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

Magdalena Bennett, Peter Bergman†

This draft: February 9, 2018

Abstract

Truancy predicts many risky behaviors and adverse outcomes. We use detailed administrative data to construct social networks based on students who miss class together. We simulate these networks and use permutation tests to show that certain students systematically coordinate their absences. Leveraging a parent-information intervention on student absences, we find spillover effects from treated students onto peers in their network; excluding these effects understates the intervention's cost effectiveness by 19%. We find scope to use these networks to implement costlier interventions more effectively. We develop an algorithm that incorporates spillovers and treatment-effect heterogeneity identified by machine-learning techniques to target interventions more efficiently given a budget constraint.

*We thank Claudia Allende, Alex Bowers, Eric Chan, Laura Giuliano, Matthew Jackson, Adam Kapor, and Bruce Sacerdote 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, or if they 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 absenteeism, 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 2016; Cook et al. 2017; Gershenson et al. 2017). 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, Check and Connect, which 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. 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 the same grade.

In this paper, we show how administrative data can be used to construct social networks around truancy, and how treatment effects spill over along these networks. 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 test if the effects of the alerts spill over to other students in the network and how these spillovers interact with the characteristics of the network. Lastly, we examine to what extent we can use the network information to target potentially costly 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

[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).

students systematically miss classes together. In fact, students are 4.7 times more likely to skip class with a specific student (whom we identify as their "closest peer") than with another of their peers. The networks also 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 show this explanation cannot fully account for the observed homophily in the network by comparing simulated moments of the data to their observed counterparts. These tests provide evidence that the observed homophily is not entirely due to contextual factors. To demonstrate the relevance of these networks, we show that the text-message alert intervention exhibits spillovers onto individuals with whom treated students have strong network ties. 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 algorithm 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 higher rates of absenteeism and have a peer with whom they miss class. We find that using the social networks to target the treatment further reduces the number of chronically absent students by an additional 15 percentage points relative to the original allocation of the treatment.

This paper contributes to a large literature on the interaction between social networks and risky youth behaviors.[3] 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 Duncan et al. (2005) and Kremer and Levy (2008) respectively. The latter two examples use the random assignment of roommates to identify peer effects and find significant effects of peer drug and alcohol consumption on own use in college. 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)

[3]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).

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.

Our network measure sits somewhere 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. 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 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. Lastly, we show that the randomly-assigned intervention exhibits meaningful 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 suggestion 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.

The rest of the paper proceeds as follows. Section 2 details the background of the original experiment and the data used for constructing the spillover analysis. Section 3 describes the social networks in each school and its measurements, whereas Section 4 shows the results for the spillover analysis. Section 5 refers to the allocation algorithm for optimal targeting, and finally, Section 6 concludes.

2 Background and Data

The experiment, which describes the original study of the direct effects of parent alerts. took place in 22 middle and high schools during the 2015-2016 school year in Kanawha County Schools (KCS), West Virginia.[4] West Virginia ranks last in bachelor degree attainment and 49th in median household income among US states and the District of Columbia.[5] 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.[6]

The district has a gradebook system for teachers. Schools record by-class attendance and teachers mark missed assignments and grades using the same web-based platform. The Bergman and Chan (2017) study used data from this 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

[4] This description closely follows that of Bergman and Chan (2017).
[5] American Community Survey one-year estimates and rankings by state can be found here.
[6] These summary statistics come from the state education website, which can be found here.

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 1. 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.

2.1 Original Experimental Design

The sample for the original experiment began with approximately 11,000 households with roughly 14,000 students who were enrolled in grades five through eleven during the end of the 2014-2015 school year. The parent or guardian of 1,137 students consented to participate in the experiment studying the effects of the alerts during the following school year, 2015-2016.

Among consenting families, random assignment was at the school-by-grade level. 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. 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. As discussed above, all parents had access to the online gradebook.

The parent alerts caused significant (40%) 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).

2.2 Data

Data for this study come from the electronic gradebook described above and baseline administrative data for students enrolled in grades 6 through 12 during the 2015-2016 school year. 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.

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.[7] 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 attendance:

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

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':

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

AA' is an $N \times N$ matrix where each cell a_{ij} represents how many times student i skipped class with student j. This number represents the strength of the tie or edge between students.[8]

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

[8] Note that the product $[A'A]_{C\times C}$ would show the number of times a class c on a specific day was skipped by 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 describe how we address this concern below.

For simplicity we focus attention on the peer in students' networks with whom they skip the most class (if there exists such a peer) and their associated characteristics and spillover effects. Given that their strongest peer 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 generally become much weaker and more imprecise when we explore weaker ties (results available upon request).

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. Table 2 and 3 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 19% 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 also present several network-level measures: clustering coefficients and degrees of centrality (Table 4). The average clustering coefficient refers to the number of closed triplets over the total number of triplets in a network (Jackson 2008).[9] Due to the fact that edges have different weights in our network, which represent the number of absences between students, we use a weighted average of the clustering coefficient using both an arithmetic or 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 4 shows estimates of clustering and centrality in the observed school-level networks.

[9] 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.

To assess the extent of homophily within the networks—whether students tend to skip class with other students who have similar characteristics to themselves—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.

Table 5 shows the results of this analysis. Across measures, the characteristics of students strongly correlate with the characteristics of their peers. GPA, gender, race, and suspensions all strongly predict these characteristics in their peers.

We benchmark these results by constructing placebo networks or randomly generated networks. These networks are constructed for each school by randomly generating absences for each student according to their overall probability of absence during the baseline period; each student's absent rate is independent of each other and uniformly distributed according to their respective baseline means. We ran 100 simulations per school to create random networks for the pre-intervention period where students randomly skipped classes based on their observed baseline probability of attendance. With these data, we generate a distribution of the measures of the network.

Table 6 shows the measures of clustering, and average degree of centrality for the nodes for the original pre-intervention network and random networks. From Table 6 we can observe that the level of clustering is very similar between the observed network and the simulated random networks, though the observed network present a larger average number of absences between students.

