Peer Effects in Active Learning

Vinicius Lima[*] Vladimir Ponczek Priscilla Tavares

August 2, 2022

Abstract

This paper studies peer effects in higher education by observing students in an environment of *active learning*, where peer interaction is a meaningful mechanism explaining performance. The identification of peer effects relies on the random assignment of students to groups and explores variation both in terms of the share of low and high-ability students in each group as well as the frequency that peers meet for group work. The main result of the paper is that the share of low and high-ability peers in a group is not able to explain by itself the existence of positive peer effects on performance, which arises only when we consider the fraction of a student's high-ability peers in the group that are also his peers in a second group. We show that when pairs of students meet in more than one group, they are more likely to establish a link, measured by students' self-reported naming of their relevant peers. Thus, increased interaction with high-ability peers impacts positively performance. However, we also show that this effect depends on whether the student himself is classified as low or high-ability.

[*]FGV/EESP, vinicius.lima@fgv.br

1 Introduction

Peer effects matter for performance in education.(Sacerdote, 2011; Epple and Romano, 2011) To some extent, policies such as the provision of educational vouchers or ability tracking rely on the possibility that a student's performance could improve in a different group. From this perspective, understanding how peer effects operate is crucial to thinking about alternative grouping policies.

In this paper, we study peer effects in a higher education institution adopting an active learning methodology, which is a pedagogical strategy that places peer interaction at the core of the learning process. In our setting, students meet in small groups for continuously evaluated classwork, providing further incentives for students to interact. To allocate students in these groups, we developed an assignment rule that ensured random variation in two dimensions: the number of peers by ability levels in each group and the frequency that specific peers meet for group work. While group composition in terms of how many low and high-ability students could affect performance, there are different channels through which this could happen. The measure on the numbers of meetings between pairs of students allows us to distinguish what drives peer effects in our setting: composition, frequency of meetings or a combination of the two factors.

We begin the analysis by showing that for peers sharing at least one group, a student is 87% more likely to report the desire of having a specific peer again in some group if they already interact in more than one group. Then, we start analyzing peer effects on performance. We find no average effect due to a change in group composition by simply replacing a low-ability student with a high-ability one. However, the absence of effect from such a change seems to result from opposing effects that becomes clear when we consider the dynamics of group interaction by using our measure of the frequency of interaction among peers. If we keep constant the fraction of frequent peers, the point estimate for the effect of replacing a low-ability student with a high-ability one is a 2.9% decrease of a standard deviation, although not statistically significant at usual levels. However, if instead we keep group composition constant and maximize the fraction of frequent peers (relative to the average fraction across students), performance would increase by 2.2% of a standard deviation.

We then show that for low-ability students, a change in group composition that replaces a low-ability student with a high-ability one and also decreases their potential interaction implies a 5% standard deviation decrease in performance while the same change in group composition that, on the contrary, increase their potential interaction does not affect performance. An analogous situation happens for high-ability students, for whom a replacement of a low-ability student with a high-ability one that decreases their potential interaction has no effect on performance, but the same change made with an increase in potential interaction seems to increase performance by 5% of a standard deviation. Finally, we construct a measure of students' connectedness within each group using the report on the desire of having specific peers again in subsequent groups. We show that improvement in the student's network position within the group could explain the positive peer effects for high-ability students.

Related literature The largely explored linear-in-means model of peer effects (Sacerdote, 2001; Zimmerman, 2003; Hanushek et al., 2003) leaves no room for reallocation policies aimed at improving efficiency: any positive effect arising from a reallocation of students that raises a group's average ability would be offset by a corresponding decrease in some other group's average ability (Hoxby, 2000). Efficiency improvement through a reallocation scheme depends on the existence of different peer effects for different subgroups of students. Carrell et al. (2009) show that the performance of low-ability students in college increases when they have more high-ability peers and there are no negative effects the other way round. Given the existence of such nonlinear peer effects, there would be scope for a policy aimed at maximizing the interaction between those types of students.

