A Preliminary Analysis of Spillovers in READI Chicago — Early Results from "Using Network Data to Measure Social Returns and Improve Targeting of Crime-Reduction Interventions" *

Ashley Craig[†] Sara B. Heller[‡] Nikhil Rao[§]

November 23, 2022

1 Introduction

This note presents preliminary estimates of spillover effects from the Rapid Employment and Development Initiative (READI) onto peers of treated individuals. The analysis here is part of a separate, broader project to estimate how interventions that reduce violence spread through participants' social networks, as well as the implications for how to target such interventions to maximize their benefits. This report provides an early preview of results for a subset of that project.

Convincingly estimating social spillovers faces two key challenges: measuring social networks and causally identifying peer effects (Angrist, 2014; Manski, 1993). In our ongoing broader project, we solve the first challenge by combining population-wide administrative police and school records in Chicago to capture different kinds of social connections: co-arrest, co-victimization, school, residential, and neighborhood networks. These measures allow us to observe a range of different social connections—who spends time together in ways that result in arrest or victimization, who takes classes together, who lives together (siblings and co-habitants), and who lives near each other.

We solve the second challenge by combining this network information with exogenous variation in crime and violence generated by four existing RCTs in Chicago: Becoming a Man (BAM), two experiments on One Summer Plus (OSC1 & OSC2), and READI (Bhatt et al., in progress; Davis and Heller, 2020; Heller, 2014; Heller et al., 2017). These four RCTs all offered a mix of cognitive behavioral therapy and other elements targeting non-academic skills (including supported employment and/or personal mentoring). And they all focused on low-income populations, predominantly consisting of Black and Hispanic juveniles & young adults, in Chicago.

Because we have administrative data going back to 1999, we can construct the social network between individuals at the time each intervention was randomly assigned. Since ties to study individuals are a baseline characteristic and treatment is randomly assigned, indirect exposure—whether someone's peer is treated—is also effectively randomly assigned. Using this variation, we can therefore test whether and how treatment changes behavior, both via direct participation and

*This project has been supported with funding from the Institute for Firearm Injury Prevention at the University of Michigan, J-PAL North America, LSA Social Science Scholarship/Research Funding, and a Faculty Research award from the Michigan Institute for Teaching and Research in Economics (MITRE). We are grateful to the Chicago Police Department for allowing the use of their data, to University of Chicago Crime Lab for their support of the project, and to the READI team for answering our questions and providing helpful feedback. Micah Baum and Julia Yates have provided excellent research assistance. All opinions and any errors are our own and do not necessarily reflect those of our funders, implementing partners, data providers, or other government agencies.

[†]University of Michigan: ashcraig@umich.edu.

[‡]University of Michigan & NBER: sbheller@umich.edu.

[§]University of Chicago & University of Michigan: nikhrao@umich.edu.

via indirect exposure to a treated peer. Separating direct effects from exposure effects for the sample that was part of an RCT allows us to unpack how social spillovers may have influenced the initial intent-to-treat (ITT) estimate, and to recover the direct effect of the RCT (i.e., treatment versus nothing) on treated individuals. Additionally, we can quantify spillovers onto first-degree peers of treated individuals who were outside the original sample frame by measuring outcomes for those who were not part of the RCT.

In this research note, we report early results from a small part of the overall project. We focus solely on one type of network—based on co-arrests during the previous 5 years—for one of the RCTs—READI. We also limit our attention to only a binary, first-degree exposure effect: the impact of having at least one peer treated (where peers are defined by having been arrested together pre-randomization).[1] Although there is still much left to learn about READI's spillover effects beyond first-degree co-arrestee peers, as well as power to be gained by adding the other RCTs, we report these early results to help with the initial interpretation of the READI RCT results (Bhatt et al., 2022).

We focus on the three main outcomes of pre-specified interest in the READI RCT: shooting and homicide arrests, shooting and homicide victimizations, and other types of serious violent-crime arrests (aggravated assault or battery, robbery, and sexual assault). The READI sample was selected to be at high risk of these outcomes, but the peers of these individuals may not be as involved in very serious violence. To capture behavior that is more common among the READI peer group, we therefore also report direct and exposure effects for broader categories of arrests: violent, property, drug, and other (everything else, which includes arrests for trespassing, vandalism, illegal use of a weapon, disorderly conduct, and so forth).

We implement the methods developed in Aronow and Samii (2017) to test for the presence of first-degree spillovers.[2] The analysis starts with measurement of the co-arrest network prior to random assignment. We then define the different ways in which individuals can be exposed to the intervention ("exposure states") given the network. These exposure states define the set of potential outcomes and treatment effects that are relevant for each individual. Finally, we use the experimental design to calculate the probability of being in each exposure state for each individual, and use that probability as a weight (or, in the appendix, a regression control) to adjust for selection into each exposure state when estimating treatment effects.

The co-arrest network we define isolates the sub-populations affected by the RCT in different ways. Of the 2,456 people in the initial READI RCT, 966 have a co-arrest tie to at least one other study member at the time of random assignment, leaving 1,490 unaffected by potential spillovers via co-arrest ties. We trim the in-network sample slightly to avoid instability that can come from very high or low likelihoods of exposure, leaving 906 individuals in the original READI RCT to identify within-RCT spillovers. The reduction in sample size relative to the initial RCT, along with the fact that the key outcomes of interest in the READI study are quite rare, highlights the limited amount of statistical power available to answer the within-RCT spillover question.

Another 6,367 people are tied to study members but not in the initial RCT themselves (6,210 after trimming). We confirm that serious violence outcomes are rarer among this group than those selected for the READI RCT—about half as common as in the READI sample—making the broader arrest categories of more central interest for this group.

Having defined the relevant populations within the social network, we can then estimate several

[1] It is of course plausible that other kinds of ties and other types of exposure shape potential outcomes. We leave this question for future work, focusing here on the question of whether first-degree co-arrestee exposure matters.

[2] Appendix A implements a related analysis strategy from Borusyak and Hull (2021), which uses a regression-based approach rather than Aronow & Samii's weighting-based approach, to assess robustness. We do not focus on this approach in the main test because of unresolved questions about how to conduct inference for estimates involving those who could be both directly treated and exposed to treated peers.

different treatment effects of interest. We start by estimating the direct effect of treatment for individuals in READI who are not connected to any other study individual.[3] There is a large and marginally significant decline in shooting and homicide arrests among these 1,490 unconnected individuals. In fact, the ITT for this subsample is slightly larger than the ITT in the main READI study, a 53 percent decline ($p = 0.097$), compared to 44 percent in the original RCT's ITT.

In terms of social spillovers among the socially-connected populations, our clearest result is that READI decreases drug-crime arrests in a way that was masked in the initial RCT. Direct treatment generates a decline of about 30 percent in these arrests among the treatment group ($p = 0.034$). But those in the control group who have treated peers also show some decline in drug arrests, while peer exposure does not generate an additional decline for the treatment group. Because spillovers decreased drug-crime arrests among controls beyond what they would have experienced if the program was never delivered, the main RCT did not pick up the effect. We also show evidence of a similarly-sized and statistically significant decline among peers of the treatment group who were not part of the study (-35 percent, $p = 0.003$).

Results for the three serious violence outcomes that are the focus of READI are less conclusive. Their rarity, combined with the smaller sample size of those who are in the the co-arrest network, means confidence intervals are quite wide. There is some suggestive evidence that spillover effects may mask a direct effect that decreases serious violent-crime arrests that are not shootings or homicides. Directionally, peer exposure decreases these arrests among controls but increases them among treatments, leading the ITT to under-state the direct impact of READI. But given the width of the confidence intervals, strong conclusions should await further analysis. In future work, we plan to explore newly-developed methods to increase statistical power, as well as incorporate measures of other social networks that may be relevant for the behavioral transmission of serious violence.

2 Background and Related Literature

There are many theories of why people might commit crimes together. Some emphasize social connections to particular people who facilitate and coordinate crime (Ballester, Calvo-Armengol, and Zenou, 2006, 2010; Diaz et al., 2018; Liu and Ye, 2012; Liu, Wang, and Eck, 2014; Lindquist and Zenou, 2014). In others, co-offending is appealing because it provides benefits such as the ability to commit more profitable crimes, or non-monetary benefits such as social support (Tremblay, 1993; Weerman, 2003). A different set of theories focuses on how social dynamics change individuals' perception of risks and benefits: Being in a group could increase feelings of anonymity or encourage risk-seeking, lowering the threshold at which people engage in a crime (Granovetter, 1978; McGloin and Piquero, 2009; O'Brien et al., 2011). Alternatively, young people may seek experienced co-offenders from whom to learn skills (McAndrew, 2000; McCarthy and Hagan, 1995), or co-offend to conform to social norms or demonstrate loyalty (Akerlof and Kranton, 2010; Patacchini and Zenou, 2009; Tankard and Paluck, 2016; Warr, 2002). And extreme behavior like shooting involvement may spread among those who spend time together engaging in risky behavior, as in contagion models (e.g., Green, Horel, and Papachristos, 2017).

The type of social influence matters for how to target crime-reduction interventions. Without peer effects, the impact of an intervention is maximized by targeting those whose own drop in crime would be largest. When peer influence does matter, the type of social influence is key to understanding optimal targeting. Economic theory has highlighted a distinction between models with complementarity versus conformity. If co-offending is a function of complementarity, an intervention that targets central individuals with many ties would therefore be most effective,

[3] People may be unconnected because they only share co-arrests with those not part of the study, were only arrested alone, or were not arrested at all prior to randomization.

because they have greater aggregate peer influence (Fortin and Boucher, 2016). However, this may not be the most effective strategy if individuals conform and peer effects are determined by the average of peer actions. On the other hand, if co-offending groups form spontaneously among friends, the peers of more connected individuals might have more easily available substitutes; this would suggest that it is less effective to target someone with more social ties.

While we know that peers with whom youth spend time—in schools, neighborhoods, and the criminal legal system—have a causal influence on their own offending decisions (e.g., Bayer, Hjalmarsson, and Pozen, 2009; Bhuller et al., 2018; Billings, Deming, and Rockoff, 2013; Billings, Deming, and Ross, 2019; Billings and Schnepel, 2017; Damm and Dustmann, 2014; Damm and Gorinas, 2016; Diaz, 2020; Drago and Galbiati, 2012; Ouss, 2011; Philippe, 2017; Stevenson, 2017), existing studies only establish that a person's social ties affect his or her own decisions about crime. They cannot do the reverse: estimate how changes in behavior spread through a whole social network. This is more demanding than simply knowing that exposure to person A has a causal effect on the behavior of person B, since we need to understand how many people are socially connected to A, which of those people are affected by treating A, and whether that varies based on who is treated or who the peer is.

Estimating these network effects is crucial to understanding the nature of social influence and the importance of different theories of criminal social interactions. It is also crucial for interpreting the results of many prior RCTs, where the ITT may mask the true role of treatment on crime by ignoring how an intervention affects the peers of treated individuals.