Estimating the same measures on the simulated data, we observe that the level of homophily in the random network is smaller compared to that in the observed networks, except for GPA, where the level of homophily is similar. This indicates that contextual factors, such as tracking students by prior performance, could drive a share of the observed homophily we found in the previous regressions. However, this is less true for student gender, race and behavior. Students in the observed networks are more likely to skip class with another student of the same gender, race, and behavior (as measured by being ever suspended), but the simulated networks show otherwise (Table 5 and 7).

In terms of discerning whether these networks are meaningful or are simply artifacts of district tracking policies and correlated shocks within networks, the results found here are somewhat mixed. Aggregate network characteristics actually match well compared to those from our randomly generated network, but certain measures of homophily substantially differ. We test the significance of the correlations derived from

the regressions more formally by using the 100 simulations to compute placebo regression coefficients. These data allow us to construct a permutation test of each regression coefficient to discern whether the empirically observed coefficient is significantly different from that found in the distribution from the simulated data. We find that we can reject the null hypothesis for all four characteristics, with a p-value < 0.01. The observed coefficients for all tested characteristics are larger than the maximum value obtained from the simulations. Figure 2 shows the distributions for the simulated coefficients, as well as the coefficient obtained from the observed data.

Additionally, we test whether students systematically coordinate their absences in a similar fashion. We examine the number of times students miss class with their strongest peer 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 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 8 shows the total number of students who have a strongest peer, and the number of those students who coordinate their absences according to our permutation test using different thresholds. Almost 50% of the students who have a strongest peer coordinate their absences (at a 90% threshold level) according to our permutation test. This share is well above what we would expect by chance.

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

Overall, we find significant evidence that students coordinate their absences as well as evidence of homophily within the networks. 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 Treatment effect spillovers from strongest peer on attendance

To assess spillovers, we use the baseline attendance data to construct school-level networks. The key characteristic that we use from these networks is an indicator for whether a student's strongest peer, which we define as the peer with the strongest tie to an individual, referred to as *Peer 1*, was treated or not. This helps answer the question: if the person you skipped the most with is treated, does this affect your attendance as well?

We estimate the following equation to examine peer effects:

$$y_i = \beta_0 + \beta_1 \text{P1Treat}_{i1} + \beta_2 \text{Peers}_i + \gamma_i X_i + \varepsilon_i \tag{1}$$

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. P1Treat_{i1} is an indicator for whether the strongest peer is treated. All regressions control for the variable Peers_i, which is the number of peers with whom student i has skipped class. This variable is important as it determines the probability of treatment.[10] All regressions also include the original strata from the treatment assignment, as well as an indicator variable for whether the student was in the original sample of the experiment and whether he or she was directly treated. The X_i include additional controls specified in the original experiment's pre-registered analysis plan. 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 with school fixed effects as these greatly improve precision. Standard errors are clustered at the school-by-grade level, which is the original unit of treatment assignment.[11]

To test for heterogeneous peer effects, we interact the P1Treat_{i1} with baseline covariates of student i and their peer. Similarly, we also examine heterogeneity by measures of centrality of the strongest peer as well, such as eigenvector centrality and the number of absences that a student shares with their strongest peer.

Lastly, if this research design is valid, we should see that P1Treat_{i1} is uncorrelated with baseline characteristics of students conditional on the Peers_i variable. Table 9 shows the result of estimating 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

[10] As a robustness check, we also incorporated flexible interactions of this variable with the P1Treat_{i1} treatment variable, and our results are extremely similar (available upon request).

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

is randomly assigned.

4.2 Results

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 peer was treated, we estimate three models that build upon each other: (1) a simple regression between P1Treat_1 and attendance controlling for the size of the network;[12] (2) we then add controls for the set of predefined covariates described above; and (3) we add fixed effects by school.

Table 10 shows the results. All three models yield a significant and positive spillover effect of treated students onto their strongest peer. The standard errors are more than 40% smaller when including fixed-effects by school and the magnitude of the effect is smaller relative to the other specifications as well. The fixed-effect specification is both more precise and more conservative, so for parsimony we focus on discussing results from this specification for the remainder of the analyses. The estimated spillover effect is 11 more attended classes. For comparison, this is 20% of the direct treatment effect of the intervention on classes attended found in Bergman and Chan (2017), and 3.6% of the mean for control students, those who were not treated nor had their closest peer treated.

We also analyze how the joint attendance between a student and their strongest peer changes during the post-intervention period. Table 11 shows the results of these analyses. If students derive utility from a joint absence or a joint attendance with a particular peer, we could observe that joint *attendance* with their strongest peer increases if that peer is treated. Students may also reallocate their attendance to be spend more time with their closest peer and less time with other, less-connected peers. Panel A shows the effect of having a student's strongest peer treated on the number of classes attended with that peer. Panel B shows the effect of having a student's strongest peer treated on the number of classes attended with all other students, excluding the strongest peer. The results show the substitution pattern described above: an increase in joint attendance with their strongest peer and a decrease in attendance with other peers. This is consistent with the idea that students derive utility from jointly attending (or missing) class with a particular peer. If spillover effects stemmed entirely from learning about the benefits and costs of attendance from peers as opposed to the experiential value of attending or missing class jointly, we may not see this substitution pattern occur.

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

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 peer measure is significant at the 10% level according to our permutation test.[13] 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.1 shows the results of spillovers for peers that are treated and coordinate their absences according to our permutation test.[14] Despite larger standard errors, we find that the point estimates are slightly larger than before. The effects are statistically significant across all specifications, however because the standard errors are much larger we cannot reject that the coefficients are larger than those found in our previous specifications above.

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 Peers_i value is more than three standard deviations away from the mean. Table A.2 shows the effects are still significant and very similar in magnitude. Removing the outliers also creates more consistency in the magnitude of the effects across specifications, which converge closer to our preferred, more precise estimate that uses school fixed effects.

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.[15] 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 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.

[13] We also constructed the same variable for a 5% threshold.

[14] We used a p-value of 10% for this definition of coordination, but results are consistent when using a lower threshold on 5% as well, but less precise.

[15] 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.

Heterogeneity

We examine how these treatment effects vary by network and demographic characteristics. In general, given the effect size and precision, it is difficult to detect significant heterogeneous effects. We analyze potential heterogeneous effects by gender, race, academic performance, suspension status, baseline absences,[16] centrality in the network.[17], and joint absences[18] Table 12 shows the results are too imprecise or small to detect significant heterogeneity. Qualitatively, the results are larger for Black students, students who were ever suspended, and students whose strongest peer is central in the network.