However, experimental evidence on such a policy shows a detrimental effect on low-ability students' performance (Carrell et al., 2013). Contrary to what the previous positive reduced-form estimates suggested, low and high-ability students did not interact as expected due to endogenous sorting of peers within the new groups. An important takeaway is that identification of the effects arising from peers' equilibrium behavior, a key parameter for optimal group design, is empirically challenging and reduced-form estimates could be misleading. (Manski, 1993; Goldsmith-Pinkham and Imbens, 2013; Lee, 2007; Bramoullé et al., 2009)

Our paper uses a credible source of exogenous variation in peer interaction within an environment where peer effects should be an important mechanism to explain performance. It contributes to an empirical literature discussing peer effects as subsidy for the design of optimal groups (Carrell et al., 2013; Booij et al., 2017; Garlick, 2018) with evidence that changing the mix of low and high-ability peers a student has in a group can only produce positive peer effects if it induces more interaction among peers. Actually, there could be negative peer effects on the performance of students at the lower end of ability distribution placed in groups with low potential for peer interaction.

To the best of our knowledge, the closest paper in terms of the analysis we do is Brady et al. (2017). In the context of the U.S. Navy Academy, all students are observed in two kinds of groups with different sizes and purposes. They identify negative peer effects from the variation in peer ability across groups at a broader level. Then, by looking at smaller and more task-oriented groups they find that positive peer effects exist due to peers in these groups that are also peers in the broader group. This is a result qualitatively similar to ours as it highlights that stronger interaction seems to matter in explaining peer effects.

The next section presents the organizational framework and the assignment mechanism of students in our context. Then we present the data and discuss the implementation of our empirical strategy. Then, we present and discuss our main results and conclude by making some final remarks.

2 Background

2.1 Organizational Framework

São Paulo School of Economics (FGV EESP) is a distinguished higher education institution in Brazil. From 2003 to 2016, the school selected up to 60 students for the undergraduate course in economics by applying a highly selective admission exam done by around 1,500 candidates. From 2017 to 2021 the number of students admitted each year gradually increased and now the school accepts up to 120 students. The annual course fee for 2019 was approximately R$ 65,000, which was roughly 3.7 times the Brazilian per capita income calculated for that year. Thus, FGV EESP's students are at the top of both the country's ability and income distribution.

In 2013, the school began to replace the traditional teaching method based on lectures with a model of *Active Learning*. Students admitted from that year on have been doing all of their coursework under the *Problem-based Learning* method. Problem-based Learning (PBL) is a pedagogical strategy in which learning evolves in a problem-solving framework. A central element of this method is the goal of engaging students in their own learning process by working with concrete problems in small groups. In FGV EESP's, each student in each discipline is allocated to a group of 12 students, on average, that will work together in sessions of two hours. In different sessions of the same discipline, group members are the same. However, across different disciplines, a given student has different pools of classmates in different groups. The total number of sessions varies by discipline.

All of these sessions occur under the supervision of a tutor and each session has two parts. In the *pre-discussion* part, a problem is presented to the group without any explicit requirement for the work on the topic. The only available tools for students at the time are what they have learned up to that point. In this phase, a student is chosen to be the leader of the session. She organizes the discussion so that the group manages to bring about all the problem's latent questions. Besides, another student organizes the contributions of each group member into a document made available to the group at the end of the session. Throughout this part, tutors must ensure that all learning goals are clear to the group. Ideally, this should be done with minimal interventions from the tutor.

Pre-discussion happens at the end of each session. At the beginning of the session, there is the *post-discussion* part, which concludes the work started in the previous pre-discussion. This way students have an interval of one day or more to make use of the pre-discussion information together with bibliographic references to get prepared for the post-discussion. During the post-discussion, students should answer all issues raised from the problem by using concepts and tools they have learned from self or group study outside class. Group study in this case is not restricted to happening among students of the same tutorial group.

In the post-discussion, besides organizing the discussion, the tutor must guide students, so they develop a comprehensive understanding of the topic under debate. When there is available time, he can also give additional information on the topic. Once students conclude the task, the tutor must evaluate individual performance and give appropriate feedback so that students know exactly how they performed. That is, in each session each student receives a grade between zero and one. At the end of the period, the mean of these tutorial grades enters the discipline's final grade as a multiplicative factor over some combination of marks from other examinations. This way, the better the performance in group work, the better the overall performance. This provides a strong incentive for students to show up (absence implies zero) and to really participate during tutorial sessions.