In addition to the reduced-form peer effect literature in the economics of crime, there is a huge literature outside of crime that estimates these kinds of network spillover effects. Development economists have tested how health and agricultural technology use spills over to geographically and socially close peers (e.g., Bhattacharya, Dupas, and Kanaya, 2013; Miguel and Kremer, 2007; Miller and Mobarak, 2014; Oster and Thornton, 2012), and how product information diffuses through networks (e.g., Banerjee et al., 2013; Beaman et al., 2021; Beaman and Magruder, 2012; Breza and Chandrasekhar, 2019; Cai, Janvry, and Sadoulet, 2015). Education economists and psychologists have examined how interventions that are randomly assigned to individuals spill over through self-reported friend networks (e.g., Bursztyn, Fujiwara, and Pallais, 2017; Dieye, Djebbari, and Barrera-Osorio, 2014), or vary with the network centrality of treated individuals, possibly due to their influence on social norms (e.g., Paluck, Shepherd, and Aronow, 2016; Shepherd and Paluck, 2015; Tankard and Paluck, 2016). Work in development economics has highlighted how interventions can change the structure of a network (Banerjee et al., 2013; Comola and Prina, 2014; Feigenberg, Field, and Pande, 2013).

We aim to bring the methods developed in these literatures to the network analysis of crime and violence. Efforts by others to estimate how crime interventions spread through networks provide reason to expect that the effects may be large, but these results are either based on non-experimental variation that could confound spillovers with pre-existing differences or common shocks (e.g., Charette and Papachristos, 2017; Dominguez, 2022; Green, Horel, and Papachristos, 2017; Lindo and Padilla-Roma, 2018; Wood and Papachristos, 2019) or are based on single studies that generate significant statistical power limitations (Abdul-Razzak et al., In Progress; Diaz, 2020). To our knowledge, our broader project will be the first to use convincing identification across multiple RCTs and multiple measures of social networks to estimate how crime changes spread through social networks.

3 Defining Potential Outcomes and Treatment Effects

The core challenge in a setting with spillovers is that the stable unit treatment value assumption (SUTVA) does not hold. This means that there are no longer simply two potential outcomes, $Y_i(T_i = 1)$ and $Y_i(T_i = 0)$, where T_i is an indicator for random assignment to treatment. Rather, an individual's potential outcomes could now potentially depend on the full vector of others' treatment assignments, $\boldsymbol{T}$, as well as their own. Note that throughout, we define treatment as random assignment to the treatment group, as opposed to actual participation in READI.[4]

In our future work, we will test a range of null hypotheses about which types of exposure matter. For this early analysis, however, we will assume that distant exposure to those far away in the network is irrelevant to potential outcomes. We allow just two types of "treatment" to matter: direct exposure (being treated one's self) and first degree indirect exposure (having at least one first-degree peer treated). Specifically, let $T_i \in \{0, 1\}$ indicate direct treatment and $E_i \in \{0, 1\}$ indicate indirect exposure.

Using this notation, we define $D_i = d_{T_i,E_i}$ as indicators for the four mutually exclusive and collectively exhaustive 'exposure states' in which each READI study member could be. First, he may be both treated and exposed to a treated peer ($d_{1,1}$). Second, he may be untreated and exposed to a treated peer ($d_{0,1}$), Third, he could be treated and not indirectly exposed ($d_{1,0}$). Finally, he may be untreated and not indirectly exposed ($d_{0,0}$). These four exposure states determine the potential outcomes for each individual, $Y_i(d_{T_i,E_i})$.

In practice, not everyone in the original RCT has all four potential outcomes. Individuals with no pre-existing social ties to other individuals in the RCT have no potential to be exposed to treated peers; they have only two relevant exposure states, ($d_{1,0}$ and $d_{0,0}$). Because this group has no peers to influence them, they are in the more standard world of 2 potential outcomes. For them, the only effect of interest is direct treatment: $E(Y|d_{1,0}) - E(Y|d_{0,0})$.

Those in the original RCT who are connected to other study members could experience all four of the exposure states, and so have four potential outcomes. We consider three different treatment effects of interest: 1) the effect of direct exposure versus nothing $E(Y_i(d_{1,0}) - Y_i(d_{0,0}))$, 2) the effect of peer exposure for treated individuals $E(Y_i(d_{1,1}) - Y_i(d_{1,0}))$, and 3) the effect of peer exposure for control individuals $E(Y_i(d_{0,1}) - Y_i(d_{0,0}))$.

As shown in Vazquez-Bare (forthcoming), the standard difference-in-means ITT effect is a weighted average of these three effects:

$$\begin{aligned} \text{ITT} = \overbrace{E(Y_i(d_{1,0}) - Y_i(d_{0,0}))}^{\textbf{Direct Treatment Effect, No Exposure}} &+ \Pr(D_i = d_{1,1}|T_i = 1)\overbrace{E[Y_i(d_{1,1}) - Y_i(d_{1,0})|T_i = 1]}^{\textbf{Exposure effect for treated}} \\ &- \Pr(D_i = d_{0,1}|T_i = 0)\underbrace{E[Y_i(d_{0,1}) - Y_i(d_{0,0})|T_i = 0]}_{\textbf{Exposure effect for control}} \end{aligned} \tag{1}$$

Equation 1 demonstrates that, in a context with spillover effects, the ITT no longer necessarily measures the direct effect of being in the READI treatment group versus nothing. It will be a misleading estimate of that direct effect to the extent that either: (i) the probability of indirect exposure varies between the treatment and control groups; or (ii) the marginal effect of direct treatment varies depending on whether one is also exposed indirectly.

Case (i) may happen by chance in a finite sample, but not in expectation. Due to random assignment, the probability of an individual being exposed is independent of their direct treatment

[4]The experiment has imperfect take-up, such that T_i is an imperfect measure of actual receipt of READI. But because the literature has not yet dealt with how to handle instrumental variable estimation in the context of network econometrics, we focus throughout on the intent-to-treat estimation of spillovers.

status:

$$\Pr(D_i = d_{1,1}|T_i = 1) = \Pr(D_i = d_{0,1}|T_i = 0) = \Pr(E_i = 1) \quad (2)$$

Thus, in expectation, whether the direct effect of treatment $E(Y_i(d_{1,0}) - Y_i(d_{0,0})$ and the ITT coincide amounts to whether the marginal effect of exposure varies by treatment status.

If exposure to a treated peer has the same effect on treated and control individuals, then the second two terms in Equation 1 are equal in expectation, and the ITT from the experiment is an unbiased estimate of the effect of direct treatment. But this will not be the case if indirect exposure has a different effect on treatment and control individuals. For example, if indirect exposure spills over onto controls with an effect of the same sign as direct treatment, but exposure has no marginal effect other treated units—as one might expect if it is a matter of information being passed on to the control group—the ITT will underestimate the direct impact of the RCT. Alternatively, the ITT would overstate the direct effect if exposed controls had the opposite reaction as those directly treated—as might happen if controls stepped into opportunities for crime that treated peers had abandoned.

In addition to having direct and exposure effects on the original study members who contribute to the ITT, the intervention may also affect the peers of READI study members who were not part of the initial RCT at all. Because this non-RCT group could not possibly have received direct treatment, they only have two relevant exposure states: untreated and exposed to a treated peer ($d_{0,1}$) and untreated and not exposed to a treated peer ($d_{0,0}$). For them, the only treatment effect of interest is the effect of exposure: $E(Y_i(d_{0,1}) - Y_i(d_{0,0}))$. This effect could be the same as the exposure effect for control study members. But it does not have to be, given that the two groups are selected and exposure effects could be heterogeneous.

In the analysis below, we estimate the treatment effects of interest for these different groups to help separate and measure the various ways READI affects the population in and around the RCT.

4 READI

The Rapid Employment and Development Initiative (READI) provided 18 months of supported work, cognitive behavioral therapy, and other social supports to men in five high-violence Chicago neighborhoods. Study members were selected to be the group of men age 18 and over at the highest risk of shooting or being shot. This was accomplished in three different ways: a predictive algorithm; outreach workers with on-the-ground experience; and screenings from those leaving prison, jail, or on parole. In total, 2,456 men were randomized to be offered READI or to be part of a control group that was free to pursue other available programming.

READI's pre-specified primary outcome of interest was an index of three proxies for serious violence involvement: counts of shooting and homicide arrests, shooting and homicide victimizations, and arrests for other serious violent-crime arrests (aggravated assault or battery, robbery, and sexual assault). The main RCT found large declines in shooting and homicide arrests that were statistically significant on their own, but crossed standard significance thresholds after adjustments for multiple testing (Bhatt et al., 2022, in progress). Point estimates on shooting and homicide victimizations were negative, while estimates for other violent-crime arrests were positive. But neither were statistically significant. When weighted by the social cost of crime, the changes in all measures of crime suggest a 50 percent decline social costs ($p = 0.03$), the benefits of which outweigh spending on the program by at least 3.4 to 1. We refer interested readers to the early results brief and upcoming main READI RCT paper for further details (Bhatt et al., 2022, in progress).

5 Data and Descriptives

We combine program data from READI ($N = 2,456$) with detailed individual-level administrative data covering the universe of Chicago Police Department (CPD) arrest and victimization records going back to 1999.[5] The arrest data include unique identifiers, based on fingerprints, that link arrests to an individual, incident identifiers to link arrests to incidents, details on the exact time and location of arrests, and the reason for the arrest. We use records from 2005 onward, since the quality and accuracy of identifiers are higher in this period.

Our primary outcome variables will be the same counts of serious violence as in the original READI analysis – counts of shooting & homicide arrests, shooting & homicide victimizations, and other serious violent-crime arrests. These three outcomes were most relevant for READI because it was aimed at stopping gun violence for individuals at the highest risk of gun violence involvement. These outcomes may not be the most relevant outcomes for peers of READI study members, however, especially those not in the study themselves. As a result, we also report the impacts on broader violent, property, drug, other, and total arrests.

Defining the co-arrest networks

Co-arrest ties are a common way to measure co-offending (e.g., van Mastrigt and Carrington (2019). This is a relationship likely to reflect joint decision-making about crime and time spent together, and non-experimental evidence suggests such ties transmit the risk of gun violence (Green, Horel, and Papachristos, 2017; Wood and Papachristos, 2019). Since the CPD records contain both incident and individual identifiers, we are able to observe all the individuals arrested together as part of the same incident.[6] We call individuals 'tied together' or peers if they were part of the same incident in the 5 years before randomization. Our focus for this note is on binary exposure (having at least one tie) via first-degree peers, though future work will test whether spillovers travel farther through more distant social ties and whether the intensive margin of exposure matters.

Identification relies on the fact that we only use social tie information prior to randomization. In the case of READI, study participants were randomized in a rolling fashion between August 2017 and March 2020. In cases where co-arrestees were randomized at different times, some social ties may be pre-randomization for one study member but use post-randomization information for another.