Lastly, we follow Athey and Imbens (2016) to identify potential heterogeneity in the spillover effects but mitigate 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.3 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 peer, we define second-degree spillovers as spillovers stemming from whether or not the strongest peer to peer 1 is treated or not. We find no significant effects.

4.2.1 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 13 shows the results for these outcomes.

In terms of GPA, we do not find strong evidence of spillover effects. We do observe a small reduction in the number of failed courses, but it is not significant at conventional levels (p-value = 0.12). The coefficient of -0.10 means that students whose closest peer was treated failed 0.08 fewer courses than students whose closest peer was not treated. We also find a marginally-significant, negative effect on dropout rates. In general, the effects on these outcomes are suggestive but too imprecise to detect reasonably-sized effects with much power. However, the directions of these effects are consistent with improved academic behaviors and outcomes as a result of the spillovers.

[16] 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.

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

[18] We define a binary variable to identify students who share above the median absences (3 absences) with their strongest peer.

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 closest peer treated, the cost per additional class attended falls by 19% to $0.17.

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. We solve for an optimal 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 consider a second object 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 closest peer. Empirically, we did not find significant spillover effects beyond a student's closest tie, so we assume that peers that have weaker ties to a student have a negligible spillover effect on that student's attendance.[19] 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

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

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.

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. 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 plus the spillovers of their treated peers onto that same student i.

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 i but not student j $(m_{ij} = z_i(1 - z_j))$. We can re-write the previous m_{ij} variable in terms of the allocation variables as following:

$$2m_{ij} \leq 1 + z_i - z_j \quad \forall\ i, j \in I,\ i \neq j$$

$$m_{ij} \geq z_i - z_j \quad \forall\ i, j \in I,\ i \neq j$$

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.

$$
\begin{aligned}
\max_{\mathbf{z},\mathbf{m}} \quad & \sum_{i\in I} a_{ii} z_i + \sum_{i\in I} \sum_{j\in I; i\neq j} a_{ij} m_{ij} & & \\
\text{s.t.} \quad & 2m_{ij} \leq 1 + z_i - z_j & & \forall i,j \in I; i \neq j \\
& m_{ij} \geq z_i - z_j & & \forall i,j \in I; i \neq j \\
& \sum_{i\in I} z_i \leq b & & \\
& z_i \in \{0,1\} & & \forall i \in I \\
& m_{ij} \in \{0,1\} & & \forall i,j \in I; i \neq j
\end{aligned}
$$

However, this individual-allocation problem is computationally difficult to solve due to the large number of students in the sample. To address this issue, we simplify the problem by defining types of students according to their relevant characteristics. This allows us to capture the heterogeneity of spillovers and treatment effects, and, at the same time, solve a large integer-programming problem in a reasonable amount of time.

To illustrate how we implement this simplification, Figure 4 shows a brief example of a reduced network with 5 students and 3 'types' of students. The directions of the arrows point to each student's closest peer. Panel 4a shows the original networks that we have, which is the basis of the previous optimization problem. To simplify the network, we group students by type, which reduces the number of nodes and edges as shown in Panels 4b and 4c. Our reduced network helps us solve the allocation problem more easily by focusing on the relevant characteristics for effects and spillover heterogeneity.

Thus, we can re-write the previous optimization problem as the following:

$$
\begin{aligned}
\max_{\mathbf{x},\mathbf{y}} \quad & \sum_{t\in T} d_t n_t x_t + \sum_{t\in T^p} \sum_{k\in T^p} s_t n_{tk} y_{tk} & & \\
\text{s.t.} \quad & & & \\
& \sum_{t\in T} n_t x_t \leq b & & \\
& \sum_{k\in T^p} n_{tk} y_{tk} = n_t x_t & & \forall t \in T^p \\
& x_t \in [0,1] & & \forall t \in T \\
& y_{tk} \in [0,1] & & \forall t \in T^p, k \in T^p
\end{aligned}
$$

where x_t defines the proportion of each student type t that is treated, n_t is the number of students of type t, n_{tk} is the number of connections from type t students to type k students (with $t,k \in T^p$, where $T^p \subset T$ represents the types of students who have a closest peer), and y_{tk} is the proportion of treated students of type t who have a closest peer of type k. 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.

To illustrate this algorithm, we consider the following groups of students: students who do not have a closest peer (i.e. have not missed class with another student), referred to as NP, and the subgroups identified by the machine-learning analysis described above: students with baseline absence rates less than 4% (Group 1) and students with baseline absence rates greater than 4% (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, $T = \{NP, P1, P2\}$. In this example, n_t represents the number of students of type t; the subset $T^p \subset T$ is defined as the types of students who have a close peer (types $P1$ and $P2$ in this example). Each of these types of students connect to other types of students, which reduces the complexity of the network by clustering individuals according to type. The types of students can be created 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.

Figure 5a shows the proportional allocation of treatment by types of students (y axis) conditional on the budget restriction b (x axis). Note that the denominator for these proportions is not the total number of students, but the number of students of a given type. If the restriction b (i.e. number of students we can treat) is less than $n_{P1,P1}$, then we only treat a portion of type $P1$ students connected to type $P1$ students. After we treat all of the students of this type, if the budget allows it, we treat type NP students.

Figure 5b 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 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 an additional 49%. If there is substantial heterogeneity in the effect of the treatment, targeting the treatment to those who would benefit the most (subject to our overall objective) can substantially improve the results of an 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.

