high-risk participants as well as among their unassigned peers through spillover effects. The program is an "induction intervention" (21, 22) that attempted to activate peer-to-peer diffusion of a desistance message by encouraging intervention participants to spread a behavioral stimulus.

For the intervention, a group comprised of law enforcement, community members, and social service agencies carried out an exercise to map violent conflicts, disputes, and episodes of gun violence in order to identify participants at high-risk of involvement in group-involved gun violence (23). These individuals were subsequently invited to attend a one-hour meeting (24, 25). At the meetings, which were held in-person at a public location such as a local park district, school, or community center, participants engaged in a four-stage process. First, law enforcement officers delivered a message that the individual is at acute risk of victimization and that further involvement in gun violence would be met with coordinated enforcement action and criminal justice sanctions. Second, community representatives emphasized that the participant is valued by their community and extended an invitation for the participant to become involved in positive community activities to avoid victimization. Third, in a process akin to perspective-taking (26), victims or relatives of victims of gunshot violence recounted the effects and trauma associated with gun victimization. Lastly, participants were offered direct links to local social service provision. Throughout the meeting, participants were encouraged to spread the desistance and community message, perspective-taking, and service provision information to their peers (27, 28, 29). See the supplementary materials for more program information. Invited individuals may or may not attend the meeting and are designated compliant or non-compliant, accordingly. There are no punitive consequences for non-compliance.

To estimate the effects of these meetings, we conducted a quasi-experimental evaluation using participation data and administrative data. We first identified the individuals who were invited to and attended a meeting (n = 1,642) and the individuals who were invited but not did not attend a meeting (n = 707) whom we call compliant seeds and non-compliant seeds, respectively.

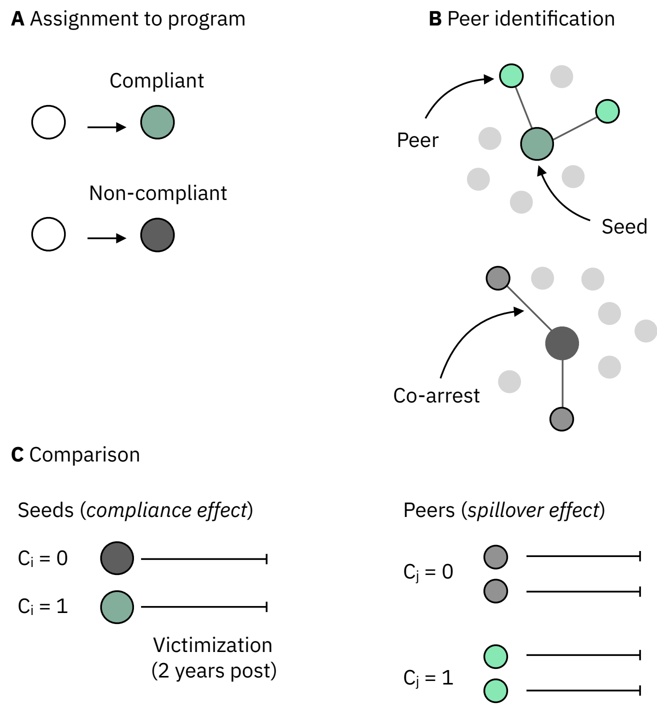

Figure 1: Design of the intervention evaluation. (A) Individuals at high-risk of gunshot victimization (seeds) are assigned to the program and select into compliance or non-compliance. (B) Each seed is identified in the Chicago co-arrest network. We apply our exposure mapping, classifying individuals adjacent to a compliant seed as compliant peers and individuals adjacent to a non-compliant seed as non-compliant peers. (C) The compliance effect is based on a comparison of victimization outcomes in the two-years after assignment for seeds in the compliant and non-compliant condition. The spillover effect is based on the equivalent comparison among peers.

We then constructed a co-arrest network (30) around each compliant and non-compliant seed, where each seed is connected to all non-invited individuals alongside whom they were arrested in the three years prior to the meeting. We define an exposure mapping (31, 32) in which unassigned individuals are designated as compliant peers ($n = 3{,}034$) if they are connected to a compliant seed–and therefore potentially exposed to the intervention message–or non-compliant peers ($n = 3{,}098$) if they are connected to a non-compliant seed. To estimate the main effect of the meeting, we compared victimization outcomes in the two-years following invitation to the program for the compliant seeds against the non-compliant seeds. To estimate the spillover effect, we compared victimization outcomes in the two-years after potential exposure to the intervention message for the compliant peers against the non-compliant peers. For each peer, potential exposure to spillover begins on the day that the seed to whom that peer is connected was invited to attend a meeting. Finally, to test the spillover effect under additional exposure to program participants, we compared victimization outcomes for peers with connections to two or more compliant seeds and peers connected to two or more non-compliant seeds.