As a result, after forming a network map using information 5 years prior to each randomization date, we purge each of these networks of ties that are not strictly orthogonal to treatment.[7] Ideally, we would retain all links between 2 people that existed before either of them were randomized. Because individuals were randomized into READI on 152 different dates, however, constructing separate pre-period networks for each date and calculating exposure probabilities within each network is computationally intensive. To simplify the computation, we place the 152 randomization dates into 41 different groups (we call these 'randomization network groups'). For the most part, these are month-year groupings of the dates, but randomization rounds with 50 or more people are assigned their own groups. We then use the earliest randomization date in a given group to

[5] Our broader project also uses Chicago Public School data to measure social networks within schools with class-taking data, as well as provide information on residential address to identify sibling, co-habitant, and neighborhood spillovers. These data contain 5.2 million unique individuals. In this early analysis, we focus on CPD arrest data.

[6] Note that when incident identifiers are missing, we assume anyone arrested at the same time and location were part of the same incident.

[7] In theory, it would be somewhat cleaner to define a single network using information prior to the earliest randomization date. While this would ensure that no ties in the network are endogenous to treatment, in practice we would lose a lot of useful information given that individuals were randomized in the study over a period of almost 3 years.

define the extent of the 5 year pre-period that we use to construct the networks. In practice, this simplification loses information on less than 1% of ties that were formed between the group's randomization date and the actual randomization date.

Using the 41 networks specific to the time of each set of randomization dates, we drop ties that could have formed as a result of a peer being treated. That is, if person A and B formed a tie by being arrested together in April 2019, with person A randomized in January 2019 and person B randomized in July 2019, the tie could have been influenced by whether A was treated. As a result, we do not count the tie for either person. There are 146 people (out of 1,112 who have co-arrests with other READI individuals in their pre-randomization network) whom we count as unconnected, even though they might have been exposed through this kind of potentially endogenous tie formation. The cost of this decision is to ignore ties that may have transmitted behavior, which introduces some mismeasurement of the network. But the benefit is that we fully purge our analysis of any possibility of endogeneous tie formation.

Describing the co-arrest network around the READI sample

Each of these 41 networks contains around 60,000-90,000 people, suggesting that the potential set of people who may be indirectly affected by READI is quite large.

Table 1 compares baseline characteristics of those present in the study but unconnected via the arrest networks (Column 1) to those who are present in the study and the arrest networks (Column 2), the basis for our spillover analysis for those in the initial RCT. Column 3 displays characteristics of those who are first-degree co-arrest peers of READI study members but not in the initial RCT themselves. This group is the basis for our non-RCT spillover analysis.

Since the intervention targeted criminally-active men with a high risk of serious violence, we see in column 1 that the number of prior arrests even among the unconnected sample is large: These individuals have been arrested around 14 times in the past.[8] Because of the way the original study sample was selected, individuals are in their mid-to-late 20s and are mostly Black individuals. Even though they are not connected to other RCT individuals, they have around 2 peers on average. The share of co-arrests row reports the proportion of someone's arrests that involves at least one other person, a measure of how 'social' crime – for this population, 13% of their arrests are with someone else.

Column 2 limits the analysis to individuals who were in the initial RCT and are connected to the other READI study members – the relevant sample for assessing spillovers within the study.[9] This group makes up 39% of the study. They tend to be younger, with relatively more Black individuals than the unconnected sample. Additionally, the study members who are part of the co-arrest network have more active criminal histories than their unconnected counterparts: They have almost 7 more prior arrests on average, have around 8 peers, and co-offend for a larger share of their prior arrests (28%). These individuals are more connected on average than individuals in the larger pre-READI co-arrest networks (which includes everyone, regardless of study membership or ties to study members) – the average number of peers in these broader networks is 2.7. However, these individuals are slightly less 'social' than average. For comparison, the average individual in the broader pre-READI networks co-offends 36% of the time.

Column 3 displays the same statistics for individuals who are not in the READI RCT themselves but are first degree peers of these study individuals, i.e., non-study people who have been arrested

[8]Note that our numbers vary very slightly from those reported in the main READI RCT, because the procedure to link individuals across administrative datasets for our broader project involves more administrative datasets. As a result, our person identifiers differ slightly in a small number of cases.

[9]Note that the sample sizes in Columns 2 and 3 are slightly different from the sample sizes used in our analyses of exposure within and outside of the study. This is due to a trimming procedure that we employ to deal extreme probability weights, which we describe below.

with one of the READI study members in the 5 years prior to their network's randomization date. Because there are over 6,000 immediate peers of study members, spillovers would not have to be large to have substantial social impacts. These peers have fewer prior arrests, consistent with fact that the READI study targeted high risk individuals. While this group seems similar to the study samples in terms of race, they tend to be slightly older than those in the in-study spillover group and no longer consist wholly of men.

Finally, having defined the pre-randomization network, we can also calculate the probability each of these groups was actually exposed to a treated peer (final row of Table 1). For the 966 men who were in the READI RCT and part of the co-arrest network, 75 percent of the treatment group and 71 percent of the control group had at least one treated peer. Within the 6,367 first-degree peers of READI study members not in the RCT, 59 percent had at least one treated peer.[10]

6 Methods

In section 3, we defined treatment effects for three distinct groups. The first is comprised of those in the original RCT but with no potential peer exposure through the co-arrest network. For them, the standard ITT effect is the only treatment effect of interest:

- Effect of direct treatment $= E[y(d_{1,0})] - E[y(d_{0,0})]$

Since this effect does not involve any peer exposure, we can estimate this quantity as in a standard RCT, and conduct inference using heteroskedasticity-robust standard errors.

The second group is comprised of individuals who were part of the original RCT and have ties to others in the RCT. For them, there are four relevant exposure states—$d_{0,1}$, $d_{0,0}$, $d_{1,1}$, and $d_{1,0}$—leading to the three components of the ITT[11] we would like to separate:

- Direct treatment with no exposure $= E[y(d_{1,0})] - E[y(d_{0,0})]$
- Exposure effect for treated individuals $= E[y(d_{1,1})] - E[y(d_{1,0})]$
- Exposure effect for untreated individuals $= E[y(d_{0,1})] - E[y(d_{0,0})]$

Lastly, for individuals who were not part of the RCT and so could never receive direct treatment, there were two relevant exposure states, $d_{0,1}$ and $d_{0,0}$, implying one treatment effect of interest:

- Effect of exposure $= E[y(d_{0,1})] - E[y(d_{0,0})]$

Our main estimation approach for these treatment effects comes from Aronow and Samii (2017). The first step is to limit the data to the population which has non-zero probability of the relevant exposure states. With this sample, we can calculate the pieces necessary to estimate each treatment effect. Because we can observe individuals in each exposure state, we can observe each $Y_i(d_{j,k})$ from the data. But the simple difference in averages of $Y_i(d_{j,k})$ would not give an unbiased estimate of the treatment effects, because there is selection into each exposure state (e.g., those with more co-arrests and connected to more people are more likely to be exposed to a treated peer). The key

[10] These probabilities vary slightly once we trim the sample for estimation. In the trimmed sample, the average exposure is 74 percent for the treatment group, 70 percent for the control group, and 58 percent for the non-RCT peers.

[11] There are other combinations of differences in exposure states that one could also estimate, e.g., the direct treatment effect for exposed individuals, $E[y(d_{1,1}) - y(d_{0,1})]$. We focus on estimation of the effects that are most related to what the original ITT estimated, which are the key for understanding the effect of treatment versus nothing and the effects of being exposed to a treated peer.

contribution from Aronow and Samii (2017) is how to account for the fact that the probability of being in a given exposure state is non-random.

Because we both observe the network structure and know each individual's probability of direct treatment from the experimental design, we can calculate the probability of being in each exposure state from the data. In practice, this involves re-assigning treatment a large number of times according to the experimental design, measuring how often each individual is either directly treated or exposed to a treated peer within each randomization, and averaging across the iterations to capture the probability that each individual is in each of the exposure states, $\pi(d_{j,k})$.

Aronow and Samii (2017) show how to use this probability, effectively as a propensity score weight, to estimate the average outcome in each exposure state:

$$\hat{y}(d_{j,k}) = \frac{1}{N} \sum_i \frac{\boldsymbol{I}(D_i = d_{j,k}) Y_i}{\pi_i(d_{j,k})} \tag{3}$$

In practice, these inverse-probability weighting-type estimators are sensitive to extreme probabilities. In our setting, we therefore discard units that have extreme probabilities of being in any of the exposure states. We define 'extreme' as being in the top or bottom 1.25% of the propensity score distribution.

The estimated causal exposure effects are then the difference between the estimated average potential outcomes in each exposure state. For the within-RCT spillovers, we limit the sample to everyone in the original READI RCT who is also in the co-arrest network and tied to another READI study member. For this sample, we can estimate:

- Direct treatment with no exposure $= \hat{y}(d_{1,0}) - \hat{y}(d_{0,0})$
- Exposure effect for treated individuals $= \hat{y}(d_{1,1}) - \hat{y}(d_{1,0})$
- Exposure effect for untreated individuals $= \hat{y}(d_{0,1}) - \hat{y}(d_{0,0})$

For the non-RCT individuals with ties to those in the RCT, we limit to the sample to those in the arrest data with ties to READI study members, but who are not part of the original RCT. For them, we are interested in:

- Effect of exposure $= \hat{y}(d_{0,1}) - \hat{y}(d_{0,0})$

We note that this estimation approach relies on the researcher correctly specifying the set of exposure states. Failure to do so introduces bias through the inclusion of effects of exposure states that had not been specified as being important, if being in those unspecified states (or their effects) are correlated with being in the exposure states included in the analysis.

6.1 Variance Estimation

Variance estimation is complicated by the fact that observations depend on each other through the network structure. We use the conservative variance estimators from Aronow and Samii (2017).[12] See Aronow and Samii (2017) for further details on the standard errors and confidence intervals.

[12] These variance estimators require accounting for the probabilities that each pair of units are in the same exposure states together, as well as in different pairs of exposure states together. Some of these probabilities can be zero in practice. Aronow and Samii (2013, 2017) provide correction terms that account for such cases that are guaranteed to have non-negative bias, and it is in this sense that these variance estimators are conservative.

6.2 Inclusion of Controls

Aronow & Samii also provide refinements to their estimator to account for residual variation in the outcomes of interest. These essentially involve using a regression to residualize the set of covariates out of outcomes of interest, separately by exposure state. Using estimated coefficients from this regression, the average fitted value across the entire sample is added back to the residualized outcomes for the given exposure state. Then, adjusted outcomes are used to estimate the average potential outcomes in each of the exposure states, and estimation proceeds as normal.

All results in this note include baseline covariates. Our covariates include indicators for race, bins of age at randomization, and indicators for missing race and age. We also include indicators for having experienced 1, 2 and 3 or more relevant prior incidents. For all the arrest outcomes, the relevant incidents include the arrests for the following crimes: serious violence, property, drug, other crimes and other violent crimes. For the victimization outcomes, we include these indicators for the following incidents: nonfatal shootings, non-violent victimizations and other violent victimizations.

6.3 An Alternative Estimator

It is also possible to estimate the exposure effects in a regression framework by controlling for $\pi_{d,k}$ rather than using it as a weight (Borusyak and Hull, 2021). While regression will weigh observations differently from the direct weighting approach described above, the idea underlying the identification and estimation of exposure probabilities is quite similar. Borusyak and Hull (2021)'s approach has several advantages: potential power gains, the ability to test for non-binary exposure effects (e.g., the number or share of peers treated), and more straightforward adjustments for baseline covariates.