2.2 Students Assignment Mechanism

We implemented the allocation of new students from the 2018 and 2019 cohorts in their first semester at the school. During this semester, they took six mandatory courses, and in every discipline each student was assigned to a tutorial group with 12 students on average. The algorithm to allocate students into tutorial groups ensured random variation across groups in two dimensions. First, it created large variation across groups considering the proportions of low and high-ability students in each group. Besides, since students take six disciplines it was possible that some of them share more than one group. Thus, some pairs of students met weekly in only one group while others, by chance, ended up meeting in more than one group. This implies that there is random variation in the frequency of meetings across pairs of students.

In principle, a simple lottery placing students across groups in a given discipline would also create variation in both dimensions. But to ensure that there would be a large suport of group composition we implemented the allocation in two steps.[1] Suppose we have a pool of students to allocate in given discipline. After classifying students as low and high-ability according to the ranking in the admission exam we performed the following two steps:

- Step 1) The algorithm randomly chooses how many students of each type will be allocated in each of the groups. Example: A group will have 6 low and 6 high-ability students, another group will have 2 low and 10 high-ability, and so on. We run this lottery without replacement (conditional on type) to minimize the chance that two groups have the same composition and then we increase variation across groups. As this step must respect constraints on group size and total students by type, it actually defines three groups plus a residual group.

- Step 2) Given the composition of groups defined previously, it randomize which students will be in each group. That is, if there is a total of 20 low-ability students and the previous step defined that a group must have 6 of them, this step defines the identity of these six students. Then, among the pool of the 14 remaining low-ability students, the algorithm draw a subset of them for the next group and so on.

Random variation in group composition The algorithm produces what we call throughout the paper *random variation in group composition*: Step 2 ensures that for a given student,

[1] The most relevant aspects of the assignment mechanism will be shown with a simple example. A detailed explanation of the allocation process is left to the appendix.

being in a group with many or few high-ability students is independent of his unobservables. The actual outcome of the allocation process is illustrated in figure 1. Our sample contains 67 groups (each point in the plot) and there is a large support of composition as defined by pairs of containing proportions of low and high-ability students in the group. Summing the proportions of low and high-ability students not always add to 100% since there are few students with no ability classifiction and students redoing the discipline (see descriptive statistics). These types of students were randomlly allocated after the definition of group composition (step 1) and the identity of students in each group (step 2).

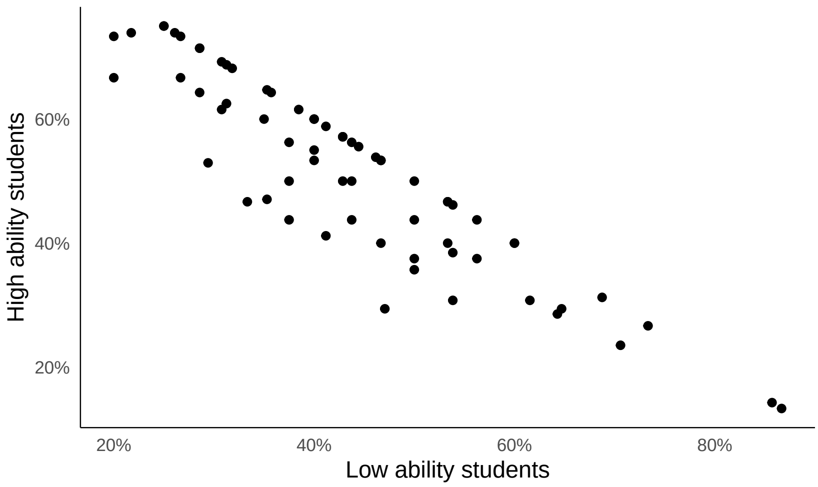

Figure 1: Random variation in terms of group composition. Each point represents a group and the position in the plot is defined by the proportion of each student type in the group as indicated by the axis.

