Better Together? Social Networks in Truancy and the Targeting of Treatment[*]

Magdalena Bennett, Peter Bergman[†]

This draft: May 3, 2019

Abstract

There is concern the risky behaviors of teenagers, such as truancy, negatively influence the behaviors of others through their social networks. We develop a strategy to use administrative data of in-class attendance to construct social networks based on students who are truant together. We simulate these networks to document that certain students systematically coordinate their absences. We validate them by showing a parent-information intervention on student absences has spillover effects from treated students onto their peers. Excluding these effects understates the intervention's cost effectiveness by 43%. We show there is potential to use networks to target interventions more efficiently given a budget constraint.

[*]We thank Claudia Allende, Josh Angrist, Andrew Barr, Alex Bowers, Eric Chan, Francois Gerard, Laura Giuliano, Jan Feld, Caroline Hoxby, Matthew Jackson, Adam Kapor, Bruce Sacerdote, Ulf Zolitz and participants at the NBER Education Program meeting, SOLE conference, and CESifo conference for their comments. All errors are our own.

[†]Teachers College, Columbia University, 525 W. 120th Street New York, New York 10027.

1 Introduction

There is concern that the risky behaviors of teenage children negatively influence the behaviors of other children through their social networks. This influence could occur if, for instance, children learn behaviors from other children, signal their behaviors to peers, or derive utility from undertaking behaviors jointly (Akerlof and Kranton 2000; Austen-Smith and Fryer 2005; Bénabou and Tirole 2011; Bursztyn et al. 2014). Such mechanisms may be particularly relevant to school truancy, which predicts a number of adverse outcomes including high school dropout, substance abuse, and criminality (Kearney 2008; Goodman 2014; Aucejo and Romano 2016; Rogers and Feller 2018; Cook et al. 2017; Gershenson et al. 2017; Billings et al. 2016). Attendance is also an important metric for schools because it is frequently tied to state funding and many state proposals use chronic absenteeism as an indicator for accountability under the new Every Student Succeeds Act.[1]

Assessing the influence of social networks on risky behaviors such as absenteeism has important implications. Many interventions that aim to attenuate these behaviors can be expensive for school districts to implement. For instance, one of the most effective interventions, Check and Connect, uses student mentors to significantly reduce student absences, costs $1,700 per child per year (Guryan et al. 2016).[2] Though difficult to assess, the benefits of this and other interventions may be understated if there are spillover effects.[3] Moreover, if these spillovers occur along a measurable network, it may be possible to target the intervention more cost effectively by incorporating the potential for spillovers. Nonetheless, this possibility is muted if the networks are expensive to estimate or imperfectly measured, for instance via labor-intensive surveys or proxying a student's social network using students in

[1] cf. This article describing proposed state policies.

[2] There have been several experiments studying Check and Connect, including studies by Sinclair et al. (1998, 2005) and Maynard et al. (2014). However, there have been more recent interventions, similar to the one that we present in this paper, that have shown positive effects at a much lower cost (e.g. Rogers and Feller (2018) intervention has a cost of $6 per additional day of attendance, while Robinson et al. (2018) had a cost of $10.69 per day).

[3] For example, Rogers and Feller (2018) find significant spillovers between siblings within a household, but do not estimate effects for other peers.

the same grade.

In this paper, we show how administrative data can be used to cheaply construct social networks around a particular behavior of interest and how treatment effects spill over along these networks. Specifically, we use student-by-class-by-day attendance data to construct networks of who misses class with whom. The strength of each tie (or edge) between students is given by the number of times they miss the same class together. We assess the features of this network and test whether students systematically miss class with other students. We then leverage the random assignment from an automated-text message alert experiment, which includes alerts to parents for each time a student misses a class, to validate whether the networks we construct are valid; if the networks are meaningful, we may see treatment effects spill over to other students in the network. Lastly, we examine to what extent we can use the network information to target potentially costly but effective attendance interventions, subject to a budget constraint, to increase their cost effectiveness.

We find that students are more likely to miss an individual class than a full day of school, and that students systematically miss classes together. In fact, students are 4.7 times more likely to skip class with a specific student (with whom they have the strongest tie in terms of joint absences) than with another of their peers. The networks also *appear* to exhibit strong homophily: students tend to miss class with other students who have GPAs, behaviors, and racial characteristics that are predictive of their own characteristics. However, the latter could be due to correlated shocks or omitted contextual factors that induce apparent homophily. We examine whether this explanation can account for the observed homophily in the network by comparing simulated moments of the data to their observed counterparts. These tests provide evidence that much of the apparent homophily is due to contextual factors; however, we do find significant homophily for gender and academic performance. To test whether these networks are driven by omitted factors, we measure whether the text-message alert intervention exhibits spillovers onto individuals with whom treated students have strong network ties. We find there are significant spillovers along our network. In contrast, Bergman and Chan (2017) find that a common alternative measure of a student's network—students in the same grade as other students—exhibits weak, statistically insignificant spillovers. Lastly, we show evidence that joint absences are, in part, due to utility derived from missing class jointly with a specific student rather than deriving general utility missing class.

We also provide a targeting approach designed to maximize the total effect of an intervention considering heterogeneous spillovers. Given a budget restriction (the overall number of students that can be treated), this algorithm shows the potential to target other, more-expensive attendance interventions cost effectively. By identifying different types of students and their connections within the network, we show that it would be more efficient to first allocate the treatment to students who have many connections as other students' strongest tie, and at the same time, have higher absenteeism rates. We find that using the social networks to target the treatment doubles the total effect on attendance in comparison to targeting the effect only based on baseline absenteeism rates.[4] We also provide the framework to conduct the same exercise under other objective functions, such as the reduction of chronically absent students.

This paper contributes to a large literature on the interaction between social networks and risky youth behaviors.[5] The influence of peers on individuals' behaviors is difficult to estimate, in part, due to the reflection problem (Manski 1993). In the context of social influence on risky behaviors, a number of papers overcome this difficulty by structurally estimating peer interactions, as in Card and Giuliano (2013) and Richards-Shubik (2015), or by using quasi-random or random variation in the assignment of peers to individuals as in Imberman et al. (2012) and Carrell and Hoekstra (2010), and Feld and Zölitz (2017), Duncan et al. (2005) and Kremer and Levy (2008) respectively. The latter two examples use

[4] Banerjee et al. (2018) show an alternative way to analyze how information flows within a network by identifying individuals that are central to the network by collecting some low-cost primary data, and using them as seeds to spread information.

[5] While we focus on social-networks effects on risky behaviors, Sacerdote (2011) reviews the broader literature on peer effects in educational contexts, such as Hoxby (2000), Sacerdote (2001), and Angrist and Lang (2004).

the random assignment of roommates to identify peer effects and find significant effects of peer drug and alcohol consumption on own use in college, while Feld and Zölitz (2017) use the random assignment of students to different sections at a university level, analyzing the effect of "peer quality" (identified by past GPA) on student performance. The authors find that even though there is a small but significant effect associated with higher performing peers, low-achieving students are actually worse off when sharing a section with high performance peers. Imberman et al. (2012) and Carrell and Hoekstra (2010) find that exposure to students who exhibit behavior problems leads to increased behavioral issues for their peers. Card and Giuliano (2013) and Richards-Shubik (2015) structurally estimate models of peer interactions around sexual initiation using self-reported networks.

Card and Giuliano (2013) write that one limitation of studies focusing on the random assignment of peers to individuals is that, because these peer relationships are formed primarily due to exogenous factors, any resulting peer effects on risky behaviors might not reflect those found in friendships that form more organically. Paluck et al. (2016) overcome this by surveying the entire student bodies of 56 schools to assess students' social networks and randomize an anti-bullying intervention. They find that highly-connected students had outsize effects on changing social norms in schools.

We develop an alternative network measure that sits between these research designs and has several advantages and disadvantages. First, in terms of the former, using administrative data on truancy has the advantage of reflecting networks based on the exhibited behavior of interest, which may be more pertinent to risky behaviors than general friendship networks, randomly assigned peers, or measures based on self-reported behaviors. Second, collecting secondary administrative data is typically lower cost than primary-data collection.[6] Finally, Marmaros and Sacerdote (2006) show that proximity and repeated interactions, which is likely to occur for students who share the same classes, are strong predictors of long-term

[6] de Paula et al. (2018) and Manresa (2016) also show how to recover social networks using panel data, identifying the entire structure of the network under a sufficient set of assumptions, but not necessarily identifying links in the network from observered joint behavior.

friendships. The disadvantages are that we place restrictions on how we define social networks by using class schedules, and we cannot be certain students are actually coordinating their absences. To test the latter, we simulate random networks under the null hypothesis that students do not coordinate their absences. We find that our observed measures of joint absences occur more frequently than what would be expected by chance under our chosen data-generating process. The latter may still have limitations due to unmodeled shocks, so we show our networks are meaningful by demonstrating that a randomly-assigned intervention exhibits spillovers along the observed networks. In this way, our paper relates to the study of peer influence in the context of a randomized intervention, as in the adoption of health and agricultural technologies (Foster and Rosenzweig 1995; Kremer and Miguel 2007; Conley and Udry 2010; Foster and Rosenzweig 2010; Duflo et al. 2011; Oster and Thornton 2012; Dupas 2014; Kim et al. 2015), the role of social interactions in retirement plan decisions (Duflo and Saez 2003), the adoption of microfinance Banerjee et al. (2013), and education technology adoption (Bergman 2016).

Another literature considers the optimal allocation of treatment assignments (cf. Bhattacharya and Dupas 2012) and the assignment of peers to individuals in the presence of potential peer effects (Bhattacharya 2009; Carrell et al. 2013; Graham et al. 2014). Carrell et al. (2013) use insights from Bhattacharya (2009) and Graham et al. (2014) to optimally assign peer groups in the United States Air Force Academy. Their findings suggest caution when optimally assigning peers; their intervention actually reduced performance, which was likely due to subsequent, endogenous peer-group formation. Our paper is related, but rather than optimally assigning peer groups, we consider the targeting of an intervention across existing peer groups. Nonetheless, caution is warranted as individuals could substitute joint absences with one friend with coordinated absences among their other friends.