Minimizing the Number of Chronically Absent Students

As described above, many state education-policymakers would like to use the rate of chronic absenteeism (often measured as the share of students missing 10% or more of their classes) in a school as a measure of accountability. This creates an incentive for schools and districts to reduce the number of chronically absent students. To that point, we re-arrange the previous optimization problem to reduce the number of students who miss more than 10% of classes. Following the same logic as before, we create different types of students according to their chronic-absentee status, but also, within those students, we identify likely to no longer be chronically absent if directly treated and those who would change their chronic-absence status by treating their closest peer, and not the student directly.[20] As an example, consider student i who would need to increase their attendance by q days to no longer be considered chronically absent. Then, if the direct effect is $d \geq q$, then student i would change status if treated. In the same fashion, if the spillover s is so that $s \geq q$, student i would no longer be considered chronically absent if their closest peer is treated. Thus, we identify seven types of students analogous to the previous types, $T = \{NP, P1, Ch-NP, P3, P4, P5, P6\}$, where NP ($Ch-NP$) represents students with no closest peers and not chronically absent (are chronically absent), and $P1$ represents students that are not chronically absent and have a closest peer. $P3$ comprises a small group of students that, even if they were directly treated as well as their closest peer, would most likely remain chronically absent. Groups $P4$, $P5$, and $P6$ represent chronically absent students that could change status if they and their closest peer were treated ($d+s>q$), if they were directly treated ($d>q$), or if they only were affected by a spillover effect ($s>q$), respectively.

Given that direct effects are larger than spillovers, the chronically absent students belonging to type $P4$ are farther from the 10% threshold than those who are in category $P6$.

For the previous case, our optimization problem is:

$$\min_{\mathbf{w},\mathbf{z}} \quad \sum_{t=T^{Ch-NP}} n_t - n_t w_t + \sum_{t=T^C} n_t - \sum_{k\in T^p} z_{tk} n_{tk}$$

s.t.

$$\begin{aligned}
\sum_{t\in T} n_t w_t &\leq b \\
\sum_{k\in T^p} n_{tk} z_{tk} &= n_t w_t && \forall t \in T^p \\
w_t &\in [0,1] && \forall t \in T \\
z_{tk} &\in [0,1] && \forall t \in T^p, k \in T^p
\end{aligned}$$

Where w_t represents the proportion of type t of students that are treated, and z_{tk} represents the treated

[20] For simplicity, we assume a constant spillover effect on students, but the extension to include heterogeneity in the effects can be easily implemented by expanding the "types" of students used in the analysis.

proportion of type t students who have a type k closest peer. For simplicity, and given that we are only focused on minimizing the number of chronically absent students, we have defined appropriate subsets of T, such as $T^{Ch-NP} = \{Ch - NP\}$, which encompasses chronically absent students with no closest peer, but that could change their status if treated. The set $T^C = \{P3, P4, P5, P6\}$ represents chronically absent students who have a closest peer, and, again T^p represents all students who have skipped class with someone.

Figure 6a shows the targeted allocation of the treatment in order to minimize the number of chronically absent students. In this case, we observe that chronically absent students are treated first with priority given to those who are farther from the threshold. Figure 6b shows the reduction of this type of students according to the number of students treated. The red line represents the number of students originally treated in the study, which shows a reduction of 604 chronically absent students from the original 3182 students *if the treatment had been targeted*. The latter is a 19% reduction in the number of chronically absent students. In the experiment with random allocation, the potential reduction was 164 students, which is a 5.2% reduction.

If school districts aim to reduce the number of chronically absent students, targeting the treatment to those students with lower attendance rates as well as their closest peers could generate fewer chronically absent students relative to random allocation, under budget constraints.

Overall, our results shows there are potentially substantial gains in the reduction of absenteeism if the intervention is targeted as described above. However, our results do not take into account the possible reformation of networks in response to a large-scale allocation of the intervention.

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.

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.

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 âĂIJacting whiteâĂİ," *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, 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.

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.

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.

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).

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.

Rogers, Todd and Avi Feller, "Reducing student absences at scale," *Unpublished paper*, 2016.

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

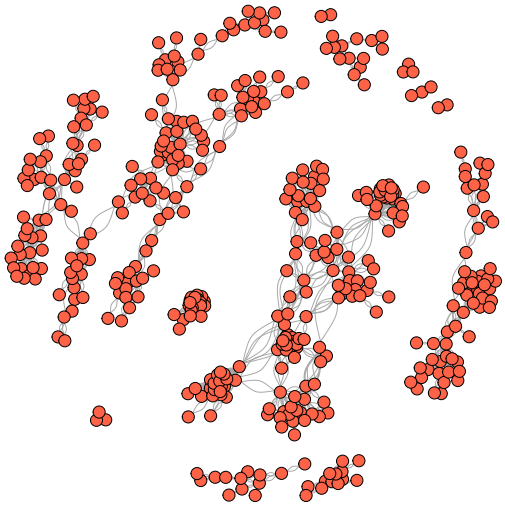


Figure 1: Social Network for School 1 (pre-intervention period)

Notes: Each circle corresponds to a student, and each interconnecting line or *edge* corresponds to the number of absences between two students.

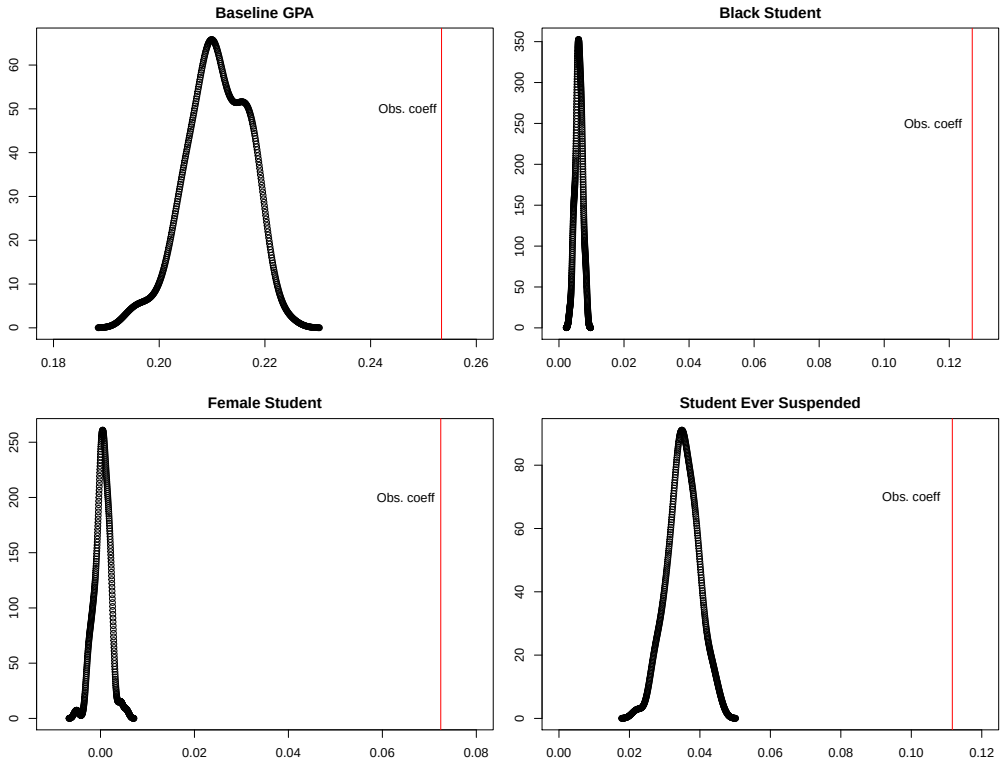


Figure 2: Distribution of simulated coefficients (100 simulations) and observed coefficient for homophily analysis