By linking program assignment data with administrative data on the networks of participants and their associates, we are able to map potential pathways of diffusion and thereby determine whether gunshot victimization is reduced among individuals assigned to the program and among individuals indirectly exposed to treatment through their network connections. One research design feature warrants further exposition. The program uses non-random assignment and selects individuals deemed to be at high-risk of victimization. As such, we could not compare those assigned to attend a meeting against a random selection from the wider Chicago population (supplementary section 7). Exploiting non-compliance with the program permits a comparison among units deemed to be similarly high-risk ex-ante. However, we cannot be certain that the choice to comply or not comply with the intervention is unconfounded with the risk of victimization. In estimating the effects of the intervention on compliant seeds, we must therefore account for the probability of selection into the compliant or non-compliant condition. Moreover, because the risk of victimization among connected peers and seeds is plausibly associated, we must also account for the probability of being connected to a compliant or non-compliant seed.

To address this problem of confounding, we estimate the effect of compliance on victimization outcomes conditional on observed covariates using Bayesian Additive Regression Trees (BART). BART has been shown to be more accurate than propensity-score based estimators in estimating treatment effects under confounding, particularly when the parametric relationship between the outcome, treatment, and confounders is unknown (33). This is in part because BART naturally allows for possible interaction effects and nonlinearities in this relationship (33, 34, 35). We then compare the BART estimate to a simple difference-in-means (DIM). We use this comparison to assess the proportion of the observed difference in victimization outcomes that is attributed to the compliance effect and spillover effect after adjusting for confounding due to selection into compliance. Finally, as an extension, we use the flexible BART approach to estimate the heterogeneity of the compliance effect. In the supplementary materials, we show further estimates of the compliance effect and spillover effect using a more conventional logistic regression model with unit re-weighting based on entropy balancing (EBAL) (36). See methods for further design and estimation details.

2. Results

2.1. Compliance effect on victimization

Compliance substantially reduced the probability of gunshot victimization in the two-years after attending an intervention meeting. The share of non-compliant individuals who were victimized in the two-years following non-attendance was 18.1%. For compliant individuals, the corresponding share was 10.6%. From the BART model, we estimate that compliance caused a median reduction in the chance of victimization of -3.2 percentage points (Fig. 2; mean estimate = -3.2, credible interval = -4.9 to -1.4). The magnitude of the reduction is substantially larger in the DIM model (estimate = -7.5, confidence interval = -10.7 to -4.3, $p < 0.001$), which does not adjust for selection into compliance and simply reflects the 7.5 percentage point difference between compliant and non-compliant units in post-intervention victimization outcomes. The reduced effect size estimate in the BART model shows that adjusting for selection into compliance accounts for part of this difference in observed victimization outcomes. The BART estimate attributes 42.6% of the difference in victimization outcomes to the effect of compliance. The estimated reduction in victimization caused by compliance is supported in the EBAL model (supplementary section 4).

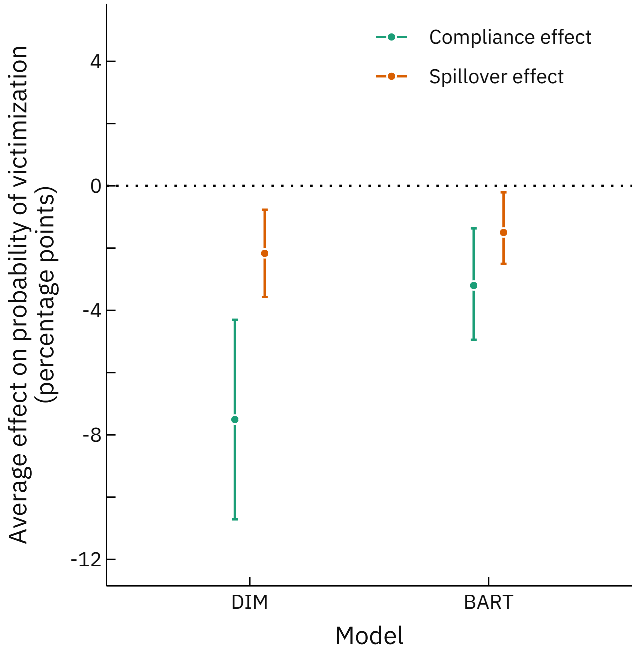

Figure 2: The estimated effect of compliance on program seeds and the estimated effect of compliance spillover on program peers. The Bayesian Additive Regression Trees (BART) model is a nonparametric estimate of the conditional average treatment effect of compliance, accounting for observed covariates. The difference-in-means (DIM) is the mean difference in the two-year probability of victimization for individuals in the compliant versus non-compliant condition. A 95% confidence interval is shown for the DIM estimate. A 95% credible interval is shown for the BART estimate.