Finally, our paper also contributes to an emerging literature on partial-day absenteeism by estimating direct and spillover effects of an intervention on class attendance in contrast to full-day attendance. Similar to Whitney and Liu (2017), we show that partial-day absences are more common than full-day absences. This makes estimating the effects of interventions on attendance at the class-level of particular relevance; relatedly, Liu and Loeb (2017) show that teachers can impact class attendance as well.

2 Background and Data

The study uses data from 22 middle and high schools during the 2015-2016 school year in Kanawha County Schools (KCS), West Virginia. West Virginia ranks last in bachelor degree attainment and 49th in median household income among US states and the District of Columbia.[7] KCS is the largest school district in West Virginia with over 28,000 enrolled students in 2016. The district's four-year graduation rate is 71% and standardized test scores are similar to statewide proficiency rates in 2016. In the school year previous to the study, 2014-2015, 44% of students received proficient-or-better scores in reading and 29% received proficient-or-better scores in math. At the state level, 45% of students were proficient or better in reading and 27% were proficient in math. 83% of district students are identified as white and 12% are identified as Black. 79% of students receive free or reduced-priced lunch compared to 71% statewide.[8]

The district uses a single gradebook system for teachers. Schools record by-class attendance and teachers mark missed assignments and grades using the same web-based platform. Bergman and Chan (2017) used data from this platform to conduct an experiment testing an automated alert system, which is described further in Section 4.

2.1 Data

The sample consists of approximately 11,000 households with roughly 14,000 students. These students were enrolled in grades five through eleven during the end of the 2014-2015 school year. Data come from the electronic gradebook described above and baseline administrative data for students enrolled in grades 6 through 12 during the following school year, 2015-2016.

[7]American Community Survey one-year estimates and rankings by state can be found here.

[8]These summary statistics come from the state education website, which can be found here and this description follows closely to that of Bergman and Chan (2017).

The administrative data record students' race and gender as well as their suspensions and English language status from the previous school year. We code baseline suspensions as an indicator for any suspension in the previous school year.

The gradebook data were available at baseline and endline, and record students' grades and class-level attendance by date. We use these data to construct measures of how many classes students attended after the intervention began as well as the number of courses they failed in the second semester of the year and their GPA. Lastly, we define retention as an indicator taking any courses post treatment. Importantly, we have these de-identified data for *all* students in the district.

3 Network Measurement and Descriptive Statistics

In this study, we define the pertinent social network as the ties between students in the same school who miss the same class on the same day.[9] The strength of the tie (or edge) between students is given by the number of times they have missed the same class together. We can formulate this network as follows. Consider a table of students' class attendance in one school over the course of the year in which students' attendance by class, by day, is indicated by a '1' or '0' as follows

Student	Class 1 day 1	Class 2 day 1	Class 1 day 2	Class 2 day 2	…
Student A	1	1	0	1	⋯
Student B	1	0	0	1	⋯
Student C	1	1	0	0	⋯
Student D	1	1	1	1	⋯
⋮	⋮	⋮	⋮	⋮	⋯

We use these data to create a matrix of student absences:

[9]We do not have information whether these students missed class together coordinately or randomly.

$$A_{N\times C} = \begin{bmatrix} 0 & 0 & 1 & 0 & \cdots \\ 0 & 1 & 1 & 0 & \cdots \\ 0 & 0 & 1 & 1 & \cdots \\ 0 & 0 & 0 & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \cdots \end{bmatrix}$$

Here, N is the number of students and C is the total number of classes times days in a year. We can then formulate a matrix of who skips class with whom by multiplying A by A':

$$S = AA'_{N\times N} = \begin{bmatrix} 0 & 0 & 1 & 0 & \cdots \\ 0 & 1 & 1 & 0 & \cdots \\ 0 & 0 & 1 & 1 & \cdots \\ 0 & 0 & 0 & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \cdots \end{bmatrix} \begin{bmatrix} 0 & 0 & 1 & 0 & \cdots \\ 0 & 1 & 1 & 0 & \cdots \\ 0 & 0 & 1 & 1 & \cdots \\ 0 & 0 & 0 & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \cdots \end{bmatrix}'$$

S is an $N \times N$ matrix where each cell s_{ij} represents how many times student i skipped class with student j. This number represents the strength of the tie or edge between students.

Figure 1 shows an example of this network for one of the schools during the pre-intervention period. In this figure each node (or vertex) represents a student, and the edges between vertexes represent the bond between students: The thicker the edge the stronger the bond, which means students have missed more classes together. In this network we can also observe a certain level of clustering, which indicates a group of students primarily skipping classes with other students in the group.

The main advantage of this network approach is that we have complete administrative attendance data at a disaggregated level, which allows us to construct all the connections between students with respect to attendance. However, we do not have information on the reason for the students' absences, which makes connections occurring from a random shock a concern. We address this concern in two ways: 1) identifying "most likely" ties by comparing the observed networks to simulated random networks under different data generating processes, and 2) leveraging the randomized experiment conducted in the KCS district to estimate spillover effects through the network. We describe the detail of these approaches below.

For simplicity we focus attention on the peer in students' networks with whom they skip the most classes (if there exists such a peer) and their associated characteristics and spillover effects, their "strongest tie". Given that their strongest tie is the one with whom a student skipped the most classes simultaneously, we expect that spillovers would be larger through this connection than through other weaker ties. All results that follow become weaker and more imprecise when we explore weaker ties (see Table A.1 in the Appendix).

3.1 Descriptive Statistics and Testing for Coordination

To analyze the social networks in each school in the absence of the intervention, we use baseline data from the beginning of August until the treatment began at the end of October to construct the networks in each school. These data are *prior* to random assignment of the intervention. Table 1 shows baseline summary statistics of the sample. Most students in KCS identify as white and 13% of students identify as Black. Additionally, 50% of students identify as female. Reflecting the student population, few students (2%) are classified as English-Language Learners, and 20% of the sample had at least one suspension in the past year. 4% of the sample was treated and 3% has a peer who was treated.

We estimate several network-level measures that we use to compare to our simulated measures: the size of the network, the clustering coefficients and the degrees of centrality. The size of the network relates to the number of nodes that are in the network, on average, which in this case corresponds to the average number of students, by school, that skip class with another student. The average clustering coefficient refers to the number of closed triplets over the total number of triplets in a network (Jackson 2008),[10] measuring how

[10] A triplet is defined by three nodes connected by two (open) or three (closed) edges. A closed triplet refers to three nodes that are directly connected by three edges.

tightly units are grouped together relative to random ties.[11] Due to the fact that edges have different weights in our network, which are represented by the number of absences between students, we use a weighted average of the clustering coefficient using both an arithmetic and geometric mean to consider the weights of a triplet (Opsahl and Panzarasa 2009).

Degrees of centrality refer to the number of edges between nodes or vertexes. In our case, the degree of centrality of a student is the number of connections with other students, while the eigenvector centrality is the measure of centrality proportional to the centrality of their neighbors (Jackson 2008). Table 2 shows estimates of clustering and centrality in the observed school-level networks.

To further describe the networks at the individual level, we assess the extent of homophily within the networks—whether students tend to skip class with other students who have similar characteristics to themselves. To do so, we regress students' own characteristics on the characteristics of the peer with whom they skip the most class. Specifically, we estimate the following:

$$\text{characteristic}_i = \beta_0 + \beta_1 \text{characteristic}_{ij} + \varepsilon_i$$

In which j is a peer of i, and j indexes the rank of this peer in terms of joint absences. For instance, j equal to 1 indicates the student with whom i has missed the most class. We focus on j equal to 1 for this analysis.

To construct p-values for our homophily estimates, we use a parametric bootstrap for a one-sided test. We construct simulated networks under a null that students skip class independent of one another. Specifically, we formulate a data generating process that matches the baseline absence rate of each student, the absence rate for each specific class[12] (if, for

[11] For example, if networks were randomly generated and students did not coordinate their absences, it would be common to see that student i skipped a class-day with student j and another class-day with student k just by chance; however, if student i, j, and k belong to the same group of friends that skip class together, then it would be common to see that those three students skipped the same day-class, and probably more than once. Then, those three students would form a "closed triplet", and would contribute to the clustering of the network.

[12] A specific class refers to a specific subject-section (e.g. Algebra I in section 1), which occurs at a specific

instance, students dislike a particular teacher), and the absence rate for a given day of the week.[13] We then randomly generated absences for each student i accordingly. We ran 100 simulations per school to create random networks for the pre-intervention period. These simulations allow us to generate distributions of statistics under a null of uncoordinated absences. For instance, if the observed statistic is in the 99th percentile of the distribution under our null, the observed statistic is significant at the 1% level.

We can use the observed aggregate network measures to compare to our aggregate network measures in the simulated networks. The extent to which these aggregate measures are close in magnitude suggests that our data generating process for the simulations is a reasonable approximation of the truth under uncoordinated absences. Table 2 and Table 3 show the average size of the network by school, measures of clustering, and average degree of centrality for the observed and random networks, respectively. The randomly generated networks are similar in terms of the size of the networks. The simulated and observed networks are also similar in their level of clustering. This provides suggestive evidence that our data generating process is reasonable, though we note several further limitations at the end of this section.

Table 4 shows the homophily in the network. The coefficients are all positive. Using the parametric bootstrap, we can reject the null hypothesis for two of the four characteristics at a 5% level: Students are significantly more likely to skip class with another student who has similar GPA or gender.[14] The observed coefficients for GPA and gender characteristics are larger than the maximum value obtained from the simulations, but we find no significant positive difference in terms of race or ever-suspended status using a one-sided test.