Notes: Distribution obtained from the coefficients of regressions using only strata fixed effects and clustered errors.

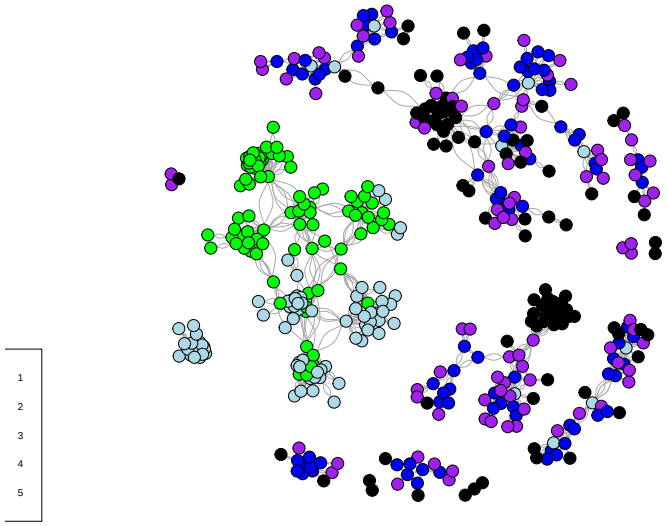
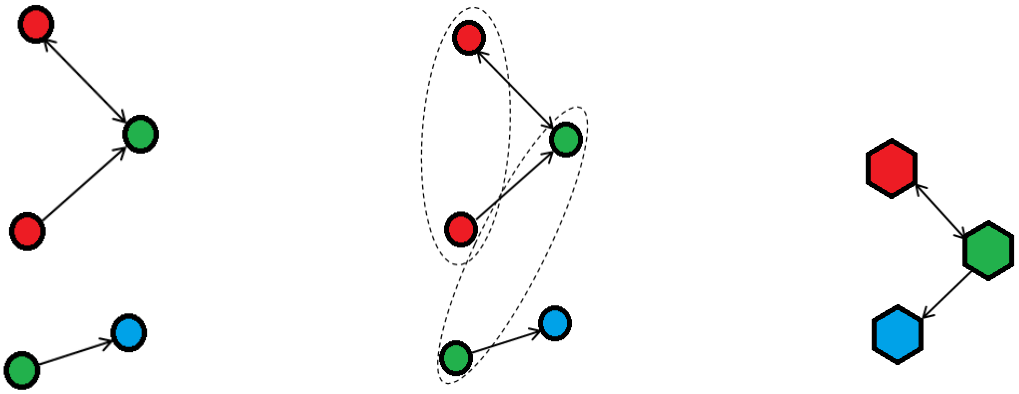
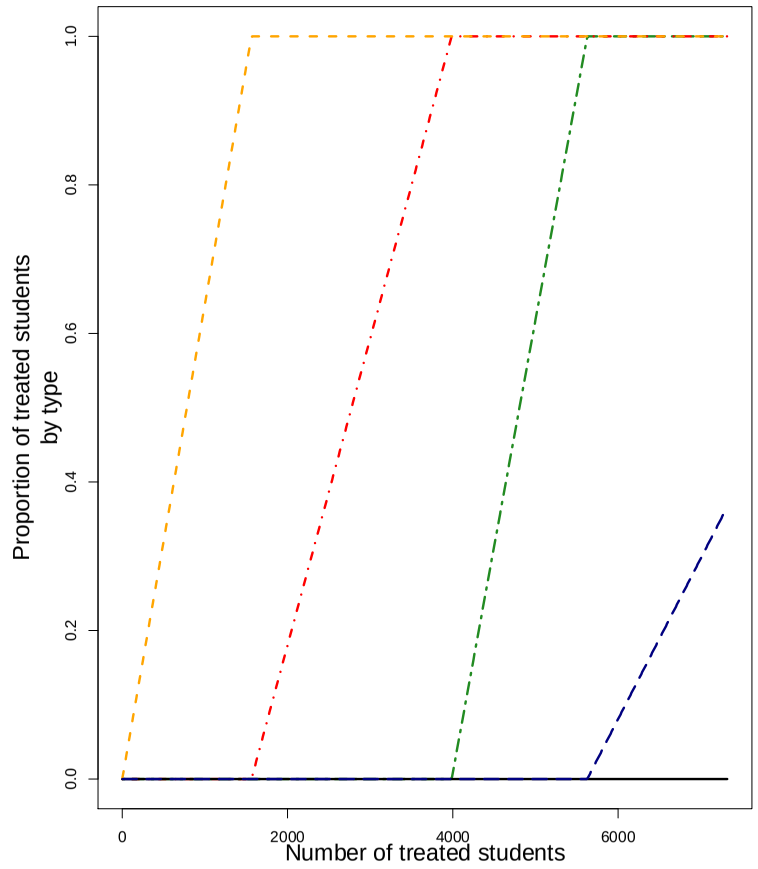


Figure 3: Social Network for School 1 (pre-intervention period) with quintiles of eigenvector degrees