However, inference is complicated by the non-independence of observations. In cases where there are only two possible exposure states, and so one treatment effect of interest, inference is straightforward. A sharp null hypothesis that there is no effect of treatment can be tested via randomization inference, accounting for the non-independence of observations. But inference is more complicated if there are more than two exposure states stemming from "multiple treatments" (e.g., direct treatment and indirect exposure both potentially matter), as is the case for our within-RCT population. Re-randomizing everyone in the study then varies two treatments at once (direct and indirect). As a result, there is no simple way to use randomization inference to test sharp null hypotheses, such as the absence of indirect exposure effects specifically, or the absence of direct effects.[13]

Appendix A provides details on the Borusyak and Hull (2021) estimation and compares these regression-based estimates to our baseline Aronow & Samii estimates. It reports inference using randomization inference for the non-RCT group that only has 2 potential outcomes. For the within-RCT spillover estimation, we do not report confidence intervals for now, since correct inference in these "multiple treatments" is an active area of econometric research. However, we compare these point estimates to the estimates derived using the Aronow & Samii methods, for which we can obtain correct, if conservative, confidence intervals.

7 Results

As laid out in Sections 3 and 6, we are interested in several different treatment effects across three populations. We start with estimates of the direct effects of READI for the 1,490 study members

[13] Some methods, as in Athey, Eckles, and Imbens (2018), fix the direct treatment status of a given set of individuals and then re-randomize the remaining individuals, but this comes with large power losses and requires a principled way of choosing which set of individuals is fixed, and which is re-randomized. In our setting, these methods appear under-powered and too sensitive to the choice of whom to fix.

who were not co-arrested with anyone else in the study prior to randomization in Figure 1. Panel A focuses on the three components of READI's primary outcome: shooting and homicide arrests, shooting and homicide victimizations, and other serious violent-crime arrests. Figure 1 shows a large and marginally significant decline in shooting and homicide arrests (-53 percent, $p = 0.097$). This result, as well as the positive but imprecise point estimate on other serious violent-crime arrests, mirrors the ITT in the main RCT. There are no significant changes in broader arrest outcomes (Panel B), with point estimates quite close to zero.

Figure 2 reports the three effects of interest for those in the READI RCT who could be both directly treated and exposed to a treated peer. The mean outcome for each non-exposed group is listed under the x-axis labels. Panel A focuses on the main outcomes of interest within the READI RCT. The standard errors are too wide to draw strong conclusions. Directionally, there is a substantively large decline in shooting and homicide arrests from direct treatment (similar to the ITT effects reported in the main RCT), but similarly-sized adverse effects from exposure to a treated peer for both treatment and control individuals. None are precise enough to interpret clearly, however.

The starkest pattern in Panel A is for other violent-crime arrests. Here, there is a 47 percent, significant ($p = 0.071$) decline for those directly treated with no peer exposure. But the exposure effects work in opposite directions: Exposed controls show a significant 56 percent decline in these arrests, while exposed treatments show a 59 percent, not-quite-significant increase. The fact that exposure decreases arrests for exposed controls but increases it for exposed treatments means that the initial ITT could be understating the net effect of treatment within the READI RCT. Substantively, this pattern could suggest that while treatment decreases violent crime for those who receive it, interaction with other treated peers may generate some offsetting adverse peer effects, while their control peers may learn from their desistance or be less likely to substitute into the deserted offending opportunities. Given the preliminary nature of the estimates, however, strong conclusions should wait for future efforts to improve the statistical power of these tests, as well as additional analysis on other types of social ties than just co-arrests.

For shooting and homicide victimizations, the pattern of the exposure effect is reversed: protective effects for exposed treatment individuals but (smaller) adverse effects on exposed controls. These estimates are quite imprecise, however, and they are sensitive to the weighting involved in the Aronow & Samii approach.[14]

In Panel B of Figure 2, we turn attention to the broader categories of arrests. Most of the point estimates are substantively small but imprecise. The clearest result is a significant 27 percent decline in drug arrests for those directly treated by READI with no peer exposure ($p = 0.034$). A smaller decline occurs when control peers are exposed to READI via their peers; their drug arrest point estimate is negative, though not statistically significant. But the exposure effects for the treated group are small and positive (though the confidence interval cannot rule out sizable negative effects). Overall, it looks as though the insignificant decline in drug arrests seen in the original RCT's ITT may mask a directly negative program effect on drug crime that spills over onto control individuals, but with little additional exposure effect for already–treated peers.

Two of the other estimates in Panel B are substantively large: a 31 percent increase in violent-crime arrests from direct treatment and an 87 percent increase in property crime for controls from exposure to a treated peer. But neither are statistically significant, even before any adjustments for the number of hypothesis tests in this graph.

Lastly, Figure 3 shows the estimated exposure effects for those not part of the original RCT. Because the READI population was selected for their high risk of these outcomes, the mean amount

[14] Appendix A shows that while the directional patterns of the estimates remain similar for both arrest outcomes, the sign of the estimates for shooting and homicide victimization actually reverses for the Boryusak & Hull regression-weighted estimation technique.

of serious violence among their unexposed peers (below the x-axis in Panel A) is considerably lower than the control means in the study itself: Shooting and homicides are about half as common among unexposed peers as in the original study sample's control group.

Partly because of the rarity of these outcomes for the non-RCT peers of READI study members, the estimates in Panel A of Figure 3 are imprecise. The point estimates suggest some adverse spillovers for shooting and homicide arrests, some protective spillovers for other kinds of violent-crime arrests, and a near-zero spillover for shooting and homicide victimization. But the standard errors are too wide to draw clear conclusions.[15]

Panel B, which reports exposure effects for outcomes that are much more common among this group, shows clearer evidence of one type of spillover. Those exposed to a treated peer have a statistically significant 35 percent decline in drug-crime arrests. This decline suggests that the direct negative effect on drug arrests for those who are assigned to READI spills over onto peers not eligible for READI. Most other crime types have point estimates close to zero.

8 Discussion

Even having solved the challenge of causally identifying peer effects within social networks, interpretation of treatment effects in the context of social spillovers requires care. Part of the challenge is being clear on what the different treatment effects are and for which populations they are relevant. Here, we are interested in the direct effect of treatment versus nothing, as well as the effect of being exposed to a treated peer. We estimate these effects for three different populations: 1) those who were in the original RCT but not tied to anyone else in the study (direct effect only), 2) those who were in the original READI RCT and had baseline co-arrest ties to other people in the study (direct and exposure effects), and 3) those with pre-randomization co-arrest ties to those in the original READI RCT but who were not in the RCT themselves (exposure effect only).

The group that is not part of the social network has direct effects very similar to those reported in the main READI RCT: large declines in shooting and homicide arrests, with no significant changes in other outcomes. Outside that group, READI RCT's focus on three rare outcomes makes it difficult to draw clear conclusions about spillovers involving serious violence. Within the READI population that shares ties with others in READI, the smaller sample size (906), along with splitting the sample to estimate separate exposure effects for treatment and control groups, contributes to the imprecision of the estimates. The pattern of point estimates suggests that the effects of READI versus nothing (neither READI nor exposure to a treated peer) is likely negative for non-shooting, non-homicide violent-crime arrests (assault and battery, robbery, and sexual assault)—an effect that was masked in the original ITT estimates because of increases among exposed treatment individuals and decreases among exposed controls. But the confidence intervals are too wide to consider this a strong conclusion. For the shooting and homicide outcomes, there is no definitive evidence of adverse spillovers, which would be the most problematic kind of spillover in terms of overturning the interpretation of the original ITTs. But the estimates are too noisy (and, as we discuss in Appendix A, too sensitive to the different weights across estimation techniques) to say for sure.

The non-RCT sample was not selected to be at elevated risk of these serious-violence outcomes. So although their co-arrest ties to READI study members may indicate that they are part of a social group with elevated risk, in practice their baseline (non-exposed) mean was about half of the READI group's. Partly as a result, estimated exposure effects for the three serious violence measures are quite noisy despite the larger sample size (6,210).

Evidence on the more common types of arrests is somewhat clearer. Here, there is a significant

[15] Appendix A shows that while there are power gains from the regression-based approach in Boryusak & Hull, they are still not enough to rule out sizable spillover effects in either direction.

decline of about 30 percent in drug-crime arrests for those in the arrest network who are directly treated with no peer exposure. That decline appears to diffuse through social ties, decreasing drug-crime arrests among both control-group and non-READI peers. There is no apparent decline for already-treated peers, indicating no additional exposure effect beyond the one created by treatment itself. The decline in control group drug arrests explains why the main RCT did not detect the decrease: by making the controls look better than they would have without exposure to treated peers, the simple ITT understated the direct effect of treatment.

Separating the direct effect of treatment versus nothing from the different types of exposure effects helps to unpack what the original READI ITT was estimating. But it is important to realize that none of our estimated effects in isolation provides the net social effect of READI. To calculate the net effect, we would need to weight each effect by how many people experienced that effect in practice, and sum everything up.

Eventually, we plan to estimate the net effect of READI. We are, however, hesitant to do so with the current estimates for several reasons. First, given the imprecision in each individual estimate, it seems crucial to assess the uncertainty in the net effect as well. This involves making more progress on inference in the regression framework, as well as figuring out how to calculate correct confidence intervals for the weighted sum of estimates (bootstrapping is not feasible in this context due to the computational intensity involved with generating each estimate). Second, this analysis is a first look at a single type of exposure: having any previously co-arrested peer offered READI. We also plan to test other exposure types (e.g., number of peers or proportion of peers) as well as other types of social ties (having been victimized together, taken many similar classes in school together, living together, and living near each other). Until we have made more progress on these next steps, we are potentially missing big pieces of the story. As a result, we hesitate to provide one single number claiming to be the net social effect before making more progress on other social spillovers.

We also note that the focus on READI and its three main outcomes here is intended to help interpret the main READI RCT. But it is only a subset of larger project, which will include three other RCTs to aid statistical power, in addition to the other social network measures. With these in hand, we also plan to test how network formation itself might be affected by treatment to capture changes in future peers in addition to existing peers. Future work will have much more to say about how intervention-driven changes in violence and other crime spread through social networks, which types of social ties appear to matter, how that helps us understand social decisions about crime, and what it implies for the optimal targeting of interventions.

References

Abdul-Razzak, Nour, Kelly Hallberg, Cristobal Pinto, and Lucia Delgado. In Progress. "Spillovers from a Youth Violence Prevention Program." .

Akerlof, George A., and Rachel E. Kranton. 2010. *Identity Economics: How Our Identities Shape Our Work, Wages, and Well-Being.* Princeton University Press.

Angrist, Joshua D. 2014. "The Perils of Peer Effects." *Labour Economics* 30(c): 98–108.

Aronow, Peter M., and Cyrus Samii. 2013. "Conservative Variance Estimation for Sampling Designs with Zero Pairwise Inclusion Probabilities." *Survey Methodology* 39: 231–241.

———. 2017. "Estimating Average Causal Effects Under General Interference, with Application to a Social Network Experiment." *The Annals of Applied Statistics* 11(4): 1912–1947.