Random variation in the frequency of meetings Consider three students i, j, k randomly selected to be in the same group in a given discipline. Since the above two-step procedure is repeated for every discipline, it is possible that in a different discipline the pair ij is again (by chance) allocated to the same group while k ended up in another group. Thus ij meets twice while ik jk meet only once. This is what we call *random variation in the frequency of meetings.* Table 1 shows that there it was possible to combine new students to form up to 9,802 pairs. However, in the actual allocation, 55% of those theoretical pairs never happen to meet, 29.6% meet in only one group and 15.4% meet in more than one group. Thus, conditional on sharing some group, one third of the pairs meet at least once again.

Randomization check By performing student allocation through the implementation of the above algorithm, it was expected that student's average ability in a group was not correlated with the number of low and high-ability students in that group. To check this was so, we regress student's (standardized) own ability on group composition, conditional on randomization controls (year, discipline and whether student is either low or high-ability). In each column of table 2, group composition is the number of low and high-ability students classified by different measures of incoming ability: final score (which determines the ranking of students), mathematics

Table 1: Number of meetings by pair of students

Number of meetings	Number of pairs	Frequency
0	5,390	55.0%
1	2,897	29.6%
2	1,099	11.2%
≥3	416	4.2%
Total	9,802	

Notes: From all possible pair of students, 55% never meet, 29.6% meet in one group and 15.4% meet in more than one group.

and writing skills.

Table 2: Regression of Own Ability on Group Composition

	Dependent variable: Own ability		
High Peers	-0.006	-0.003	-0.005
	(0.008)	(0.008)	(0.010)
Low Peers	0.001	-0.004	-0.002
	(0.009)	(0.011)	(0.010)
Retained Peers	0.078*	0.073+	0.074+
	(0.039)	(0.039)	(0.038)
Other Peers	0.002	0.002	0.002
	(0.016)	(0.016)	(0.016)
Peers classified by	Final score	Math	Writing
Students	178	178	178
Observations	994	994	994
Groups	67	67	67

Notes: *Own ability* refers to the score in the admission exam ability measures standardized among new students by year. Peers variables count the number of students of each type in each group considering their classification according to the ability indicated in the bottom of the table. Standard errors clustered by student in parentheses and p-values in brackets. $^{+}$ $p < 0.1$, * $p < 0.05$

Point estimates for on Low and High Peers are virtually zero, as well as for the coefficients on the number of peers not classified by ability (*Other peers*). The only significant results appear for the coefficients on the number of peers doing the course for a second time (*Retained peers*). However, we do not believe this indicates any problem with the allocation. There are only eight groups out of 67 in which there is one such student plus two groups in which there are two such students. The average standardized final score in the latter two groups is 0.26 while it is zero for the remaining groups. Thus, it is likely that the estimate is really due to chance.

As a final remark, it is worth noticing some aspects of the workflow to implement the allocation. After doing the above procedure we provided to the school staff a list indicating the groups of each student in each discipline. Between this information and the beginning of classes, some students usually accepted offers from other institutions before knowing their allocation.

These are usually high-ability students that end up being replaced by low-ability students. For those who started the course, there was perfect compliance with the allocation. One last important aspect is that the assignment of tutors to each group was done independently and in advance of the assignment of students.

3 Data

Our analysis use administrative data from the first semester of 2018 and 2019 academic years. All variables in our data are built from information on students' performance in the admission exam and in all disciplines they take during their first semester and from the allocation of students implemented according to our guidelines. In this section we define the basic variables we use and present descriptive statistics.

3.1 Variables

Academic performance The final measure of performance in a given discipline may take into account different kinds of activities. Each discipline combines exams, problem sets or other evaluations with varying weights. To have a more homogeneous measure of performance we use scores of the first exam taken in each discipline (standardized by discipline). Students take all exams during a week two months after the beginning of the course. There is a discipline (Probability) that starts only after the round of first exams. For this discipline we consider its final and single exam.

Students classification Students are classified according to their predetermined ability in mathematics and writing skills measured by the admission exam. For each student, we construct dummy variables indicating whether she belongs to the top or bottom 50% of his cohort's math or writing score distribution. This is what defines low and high-ability students in the analysis. Whenever we need to look at subgroups of students, we do that based on the math classification. Data on previous cohorts show evidence that this is the best predictor of GPA. To classify a given student's set of peers, we will sometimes use their classification based on writing skills.