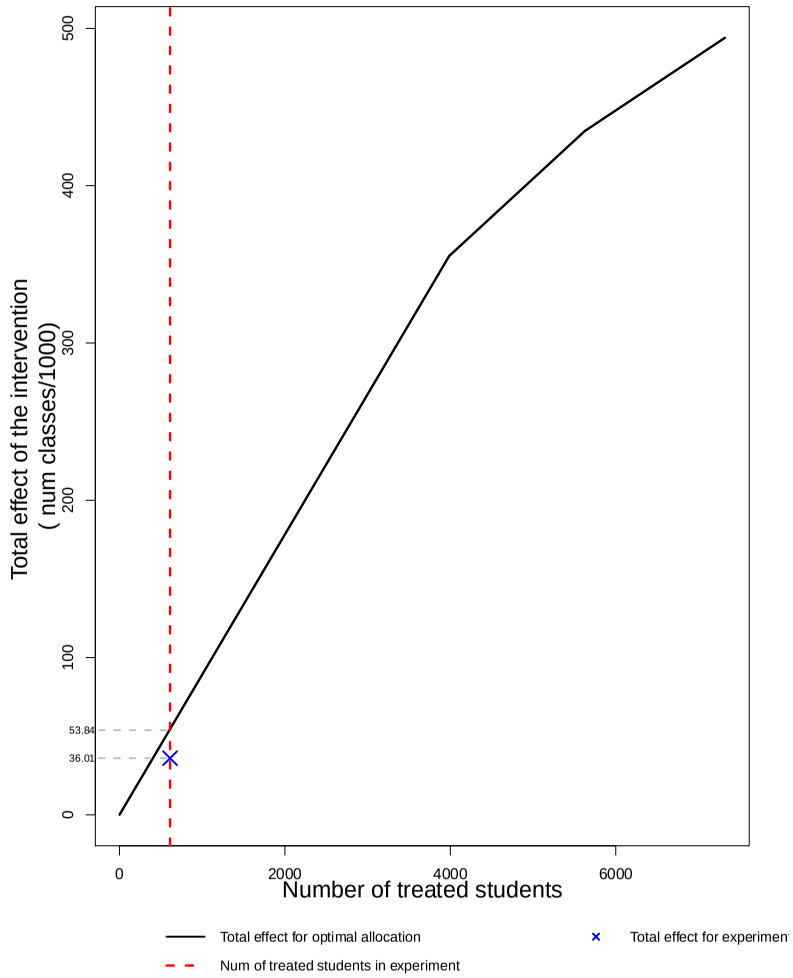


(a) 5 students and 3 types (b) Clustering by type of student (c) Reduced social network (by types)

Figure 4: Example of Network by Type



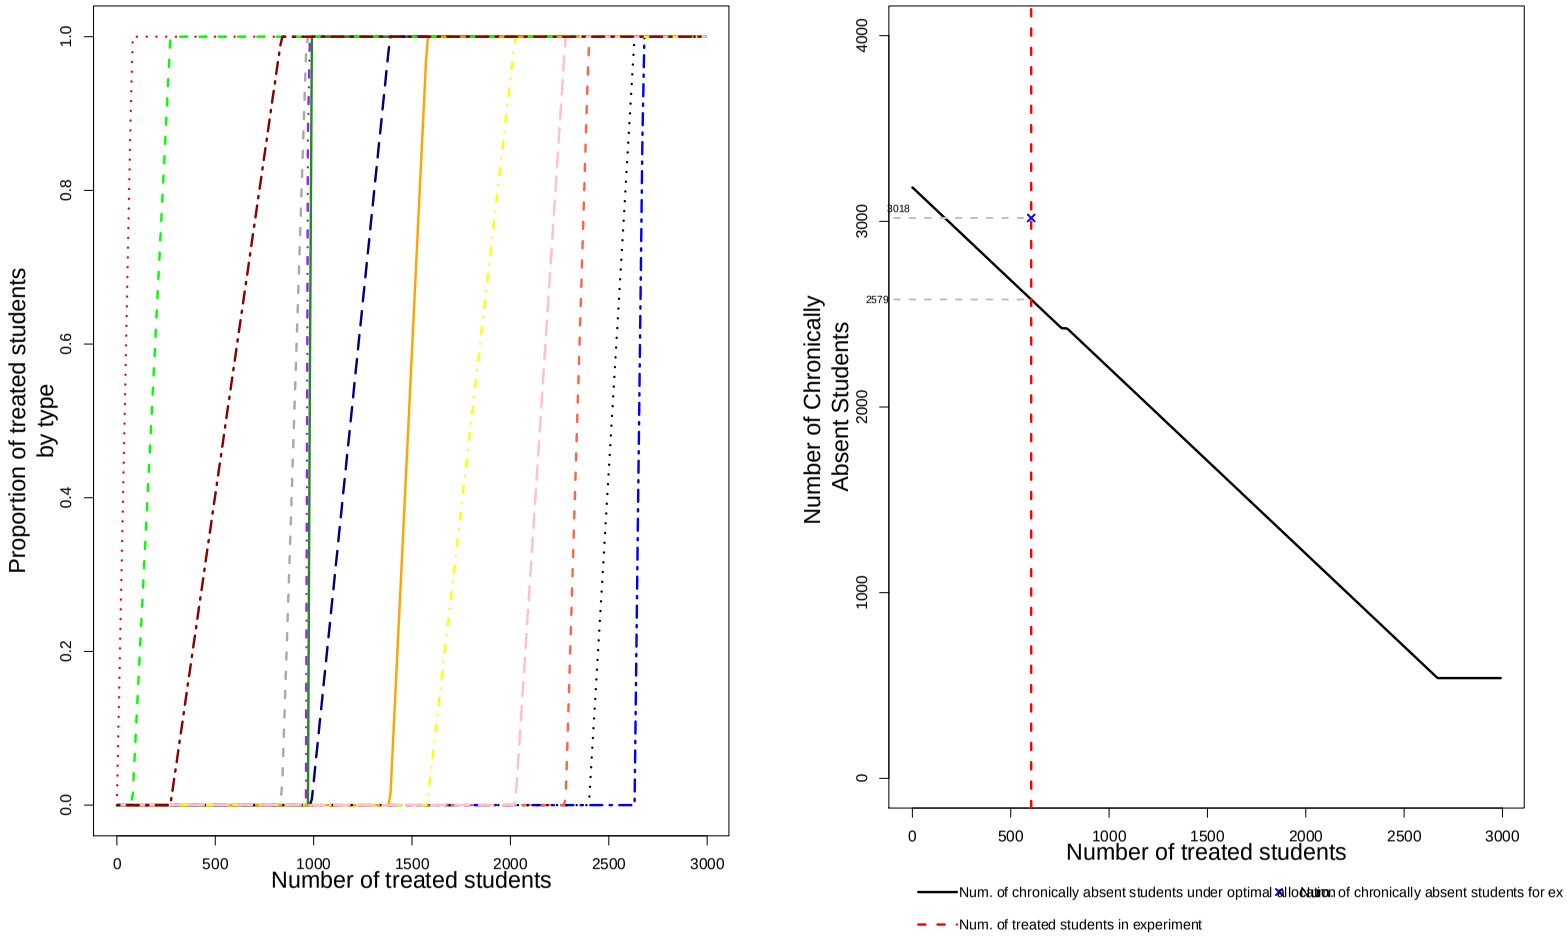
(a) Optimal allocation by number of students treated for maximization of treatment effect