Athey, Susan, Dean Eckles, and Guido W. Imbens. 2018. "Exact P-Values for Network Interference." *Journal of the American Statistical Association* 113(521): 230–240.

Ballester, Coralio, Antoni Calvo-Armengol, and Yves Zenou. 2006. "Who's Who in Networks. Wanted: The Key Player." *Econometrica* 74(5): 1403–1417.

———. 2010. "Delinquent Networks." *Journal of the European Economic Association* 8(1).

Banerjee, Abhijit, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson. 2013. "The Diffusion of Microfinance." *Science* 341: 363–370.

Bayer, Patrick, Randi Hjalmarsson, and David Pozen. 2009. "Building Criminal Capital Behind Bars: Peer Effects in Juvenile Corrections." *The Quarterly Journal of Economics* 124(1): 105–147.

Beaman, Lori, Sylvan Herskowitz, Niall Keleher, and Jeremy Magruder. 2021. "Stay in the Game: A Randomized Controlled Trial of a Sports and Life Skills Program for Vulnerable Youth in Liberia." *Economic Development and Cultural Change* 70(1): 129–158.

Beaman, Lori, and Jeremy Magruder. 2012. "Who Gets the Job Referral? Evidence from a Social Networks Experiment." *American Economic Review* 102(7): 3574–3593.

Bhatt, Monica, Marianne Bertrand, Chris Blattman, Sara B. Heller, and Max Kapustin. 2022. "READI Chicago: A Community-Based Approach to Reducing Gun Violence." .

———. in progress. "Predicting and Reducing Gun Violence in Chicago: Experimental Results from the READI Program." .

Bhattacharya, Debopam, Pascaline Dupas, and Shin Kanaya. 2013. "Estimating the Impact of Means-tested Subsidies under Treatment Externalities with Application to Anti-Malarial Bednets." *NBER Working Paper No. 18833* .

Bhuller, Manudeep, Gordon B. Dahl, Katrine V. Loken, and Magne Mogstad. 2018. "Incarceration Spillovers in Criminal and Family Networks." *NBER Working Paper No. 24878* .

Billings, Stephen B., David J. Deming, and Jonah Rockoff. 2013. "School Segregation, Educational Attainment, and Crime: Evidence from the End of Busing in Charlotte-Mecklenburg." *The Quarterly Journal of Economics* 129(1): 435–476.

Billings, Stephen B., David J. Deming, and Stephen L. Ross. 2019. "Partners in Crime." *American Economics Journal: Applied Economics* 11(1): 126–150.

Billings, Stephen B., and Keven T. Schnepel. 2017. "Life After Lead: Effects of Early Interventions for Children Exposed to Lead." *American Economic Journal: Applied Economics* 10(3): 315–344.

Borusyak, Kirill, and Peter Hull. 2021. "Non-Random Exposure to Exogenous Shocks." .

Breza, Emily, and Arun G. Chandrasekhar. 2019. "Social Networks, Reputation, and Commitment: Evidence from a Savings Monitors Experiment." *Econometrica* 87(1): 175–216.

Bursztyn, Leonardo, Thomas Fujiwara, and Amanda Pallais. 2017. "'Acting Wife': Marriage Market Incentives and Labor Market Investments." *American Economic Review* 107(11): 3288–3319.

Cai, Jing, Alain De Janvry, and Elisabeth Sadoulet. 2015. "Social Networks and the Decision to Insure." *American Economic Journal: Applied Economics* 7(2): 81–108.

Charette, Yanick, and Andrew V. Papachristos. 2017. "The Network Dynamics of Co-Offending Careers." *Social Networks* 51: 3–13.

Comola, Margherita, and Silvia Prina. 2014. "Do Interventions Change the Network? A Panel Peer-Effect Model Accounting for Endogenous Network Changes." *IZA Discussion Papers, No. 8641* .

Damm, Anna Piil, and Christian Dustmann. 2014. "Does Growing Up in a High Crime Neighborhood Affect Youth Criminal Behavior?" *American Economic Review* 104(6): 1806–1832.

Damm, Anna Piil, and Cedric Gorinas. 2016. "Prison as a Criminal School: Peer Effects and Criminal Learning Behind Bars." *The Rockwool Foundation, Study paper no. 105* .

Davis, Jonathan M.V., and Sara B. Heller. 2020. "Rethinking the Benefits of Youth Employment Programs: The Heterogeneous Effects of Summer Jobs." *The Review of Economics and Statistics* 102(4): 664–677.

Diaz, Lelys Ileana Dinarte. 2020. "Peer Effects on Violence: Experimental Evidence from El Salvador." *The World Bank (WPS9187)* .

Diaz, Sandra, Unai Pascual, Marie Stenseke, Berta Martín-López, Robert T. Watson, Zsolt Molnár, Rosemary Hill, Kai M. A. Chan, Ivar A. Baste, Kate A. Brauman, Stephen Polasky, Andrew Church, Mark Lonsdale, Anne Larigauderie, Paul W. Leadley, Alexander P. E. van Oudenhoven, Felice van der Plaat, Matthias Schröter, Sandra Lavorel, Yildiz Aumeeruddy-Thomas, Elena Bukvareva, Kirsten Davies, Sebsebe Demissew, Gunay Erpul, Pierre Failler, Carlos A. Guerra, Chad L. Hewitt, Hans Keune, Sarah Lindley, and Yoshihisa Shirayama. 2018. "Assessing Nature's Contribution to People." *Science* 359(6373): 270–272.

Dieye, Rokhaya, Haviva Djebbari, and Felipe Barrera-Osorio. 2014. "Accounting for Peer Effects in Treatment Response." *AMSE Working Papers 1435* .

Dominguez, Magdalena. 2022. "Sweeping Up Gangs: The Effects of Tough-on-Crime Policies from a Network Approach." *Working Paper* .

Drago, Francesco, and Roberto Galbiati. 2012. "Indirect Effects of a Policy Altering Criminal Behavior: Evidence from the Italian Prison Experiment." *American Economic Journal: Applied Economics* 4(2): 199–218.

Feigenberg, Benjamin, Erica Field, and Rohini Pande. 2013. "The Economic Returns to Social Interaction: Experimental Evidence from Microfinance." *The Review of Economic Studies* 80(4): 1459–1483.

Fortin, Bernard, and Vincent Boucher. 2016. *The Oxford Handbook of the Economics of Networks*, chap. Some Challenges in the Empirics of the Effects of Networks.

Granovetter, Mark. 1978. "Threshold Models of Collective Behavior." *American Journal of Sociology* 83(May): 489–515.

Green, Ben, Thibaut Horel, and Andrew V. Papachristos. 2017. "Modeling Contagion Through Social Networks to Explain and Predict Gunshot Violence in Chicago, 2006 to 2014." *JAMA Internal Medicine* 177(3): 326–333.

Heller, Sara B. 2014. "Summer jobs reduce violence among disadvantaged youth." *Science* 346: 1219–1223.

Heller, Sara B., Anuj K. Shah, Jonathan Guryan, Jens Ludwig, Sendhil Mullainathan, and Harold A. Pollack. 2017. "Thinking, Fast and Slow? Some Field Experiments to Reduce Crime and Dropout in Chicago." *The Quarterly Journal of Economics* 132(1): 1–54.

Lindo, Jason M., and Maria Padilla-Roma. 2018. "Kingpin Approaches to Fighting Crime and Community Violence: Evidence from Mexico's Drug War." *Journal of Health Economics* 58: 253–268.

Lindquist, Matthew J., and Yves Zenou. 2014. "Key Players in Co-Offending Networks." *CEPR Discussion Paper No. DP9889* .

Liu, Lin, Ninghua Wang, and John E. Eck. 2014. "Analyzing Crime Displacement with a Simulation Approach." *Environmental and Planning B Planning and Design* 41(2).

Liu, Lin, and Xinyue Ye. 2012. "Spatial Crime Analysis and Modeling." *Annals of GIS* 18(3).

Manski, Charles F. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *The Review of Economic Studies* 60(3): 531–542.

McAndrew, Duncan. 2000. *The Social Psychology of Crime: Groups, Teams and Networks*, chap. The Structural Analysis of Criminal Networks. Ashgate Publishing.

McCarthy, Bill, and John Hagan. 1995. "Getting into Street Crime: The Structure and Process of Criminal Embeddedness." *Social Science Research* 24(1): 63–95.

McGloin, Jean Marie, and Alex R. Piquero. 2009. "On the Relationship between Co-Offending Network Redundancy and Offering Versatility." *Journal of Research in Crime and Delinquency* 47(1): 63–90.

Miguel, Edward, and Michael Kremer. 2007. "The Illusion of Sustainability." *The Quarterly Journal of Economics* 122(3): 1007–1065.

Miller, Grant A., and Mushfiq Mobarak. 2014. "Learning About New Technologies Through Social Networks: Experimental Evidence on Nontraditional Stoves in Bangladesh." *Marketing Science* 34(4): 473–626.

O'Brien, Lia, Dustin Albert, Jason Chein, and Laurence Steinberg. 2011. "Adolescents Prefer More Immediate Rewards When in the Presence of their Peers." *Journal of Research on Adolescence* 21(4): 747–753.

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).

Ouss, Aurelie. 2011. "Prison as a School of Crime: Evidence from Cell-Level Interactions." *SSRN Electronic Journal* .

Paluck, Elizabeth Levy, Hana Shepherd, and Peter M. Aronow. 2016. "Changing Climates of Conflict: A Social Network Experiment in 56 Schools." *PNAS* 113(3): 566–571.

Patacchini, Eleonara, and Yves Zenou. 2009. "Juvenile Delinquency and Conformism." *CEPR Discussion Paper No. DP7565* .

Philippe, Arnaud. 2017. "Incarcerate One to Calm the Others? Spillover Effects of Incarceration Among Criminal Groups." *TSE Working Paper n. 17-840* .

Shepherd, Hana, and Elizabeth Levy Paluck. 2015. "Stopping the Drama: Gendered Influence in a Network Field Experiment." *Social Psychology Quarterly* 78(2): 173–193.

Stevenson, Megan. 2017. "Breaking Bad: Mechanisms of Social Influence and the Path to Criminality in Juvenile Jails." *Review of Economics and Statistics* 99(5): 824–838.

Tankard, Margaret E., and Elizabeth Levy Paluck. 2016. "Norm Perception as a Vehicle for Social Change." *Social Issues and Policy Review* 10(1): 181–211.

Tremblay, Pierre. 1993. "Searching for Suitable Co-Offenders." *Routine Activity and Rational Choice: Advances in Criminological Theory* 5: 17–36.

van Mastrigt, Sarah B., and Peter Carrington. 2019. *The Oxford Handbook of Developmental and Life-Course Criminology*, chap. Co-Offending.

Vazquez-Bare, Gonzalo. forthcoming. "Identification and Estimation of Spillover Effects in Randomized Experiments." *Journal of Econometrics* .

Warr, Mark. 2002. *Companions in Crime: The Social Aspects of Criminal Conduct*. Cambridge University Press.

Weerman, Frank M. 2003. "Co-Offending as a Social Exchange: Explaining Characteristics of Co-Offending." *The British Journal of Criminology* 43(2): 398–416.