Figure 2 shows the distributions for the simulated coefficients as well as the coefficient obtained from the observed data. A number of the simulated coefficients are positive as well. This implies that if one naively regressed students' characteristics against those of the peer

time.

[13]Predicted probabilities p_{ijt} are obtained using a linear probability model including fixed effects by student, class id, and day of the week.

[14]Correlations of characteristics between students and their peer with the strongest tie at the middle school and high school level are very similar, with the exception of baseline GPA, which is 0.17 vs 0.28 for middle school and high school, respectively. Results are available upon request.

with whom they have the strongest tie many coefficients would be significant even if students were actually skipping class independent of one another. This likely reflects district tracking policies and correlated shocks within networks. Note that the simulated coefficient for gender is close to zero and slightly negative however, while the observed is larger in magnitude and positive.

We use the simulated networks to identify whether students systematically coordinate their absences as well. We compare the number of times students miss class with the peer with whom they share the strongest tie and compare this to the distribution of absences for this student pair in the simulated data. For each simulated network, we constructed the joint absences for each pair of students, which gives us a distribution under the null hypothesis of uncoordinated absences holding each students' individual absence rate by class and day of the week constant. We then calculate the p-value for the test that absences are uncoordinated based on the observed number of classes that student pairs skip together compared to that found under the null distribution. If students skip class randomly and do not coordinate their absences, we should not be able to reject the null hypothesis. However, if student i coordinates their absences with student j, then the observed joint absences would be on the right tail of the distribution, allowing us to reject the null hypothesis for that particular student pair.

Table 5 shows the total number of students who have missed a class with at least one other student, and the number of those students who coordinate their absences according to our parametric bootstrap approach using different thresholds. Almost 50% of the students who have a strongest tie coordinate their absences (at a 90% threshold level) compared to our simulated networks and, on average, skip 6.5 more classes with their strongest tie than what is predicted by randomly simulated networks. This share is well above what we would expect by chance, given our data generating process.

Lastly, to get a sense of the magnitude of coordination, we compare the share of a student's absences with their strongest tie to the average share of absences across their other peers. Students are absent 4.7 times more often with their the strongest tie relative to the average absences across their other peers.

Overall, we find significant evidence that students coordinate their absences as well as evidence of homophily within the networks under our null distribution assumption. However, there are several important limitations concerning the previous analysis. For instance, we cannot observe all contextual aspects, such as the time of the day students skip class or whether other events might account for coordinated absences (e.g. sport events or class activities). It could also be the case that students that share similar characteristics tend to skip the same classes without coordination (e.g. neighborhood effects). To provide an additional robustness test to some of these concerns, we conducted a second set of simulations for a null distribution matching the absence rate for each specific class, day of the week, and specific characteristics of the students (i.e. gender, English learner status, baseline GPA, and neighborhood). Results show similar patterns as described by the first set of simulations. Additionally, to further assess the importance of networks and attendance, in the following section we analyze whether treatment effects from the alert intervention spill over onto peers within a student's baseline attendance network.

4 Network Spillovers

4.1 Previous Experimental Design

Bergman and Chan (2017) used data from the online gradebook platform to create and test a text-message alert system to inform parents about their child's academic progress. That study tested three types of parent alerts: Low-grade alerts, missed assignment alerts, and by-class attendance alerts. On Mondays, parents received a text-message alert on the number of assignments their child was missing (if any) for each course during the past week. These assignments included homework, classwork, projects, essays, missing exams, tests, and quizzes. On Wednesdays, parents received an alert for any class their child had missed the previous week. Lastly, and normally on the last Friday of each month, parents received an alert if their child had a cumulative average below 70% in any course during the current marking period. Each alert was sent at 4:00 P.M. local time and the text of each alert is provided in Table 6. The text messages also included a link to the website domain of the parent portal, where the parent could obtain specific information on class assignments and absences if necessary.

Within the same sample used in the network analysis, Bergman and Chan consented 1,137 students' parents or guardians to participate in the experiment studying the effects of the alerts. Among consenting families, random assignment was clustered at the school-by-grade level.[15] The intervention began in late October 2015 and continued through the remainder of the school year.

Parents in the control group received the default level of information that the schools and teachers provided. This included report cards that are sent home after each marking period every six to nine weeks along with parent-teacher conferences and any phone calls home from teachers. All parents had access to the online gradebook. Figure 3 shows a diagram of the experiment randomization. Note that we have administrative data on *all* students—those who consented and those who did not—which allows us to identify networks and spillover effects from consenting, treated students, onto all other students.

The parent alerts caused significant (39%) reductions in course failures and increases (17%) in by-class attendance. For further details on the experiment and the direct effects of the intervention, see Bergman and Chan (2017).

4.2 Treatment effect spillovers from strongest tie on attendance

To assess spillovers, we use the baseline attendance data—data collected prior to any intervention implementation or random assignment—to construct school-level networks. Note, we have data on all students; both those who participated in the experiment and those who did not. The key characteristic that we use from these networks is an indicator for whether

[15]The data were collapsed at the grade-by-school level and randomization was subsequently stratified by indicators for below-median grade point average (GPA) and middle versus high school grades.

a student's strongest tie, which we define as the peer with whom the student has the most joint absences, referred to as *Peer 1*, was treated or not. This helps answer the question: if the person you skipped the most classes with is treated, does this affect your attendance as well?

We estimate the following equation to examine peer effects:

$$\mathrm{y}_i = \beta_0 + \beta_1 \mathrm{P1Treat}_i + \beta_2 \mathrm{Treat}_i + \beta_3 \mathrm{Peers}_i + \beta_4 \mathrm{D1}_i + \beta_5 \mathrm{sample}_i + \beta_6 \mathrm{sample1}_i + \gamma_i X_i + \varepsilon_i \quad (1)$$

In this equation, the key outcome of interest, y_i, is the number of classes attended after the intervention began, though we also check for effects on other gradebook outcomes such as course failures and GPA. $\mathrm{P1Treat}_i$ is an indicator for whether the strongest tie is treated, and Treat_i is a binary variable for whether student i was directly treated. All regressions control for the variable Peers_i, which is the number of peers with whom student i has skipped class, as well as an indicator $\mathrm{D1}_i$ which accounts for whether the student has skipped class with at least one other student. Additionally, we include indicator variables for whether the student and their strongest tie were in the original experimental sample (sample_i and $\mathrm{sample1}_i$, respectively). These variables are key, as treatment is randomly assigned, but having a peer in the original experimental sample is not.[16] All regressions also include the original strata from the treatment assignment. The X_i include additional controls specified in Bergman and Chan's pre-registered analysis plan for the student and his or her strongest tie. These variables are indicators for race, gender, suspension in the past year and IEP status, as well as baseline attendance and GPA. Our preferred specification shows results controlling for covariates as these greatly improve precision. P-values adjust for the clustered design at the school-by-grade level, which is the original unit of treatment assignment.[17]

To test for heterogeneous peer effects, we interact the $\mathrm{P1Treat}_i$ with baseline academic

[16] As a robustness check, we also incorporated flexible interactions of Peers_i variable with the $\mathrm{P1Treat}_i$ treatment variable, and our results are extremely similar.

[17] Moreover, students' own grade level is a near one-to-one predictor of their strongest tie's grade level; the coefficient on a regression of own grade level on their strongest tie's grade level is 0.96.

and demographic covariates for student i. Similarly, we also examine heterogeneity by measures of centrality of the strongest tie as well, such as eigenvector centrality and the number of students for whom Peer1_i is also the strongest tie. This last measure would give us an indication of how many students are connected to student i's strongest tie.

One important challenge in our analyses is adjusting inference for the small number of network clusters (i.e. 22 schools), particularly given the small share of treated students within each cluster, which may bias cluster-robust standard errors. Accurate inference often relies on the number of clusters going to infinity. For that reason, we use a randomization inference approach to obtain accurate p-values for our spillover effects under sharp null hypothesis. Following Athey et al. (2018) (AEI), we randomly choose half of the clusters to be our "focal" units, and then randomly assign the non-focal clusters to either treatment or control maintaining the observed number of clusters in each group according to the observed experiment.[18] Then, as in AEI, we construct the observed test statistic of choice, in this case a T-score, as the covariance between the P1Treat variable for our focal group and the residual of our preferred specification excluding that variable. The idea is that if there is little spillover effect, then this test statistic would be closer to 0. [19] Finally, we compare our observed test statistic (T-score) with the distribution of test statistics under the null of no spillovers using 2,000 draws.[20].

Given the design of the experiment, identification for the spillover effect will come from students who have a strongest tie in another cluster (grade) or that did not consent to participate in the original experiment. To test the validity of the identification strategy, we should see that P1Treat_i is uncorrelated with baseline characteristics of students conditional on the Peers_i variable if this research design is valid. Table 7 shows the result of estimating

[18] The allocation of clusters to the focal group was random, based on the number of network edges between clusters in an adaptation of a *2-net* approach. See Athey et al. (2018) for more details.

[19] While Athey et al. (2018) find this test statistic is more powerful for detecting treatment effects, our results hold if we simply use the regression coefficients as the test statistic; we find the power gains are quite small in our context (results available upon request).

[20] For estimating the p-values for the direct effect under sharp nulls, we followed a standard randomization inference approach for testing sharp null hypothesis on treatment effects detailed in AEI.

equation (1) with baseline covariates as the dependent variable. The magnitudes are all small and statistically insignificant, particularly around baseline absence measures, which provides reassurance that peer treatment status is randomly assigned.