(b) Overall intervention effect under optimal allocation by number of student treated

Figure 5: Optimal Allocation Results for Maximization of Treatment Effect

Notes: NP is the group of students who do not have a closest peer; Pij refers to students of type i ($i = \{1, 2\}$, where 1: Students with low baseline absenteeism rates (< 0.044), 2: Students with higher baseline absenteeism rates (≥ 0.044)) whose closest peer is type j ($j = \{1, 2\}$). We assign at most $N/2$ students to treatment to reduce the potential for non-linear, general-equilibrium effects. Some types t of students are not shown in the graph given that they are only treated when $b > N/2$.



(a) Optimal allocation by number of students treated for minimizing chronically absent students

(b) Number of chronic absenteeism under optimal allocation by number of students treated

Figure 6: Optimal Allocation Results for Minimization of Chronic Absenteeism

Notes: $Ch-NP$ refers to chronically absent (CA) students who do not have a closest peer. Group Pij refers to students of type i whose closest peer is of type j, where the types are defined as following: 1, not chronically absent students; 3, CA students whose days to no longer be CA q are greater than the sum of the direct effect d and spillover s; 4, CA students who would not be CA if treated and closest peer was treated too ($d < q \leq d+s$); 5, CA students that would change status if treated ($s < q \leq d$); and 6, CA students that would change status if closest peer was treated ($0 < q \leq s$).We assign at most $N/2$ students to treatment to reduce the potential for non-linear, general-equilibrium effects. Some types t of students are not shown in the graph given that they are only treated when $b > N/2$.

8 Tables

Table 1: 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 figure shows the script for each of the three types of alerts sent via text messages: low class average, absence, and missing assignments.

Table 2: Treated students and sample size

Variable	Percentage
Treated students	0.04
Students with closest peer	0.87
Students with closest peer treated	0.03
Observations	14,653

Notes: Closest peer is constructed using the by-class attendance data.

Table 3: Baseline characteristics for the sample

Variable	Mean	Observations
Female	0.50	13,641
Black	0.13	13,641
English Language Learner	0.02	12,955
IEP	0.15	12,955
Baseline GPA	2.54	14,653
Ever suspended last year	0.19	14,586
Percent of classes absent pre-intervention	0.14	14,621
Percent of days missed (2014-2015)	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.

Table 4: Observed network features

Characteristic	Observed
Average size	664.6
Weighted clustering coefficient	
Arithmetic mean	0.41
Geometric mean	0.41
Centrality	
Average number of edges	36.40
Eigenvector centrality	0.12

Notes: These statistics are constructed using by class attendance data.

Table 5: Homophilly in the observed network

Baseline Variable	GPA	Female	Black	Suspended
GPA peer	0.25***			
	(0.02)			
Female peer		0.07***		
		(0.01)		
Black peer			0.13***	
			(0.02)	
Suspended				0.11***
				(0.01)
Observations	11,666	10,483	11,873	11,669

Notes: Standard errors in parenthesis. Specification of the models include fixed effects by strata. Peer represents the person with whom a student missed the most class and is constructed using by class attendance data. GPA is from the gradebook data. Demographic information and suspensions are from school administrative data.
* Significant at 10%; ** significant at 5%; *** significant at 1%

Table 6: Simulated (random) network features

Characteristic	Random
Average size	664.6
Weighted clustering coefficient	
Arithmetic mean	0.40
Geometric mean	0.41
Centrality	
Average number of edges	32.41
Eigenvector centrality	0.12

Notes: These statistics are constructed using simulated network data, with 100 simulations per schools.

Table 7: Homophilly in the simulated network

Baseline Variable	GPA	Female	Black	Suspended
GPA peer	0.211***			
	(0.005)			
Female peer		0.000		
		(0.001)		
Black peer			0.006***	
			(0.001)	
Suspended				0.035***
				(0.003)
Observations	12,186	12,186	12,186	12,186

Notes: Standard errors in parenthesis. Specification of the models include fixed effects by strata. Regressions are estimated using the simulated network data, with 100 simulations per schools.
* Significant at 10%; ** significant at 5%; *** significant at 1%

Table 8: Students with coordinated absences for different thresholds

Variable	N Total	Number of Peer1 Treated
Students in sample (Total)	14,653	433
Students with a closest peer	12,740	433
Students with coordinated absences:		
90% threshold	6,336	263
95% threshold	4,675	187
99% threshold	2,131	79

Notes: Results obtained from permutation test using 100 simulations of random networks. Data are from the by-class attendance data.