Wood, George, and Andrew V. Papachristos. 2019. "Reducing Gunshot Victimization in High-Risk Social Networks Through Direct and Spillover Effects." *Nature Human Behaviour* 3: 1164–1170.

9 Tables and Figures

Table 1: Characteristics of study individuals and their first-degree co-arrest peers

	Unexposed study sample	Within-study exposure sample	Out of study first-degree peers
N	1490	966	6367
Age	27	23.7	26.4
Black	0.935	0.997	0.945
Hispanic	0.039	0.003	0.042
Male	1	1	0.9
Prior Arrests	14.05	20.75	13.27
Degree	2.15	8.15	6.11
Share Co-Arrests	0.13	0.28	0.36
Pr(Exposure)	–	T=0.75; C=0.71	0.59

Figure 1: Direct effects of READI for those in the RCT but not the network

(a) Main READI outcomes

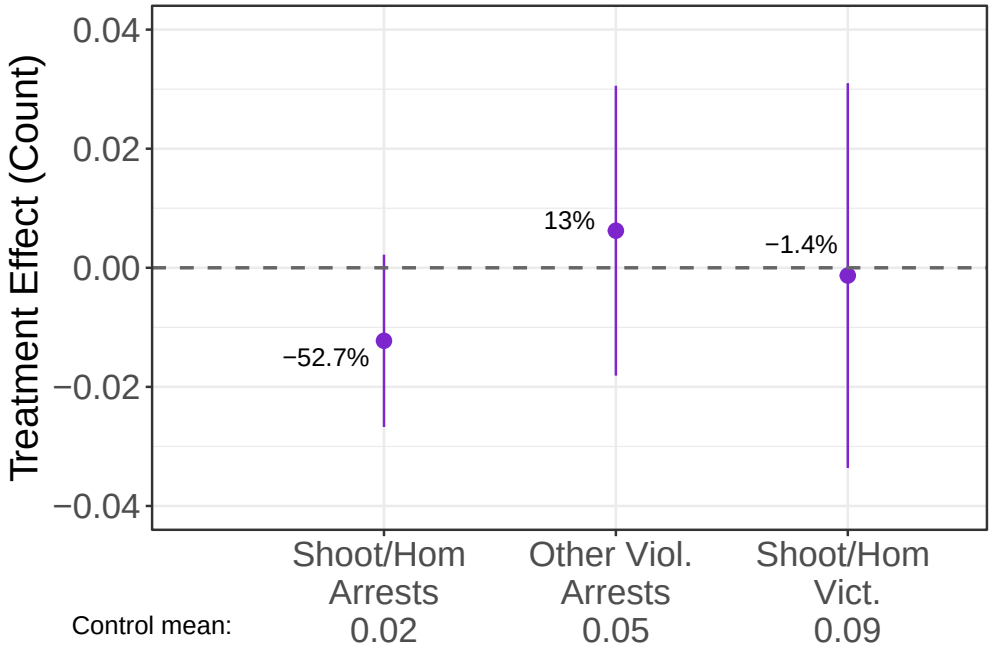

(b) Broad arrest categories

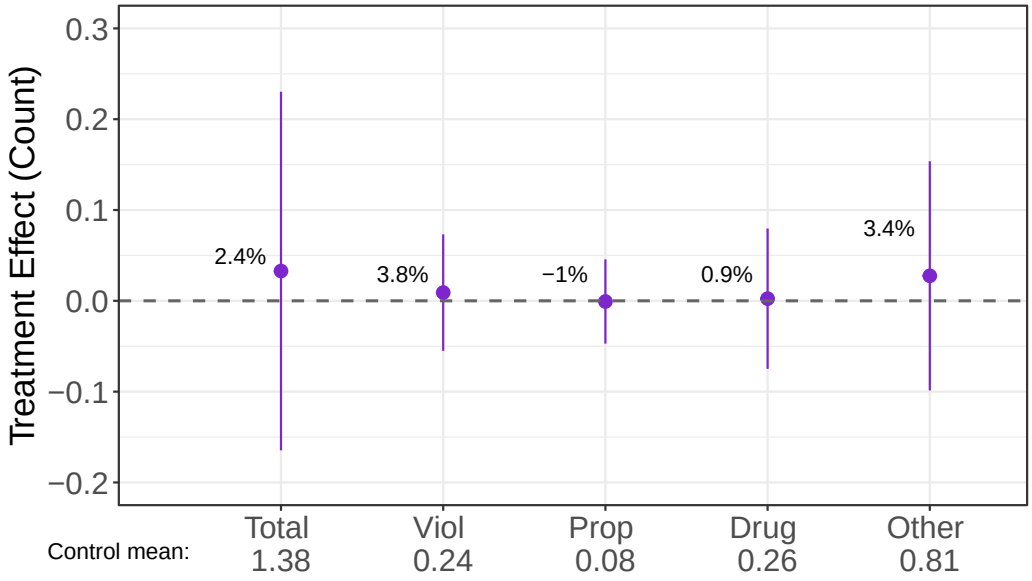

Notes: This figure displays direct treatment effects for READI study members who are not present in the co-arrest network ($N = 1,490$). Estimates are covariate-adjusted. 95% confidence intervals are reported using robust standard errors. Average outcomes for unexposed individuals are below each x-axis label, with the implied proportional change labeled on the graph.

Figure 2: Direct and exposure effects of READI for those in the RCT and in the network

(a) Main READI outcomes

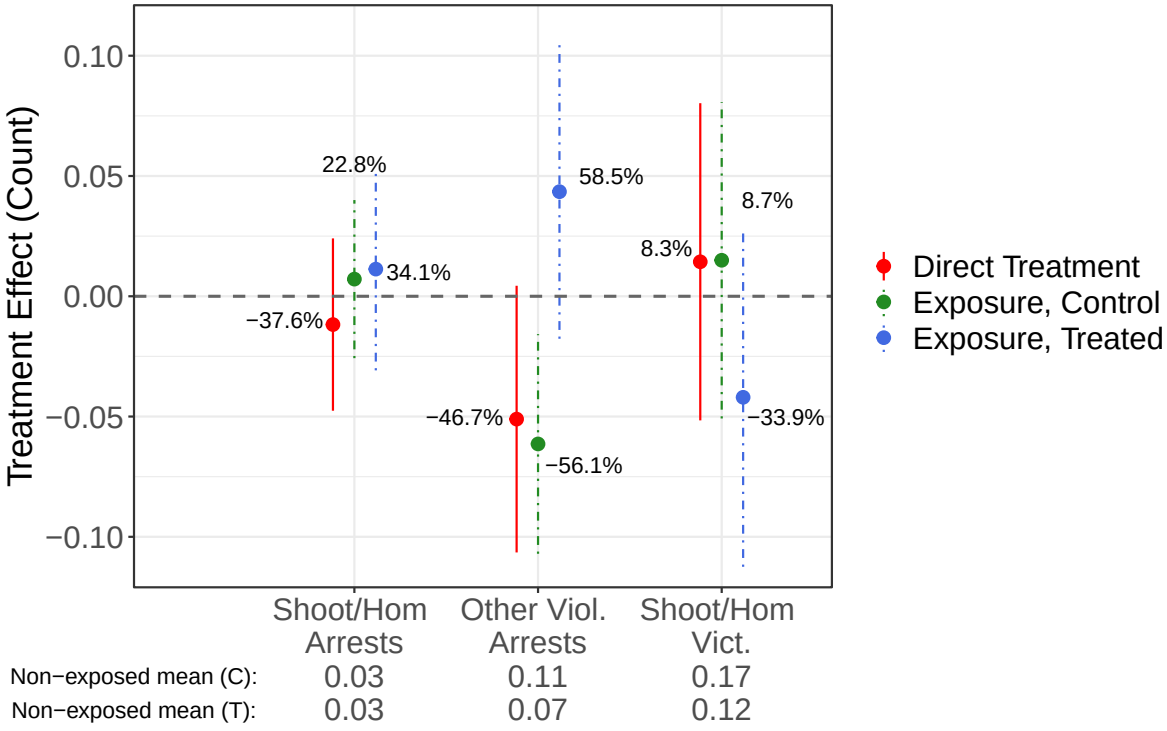

(b) Broad arrest categories

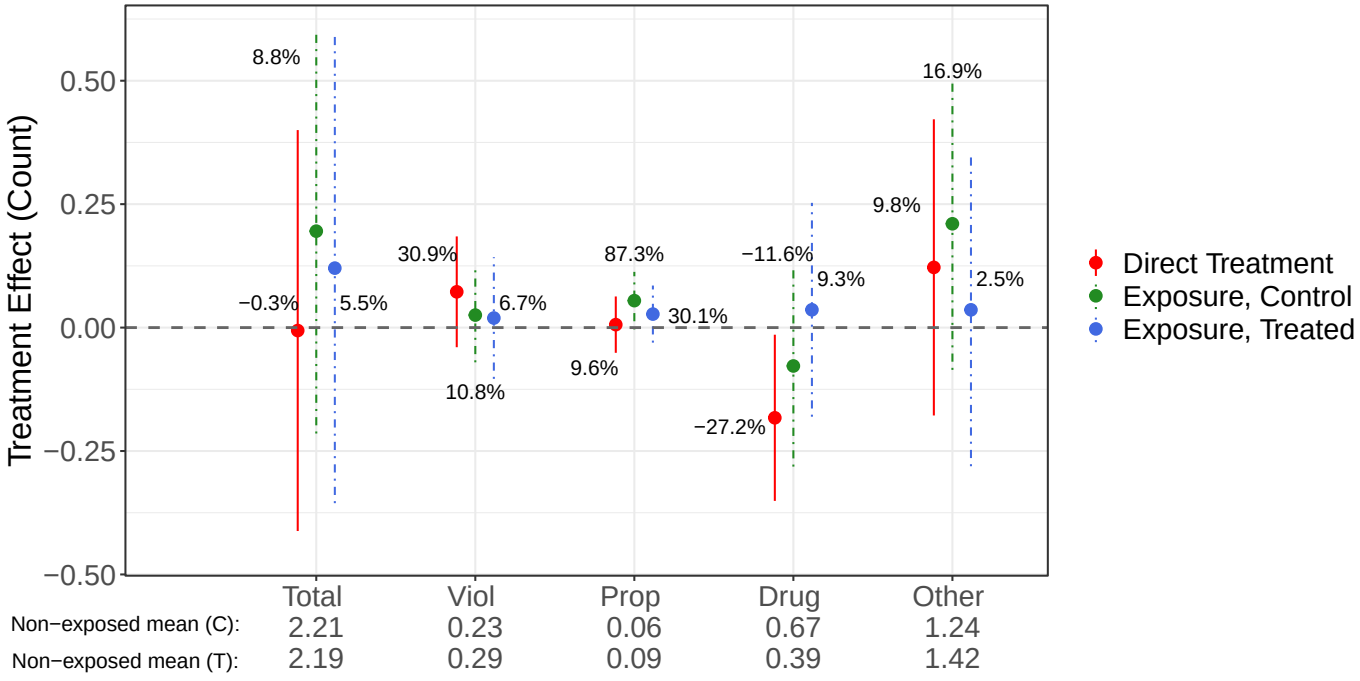

Notes: This figure displays average direct treatment effects and marginal exposure effects, separately by direct treatment status, for READI study members who are also tied to other READI members via the co-arrest network ($N = 906$). Covariate-adjusted point estimates and conservative confidence intervals calculated with the methods in Aronow & Samii (2017). Exposure is binary – a unit is exposed if they have any treated peers. Average outcomes for unexposed individuals are below each x-axis label (C and T separately), with the implied proportional change labeled on the graph.