4.3 Results

4.3.1 Attendance

Given the networks are constructed based on class absences, we focus on whether there are spillover effects of the treatment on students' by-class attendance and the robustness of these effects to different measures of peers. To test whether there are spillover effects on attendance for students whose strongest tie was treated, we estimate two models that build upon each other: (1) a simple regression between P1Treat_1 and attendance controlling for the size of the network and whether the students and his strongest tie was in the original experiment sample and (2) we then add controls for the set of predefined covariates described above.[21]

Table 8 shows the results, using randomization inference for estimating p-values. Both models yield stable and positive spillover effect of treated students onto their strongest tie. The estimated spillover effect is 24 more attended classes (note that, as described previously, classes are not the same as days of school). We can reject the sharp null hypothesis of no spillover effect with a p-value < 0.01 (see also Figure 4a and 4b). This is 72% of the direct treatment effect (ITT) of the intervention on classes attended found in Bergman and Chan (2017), and 8% of the mean for control students for those who were not treated and did not have their strongest tie treated.

As a reference point for these magnitudes, several papers have studied the impacts of parent interventions on attendance while proxying for study participants' peers using students in the same classroom or school (as opposed to a more direct attempt to measure students' networks). Avvisati et al. (2013) examined the spillover effects of a parent-engagement intervention in which some students' parents were treated in a classroom while others were

[21] The size of the network is defined as the number of peers with whom a student skipped classes.

not, and then compared these two groups of students to control-group classrooms where no students' parents were treated. In that setting, students remained with the same group of students throughout the day. They found spillover effects on attendance roughly 50% of the size of the direct (ITT) effect of the intervention. Cunha et al. (2017) conducted a parent-texting intervention in Brazil—messaging parents about their math assignments and absences—randomized to detect within-school spillovers. They find spillover effects on GPA and attendance nearly 100% of the size of the direct effects when comparing untreated students in treated schools to students in schools where no one was treated.

We also analyze how the joint attendance between a student and their strongest tie changes during the post-intervention period, given that there could be potential substitution effects in absences with other peers. Table 9 shows the results of these analyses. If students derive utility from a joint absence or a joint attendance with a particular peer, we might observe that joint *attendance* with their strongest tie increases if that peer is treated. That is, if a strongest tie is treated and students value spending time with that peer, we should see that attendance in classes *with that peer* increases. Panel A shows the effect of having a student's strongest tie treated on the number of classes attended with that peer. Panel B shows the effect of having a student's strongest tie treated on the number of classes attended with all other students, excluding the strongest tie. The results show that most of the increase in attendance occurs through attendance with the strongest tie, which is consistent with the idea that students derive utility from spending time with a particular peer. However, the results in Panel A are imprecise and we do not have enough precision to rule out a zero effect.

4.3.2 GPA, Course Failures and Retention

Finally, we also analyze whether spillovers from treated peers extend beyond attendance to other measures of academic performance: GPA, number of failed courses, and dropout. Table 10 shows the results for these outcomes.

For inference, we again used the randomization inference approach of Athey et al. (2018)

as we did for the attendance outcomes. We do not find evidence of a significant effect on GPA, failed courses or retention, though we may be under-powered to detect effects on academic outcomes other than attendance.

4.3.3 Heterogeneity

We examine how these treatment effects vary by network and demographic characteristics. We explore heterogeneous effects by gender, race, academic performance, suspension status, baseline absences,[22] centrality in the network of their peer,[23] and whether their strongest tie is also the strongest tie to more than 2 students. Table 11 shows the results for heterogeneous effects. We find that spillovers seem to be larger for those students who have ever been suspended the previous year. We test many interactions however, so these results should be interpreted as exploratory.

Lastly, we follow Athey and Imbens (2016) to identify potential heterogeneity in the spillover effects while mitigating the threat of "data mining." Athey and Imbens (2016) use machine learning techniques to identify groups of students within the data who experience differential spillovers. Table A.2 shows the results for the spillover effects under heterogeneity. We find that one subgroup, those with a low baseline absence rate, have a higher spillover effect than others in the sample (significant at the 10% level).

In results not shown, we also analyzed whether there are second-degree spillovers. If peer 1 is a student's strongest tie, we define second-degree spillovers as spillovers stemming from whether or not the strongest tie to peer 1 is treated or not. We find no significant effects.

4.3.4 Robustness

In the appendix, we consider the robustness of our spillover measure to a stricter definition of coordinated absences. We generate an indicator that equals one if the strongest tie measure

[22] A student is considered to have a high percentage of baseline absences if it is higher than the median for all students with at least one absence.

[23] We define a student as "central" if, according to their eigenvector centrality, they are on the top 50% of the distribution.

is significant at the 10% level according to our parametric bootstrap.[24] While more than 40% of the sample have joint absences that pass this test, very few of these students are treated, so the estimates are much less precise than before. Table A.3 shows the results of spillovers for peers that are treated and coordinate their absences according to our bootstrap sampling. Despite larger standard errors, we find that the point estimates are slightly larger than before, though the results are not statistically significant at conventional levels.

Lastly, if a handful of students miss many classes, this may generate some extreme baseline values for the Peers_i variable, which is the number of individuals with whom a student missed class. We examine the robustness of our effects by dropping observations whose baseline, pre-random assignment Peers_i value is more than three standard deviations away from the mean. Table A.4 shows the effects are still significant and very similar in magnitude.

These results contrast to results using another measure commonly used as a proxy for networks in clustered-randomized controlled trials. This alternative strategy for measuring spillovers compares the untreated students in a treated cluster to another cluster that is completely untreated. This occurs, for instance, if fractions of a classroom are treated and some classrooms are untreated. This strategy can be effective in settings where students may change classrooms during the day but they do so with the same group of students.[25] In the United States, many high-school and middle-school students change classrooms and the student composition of the classrooms may change as well. Students do tend to remain within grades however, and Bergman and Chan (2017) analyze within-grade spillovers using this design, but find no evidence spillover effects.

These findings show that our network measures are meaningful not only in the sense that students coordinate absences, as shown in the previous section, but also because there are meaningful spillover effects that occur along this network as well.

[24]We also constructed the same variable for a 50%, 40%, 30%, and 20% threshold.

[25]For instance, Avvisati et al. (2013) use this measure in an experiment aimed to involve parents in their education, and they find evidence of spillover effects.

5 Targeting of treatment and Cost Effectiveness

We show the implications of these spillovers for targeting the intervention and a basic accounting exercise to measure cost effectiveness. For the latter, the cost of the learning management software, training, and text messages is $7 per student. Without accounting for spillovers, the cost per additional class attended given the number of students treated and the intent to treat effects found in Bergman and Chan (2017) is $0.21. Incorporating the average spillover effect given the number of students with their strongest tie treated, the cost per additional class attended falls by 43% to $0.12.

We next assess the extent which we can leverage the social network information to target the treatment more cost effectively. This is primarily a conceptual exercise, as this particular treatment has low marginal cost, but other evidence-based absence interventions cited above (e.g. Check and Connect) cost thousands of dollars per treated student. Nonetheless, even for low cost interventions such as this one or the one presented in Rogers and Feller (2018), it could improve the targeting of the treatment by maximizing the number of students the treatment could reach directly and indirectly. We solve for a targeted allocation of the intervention given the direct effects and spillovers previously estimated, subject to a cost restriction. We represent the budget restriction as a maximum number of students that can be treated. Given that most policy-relevant interventions are subject to a budget constraint, we aim to find the maximum impact on class attendance subject to the number of possible students who can be treated. We could also consider objectives as well in which we aim minimize the number of chronically absent students, which could be an appealing objective for schools and policymakers.

We make the following assumptions to simplify this problem:

1. No general equilibrium effects. We assume that the direct and spillover effects do not change with the share of treated students. This assumption is plausible when small shares of students are treated, but could be violated when the proportion of treated

students increases.

2. Homogeneous effects within types of students. To simplify our model and make it computationally feasible, we consider heterogeneous effects on a particular set of characteristics (specified below), and assume the effects and spillovers are constant with respect to other individual and school characteristics not considered in the optimization model. This assumption may not hold in all settings, and other relevant characteristics should be included in the analysis depending on the context, increasing the types' of students considered.

3. Spillover effects only occur through a student's strongest tie. Empirically, we did not find significant spillover effects beyond a student's strongest tie, so we assume that peers that have weaker ties to a student have a negligible spillover effect on that student's attendance.[26] Formally, let $\mathbf{PTreat_i}$ be a vector indicating the treatment statuses of student i's J peers, where J is ordered by the strength of ties to student i such that $j = 1$ indicates with whom student i skips the most class. We assume that

$$\mathbf{PTreat_i} = \text{PTreat}_{i1}$$

4. No school-boundary considerations. For simplicity, we present the optimization model for the total population of the 22 schools in our sample, without considering allocation restrictions within schools. Adding restrictions within schools for student allocations is straightforward and can be implemented by either solving the same model for each school, or adding a school variable interacted with the type of student characteristic.

5.1 Maximizing the Effect of the Intervention

We begin by setting up the objective function as the maximization of the effect of the intervention on class attendance, irrespective of the distribution of this effect across students.

[26] We previously defined ties between peers as the number of joint absences.

Let $I = \{1, 2, .., n\}$ be the set of students that could potentially be treated, and, as defined by Sviatschi (2017), let A be a $n \times n$ matrix for effects and spillovers. In this case, each off-diagonal cell (i, j), where $i \neq j$, contains the spillover effect of student i on student j, and the elements on the diagonal, (i, i), represent the direct effect of the treatment of student i[27].

In order to set up our maximization problem, we need to define our strongest tie matrix D as well. This network will contain an indicator $d_{ij} = 1$ if student i's strongest tie is student j, and 0 otherwise[28]. The definition of this matrix can be easily obtained from the social network matrix S previously defined, by taking an indicator function to each row vector and applying it to the maximum number of joint absences.

In this case, the decision variables are z_i $(i \in I)$ and m_{ij} $(i, j \in I;\ i \neq j)$, where z_i is a binary allocation variable that indicates whether the treatment is assigned to student i, and m_{ij} is an indicator variable for treatment assigned to student j, and he/she is student's i strongest tie $(m_{ij} = z_j \cdot d_{ij})$.