Number of peers For a given student, this variable simply counts the number of other students in the group falling in each ability category (low, high). Most of our analysis will consider peers classified by math ability, although we use the writing ability classification in some exercises. There are two other types of students we count separately: those doing a discipline for a second time (*retained peers*) and few that did not use school's admission exam (*other peers*). These types were randomly allocated in the group and appear only as controls in the regressions.

Frequent peers As discussed before, some pairs of students take more than one discipline together. Thus, for a given student, we identify how many peers in her group are also present in at least some other group. This is done separately for each ability category. Then, we construct a variable containing the fraction of peers in each category that are also peers in a different group. That is, suppose a student has 4 low-ability peers in some group, and 3 of these peers also appear in some other groups. We label this variable as "Frequent Low Peers (%)" and in this example, it is equal to $3/4 = 75\%$.

Peers in other groups For identification purposes, we will have to control students' total exposition to peers of different ability levels (more on this later). Suppose a given student is allocated to groups (a, b, c) and that the distribution of his high-ability peers across the groups is given by $(2, 3, 4)$, respectively. This student is exposed to a total of $2+3+4 = 9$ high-ability peers (it does not matter peers' identity here). In groups (a, b, c), the variable "High Peers in other groups" are $(9, 9, 9) - (2, 3, 4) = (7, 6, 5)$, respectively.

Reported link We applied a survey asking students to name the most relevant peers for group work in each group. Also, we asked them to indicate peers they wish to be part of their groups in subsequent courses (independent of sharing some group at the time). We used this information to construct variables at the dyad level indicating whether student i reported peer j both conditional (j is in some of the i's groups) and unconditionally (j need not be in any of i's group). About two thirds of the students answered the questionnaire and table 6 in the appendix shows that women were more likely to answer the survey but answering is not related to ability.

3.2 Descriptive statistics

In table 3 we provide some basic statistics that might be useful as a benchmarking when discussing some of our results. Although table 2 shows that we allocated 178 students in 67 groups, we actually observe 135 students in these groups, resulting in a total of 791 observations. These are the students that effectively started the course and took at least the first exam. The median group is composed by 6 low and 6 high-ability students, and students have, on average, half of the peers in each group appearing in at least some other group.

4 Empirical Strategy

The empirical strategy of the paper explores the random allocation that placed students into tutorial sessions of different disciplines as the source of identification. As discussed previously, the allocation is done at the discipline level and our main results compare students within a given discipline by using the random variation of peer-related variables across groups as the source of identification.

In terms of inference, it is possible that unobservables are correlated for the same student across groups or for different students in the same group due to common shocks. Thus, our inference procedure acknowledges that for students i and j in groups g of d and g' of d' it would

Table 3: Descriptive statistics

Variables	Mean	SD	Min	Median	Max	N
Student's predetermined ability						
Final score	5.29	0.55	4.18	5.22	6.71	135
Mathematics	5.58	0.97	3.72	5.53	7.90	135
Writing	5.59	0.78	3.21	5.93	6.77	135
Number of peers by group						
Group size	12.58	2.39	8.00	13.00	16.00	67
High-ability	6.37	2.32	1.00	6.00	11.00	67
Low-ability	5.43	1.71	2.00	6.00	9.00	67
Fraction of peers with more than one interaction						
High-ability	0.48	0.26	0.00	0.50	1.00	791
Low-ability	0.51	0.27	0.00	0.50	1.00	791
Grade	6.78	2.13	0.00	7.00	10.00	791

Notes: *Student's predetermined ability* refers to to the the admission exam ability measures and ranges from 0 to 10. *Number of peers* counts the number of students in each groups considering their classification according to the math ability. *Fraction of peers with more than on interaction* contains the proportion of peers (by type) in a group that also appear in some other group.

be possible that $E\left[\varepsilon_{i,g,d}\varepsilon_{j,g',d'}\right] \neq 0$ if $i = j$ or $g = g'$ (conditional on observables). To deal with this issue we follow Cameron et al. (2011) by estimating a two-way clustered standard error at both student and group levels.[2]