2.2. Spillover effect on victimization

Analysis of peer outcomes showed that units with a social connection to a compliant seed had a lower rate of victimization than units with a connection to a non-compliant seed. In the two years following potential exposure to program spillover effects, the share of non-compliant peers victimized was 9.7% compared to 7.5% among compliant peers. From the BART model, we estimate that compliance spillover caused a median reduction in the probability of victimization of 1.5 percentage points (Fig. 2; mean estimate = -1.4, credible interval = -2.5 to -0.2). The magnitude of the spillover effect is approximately 0.47 as large as the primary compliance effect. The effect size in the unadjusted DIM model (estimate = -2.2, confidence interval = -3.6 to -0.8, $p = 0.0024$) reflects the raw 2.2 percentage points difference in victimization outcomes. The effect size is tempered in the BART model, which attributed 69% of the difference in outcomes to the effects of intervention spillover. The attribution of some of the difference in victimization to pre-intervention differences between compliant and non-compliant peers, rather than to spillover effects, implies that differential selection into compliance among seeds is informative for characterizing the risk of peer victimization. That is, both non-compliant seeds and the peers of non-compliant seeds appear to have a higher baseline risk of victimization than their compliant counterparts. Notably, although similar in magnitude to the BART estimate, there is considerably greater uncertainty around the spillover effect estimates (estimate = -1.2, confidence interval = -2.7 to 0.3, $p = 0.1178$) in the EBAL model (supplementary section 4).

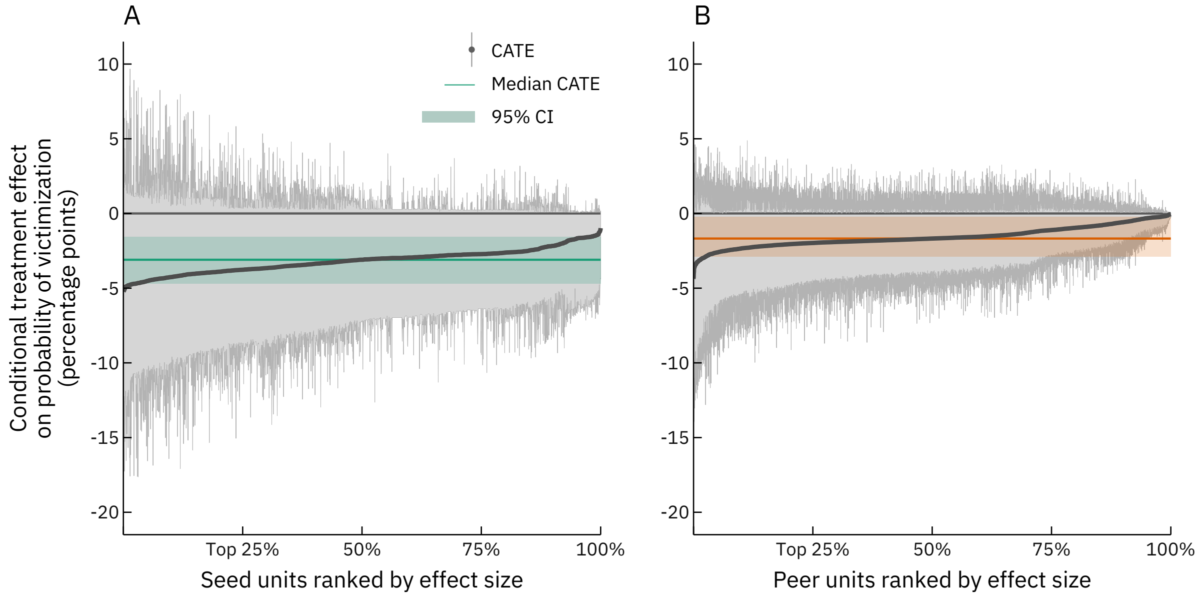

Figure 3: Conditional Average Treatment Effect (CATE) of (A) compliance on the probability of victimization for program seeds and (B) compliance spillover on the probability of victimization for peers. For each individual, the CATE is estimated as a function of their covariate profile using the fitted BART models for the compliance and spillover effect. 95% credible intervals are shown for each CATE. The median CATE and associated 95% credible interval is also shown.

2.3. Heterogeneity of compliance and spillover effects

To assess the extent to which the compliance effect and spillover effect resulted in similar reductions in the risk of victimization among all compliant seeds and compliant peers, we used the BART models to estimate the effect for each seed and peer as a function of their baseline covariate profile. This is a natural extension of BART, which allows for variation in the treatment effect across covariates (33, 35). The primary effect of compliance on the probability of victimization ranges from -5.8 to -1.0 percentage points, at the extremes (Fig. 3A). Overall, however, our analysis shows little heterogeneity in the compliance effect, with 46.1% of units falling within ± 0.5 percentage points of the median estimated compliance effect and 74.9% of units falling within ± 1 percentage point. Similarly, the spillover effect of compliance is largely homogenous (Fig. 3B), ranging from -3.6 to -0.0 with 61.7% and 87.0% falling within ± 0.5 and ± 1 percentage points of the median spillover effect, respectively. Thus, there is little evidence to suggest that the primary compliance effect and spillover effect are moderated by the covariate profile of compliant seeds or peers to a degree that would warrant restructuring of the intervention for particular subgroups of assigned individuals.