The objective function we maximize corresponds to the function that optimizes the allocation of the treatment given the direct effects and spillovers we estimated, subject to the budget restriction of treating at most b students as follows:

$$
\begin{aligned}
\max_{\mathbf{z},\mathbf{m}} \quad & \sum_{i \in I} a_{ii} z_i + \sum_{i \in I} \sum_{j \in I; i \neq j} a_{ij} m_{ij} & \\
\text{s.t.} \quad & m_{ij} = z_j \cdot d_{ij} & \forall i, j \in I; i \neq j \\
& \sum_{i \in I} z_i \leq b & \\
& z_i \in \{0, 1\} & \forall i \in I \\
& m_{ij} \in \{0, 1\} & \forall i, j \in I; i \neq j
\end{aligned}
$$

In order to estimate the effects matrix A and the potential heterogeneity in spillover effects, we used a machine learning approach proposed by Athey and Imbens (2016) to

[27] One key difference with Sviatschi (2017) is that we do allow for spillovers onto students that are also directly treated

[28] Note that because we are only choosing one strongest tie, $\sum_{i \in I} D_{i\cdot} = 1$

identify different "types" of students according to baseline characteristics. We define $t = 1, ..., T$ as the different types of students that are identified through this approach.

Then, the total effect exerted by student i (who is type t), would be:

$$E_i = a_{ii}z_i + \sum_{j \in I; i \neq j} a_{ij}m_{ij}$$

$$E_i = d_t z_i + \sum_{t \in T} s_t n_{it} z_i = z_i(d_t + \sum_{t \in T} s_t n_{it})$$

Where d_t is the direct effect for treated students of type t, s_t is the spillover effect that a student of type t would experience if his/her strongest tie was treated, and n_{it} is the total number of students type t for whom student i is a strongest tie, meaning $n_{it} = \sum_{j \in t, j \neq i} d_{ij}$.

Thus, it is easy to see that we can re-write the previous optimization problem as following:

$$\begin{aligned}
\max_{\mathbf{z}} \quad & \sum_{t \in T} \sum_{i \in I_t} z_i(d_t + \sum_{k \in T} s_k n_{ik}) \\
\text{s.t.} \quad & \\
& \sum_{t \in T} \sum_{i \in I_t} z_i \quad \leq \quad b \\
& \qquad\qquad z_i \quad \in \quad [0, 1] \qquad\qquad \forall i \in I
\end{aligned}$$

The budget restriction is given by b, which represents the maximum number of students that can be treated given our budget constraint. Vectors d and s represent the direct effects and spillovers for each type of student. Thus, the objective function sums the direct effect of the treatment in terms of the number of present days with the spillover effects that the treatment has on other students, considering the number of students who have the most connections.

To illustrate this algorithm, we consider the following groups of students: students who do not have a strongest tie (i.e. have not missed class with another student), referred to as NP, and the subgroups identified by the machine-learning analysis proposed by Athey and Imbens (2016) described above: students with baseline absence rates greater than 4.5% (Group 1) and students with baseline absence rates lower than 4.5% (Group 2). This analysis could easily be extended to more groups if they are relevant in other contexts. These characteristics define three different types of students, $T = \{NP, P1, P2\}$. Each of these types of students connect to other types of students. The types of students can be identified considering other relevant characteristics that affect the magnitude of the effects or spillovers, or students whom the district would otherwise like to target. The same logic above would apply as well.

For our particular example, we define our objective function as the maximization of the overall effect on our population. Given that all types of students have the same direct effect and that students with lower baseline absenteeism rates have higher spillover effects, we first treat students in Group 1. From this specific group, though, we also want to target first those students who have the most connections, meaning that they are other students' strongest tie, and, thus, that would exert the largest overall effects.

Figure 5 shows the total effect of the treatment when targeted, as well as the total effect of the intervention (given our model) for the observed experiment. We compare this optimal targeting strategy to a commonly used rule of thumb, such as targeting students with the highest absenteeism rates[29]. We can see that when the intervention is targeted to 604 students (the number of students treated in our sample), targeting the intervention increases the overall effect by almost 5 times. Compared to an alternative targeting mechanism, such as targeting by absenteeism rates, the optimal allocation by effects considering the network structure would double the total effect of the intervention.

This analysis can be further extended to different types of students following the same optimization algorithm and given other relevant characteristics that might affect the magnitude of the spillovers, direct effects, or otherwise prioritized students.

The same analysis can be easily adapted to other objectives, such as the reduction of chronically absent students. In that case, further assumptions should be made about the stability of the absenteeism rates in the post-intervention period, as well as considering how

[29] This targeting scheme performs better than random allocation, particularly because students with high absenteeism rates also tend to have more connections to other students, thus exerting more spillovers.

"far" from the chronically absent threshold the student is. For instance, some students that are considered chronically absent might change status if they are close enough to the threshold and their strongest tie is treated, without treating them directly. Others, for instance, might need to be treated directly as well as their strongest tie.

6 Conclusion

In this paper, we demonstrate a straightforward way to estimate meaningful social networks around a student's risky behavior, truancy. Our method is based on detailed, student-by-class-by-day attendance information for every student, which we use to construct a matrix of attendance for every student within a school.

Our study has several advantages and limitations compared to previous research studying social network effects on risky behaviors. First, one branch of the literature uses the exogenous assignment of peer groups to identify peer effects on risky behaviors, as opposed to naturally occurring friendships. In contrast, Card and Giuliano (2013) use self-reported friendship networks and estimate a structural model to discern peer effects on risky behaviors. Our paper sits somewhat in between these strategies. We use administrative data on the risky behavior, truancy, to construct social networks and combine that with a randomly assigned text-messaging intervention to identify social network effects. This strategy has the advantage of being low cost because it relies on existing data and it is less subject to bias arising from self-reporting. However, the disadvantage is that we cannot be certain students are actually coordinating their absences. We overcome the latter by simulating random networks under the null hypothesis that students do not coordinate their absences. We find that our observed measures of joint absences occur more frequently than what would be expected by chance under our chosen data-generating process. Lastly, we show that a randomly-assigned attendance intervention exhibits meaningful spillovers along our estimated networks.

It is important to note that the social networks we identify through behavior pertain specifically to truancy, which might be different to networks related to other behaviors. Thus, we might have limited scope to detect spillover effects onto other outcomes, especially if the intervention is mainly operating through attendance (Rogers and Feller 2018). Nonetheless, given the importance of school attendance, we believe the method we present is still useful to target interventions aimed at reducing students' absenteeism.

We show that these spillover effects are meaningful in terms of their implications for measuring cost effectiveness and targeting to improve efficiency. By accounting for the spillovers, the intervention is 19% more cost effective than not accounting for these effects. Moreover, we show that leveraging baseline network information could help target the intervention and increase its overall effectiveness under different objective functions. Future research could test this targeting via a randomized controlled trial.

References

Akerlof, George A and Rachel E Kranton, "Economics and identity," *The Quarterly Journal of Economics*, 2000, *115* (3), 715–753.

Angrist, Joshua D and Kevin Lang, "Does School Integration Generate Peer Effects? Evidence from Boston's Metco Program," *American Economic Review*, 2004, *94* (5), 1613–1634.

Athey, S. and G. Imbens, "Recursive partitioning for heterogeneous causal effects," *Sackler Colloquium on Drawing Causal Inference from Big Data - Colloquium Paper*, 2016.

_, D. Eckles, and G. Imbens, "Exact p-values for network interference," *Journal of the American Statistical Association*, 2018, pp. 230–240.

Aucejo, Esteban M and Teresa Foy Romano, "Assessing the effect of school days and absences on test score performance," *Economics of Education Review*, 2016, *55*, 70–87.

Austen-Smith, David and Roland G Fryer, "An economic analysis of "acting white"," *The Quarterly Journal of Economics*, 2005, *120* (2), 551–583.

Avvisati, Francesco, Marc Gurgand, Nina Guyon, and Eric Maurin, "Getting parents involved: A field experiment in deprived schools," *Review of Economic Studies*, 2013, *81* (1), 57–83.

Banerjee, A. V, A. Chandrasekhar, E. Duflo, and M. Jackson, "Using Gossips to Spread Information: Theory and Evidence from Two Randomized Controlled Trials," *The Review of Economic Studies (Forthcoming)*, 2018.

Banerjee, Abhijit, Arun G Chandrasekhar, Esther Duflo, and Matthew O Jackson, "The diffusion of microfinance," *Science*, 2013, *341* (6144), 1236498.

Bénabou, Roland and Jean Tirole, "Identity, morals, and taboos: Beliefs as assets," *The Quarterly Journal of Economics*, 2011, *126* (2), 805–855.

Bergman, Peter, "Technology Adoption in Education: Usage, Spillovers and Student Achievement," *Columbia University Teachers College Working Paper*, 2016.

_ **and Eric W Chan**, "Leveraging Technology to Engage Parents at Scale: Evidence from a Randomized Controlled Trial," Technical Report 2017.

Bhattacharya, Debopam, "Inferring optimal peer assignment from experimental data," *Journal of the American Statistical Association*, 2009, *104* (486), 486–500.

_ **and Pascaline Dupas**, "Inferring welfare maximizing treatment assignment under budget constraints," *Journal of Econometrics*, 2012, *167* (1), 168–196.

Billings, S., D. Deming, and S. Ross, "Partners in Crime: Schools, Neighborhoods and the Formation of Criminal Networks," *NBER Working Paper*, 2016, (21962).

Bursztyn, Leonardo, Florian Ederer, Bruno Ferman, and Noam Yuchtman, "Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decisions," *Econometrica*, 2014, *82* (4), 1273–1301.

Card, David and Laura Giuliano, "Peer effects and multiple equilibria in the risky behavior of friends," *Review of Economics and Statistics*, 2013, *95* (4), 1130–1149.