Figure 3: Exposure effects of READI for non-RCT peers

(a) Main READI outcomes

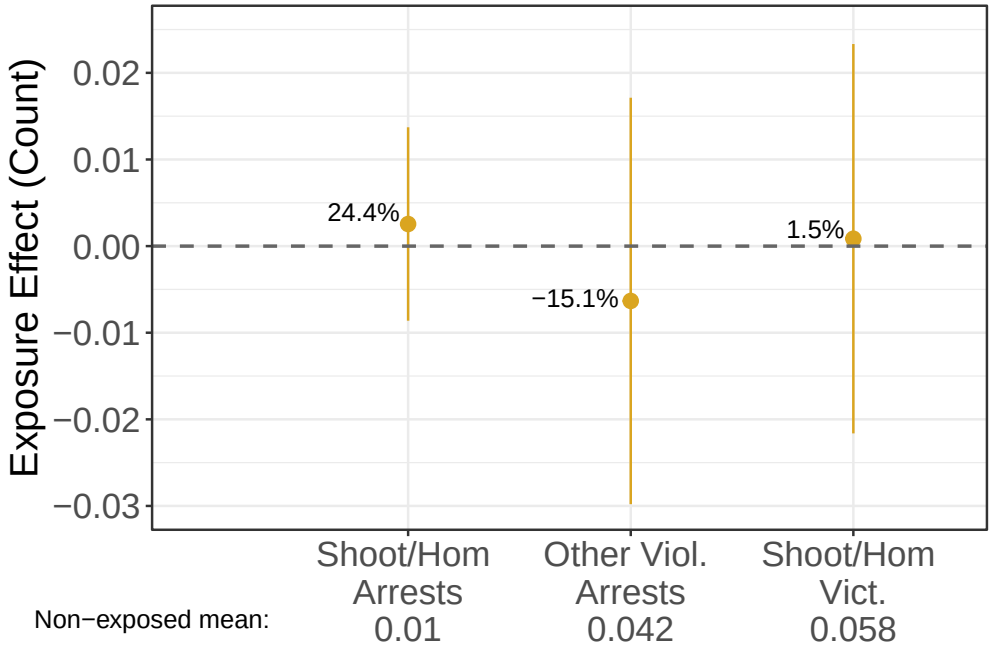

(b) Broad arrest categories

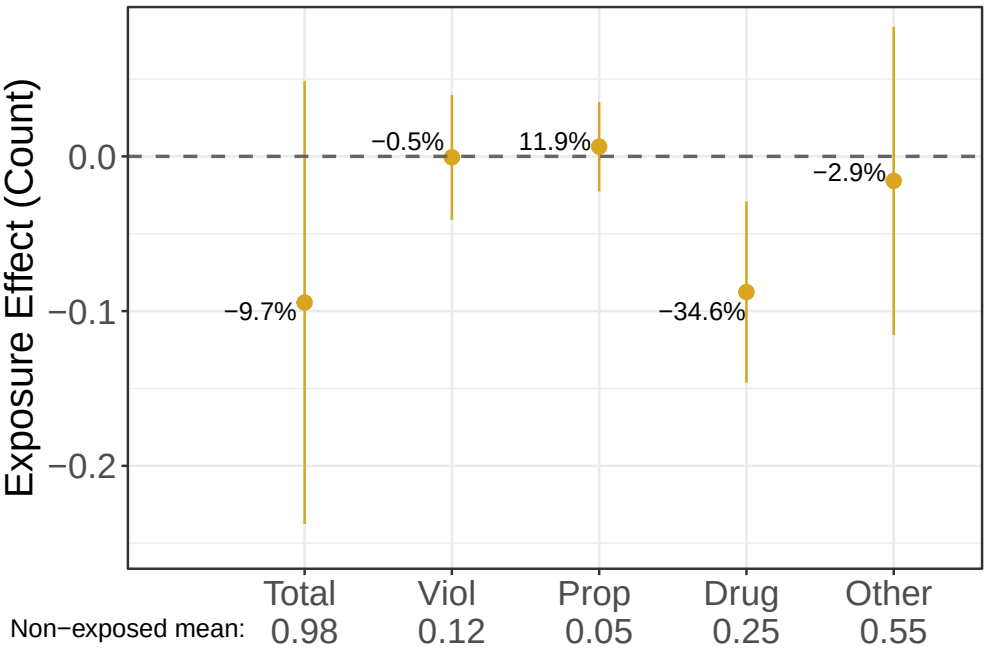

Notes: This figure displays average marginal exposure effects for individuals who are not in READI themselves but are tied to READI members via the co-arrest network ($N = 6,210$). Covariate-adjusted point estimates and conservative confidence intervals calculated with the methods in Aronow & Samii (2017). Exposure is binary – a unit is exposed if they have any treated peers. Average outcomes for unexposed individuals are below each x-axis label, with the implied proportional change labeled on the graph.

A Alternative Methods

A.1 Regression-based estimation and inference

In this section, we implement the Borusyak & Hull regression-based approach to estimation discussed in Section 6. This approach estimates exposure effects using a regression and controlling for $\pi_i(d_{j,k})$, rather than directly weighting, as in the Aronow and Samii (2017) approach. We use the same trimming procedure as above, so the two methods are estimated on the same sample. For the simpler case of the non-RCT peers of the READI sample, the regression is:

$$y_i = \alpha + \beta_1 Exposure_i + \beta_2 \pi_i(d_{0,1}) + \gamma X_i + \varepsilon_i \tag{4}$$

where the sample is only the first-degree peers of those in the READI study who were not also in the study themselves; $Exposure_i$ is an indicator for whether at least one peer was treated; and $\pi_i(d_{0,1})$ is the probability at least one peer was indirectly exposed. Since there is only one treatment effect of interest here, we can conduct hypothesis testing on β_1 with randomization inference. We report p-values from this inference procedure.

When we shift attention to within-RCT spillovers, both the regression and inference are less straightforward. Because there are multiple types of exposure and treatment simultaneously occurring here (direct treatment, exposure among the directly treated, and exposure among the controls), we need to calculate and control for the probability of each type of exposure. We then estimate Equation 5:

$$\begin{aligned} y_i = \alpha + \beta_1 Treat_i + \beta_2 Treat_i \times Exposure_i + \beta_3 Control_i \times Exposure_i \\ + \delta_1 \pi_i(d_{1,\cdot}) + \delta_2 \pi_i(d_{1,1}) + \delta_3 \pi_i(d_{0,1}) + \gamma X_i + \varepsilon_i \end{aligned} \tag{5}$$

and recover the three objects of interest: the ITT estimate of direct treatment among the unexposed sample (β_1), the exposure effect among the treated individuals (β_2), and the exposure effect among the control individuals (β_3).

In this setting with simultaneously-varying multiple treatments, randomization inference is not straightforward. As Borusyak and Hull (2021) discuss, there is currently no consensus on the best way to solve this inference problem.[16] However, in order to understand the sensitivity of our estimates to the choice of estimator, we compare the point estimates obtained using the Borusyak & Hull method to those using the Aronow & Samii described above, for which we can obtain correct confidence intervals.

Note that even in infinite samples we would not expect the Borusyak & Hull estimates to be exactly the same as those in the previous section, since Aronow & Samii reweight observations by their probability of exposure whereas Borusyak & Hull control for the probability of exposure in a regression. The former will put more weight on observations with lower exposure probabilities, while the latter will put more weight on exposure probabilities that are more common in the data. So the two approaches represent slightly different estimators, which may differ under heterogeneous treatment effects.

A.2 Regression-based results

Figures 4 and 5 report results for the group of individuals who are both in the original READI RCT and tied to another READI study member. To prevent over-crowding the graphs, we separate the direct effects of treatment (Figure 4) and the exposure effect for treatment and control individuals

[16] This is why we emphasize the Aronow & Samii method in the main text.

(Figure 5). Both Figures compare treatment effects using the Aronow & Samii estimates (AS) to those obtained using the Borusyak & Hull method (BH).

Panel A of Figure 4 suggests that the point estimates on the direct effects of READI are somewhat sensitive to reweighting; they all differ somewhat between AS and BH estimation. The sensitivity both here and below are likely in part because the outcomes are relatively rare and in part because the sample size that is part of the network sample is much smaller than in the original study (906 versus 2,456). Additionally, the main RCT documents significant heterogeneous effects. If effects differ greatly across individuals, the differences in who gets more weight in AS versus BH may matter. Panel B shows that the more common arrest categories have estimates that vary less across estimation strategies, with consistent evidence of a decline in drug arrests among those who are directly treated.

Figure 5 turns to exposure effects by treatment group. There are small adverse point estimates for exposure effects on shooting and homicide arrests. Shooting and homicide victimizations are particularly sensitive to the differences in weighting across estimation techniques. The most consistent pattern is that exposure to treated peers increases other violent-crime arrests among the treatment group but decreases it among controls, which would imply that the original ITT likely under-stated the effect of direct treatment. Again, uncertainty around the estimates means that this pattern is far from a clear result.

Panel B shows that, as with direct effects, there is less sensitivity across methods for the more common arrest types. Most point estimates are quite small. Directionally, the treatment-control difference in exposure effects for drug-crime arrests suggests the ITT may under-state the direct effect while the ITT on other arrests may over-state it. But here again, the sample is too small to draw strong conclusions.

Tables 2 and 3 report results for the effect of exposure on non-RCT peers, separately by the main READI outcomes and broader arrest categories. To aid comparisons to the previous set of results, the first two rows in each table reproduce the Aronow & Samii estimates from Figure 3 and their corresponding p-values. The second two rows of each table display the analogous estimates using Borusyak & Hull and the randomization inference p-value for the regression-based estimates.

Table 2 shows, for the non-RCT peers, the effect of being exposed to someone in the READI treatment group on the three main READI outcomes. In general, the results are quite similar to the previous section. Even with the power gains from this method, there are still no statistically significant exposure effects and the p-values are quite large. The BH point estimate for other violent-crime arrests is only a third as large as the AS estimate, suggesting some sensitivity to the difference in weighting.

Table 3 shows the exposure effect for the broader crime categories, which as the non-exposed means show, are generally far more common among this population. Results are quite similar across estimation methods: we still see the large decline in drug-crime arrests for those with a peer in READI, with no clear indications of other types of spillovers.