4.1 Regressions at the dyad level

To test the hypothesis that meeting more frequently makes two students more likely to estabilish a link between we will use two different approaches. Using the sample restricted to individuals i that answered the survey and peers p with which they were allocated in at least one group, we estimate

$$y_{ipg} = \alpha_i + \beta \text{frequent\_peer}_{ip} + \mathbf{x}_p'\theta + \eta_g + \varepsilon_{ipg} \tag{1}$$

where y_{ipg} is an indicator on whether i reported p as a relevant peer in group g, $\text{frequent\_peer}_{ip}$ is an indicator of wheter p appears in at least some other group g', $\mathbf{x}_p$ a vector with characteristics of p, and the estimates within-student variation (α_i) conditional on the group (η_g). Thus, the estimate for β gives the average effect of p being a frequent peer on the probability of being reported as a relevant peer by i, *conditional* on having at least one group interaction. In this exercise, an observation is a pair ip in every group g in which they meet.

Another approach considers all survey respondents i and every potential peer p, irrespective of p's response to the survey and of wheter they share some group. In this case, we estimate

$$y_{ip} = \alpha_i + \beta \text{frequent\_peer}_{ip} + \gamma \text{interact\_once}_{ip} + \mathbf{x}_p' + \varepsilon_{ip} \tag{2}$$

[2]The inference assessment procedure proposed by Ferman (2019) does not indicate over-rejection problems.

where y_{ip} is an indicator of whether i chose p to be in some of his group in subsequent periods, frequent_peer$_{ip}$ and $\mathbf{x}'_p$ as before, but now, as it is possible that i and p does not share any group, there is the indicator interact_once$_{ip}$ to inform if they share at least one group. In this case, the estimate for γ gives the average effect of meeting on the probability of indicating a peer p, and the estimate for β gives the additional effect of meeting more than once. Here, the sample contains both ip and pi, $i \neq p$, but only a single observation for each case. Inference in both cases acknowledges the existence of potential correlation in own unobservables for both students and peers. So, standard errors are clustered at both student and peer dimensions.

4.2 Peer effects specification

A basic regression to estimate peer effects in our context is

$$y_{i,g,d} = \beta \text{high\_peers}_{i,g,d} + \mathbf{z}'_{g,d}\pi + \mathbf{x}'_{i,g,d}\theta + \eta_d + \varepsilon_{i,g,d} \tag{3}$$

where $y_{i,d,g}$ is performance of student i allocated to group g in discipline d, high_peers$_{g,d}$ is the number of high-ability peers he has in the group, $\mathbf{z}_{g,d}$ controls for further peer variables (other/retained) and group size, $\mathbf{x}_{i,g,d}$ contains student's ability measures (to increase precision) and randomization controls, η_d is a discipline fixed-effect, and $\varepsilon_{i,g,d}$ is a random shock. An estimate for β would potentially give the effect of having one more high-ability peer in replacement of a low-ability one (since group size and composition are fixed). However, there is a drawback to this strategy.

Different papers on peer effects raise the relevant question about whether the whole group in which students interact is the reference group that really matters. That is, although a student could interact with every person in the group, it might be the case that only a subset of these people forms the group we would like to define as *peers*. Our setting gives the opportunity to refine the analysis in this dimension by exploring the random variation in the frequency of meetings. Under the (testable) hypothesis that meeting more frequently makes two students more likely to establish a link between them, we argue that this kind of frequent peer is part of a more relevant reference group. Thus, we augment equation 3 as

$$\begin{aligned} y_{i,g,d} = {} & \beta \text{high\_peers}_{i,g,d} \\ & + \gamma_1 \Big(\text{share\_freq\_high\_peers}_{i,g,d} \times \text{high\_peers}_{i,g,d} \Big) \\ & + \gamma_2 \Big(\text{share\_freq\_low\_peers}_{i,g,d} \times \text{low\_peers}_{i,g,d} \Big) \\ & + \mathbf{z}'_{g,d}\pi + \mathbf{x}'_{i,g,d}\theta + \eta_d + \varepsilon_{i,g,d} \end{aligned} \tag{4}$$

where share_freq_high_peers$_{i,g,d}$ and share_freq_low_peers$_{i,g,d}$ are the variables we labelled before as *frequent peers*, computing the share of i's peers in group g of discipline d who appears in at least some other group g' of d'. Then, conditional on group composition, an estimate for γ_1 gives the average effect of having *all* high peers in g also in some other g' compared to having no high peers in g appearing somewhere else (analogously for γ_2).