Carrell, Scott E. and Mark L. Hoekstra, "Externalities in the Classroom: How Children Exposed to Domestic Violence Affect Everyone's Kids," *American Economic Journal: Applied Economics*, January 2010, *2* (1), 211–28.

Carrell, Scott E, Bruce I Sacerdote, and James E West, "From natural variation to optimal policy? The importance of endogenous peer group formation," *Econometrica*, 2013, *81* (3), 855–882.

Conley, Timothy G and Christopher R Udry, "Learning about a new technology: Pineapple in Ghana," *The American Economic Review*, 2010, pp. 35–69.

Cook, Philip J, Kenneth Dodge, Elizabeth Gifford, and Amy Schulting, "Preventing primary school absenteeism," *Children and Youth Services Review*, 2017.

Cunha, Nina, Guilherme Lichand, Ricardo Madeira, and Eric Bettinger, "What Is It About Communicating With Parents," Technical Report, Stanford CEPA Working Paper 2017.

de Paula, A., I. Rasul, and P. Souza, "Recovering Social Networks from Panel Data: Identification, Simulations and an Application," *CEMMAP Working Paper*, 2018, *CWP17/18.*

Duflo, Esther and Emmanuel Saez, "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment*," *The Quarterly journal of economics*, 2003, *118* (3), 815–842.

_ , Michael Kremer, and Jonathan Robinson, "Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya," *American Economic Review*, 2011, *101*, 2350–2390.

Duncan, Greg J, Johanne Boisjoly, Michael Kremer, Dan M Levy, and Jacque Eccles, "Peer effects in drug use and sex among college students," *Journal of abnormal child psychology*, 2005, *33* (3), 375–385.

Dupas, Pascaline, "SHORT-RUN SUBSIDIES AND LONG-RUN ADOPTION OF NEW HEALTH PRODUCTS: EVIDENCE FROM A FIELD EXPERIMENT," *Econometrica: journal of the Econometric Society*, 2014, *82* (1), 197.

Feld, J. and U. Zölitz, "Understanding Peer Effects: On the Nature, Estimation, and Channels of Peer Effects," *Journal of Labor Economics*, 2017, *35* (2), 387–428.

Foster, Andrew D and Mark R Rosenzweig, "Learning by doing and learning from others: Human capital and technical change in agriculture," *Journal of political Economy*, 1995, pp. 1176–1209.

_ **and** _ , "Microeconomics of technology adoption," *Annual Review of Economics*, 2010, *2.*

Gershenson, Seth, Alison Jacknowitz, and Andrew Brannegan, "Are student absences worth the worry in US primary schools?," *Education Finance and Policy*, 2017.

Goodman, Joshua, "Flaking out: Student absences and snow days as disruptions of instructional time," Technical Report, National Bureau of Economic Research 2014.

Graham, Bryan S, Guido W Imbens, and Geert Ridder, "Complementarity and aggregate implications of assortative matching: A nonparametric analysis," *Quantitative Economics*, 2014, *5* (1), 29–66.

Guryan, Jonathan, Sandra Christenson, Amy Claessens, Mimi Engel, Ijun Lai, Jens Ludwig, Ashley Cureton Turner, and Mary Clair Turner, "The Effect of Mentoring on School Attendance and Academic Outcomes: A Randomized Evaluation of the Check & Connect Program." PhD dissertation, Northwestern University 2016.

Hoxby, Caroline, "Peer effects in the classroom: Learning from gender and race variation," Technical Report, National Bureau of Economic Research 2000.

Imberman, Scott A, Adriana D Kugler, and Bruce I Sacerdote, "Katrina's children: Evidence on the structure of peer effects from hurricane evacuees," *The American Economic Review*, 2012, *102* (5), 2048–2082.

Jackson, M., *Social and Economic Networks*, Princeton University Press, 2008.

Kearney, Christopher A, "School absenteeism and school refusal behavior in youth: A contemporary review," *Clinical psychology review*, 2008, *28* (3), 451–471.

Kim, David A, Alison R Hwong, Derek Stafford, D Alex Hughes, A James O'Malley, James H Fowler, and Nicholas A Christakis, "Social network targeting to maximise population behaviour change: a cluster randomised controlled trial," *The Lancet*, 2015, *386* (9989), 145–153.

Kremer, Michael and Dan Levy, "Peer effects and alcohol use among college students," *The Journal of Economic Perspectives*, 2008, *22* (3), 189–189.

_ **and Edward Miguel**, "The Illusion of Sustainability*," *The Quarterly journal of economics*, 2007, *122* (3), 1007–1065.

Liu, J. and S. Loeb, "Engaging Teachers: Measuring the Impact of Teachers on Student Attendance in Secondary School," *CEPA Working Paper*, 2017, (17-1).

Manresa, E., "Estimating the Structure of Social Interactions Using Panel Data," *Working Paper, MIT Sloan*, 2016.

Manski, Charles F, "Identification of endogenous social effects: The reflection problem," *The review of economic studies*, 1993, *60* (3), 531–542.

Marmaros, David and Bruce Sacerdote, "How do friendships form?," *The Quarterly Journal of Economics*, 2006, *121* (1), 79–119.

Maynard, Brandy R, Elizabeth K Kjellstrand, and Aaron M Thompson, "Effects of Check and Connect on attendance, behavior, and academics: A randomized effectiveness trial," *Research on Social Work Practice*, 2014, *24* (3), 296–309.

Opsahl, T. and P. Panzarasa, "Clustering in weighted networks," *Social Networks*, 2009, *31* (2), 155–163.

Oster, Emily and Rebecca Thornton, "Determinants Of Technology Adoption: Peer Effects In Menstrual Cup Take-Up," *Journal of the European Economic Association*, 2012, *10* (6), 1263–1293.

Paluck, Elizabeth Levy, Hana Shepherd, and Peter M Aronow, "Changing climates of conflict: A social network experiment in 56 schools," *Proceedings of the National Academy of Sciences*, 2016, *113* (3), 566–571.

Richards-Shubik, Seth, "Peer effects in sexual initiation: Separating demand and supply mechanisms," *Quantitative Economics*, 2015, *6* (3), 663–702.

Robinson, C., M. Lee, E. Dearing, and T. Rogers, "Reducing Student Absenteeism in the Early Grades by Targeting Parental Beliefs," *American Educational Research Journal*, 2018, *55* (6), 1163–1192.

Rogers, Todd and Avi Feller, "Reducing student absences at scale by targeting parents' misbeliefs," *Nature Human Behaviour*, 2018.

Sacerdote, Bruce, "Peer Effects with Random Assignment: Results for Dartmouth Roommates," *The Quarterly Journal of Economics*, 2001, *116* (2), 681–704.

_ , "Peer effects in education: How might they work, how big are they and how much do we know thus far," *Handbook of the Economics of Education*, 2011, *3* (3), 249–277.

Sinclair, Mary F, Sandra L Christenson, and Martha L Thurlow, "Promoting school completion of urban secondary youth with emotional or behavioral disabilities," *Exceptional Children*, 2005, *71* (4), 465–482.

_ , _ , David L Evelo, and Christine M Hurley, "Dropout prevention for youth with disabilities: Efficacy of a sustained school engagement procedure," *Exceptional Children*, 1998, *65* (1), 7.

Sviatschi, M., "Making a Narco: Childhood Exposure to Illegal Labor Market and Criminal Life Paths," *Job Market Paper, Columbia University*, 2017.

Whitney, C. and J. Liu, "What We're Missing: A Descriptive Analysis of Part-Day Absenteeism in Secondary School," *AERA Open*, 2017, *3* (2), 1–17.

7 Figures

Figure 1: Social Network for School 1 (pre-intervention period)

Notes: Each node (circle) corresponds to a student, and each interconnecting line or *edge* corresponds to the number of absences between two students on the same day and same class during the pre-intervention period. Attendance information is from the gradebook data

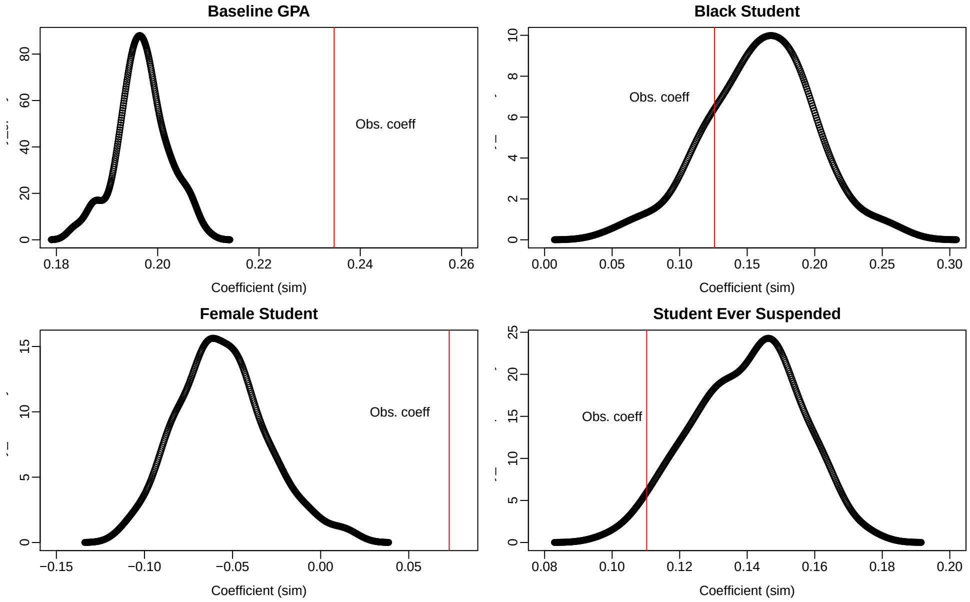

Figure 2: Distribution of simulated coefficients (100 simulations per school) and observed coefficient for homophily analysis

Notes: Distribution of simulated coefficients obtained from the coefficients of regressions using only strata fixed effects and clustered standard errors. Demographic information and suspension data are from district administrative data. Attendance and GPA are from the gradebook data.