2.4. Spillover effect of multiple exposures

Our results so far show that compliance reduces the probability of victimization among seeds and that compliance spillover effects reduce the probability of victimization among peers. As a further test of the spillover effects of compliance, we compare victimization outcomes among the subset of peers with exposure to two compliant seeds ($n = 200$) or two non-compliant seeds ($n = 350$). In the two years following exposure, 12% of peers with connections to two compliant seeds and 10.6% of peers connected to two non-compliant seeds were victimized. Thus, the victimization rate was higher among peers with a double exposure to compliant seeds than peers with a double exposure to non-compliant seeds. Additionally, the incidence of victimization was higher among peers with exposure to two seeds than among peers with a single exposure. Peers with two or more exposures exhibit a different pre-intervention covariate profile in terms of age, gender, and gang membership as well as having a higher incidence of victimization and arrest among their co-arrestees than peers with a single exposure (supplementary section 6). From BART Model 1, we estimate that peers with a double compliant exposure had a median -0.1 percentage points lower probability of victimization (mean = 0.0, credible interval = -0.6 to 1.8). However, the credible interval for the estimated effect includes zero with considerable posterior mass at both positive and negative values (supplementary Fig. S3). We therefore have insufficient information to ascertain whether double exposure to compliant seeds increases, reduces, or does not affect the probability of victimization relative to double exposure to non-compliant seeds.

2.5. Expected victimizations under the counterfactual

The parameter estimates support the idea that the violence reduction program reduced the probability of victimization among both seeds and peers. However, the magnitude of the reduction among peers that is attributable to spillover effects is small in absolute terms: a median reduction in the probability of victimization of 1.5 percentage points (Fig. 2) with 25.7% of peers experiencing a spillover effect that is smaller in magnitude (i.e. closer to zero) than 1 percentage point (Fig. 2). Moreover, the p-value for the spillover effect estimate is greater than 0.05 in the EBAL model, which imposes a stronger parametric form on the relationship between victimization, compliance and the observed covariates than the BART model (supplementary section 4). To put the magnitude of this effect into context, 174 of the 1,642 seeds who complied with the intervention were victimized in the two years after compliance. In the counterfactual condition in which these seeds did not comply, our BART effect estimate implies that 227 (lower bound = 196, upper bound = 255) might have been victimized. For the compliance spillover effect, 228 of the 3,034 peers of compliant seeds were victimized in the two years after potential exposure to the program. In the counterfactual condition in which those peers were connected to non-compliant seeds, our BART effect estimate implies that 273 (lower bound = 234, upper bound = 304) might have been victimized by gunshot violence. Our estimates therefore indicate that 53 victimizations were averted as a product of the primary effect and a further 45 through spillover effects, albeit with considerable uncertainty around these estimates. In summary, these findings imply that, by holding call-in meetings with 1,642 individuals in high-risk social networks for whom we could identify 3,034 peers, the program contributed to approximately 98 fewer gunshot victimizations.

3. Discussion

Gun violence prevention, intervention, and treatment programs are increasingly directing their efforts toward the small networks and geographic areas disproportionately impacted by gunshot victimization and associated trauma. Although a networked-logic drives many such programs, most evaluations to date have analyzed changes in aggregate rates of victimization or crime before and after program implementation, paying little attention to the direct effects on individuals who are assigned to such programs let alone the spillover effects that are theoretically deemed to be responsible for reductions in gun violence.

This study analyzed a field-intervention in Chicago that focused on high-risk individuals whom were actively involved in on-going disputes. Using a quasi-experimental design and data on co-arrest network ties, we evaluated the direct effect on those participating in the intervention as well as the spillover effect on their unassigned co-arrest associates. Our findings show that participation in the intervention reduced gunshot victimization by 3.2 percentage points over two years. At the same time, potential spillover reduced victimization by 1.5 percentage points among the unassigned associates of intervention participants. However, we did not detect a spillover effect among the subset of peers who were connected to two or more compliant intervention participants. While the absence of a detectable effect may be due to the relatively small number of peers in this subset, it could also be due to a difference between individuals with exposure to two or more seeds and individuals with a single exposure. For example, individuals with greater seed exposure may be involved in higher risk networks, as indicated by the greater incidence of victimization among peers with two exposures compared to peers with one, and may be less amenable to a behavioral shift in response to the intervention (supplementary section 6).