Table 9: Balance between students with a treated peer and without a treated peer

Variable	Control Mean	Treatment-Control Difference	P-value	Observations
Female	0.43	0.03	0.316	13,641
Black	0.16	0.00	0.856	13,641
English Language Learner	0.01	-0.00	0.579	12,955
IEP	0.23	-0.03	0.052*	12,955
Baseline GPA	2.18	0.04	0.359	13,641
Ever suspended last year	0.15	0.03	0.281	13,641
Percent of classes absent	0.13	-0.01	0.534	13,609
Percent of days missed	0.05	-0.00	0.911	13,641

Notes: Specification of the models include fixed effects by strata and clustered standard errors. 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 attendance

Variable	Classes present	Classes present	Classes present
P1Treat	17.44**	15.85*	11.01**
	(8.32)	(8.10)	(4.97)
Controls	Strata	All	All
Fixed Effects	No	No	Yes
Observations	12776	12774	12774

Notes: Cluster-robust standard errors shown in parenthesis. The model from column (1) includes includes fixed effects for strata, an indicator for *sample* and *treatment*, and no additional controls. Column (2) includes fixed effects by strata and controls by *GPA baseline, IEP, ELL, missed days (previous year), gender, black,* and *ever suspended*, in addition to a *sample* and *treatment* indicator. Column (3) includes all the previous controls plus fixed effects by school. 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: Effect of treated peer on joint attendance (classes) post-intervention

	Classes present with P1		
Variable	(1)	(2)	(3)
P1Treat	25.52***	25.30***	26.69***
	(6.53)	(6.40)	(6.29)
Controls	Strata	All	All
Fixed Effects	No	No	Yes
Observations	11329	11325	11325

	Classes present *not* with P1		
Variable	(1)	(2)	(3)
P1Treat	-16.90	-18.73*	-19.72***
	(10.21)	(10.07)	(7.36)
Controls	Strata	All	All
Fixed Effects	No	No	Yes
Observations	11329	11325	11325

Notes: Top panel shows effect of having Peer 1 treated on joint attendance; bottom panel shows effect of having Peer 1 treated on attendance with other peers that are *not* Peer 1. Cluster-robust standard errors shown in parenthesis. The model from column (1) includes includes fixed effects for strata, an indicator variable for *sample* and *treatment*, and no added controls. Column (2) includes fixed effects by strata and controls by *GPA baseline, IEP, ELL, missed days (previous year), gender, black,* and *ever suspended*, in addition to a *sample* and *treatment* indicator. Column (3) includes all the previous controls plus fixed effects by school. 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 12: Heterogeneity of effect of treated peer on attendance

Variable	Classes present (1)	Classes present (2)	Classes present (3)	Classes present (4)	Classes present (5)	Classes present (6)	Classes present (7)
Peer Treat	11.61	8.46	21.65**	6.45	15.44***	5.83	27.45***
	(8.01)	(5.46)	(9.09)	(6.50)	(5.26)	(7.39)	(9.28)
Peer Treat × female	-1.29						
	(11.27)						
Peer Treat × black		20.94					
		(18.62)					
Peer Treat × below med GPA			-14.34				
			(13.43)				
Peer Treat × suspended				23.14			
				(15.42)			
Peer Treat × below med absences					-6.47		
					(11.54)		
Peer Treat × central peer						11.62	
						(9.46)	
P1 Treat × > 3 abs together							-13.02
							(16.99)
Controls	All	All	All	All	All	All	All
Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	12774	12774	12774	12774	12774	12774	12774

Notes: Cluster-robust standard errors shown in parenthesis. All models include fixed effects by strata and controls by *GPA baseline, IEP, ELL, missed days (previous year), gender, black,* and *ever suspended*, as well as fixed effects by schools and a *sample* and *treatment* indicator. 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 13: Effect of treated peer on other outcomes

	GPA		
P1Treat	-0.02	-0.02	-0.02
	(0.04)	(0.04)	(0.03)
Controls	Strata	All	All
Fixed Effects	No	No	Yes
Observations	12146	12146	12146
	Courses failed		
P1Treat	-0.08	-0.09	-0.08
	(0.05)	(0.06)	(0.05)
Controls	Strata	All	All
Fixed Effects	No	No	Yes
Observations	12146	12146	12146
	Drop out		
P1Treat	-0.02	-0.02	-0.02*
	(0.01)	(0.01)	(0.01)
Controls	Strata	All	All
Fixed Effects	No	No	Yes
Observations	13609	13598	13598

Notes: Clustered standard errors shown in parenthesis (). The model from column (1) includes includes fixed effects for strata, an indicator for *sample* and *treatment*, and no added controls. Column (2) includes fixed effects by strata and controls by *GPA baseline, IEP, ELL, missed days (previous year), gender, black,* and *ever suspended*, in addition to a *sample* and *treatment* indicator. Column (3) includes all the previous controls plus fixed effects by school. 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 peer for peers that coordinate absences according to permutation test

Variable	Days present	Days present	Days present
P1Treat (CA)	23.28*	19.87	17.01**
	(12.65)	(12.86)	(7.97)
Controls	Strata	All	All
Fixed Effects	No	No	Yes
Observations	12808	12806	12806

Notes: Cluster-robust standard errors shown in parenthesis. The model from column (1) includes includes fixed effects for strata, an indicator for *sample* and *treatment*, and no additional controls. Column (2) includes fixed effects by strata and controls by *GPA baseline, IEP, ELL, missed days (previous year), gender, black,* and *ever suspended*, in addition to a *sample* and *treatment* indicator. Column (3) includes all the previous controls plus fixed effects by school. $PTreat(CA)$ is the interaction between the closest peer being treated and the students coordinating their absences according to our permutation test, using a 90th-percentile threshold with respect to simulations (described in the text) to define coordination. 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.2: Effect of treated peer on attendance removing outliers (>3 SD)

Variable	Days present	Days present	Days present
P1Treat	15.25*	14.02*	10.52**
	(8.14)	(7.89)	(5.15)
Controls	Strata	All	All
Fixed Effects	No	No	Yes
Observations	12458	12456	12456

Notes: Cluster-robust standard errors shown in parenthesis. Outliers with joint absences larger than 3 standard deviations from the mean are removed from the sample. The model from column (1) includes includes fixed effects for strata, an indicator for *sample* and *treatment*, and no added controls. Column (2) includes fixed effects by strata and controls by *GPA baseline, IEP, ELL, missed days (previous year), gender, black,* and *ever suspended*, in addition to a *sample* and *treatment* indicator. Column (3) includes all the previous controls plus fixed effects by school. 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.3: Direct and spillover effects for optimal allocation problem

Variable	Effect
Treatment	48.46**
P1Treat × Fraction of absences < 0.04	40.67*
P1Treat × Fraction of absences ≥ 0.04 × GPA baseline < 2.90	0.05
P1Treat × Fraction of absences ≥ 0.04 × GPA baseline ≥ 2.90	-34.74
Controls	All

Notes: Treatment coefficient is obtained from the original experiment (Bergman and Chan 2017). 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%