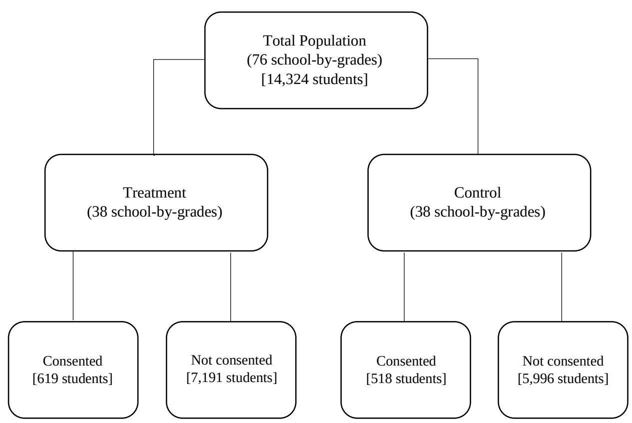

Figure 3: Randomization diagram for experiment at the cluster level (i.e. school-by-grade)

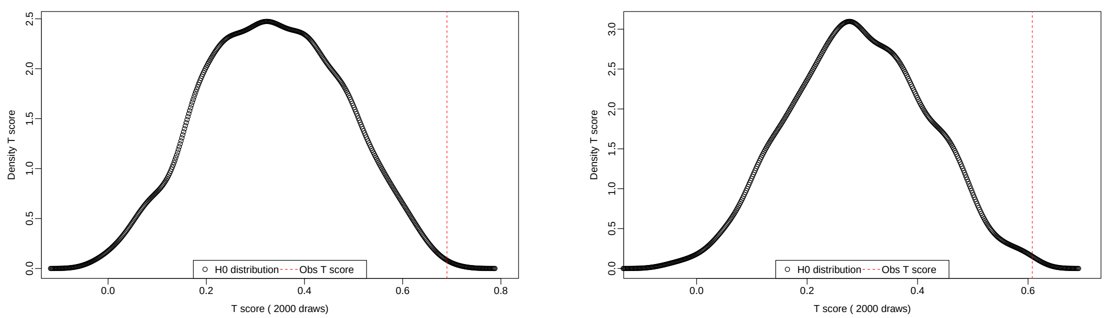

(a) Observed T-score (covariance between regression residual and P1Treat variable) and null distribution of T-score (2,000 draws) for model with no controls.

(b) Observed T-score (covariance between regression residual and P1Treat variable) and null distribution of T-score (2,000 draws) for preferred model with controls.

Figure 4: Randomization Inference Approach for Sharp Null Hypothesis of No Spillover Effect

Notes: Both models include controls by strata, indicators for sample (for both the student and his strongest tie), number of peers, and an indicator for whether the student had a tie to another student or not. Additional controls for preferred model include *GPA baseline, IEP, ELL, missed days (previous year), gender, black,* and *ever suspended* for both the student and his/her strongest tie. Demographic information, IEP status, English Language Learner status and suspension data are from district administrative data. Attendance and GPA are from the gradebook data.

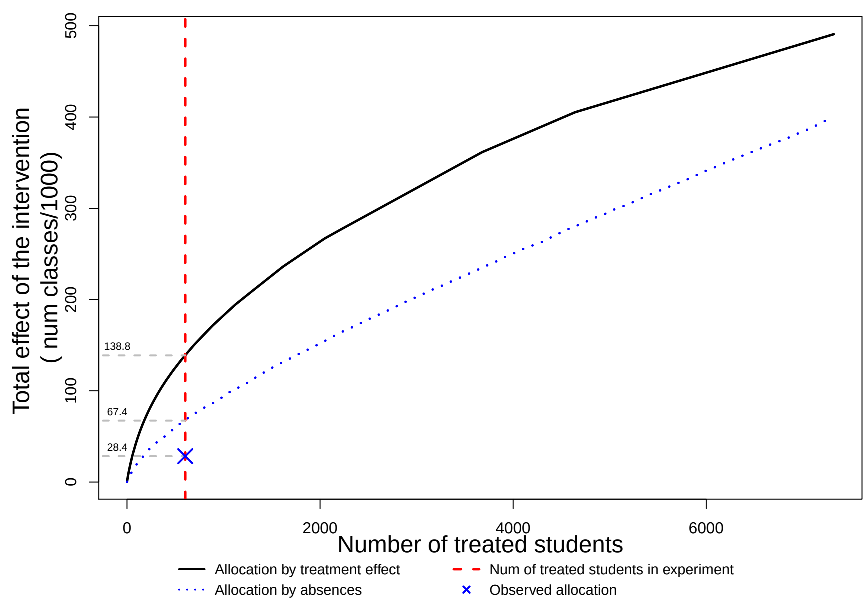

Figure 5: Total effect of the intervention effect under different allocation mechanisms by number of students treated

Notes: Allocation by treatment effect corresponds to the allocation that maximizes total treatment effect by considering spillover onto strongest ties. Allocation by absences corresponds to the allocation of treatment given the number of absences a student has (i.e. students with more absences get treated first).

8 Tables

Table 1: Baseline characteristics for the sample

Variable	Mean	Observations
Female	0.50	13,641
Black	0.13	13,641
KCS: Eng Lang Learner	0.02	13,609
IEP	0.15	13,609
Baseline GPA	2.64	14,653
Ever suspended last year	0.20	13,641
Percent of days absent pre-intervention	0.10	14,621
Percent of classes absent pre-intervention	0.14	14,621
Percent of Days Missed	0.06	14,653

Notes: Mean characteristics consider only non-missing observations. Demographic information, IEP status, English Language Learner status and suspension data are from district administrative data. Attendance and GPA are from the gradebook data.
Observations consider school-students.

Table 2: Observed network features

Characteristic	Mean	1st Quartile	Median	3rd Quartile
Average size	606.6	363.0	509.0	759.0
Weighted clustering coefficient				
Arithmetic mean	0.41	0.30	0.43	0.48
Geometric mean	0.41	0.29	0.42	0.48
Centrality				
Average number of edges	39.38	13.55	21.56	65.12
Eigenvector centrality	0.14	0.07	0.11	0.22

Notes: Statistics are constructed using by observed by-class attendance obtained from the gradebook data. Average size of the network measures the mean number of nodes (i.e. students) within each school network. Weighted clustering coefficient estimates the number of closed triplets over the total number of triplets in a network (Opsahl and Panzarasa 2009) weighted by the strength of the edges, using different either an arithmetic or geometric mean. Number of edges corresponds to the number of connections in the networks, including their weight. Eigenvector centrality determines how central a node (i.e. student) is within a network, considering not only the number of connections to other nodes, but also considering its position within the network.

Table 3: Average simulated (random) network features

Characteristic	Mean	1st Quartile	Median	3rd Quartile
Average size	608.8	364.0	494.7	780.6
Weighted clustering coefficient				
Arithmetic mean	0.40	0.30	0.42	0.48
Geometric mean	0.40	0.29	0.42	0.48
Centrality				
Average number of edges	38.47	14.53	23.26	65.30
Eigenvector centrality	0.13	0.07	0.11	0.21

Notes: These statistics are constructed using simulated network data, with 100 simulations per school. Statistics are constructed using by simulated by-class attendance obtained from a data generating process that matches the mean absence rate for students and class id, as well as day of the week. Average size of the network measures the mean number of nodes (i.e. students) within each school network. Weighted clustering coefficient estimates the number of closed triplets over the total number of triplets in a network (Opsahl and Panzarasa 2009) weighted by the strength of the edged, using either an arithmetic or geometric mean. Number of edges corresponds to the average number of connections in the networks, including their weight. Eigenvector centrality determines how central a node (i.e. student) is within a network, considering not only the number of connections to other nodes, but also considering its position within the network.

Table 4: Homophily in the observed network

Baseline Variable	GPA	Female	Black	Ever suspended
Peer1 GPA	0.23***			
	[<0.01]			
Peer1 female		0.07***		
		[<0.01]		
Peer1 black			0.13	
			[0.81]	
Peer1 ever suspended				0.11
				[0.96]
Observations	11654	10474	11861	11657

Notes: P-values from parametric bootstrap in squared parenthesis for one-sided test. Specification of the models include fixed effects by strata. $Peer1$ represents the person with whom a student missed the most class and is constructed using by class attendance data. Demographic information and suspensions are from school administrative data. Attendance and GPA are from the gradebook data.
Results are similar for middle and high school levels.
* Significant at 10%; ** significant at 5%; *** significant at 1%

Table 5: Students with coordinated absences for different thresholds

Variable	Total	% of Sample
Observations in sample (Total)	14,653	100%
Observations with a closest peer	12,740	87%
Observations with coordinated absences		
90% threshold	5,777	39%
95% threshold	4,145	28%
99% threshold	1,821	12%

Notes: Results obtained from parametric bootstrap using 100 simulations of random networks. By-class attendance data in the pre-intervention period is from the gradebook data.
Observations correspond to school-students.

Table 6: Text messages sent to parents

Alert	**Frequency**	**Message**
Low Class Average Alert	monthly	“Parent Alert: [Student Name] has a [X]% average in [Class Name]. For more information, log in to [domain]”
Absence Alert	weekly	“Parent Alert: [Student Name] has [X] absence(s) in [Class Name]. For more information, log in to [domain]”
Missing Assignment Alert	weekly	“Parent Alert: [Student Name] has [X] missing assignment(s) in [Class Name]. For more information, log in to [domain]”

Notes: This table shows the script for each of the three types of alerts sent via text messages: low class average, absence, and missing assignments.