Unlike many public health interventions (37), gun violence reduction programs seldom utilize formal network metrics for the selection of participants. Consistent with several recent evaluations of networked interventions (31, 38, 39, 40), our findings suggest that engaging the social networks of intervention participants can yield strong direct effects as well as a potential amplification of program effects via treatment spillover. While the intervention directly treated 1,642 individuals, the structure of the co-arrest network expanded the potential reach of the program to an additional 3,034 indirectly treated individuals—a 1.8-fold expansion of the affected population. In total, our findings imply that direct and spillover effects resulted in approximately 98 fewer gunshot victimizations over two years.

Programs like the one studied here are not a panacea for gun violence. The reduction in victimizations caused by the program equates to a roughly 1.5% reduction in city-wide gun victimization.[1] However, the results suggest that violence reduction programs can have a substantial impact on the incidence of victimization among targeted high-risk individuals and their peers. The program is scalable and, importantly, it minimizes traditional law enforcement responses that can have a negative impact on communities (41), especially incarceration. Our evidence suggests that gun violence interventions might amplify programmatic effects if designed more explicitly around maximizing network diffusion (21, 42, 43). For example, street outreach (13) or hospital intervention programs (11) might use formal network analytics to guide their efforts and place workers into those parts of a network or neighborhood experiencing acute rates of gun violence.

Field-interventions aimed at high-risk social networks have become a key policy tool for reducing gunshot violence. Studying peer effects provides an avenue for enhancing the efficacy of these interventions and improving our understanding of the role of social influence in the emergence and perpetuation of gun violence and violent conflict.

[1] This percentage reduction is calculated based on estimates of the mean two-year incidence of gunshot victimization in Chicago, IL during 2010–2016.

4. Materials and Methods

4.1. Network exposure mapping

To measure exposure to the program through social ties, we used data on 868,607 arrests in the jurisdiction of the Chicago Police Department (CPD) from 2007—2017. We created a co-arrest network in which an edge is present between two individuals if they had been arrested together. We located each program seed in the network and classified adjacent units as peers if those units were arrested alongside the seed at least once in the three-year period prior to the date that the seed was assigned to an intervention meeting. We define the exposure mapping $f(z, g)$ where $z_i = \{0, 1\}$ indicates whether a unit was assigned to the program and $g_i = \{0, 1\}$ indicates whether a unit has a tie to a seed (32, 44). For units with $z_i = 1$, i.e. the seeds, we denote a compliance indicator C_i which is 1 if the unit was compliant and 0 if the unit was non-compliant. For units with $z_i = 0, g_i = 1$, i.e. the peers, we denote a compliance indicator C_j which is 1 if the seed to whom the peer is connected was compliant and 0 if the seed was non-compliant.

For the estimate of the effect of additional spillover exposure on victimization, we performed a subset analysis in which we identified 550 peers with exposure to two seeds $Z_i = 0, g_i \geq 2$. For units connected to three or more seeds, we considered only the first two exposures to maintain a suitably large n for estimation. We denote the compliance indicator C_j which is 1 if the peer was connected to two compliant seeds and 0 if the peer was connected to two non-compliant seeds.

4.2. Identification strategy

Identifying the effect of the intervention is challenging due to the non-randomized assignment protocol. As a strategy for overcoming this challenge, we exploited compliance and non-compliance with the program[2] in order to identify the effect of compliance with the program among assigned seeds who, by virtue of their assignment, were deemed by the administrators of the program to have a similar baseline risk of victimization ex-ante (46). To identify the effect of compliance with the program, we compared two-year victimization outcomes among seeds $Z_i = 1$ with $C_i = 1$ and $C_i = 0$. To identify the effect of compliance spillover, we compared two-year victimization outcomes among peers $Z_i = 0, g_i = 1$, with $C_j = 1$ and $C_j = 0$. We denote the potential outcome $Y_i(1)$ which is the outcome that would be observed if unit i was in the compliant condition $C_i = 1$ and $Y_i(0)$ if unit i was in the non-compliant condition $C_i = 0$.

To estimate the compliance effect we rely on the conditional ignorability assumption $Y_i(0), Y_i(1) \perp C_i | X_i$ (47). That is, we assume independence between the potential outcomes and compliance status conditional on observed covariates X_i. We defined the effect of compliance, τ_c, as the average of the conditional expectation,

$$\tau_c = \mathrm{E}\Big[\, \mathrm{E}[\, Y(1) \mid C_i = 1, X_i] - \mathrm{E}[\, Y(0) \mid C_i = 0, X_i] \Big] \quad \forall\, Z_i = 1, \tag{1}$$

For the spillover effect, we made the analogous assumption $Y_i(0), Y_i(1) \perp C_j | X_i$, i.e. independence between the potential outcomes and the compliance status of the seed j to whom the peer unit i is connected conditional on the peer's observed covariates X_i. We defined the spillover