B Appendix Tables and Figures

Table 2: Comparing methods to estimate exposure effects of READI, Non-RCT Peers (Main READI Outcomes)

	Shoot/Hom Arr	Other Viol Arr	Shoot/Hom Vict
Aronow & Samii (AS)	0.00255	-0.00633	0.000853
AS p-value	0.654	0.597	0.941
Borusyak & Hull (BH)	0.0033	-0.00138	-0.00105
BH p-value	0.291	0.831	0.875
Non-exposed mean	0.0105	0.0418	0.0581

Notes: This table compares the Aronow & Samii (AS) estimates to the regression-based estimates proposed by Borusyak & Hull (BH). Randomization inference p-values for the Borusyak & Hull estimates are displayed as well. The sample is individuals who are not in READI themselves but are tied to READI members via the co-arrest network ($N = 6,210$). Exposure is binary – a unit is exposed if they have any treated peers. and Regressions include baseline covariates and AS estimates and standard errors are computed as described in Figure 3. Average outcomes for unexposed individuals are present in the final row.

Table 3: Comparing methods to estimate exposure effects of READI, Non-RCT Peers (Broad Arrest Categories)

	Total	Viol	Prop	Drugs	Other
Aronow & Samii (AS)	-0.0944	-0.000609	0.0063	-0.0877	-0.0158
AS p-value	0.196	0.976	0.669	0.00335	0.756
Borusyak & Hull (BH)	-0.0852	0.0059	0.00797	-0.0798	-0.0216
BH p-value	0.0632	0.625	0.369	6.67e-05	0.479
Non-exposed mean	0.976	0.124	0.053	0.254	0.545

Notes: This table compares the Aronow & Samii (AS) estimates to the regression-based estimates proposed by Borusyak & Hull (BH). Randomization inference p-values for the Borusyak & Hull estimates are displayed as well. The sample is individuals who are not in READI themselves but are tied to READI members via the co-arrest network ($N = 6,210$). Exposure is binary – a unit is exposed if they have any treated peers. and Regressions include baseline covariates and AS estimates and standard errors are computed as described in Figure 3. Average outcomes for unexposed individuals are present in the final row.

Figure 4: Comparing methods to estimate direct effects of READI, RCT Exposure Sample

(a) Main READI outcomes

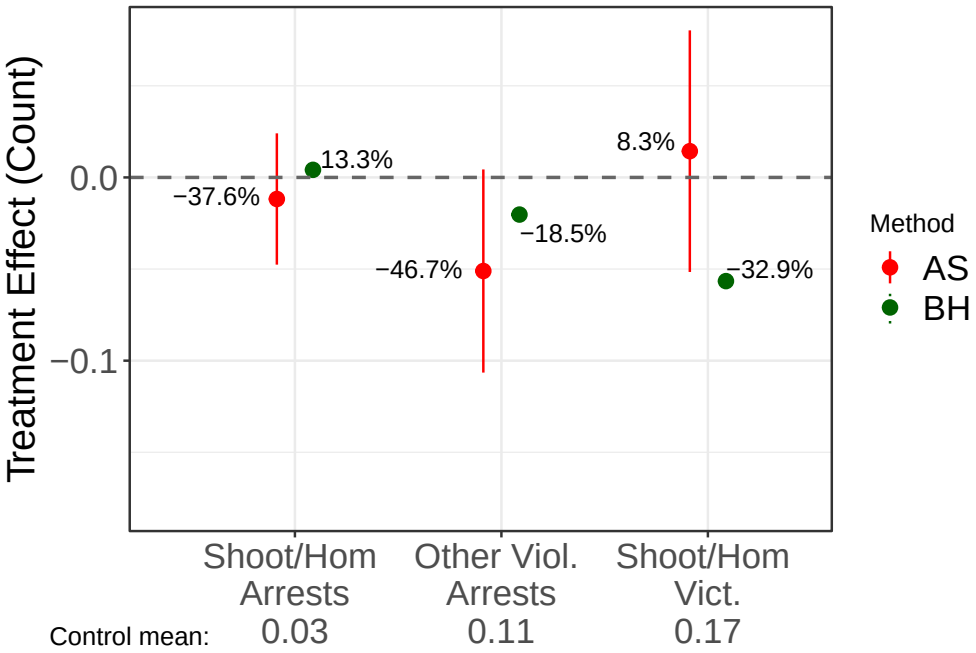

(b) Broad arrest categories

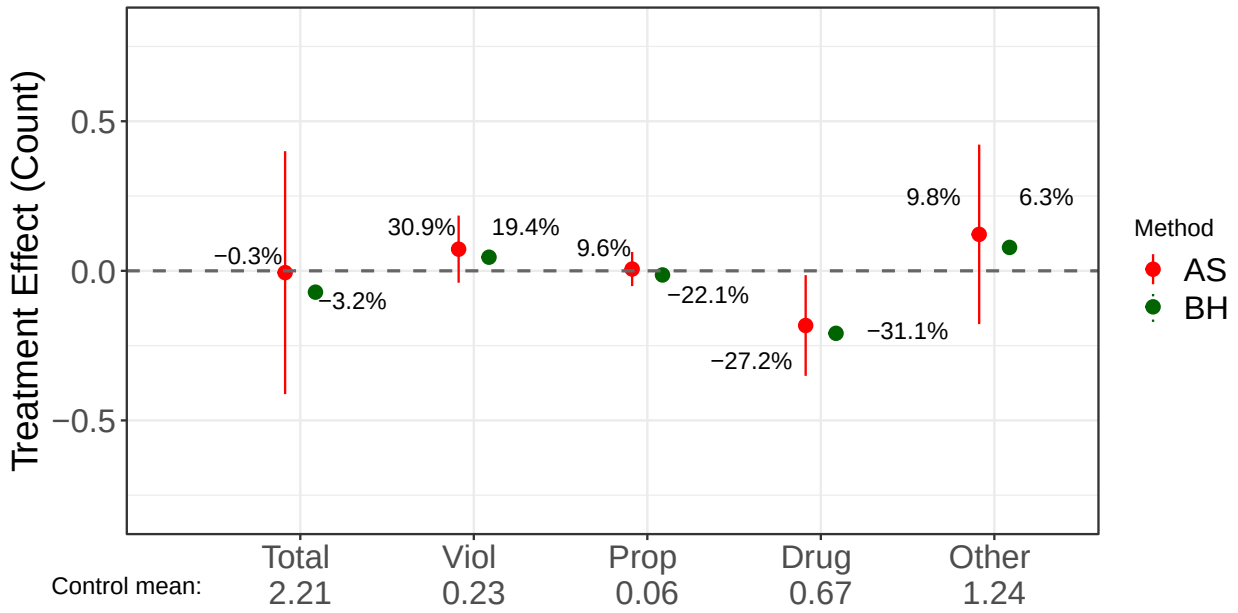

Notes: This figure compares the Aronow & Samii (AS) estimate to the regression-based estimates proposed by Borusyak & Hull (BH). The sample is READI study members who are also tied to other READI members via the co-arrest network ($N = 906$). Direct treatment is binary, and regressions include baseline covariates. AS estimates and standard errors are computed as described in Figure 2.

Figure 5: Comparing methods to estimate exposure effects of READI, RCT Exposure Sample

(a) Main READI outcomes

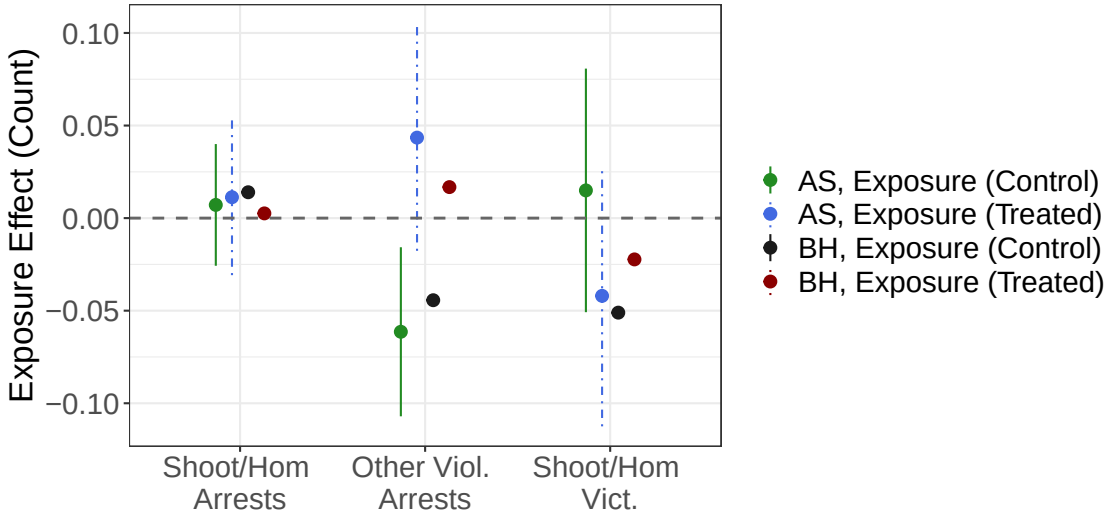

(b) Broad arrest categories

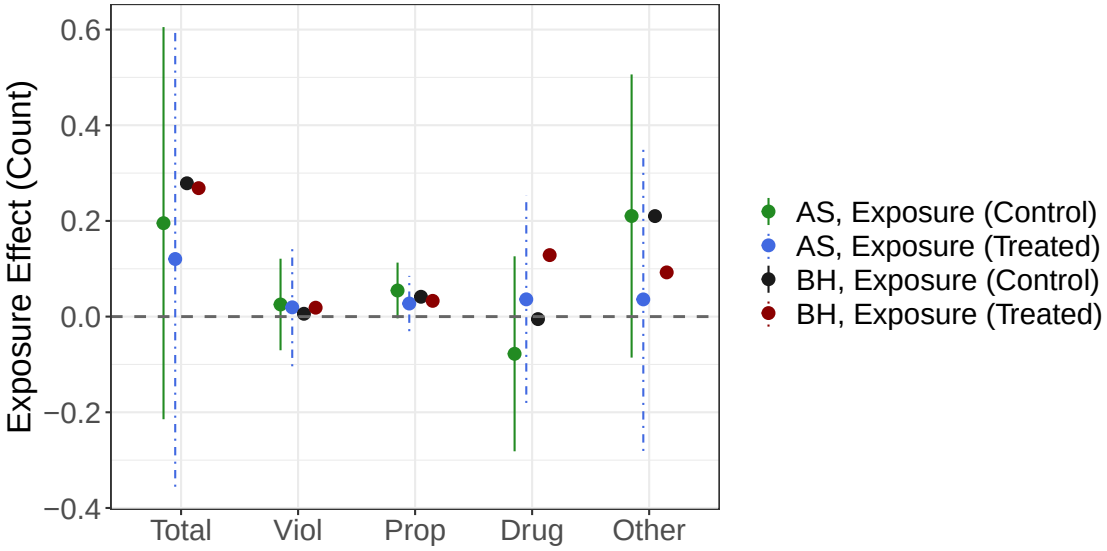

Notes: This figure compares the Aronow & Samii (AS) estimate to the regression-based estimates proposed by Borusyak & Hull (BH). The sample is READI study members who are also tied to other READI members via the co-arrest network ($N = 906$). Direct treatment and exposure are both binary, and regressions include baseline covariates. AS estimates and standard errors are computed as described in Figure 2.