Table 7: Balance between students with a treated peer and without a treated peer

Variable	Control Mean	Treat-Control Difference	P-value	Obs
Female	0.50	0.05	0.233	13,596
Black	0.13	0.00	0.975	13,596
KCS: Eng Lang Learner	0.02	-0.02	0.142	13,596
IEP	0.15	-0.05	0.080*	13,596
Baseline GPA	2.83	0.01	0.855	13,596
Ever suspended last year	0.20	-0.00	0.984	13,596
Percent of days absent pre-intervention	0.10	-0.00	0.955	13,596
Percent of classes absent pre-intervention	0.14	-0.00	0.827	13,596
Percent of Days Missed	0.07	0.00	0.731	13,596

Notes: Specification of the models includes fixed effects for strata, an indicator for *treatment* and *sample* for the student and the strongest tie, a control for number of peers and whether the student has a tie with another student. In the regressions, demographic information, IEP status, English Language Learner status and suspension data are from district administrative data. Attendance and GPA are from the gradebook data. Observations correspond to school-students.
* Significant at 10%; ** significant at 5%; *** significant at 1%

Table 8: Effect of treated peer on attendance

Variable	Classes present	Classes present
Treat	29.80	31.39*
	[0.12]	[0.06]
P1Treat	21.76***	24.16***
	[<0.01]	[<0.01]
Control mean	310.68	310.68
Controls	Strata	All
Observations	12766	12764

Notes: P-values from randomization inference following Athey et al. (2018) shown in squared parenthesis. The model from column (1) includes fixed effects for strata, an indicator for *treatment* and *sample* for the student and the strongest tie, a control for number of peers and whether the student has a tie with another student, with no additional controls. Column (2) includes all the variables from column (1) in addition to controls by *GPA baseline, IEP, ELL, missed days (previous year), gender, black,* and *ever suspended* (for both the student and his strongest tie). In the regressions, demographic information, IEP status, English Language Learner status and suspension data are from district administrative data. Attendance and GPA are from the gradebook data.
* Significant at 10%; ** significant at 5%; *** significant at 1%

Table 9: Effect of treated peer on joint attendance (classes) post-intervention

Variable	Classes present with P1 (1)	(2)
P1Treat	15.80	16.95
	[0.802]	[0.749]
Control mean	46.74	46.75
Controls	Strata	All
Observations	12766	12764

Variable	Classes present *not* with P1 (1)	(2)
P1Treat	5.96***	7.21***
	[<0.01]	[<0.01]
Control mean	263.93	263.93
Controls	Strata	All
Observations	12766	12764

Notes: Top panel shows effect of having Peer 1 on joint attendance; bottom panel shows effect of having Peer 1 treated on attendance with other peers that are *not* Peer 1. P-values from randomization inference following Athey et al. (2018) are shown in square parenthesis. The model from column (1) includes fixed effects for strata, an indicator for *treatment* and *sample* for the student and the strongest tie, a control for number of peers and whether the student has a tie with another student, with no additional controls. Column (2) includes all the variables from column (1) in addition to controls by *GPA baseline, IEP, ELL, missed days (previous year), gender, black,* and *ever suspended* (for both the student and his strongest tie). In the regressions, demographic information, IEP status, English Language Learner status and suspension data are from district administrative data. Attendance and GPA are from the gradebook data.

* Significant at 10%; ** significant at 5%; *** significant at 1%

Table 10: Effect of treated peer on other outcomes

	GPA	
P1Treat	-0.01	-0.02
	[0.328]	[0.363]
Control mean	2.75	2.75
Controls	Strata	All
Observations	12139	12139
	Courses failed	
P1Treat	-0.11	-0.07
	[0.581]	[0.481]
Control mean	0.91	0.91
Controls	Strata	All
Observations	12139	12139
	Drop out	
P1Treat	-0.03	-0.00
	[0.633]	[0.583]
Control mean	0.11	0.11
Controls	Strata	All
Observations	13596	13583

Notes: P-values from randomization inference following Athey et al. (2018) in squared parenthesis. The model from column (1) includes fixed effects for strata, an indicator for *treatment* and *sample* for the student and the strongest tie, a control for number of peers and whether the student has a tie with another student, with no additional controls. Column (2) includes all the variables from column (1) in addition to controls by *GPA baseline, IEP, ELL, missed days (previous year), gender, black,* and *ever suspended* (for both the student and his strongest tie). In the regressions, demographic information, IEP status, English Language Learner status and suspension data are from district administrative data. Attendance and GPA are from the gradebook data.

* Significant at 10%; ** significant at 5%; *** significant at 1%

Table 11: Heterogeneity of effect of treated peer on attendance

Variable	Classes present	Classes present	Classes present	Classes present	Classes present	Classes present	Classes present
P1Treat (CA)	32.29*	25.50*	38.88**	26.36**	38.83**	20.87	36.72
	[0.062]	[0.082]	[0.03]	[0.034]	[0.042]	[0.64]	[0.5]
P1Treat (CA) × female	-8.61						
	[0.178]						
P1Treat (CA) × black		22.78***					
		[0.004]					
P1Treat (CA) × below med GPA			-18.90				
			[0.450])				
P1Treat (CA) × suspended				8.44***			
				[<0.01]			
P1Treat (CA) × below median absences					-24.57		
					[0.104]		
P1Treat (CA) × Peer1 is eigen-central						17.82	
						[0.416]	
P1Treat (CA) × Peer1 is closest peer to >2 students							-19.28
							[0.76]
Control mean	310.68	310.68	310.68	310.68	310.68	310.68	310.68
Controls	All	All	All	All	All	All	All
Observations	12,764	12,764	12,762	12,764	12,764	12,764	12,764

Notes: P-values from randomization inference following Athey et al. (2018) in square parenthesis. All models include fixed effects for strata, an indicator for *treatment* and *sample* for the student and the strongest tie, a control for number of peers and whether the student has a tie with another student, and additional controls by *GPA baseline, IEP, ELL, missed days (previous year), gender, black,* and *ever suspended* for the students and his strongest tie. In the regressions, demographic information, IEP status, English Language Learner status and suspension data are from district administrative data. Attendance and GPA are from the gradebook data.

* Significant at 10%; ** significant at 5%; *** significant at 1%

Appendix A

Table A.1: Effects of treated peers (1st, 2nd, 3rd, 4th, and 5th) on attendance

Variable Classes present	Classes present
P1Treat	22.88**
	(11.14)
P2Treat	19.52
	(12.34)
P3Treat	4.94
	(15.98)
P4Treat	5.66
	(14.11)
P5Treat	10.27
	(17.26)
Control mean	
Controls	All
Observations	

Notes: PiTreat is a binary variable that reflects whether Peer$i$ was treated and Peer$j$ was not, for all $j \neq i$. Clustered-robust SE in parenthesis. The preferred model includes fixed effects for strata, an indicator for *treatment* and *sample* for the student and the strongest tie, a control for number of peers and whether the student has a tie with another student (up to 5), in addition to controls by *GPA baseline, IEP, ELL, missed days (previous year), gender, black,* and *ever suspended* (for both the student and his 5 strongest ties). Demographic information, IEP status, English Language Learner status and suspension data are from district administrative data. Attendance and GPA are from the gradebook data. For robustness, outliers are excluded from the regression (joint absences between student i and student j are greater than 3 SD).
* Significant at 10%; ** significant at 5%; *** significant at 1%

Table A.2: Direct and spillover effects for optimal allocation problem

Variable	Effect
Treatment	31.91*
P1Treat $\times$ Fraction of absences ≥ 0.045	13.52**
P1Treat $\times$ Fraction of absences $< 0.045 \times$ GPA baseline ≥ 3.4	-20.55
P1Treat $\times$ Fraction of absences $< 0.045 \times$ GPA baseline < 3.4	79.40
Controls	All

Notes: Treatment coefficient is obtained from the prefered specification. Spillover coefficients are obtained using Athey and Imbens (2016) approach using the following covariates: *GPA baseline, IEP, ELL, missed days (previous year), gender, black,* and *ever suspended*.

* Significant at 10%; ** significant at 5%; *** significant at 1%

Table A.3: Effects of treated peer for peers that coordinate absences according to parametric bootstrap

Variable	Classes present	Classes present
Treat	27.61	32.01*
	[0.12]	[0.05]
P1Treat (CA)	25.59	28.44
	[0.44]	[0.14]
Controls	Strata	All
Observations	12766	12764

Notes: P-values from randomization inference following Athey et al. (2018) shown in squared parenthesis. The model from column (1) includes fixed effects for strata, an indicator for *treatment* and *sample* for the student and the strongest tie (90% threshold for coordinated absences), a control for number of peers they coordinate absences with and whether the student has a tie with another student with whom they coordinated absences, with no additional controls. Column (2) includes all the variables from column (1) in addition to controls by *GPA baseline, IEP, ELL, missed days (previous year), gender, black,* and *ever suspended* (for both the student and his strongest tie). In the regressions, demographic information, IEP status, English Language Learner status and suspension data are from district administrative data. Attendance and GPA are from the gradebook data.

* Significant at 10%; ** significant at 5%; *** significant at 1%

Table A.4: Effect of treated peer on attendance removing outliers (>3 SD)

Variable	Classes present	Classes present
Treat	24.35	27.87*
	[0.14]	[0.08]
P1Treat	18.32*	22.92**
	[0.07]	[0.02]
Controls	Strata	All
Observations	12449	12447

Notes: P-values from randomization inference following Athey et al. (2018) shown in squared parenthesis. Outliers with joint absences larger than 3 standard deviations from the mean are removed from the sample. The model from column (1) includes fixed effects for strata, an indicator for *treatment* and *sample* for the student and the strongest tie, a control for number of peers and whether the student has a tie with another student, with no additional controls. Column (2) includes all the variables from column (1) in addition to controls by *GPA baseline, IEP, ELL, missed days (previous year), gender, black,* and *ever suspended* (for both the student and his strongest tie). In the regressions, demographic information, IEP status, English Language Learner status and suspension data are from district administrative data. Attendance and GPA are from the gradebook data.
* Significant at 10%; ** significant at 5%; *** significant at 1%