[2] A recent study finds that non-compliers and control units assigned by randomization had nearly identical post-assignment outcomes in a similar deterrence-based intervention in St. Louis (45).

effect of compliance, $\tau_s = 1$, as the average of the conditional expectation,

$$\tau_{s=1} = \mathrm{E}\Big[\, \mathrm{E}[\, Y(1) \mid C_j = 1, X_i] - \mathrm{E}[\, Y(0) \mid C_j = 0, X_i] \Big] \quad \forall\, Z_i = 0, g_i = 1 \tag{2}$$

Finally, we defined the spillover effect of double exposure to compliance $\tau_{s=2}$ as in Eq. 2, substituting $Z_i = 0, g_i \geq 2$ in order to restrict analysis to the subset of peers with connections to at least two seeds.

4.3. Suitability of the observed covariates

The identification strategy is valid if the observed covariates X_i provide an admissible back-door adjustment from the victimization outcome to compliance status (48). The back-door adjustment is admissible if there does not exist unmeasured confounders of victimization and compliance. In practice, it is not possible to demonstrate this admissibility. As our outcome is gunshot victimization, we included several covariates related to gunshot violence (49) and violence more generally, all measured pre-intervention. These include gunshot victimizations; arrests on violence-related charges, including first-degree assault and second-degree assault; and arrests on weapons-related charges, including carrying weapons without a permit and ownership of illegal firearms. To account for the risk-level in the co-arrest network surrounding each unit, we included counts of victimizations, arrests, arrests on violence-related charges, and arrests on weapons-related charges among adjacent units. To account for the connectivity of each unit, we also included degree (i.e. the number of co-arrestees). Finally, we include age, race, gender, and gang status, which are related to the incidence of gunshot victimization (50, 51, 52, 53). We assume this broad array of covariates adequately accounts for confounding due to selection into the compliant or non-compliant conditions. However, we cannot rule out unobserved factors that could bias the estimates, although our adjustment may partly reduce bias due to any unmeasured confounders that are correlated with the observed covariates (54, 55).

4.4. Estimation strategy

To estimate the compliance effect τ_C and spillover effect of compliance $\tau_{s=1}$ we used two models: (1) Bayesian Additive Regression Trees (BART) and (2) a simple difference-in-means (DIM).

BART is a sum-of-trees model that allows a flexible relationship between the victimization outcome Y_i, compliance indicator C_i (or C_j in the case of spillover), and observed covariates X_i. We use this nonparametric approach to avoid imposing a structural form on the relationship between the compliance indicator and the observed covariates, avoiding researcher-imposed choices regarding the ways in which the covariates might be associated with selection into compliance and instead allowing automatic detection of interactions and nonlinearities in this relationship (35). We fit the BART model using the dbarts R package (34) with the default settings (ntrees = 200, $\alpha = 0.95, \beta = 2$), 1000 burn-in MCMC iterations, and 5,000 posterior samples.

We contrast this approach with a simple difference-in-means (DIM),

$$Y_i = \alpha + C_i \quad \forall\, Z_i = 1, \tag{3}$$

where α is an intercept and C_i is the compliance indicator. Similarly, for the spillover effect we estimate,

$$Y_i = \alpha + C_j \quad \forall\, Z_i = 0, g_i = 1, \tag{4}$$

where C_j for the indicator for the compliance status of the seed to whom the peer is connected. This model simply compares the average probability of victimization within the compliant and non-compliant conditions without adjustment for confounding. By contrasting the results from BART with the DIM model, we are able to gauge the extent to which the relationship between the victimization outcomes Y_i and compliance status C_i (C_j for spillover) is moderated by the covariates X_i.

To further assess the effect estimates, we compared the results from BART with the results of a model which weights units based on covariate balancing propensity scores. See the supplementary materials for further details.

4.5. Estimating heterogenous effects

The BART model naturally detects interactions between the compliance indicator and observed covariates (33). To estimate the heterogeneity of the compliance and spillover effects, we used the fitted BART model to estimate the conditional average treatment effect (CATE) of compliance on each unit as a function of that unit's covariate profile. For the main effect we duplicate the seed covariate data into two $N \times K$ matrices, where N is the number of seeds and K is the number of covariates. We set the compliance indicator to 1 in the first matrix and 0 in the second. We then take 1,000 posterior draws for each matrix from the fitted model and subtract the second matrix from the first. The CATE is calculated by taking the mean of the rows and the 95% uncertainty bounds are calculated by taking the 0.025 and 0.975 quantiles (35). We repeat this process for the spillover effect, substituting in the peer covariate data and the fitted spillover BART model.

References

[1] G J Wintemute. The epidemiology of firearm violence in the twenty-first century united states. *Annual Review of Public Health*, 36:5–19, 2015.

[2] CDC. Web-based injury statistics query and reporting system (wisqars). *Natl. Cent. Injury Prev. Control, CDC, http://www.cdc.gov/injury/wisqars/index.html*, 2019.

[3] R D Peterson and L J Krivo. *Divergent Social Worlds: Neighborhood Crime and the Racial-Spatial Divide*. Russell Sage, New York, 2010.

[4] S Harper, J Lynch, S Burris, and G Davey Smith. Trends in the black-white life expectancy gap in the united states, 1983-2003. *JAMA*, 297(11):1224–1232, 2007.

[5] P Sharkey. The acute effect of local homicides on children's cognitive performance. *Proceedings of the National Academy of Sciences*, 107(26):11733–11738, 2010.

[6] M Tracy, A A Braga, and A. V. Papachristos. The transmission of gun and other weapon-involved violence within social networks. *Epidemiological Review*, 38(1):70–86, 2016.

[7] A V Papachristos, C Wildeman, and E Roberto. Tragic, but not random: the social contagion of nonfatal gunshot injuries. *Social Science & Medicine*, 125:139–150, 2015.

[8] A V Papachristos, A A Braga, and D Hureau. Social networks and the risk of gunshot injury. *Journal of Urban Health*, 89(6):992–1003, 2012.

[9] B Green, Horel T, and A V Papachristos. Modeling contagion through social networks to explain and predict gunshot violence in chicago, 2006 to 2014. *JAMA Intern Med*, 177(3):326–333, 2017.

[10] C Cooper, D M Eslinger, and P D Stolley. Hospital-based violence intervention programs work. *Journal of Trauma and Acute Care Surgery*, 61(3):534–540, 2006.

[11] J Purtle, J A Rich, J A Fein, T James, and T J Corbin. Hospital-based violence prevention: progress and opportunities. *Annals of Internal Medicine*, 163(9):715–717, 2015.

[12] J A Butts, C G Roman, L Bostwick, and J R Porter. Cure violence: a public health model to reduce gun violence. *Annual Review of Public Health*, 36:39–53, 2015.

[13] J M Whitehill, D W Webster, S Frattaroli, and E M Parker. Interrupting violence: how the ceasefire program prevents imminent gun violence through conflict mediation. *Journal of Urban Health*, 91(1):84–95, 1994.

[14] V Crandall and S L Wong. *Group Violence Reduction Strategy: Call-in Preparation and Execution*. U.S.D.o.J. The Office of Community Oriented Policing Strategies, Washington, DC, 2012.

[15] A A Braga and D Weisburd. The effects of focused deterrence strategies on crime: a systematic review and meta-analysis of the empirical evidence. *Journal of Research in Crime and Delinquency*, 49(3):323–358, 2012.

[16] A A Braga, A V Papachristos, and D Hureau. The effects of hot spots policing on crime: An updated systematic review and meta-analysis. *Justice Quarterly*, 31(4):633–663, 2014.

[17] A A Braga, D Weisburd, and B Turchan. Focused deterrence strategies and crime control: An updated systematic review and meta-analysis of the empirical evidence. *Criminology & Public Policy*, 17(1):205–250, 2018.

[18] J J Heckman. The scientific model of causality. *Sociological Methodology*, 35(1):1–97, 2006.

[19] S N Durlauf, S Navarro, and D A Rivers. Understanding aggregate crime regressions. *Journal of Econometrics*, 158:306–317, 2010.

[20] J Gravel and G E Tita. With great methods comes great responsibilities. *Criminology & Public Policy*, 14:559–572, 2015.

[21] T W Valente. Network interventions. *Science*, 337(6090):49–53, 2012.

[22] A Proestakis, E P di Sorrentino, H E Brown, E van Sluijs, A Mani, S Caldeira, and B Herrmann. Network interventions for changing physical activity behaviour in preadolescents. *Nature Human Behaviour*, 2(10):778–787, 2018.

[23] D M Kennedy, A A Braga, and A M Piehl. *Crime Mapping and Crime Prevention*, chapter The (Un)Known Universe: Mapping Gangs and Gang Violence in Boston. Criminal Justice Press, Monsey, NY, 1997.

[24] A V Papachristos and D S Kirk. Changing the street dynamic: evaluating chicago's group violence reduction strategy. *Criminology & Public Policy*, 14(3):525–558, 2015.

[25] National Network for Safe Communities. *Group violence intervention: An implementation guide.* U.S. Department of Justice, Office of Community Oriented Policing Services, Washington, DC, 2016.

[26] D Broockman and J Kalla. Durably reducing transphobia: A field experiment on door-to-door canvassing. *Science*, 352(6282):220–224, 2016.

[27] Clarke R V and D Weisburd. Diffusion of crime control benefits: observations on the reverse of displacement. *Crime Prev. Stud.*, 2:165–184, 1994.

[28] D M Kennedy, A M Piehl, and A A Braga. Youth violence in boston: gun markets, serious youth offenders, and a use-reduction strategy. *Law Contemp. Probl.*, 59:147—196, 1996.

[29] A A Braga and D Weisburd. Focused deterrence and the prevention of violent gun injuries: practice, theoretical principles, and scientific evidence. *Annual Review of Public Health*, 36:55–68, 2015.

[30] Y Charette and A V Papachristos. The network dynamics of co-offending careers. *Social Networks*, 51:3–13, 2017.

[31] E L Paluck, H Shepherd, and P M Aronow. Changing climates of conflict: A social network experiment in 56 schools. *Proceedings of the National Acdemy of Sciences*, 113(3):566–571, 2016.

[32] P M Aronow and C Samii. Estimating average causal effects under general interference, with application to a social network experiment. *Ann. Appl. Stat.*, 11(4):1912–1947, 2017.

[33] J L Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.

[34] H A Chipman, E I George, and R E McCulloch. Bart: Bayesian additive regression trees. *The Annals of Applied Statistics*, 4(1):266–298, 2010.

[35] D P Green and H L Kern. Modeling heterogenous treatment effects in survey experiments with bayesian additive regression trees. *Public Opinion Quarterly*, 76:491–511, 2012.

[36] J Hainmueller. Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political Analysis*, 20:25–46, 2012.

[37] L A Palinkas, I W Holloway, E Rice, C Hendricks Brown, T W Valente, and P Chamberlain. Influence network linkages across implementation strategy conditions in a randomized controlled trial of two strategies for scaling up evidence-based practices in public youth-serving systems. *Implementation Science*, 8(1):133, 2013.

[38] N A Christakis and J H Fowler. Social network sensors for early detection of contagious outbreaks. *PLoS One*, 5(9):e12948, 2010.

[39] D Centola. The spread of behavior in an online social network experiment. *Science*, 329(5996):1194–1197, 2010.

[40] D A Kim, A R Hwong, D Stafford, D A Hughes, A J O'Malley, J H Fowler, and N A Christakis. Social network targeting to maximise population behavior change: a cluster randomised controlled trial. *The Lancet*, 386(9989):145–153, 2015.

[41] J Legewie and J Fagan. Aggressive policing and the educational performance of minority youth, 2018.

[42] T W Valente. Social network thresholds in the diffusion of innovations. *Social Networks*, 18(1):69–89, 1996.

[43] E M Rogers. Diffusion of preventative innovations. *Addictive Behaviors*, 27(6):989–993, 2002.

[44] L Forastiere, E M Airoldi, and F Mealli. Identification and estimation of treatment and interference effects in observational studies on networks. *arXiv:1609.06245v4*, 2018.

[45] B Hamilton, R Rosenfeld, and A Levin. Opting out of treatment: self-selection bias in a randomized controlled study of a focused deterrence notification meeting. *Journal of Experimental Criminology*, 14(1):1–17, 2018.

[46] P R Rosenbaum. Choice as an alternative to control in observational studies. *Statistical Science*, 14(3):259–304, 1999.

[47] D B Rubin. Bayesian inference for causal effects: The role of randomization. *The Annals of Statistics*, 6:34–58, 1978.

[48] J Pearl. Causal inference in statistics: An overview. *Statistics Surveys*, 3:96–146, 2009.

[49] K J Steinman and M A Zimmerman. Episodic and persistent gun-carrying among urban african-american adolescents. *J Adolesc Health*, 35(5):356–364, 2003.

[50] P J Cook and J H Laub. After the epidemic: recent trends in youth violence in the united states. *Crime and Justice*, 29:1–37, 2002.

[51] R D Peterson and L J Krivo. Macrostructural analyses of race, ethnicity, and violent crime: recent lessons and new directions for research. *Annu Rev Sociol*, 31(1):331–356, 2005.

[52] R Jones-Webb and M Wall. Neighborhood racial/ethnic concentration, social disadvantage, and homicide risk: an ecological analysis of 10 u.s. cities. *J Urban Health*, 8(5):662–676, 2008.

[53] A V Papachristos, A A Braga, Piza E, and L S Grossman. The company you keep? the spillover effects of gang membership on individual gunshot victimization in a co-offending network. *Criminology*, 53(4):624–649, 2015.

[54] Z Fewell, G Davey Smith, and J A C Sterne. The impact of residual and unmeasured confounding in epidemiologic studies: A simulation study. *American Journal of Epidemiology*, 166(6):646–655, 2007.

[55] R H H Groenwold, J A C Sterne, D A Lawlor, K G M Moons, A W Hoes, and K Tilling. Sensitivity analysis for the effects of multiple unmeasured confounders. *Annals of Epidemiology*, 26(9):605–601, 2016.