However, given a set of high peers a student i has in group g, there is an increasing chance that some of these peers fall in the category of frequent peers as the number of high peers she encounters in every other group rises. From an identification perspective, not controlling for this total exposition to high-ability peers (the same applies for low-ability) could inflate the estimates for γ_1 and γ_2: these estimates would capture not only the effect of having more relevant peers in a group g, but also a feedback effect arising from the interaction in other groups. Thus, our main specification becomes

$$\begin{aligned} y_{i,g,d} = {} & \beta \text{high\_peers}_{i,g,d} \\ & + \gamma_1 \Big(\text{share\_freq\_high\_peers}_{i,g,d} \times \text{high\_peers}_{i,g,d} \Big) \\ & + \gamma_2 \Big(\text{share\_freq\_low\_peers}_{i,g,d} \times \text{low\_peers}_{i,g,d} \Big) \\ & + \delta_1 \text{total\_high\_peers}_{i,-g} + \delta_2 \text{total\_low\_peers}_{i,-g} \\ & + \mathbf{z}'_{g,d} \pi + \mathbf{x}'_{i,g,d} \theta + \eta_d + \varepsilon_{i,g,d} \end{aligned} \tag{5}$$

where total_high_peers$_{i,-g}$ and total_low_peers$_{i,-g}$ measure the total number of high and low-ability peers that i has in all other groups different from g, indicated by $-g$.

5 Results

In this section, we present and discuss three sets of results. Firstly, we show that pairs of students meeting more frequently are indeed more likely to establish a link, which we proxy by two different kinds of peer interaction variable: peer report conditional on sharing some group and unconditional peer report. Then, we show that a basic peer effects regression does not identify any effect on performance, while a positive effect arises when we consider a specification able to distinguish a set of more relevant peers. Finally, we show that, for at least a subset of the students, this positive effect might be explained by an improvement in their network position within groups.

5.1 Effects on peer reporting

We begin by discussing results on table 4, which presents estimates for equations 1 and 2. Columns 1 and 2 present estimates for equation 1 with and without controls, respectively. In these cases, *control mean* shows that from the pool of group mates with which students interact only once, there is a 27.5% of chance that some of these mates are named as a relevant peer in the group. However, the estimate on *frequent peer* shows there is a 3.8 pp increase in this chance when the pair of students has at least one more opportunity to meet, what means a 13.8% effect. One potential concern with this estimate could be that by meeting more frequently with someone, a student can learn about the characteristics of this frequent peer. Then, reporting this person as a relevant frequent peer in a group could just reflect that more skilled peers are being recognized. To check this is not the case, results of column 2 controls for peer ability and gender. Although these characteristics seem to matter for reporting a peer as relevant, the effect of a random increase in the frequency of meeting is actually slightly increased to be 4 pp.

Table 4: Effect of frequent interaction on peer reporting

Dependent variable: Reported peer				
	Conditional		Unconditional	
	(1)	(2)	(3)	(4)
Frequent peer	0.038*	0.040*	0.067*	0.067*
	(0.017)	(0.016)	(0.014)	(0.013)
Peer interact at least once			0.038*	0.038*
			(0.007)	(0.007)
Peer's writing ability		0.036$^+$		0.006
		(0.020)		(0.006)
Peer's math ability		0.058*		0.021*
		(0.020)		(0.006)
Peer is a woman		0.023		0.028$^+$
		(0.048)		(0.017)
Control mean	0.275	0.275	0.039	0.039
Observations	7343	7343	7852	7852

Notes: The dependent variable is an indicator that i reported j as relevant peer (i) *conditional* on i and j being in the same group ; and (ii) *unconditionally*, which means i and j could have never met for class work. Each observation in case (i) is a pair i and j in each group they share. In case (ii), each observation is every potential pair ij. Standard errors clustered by student and peer in parentheses. $^+$ $p < 0.1$, * $p < 0.05$

Now, we turn to columns 3 and 4, which present estimates for equation 2. This time, *control mean* reports that there is a 3.9% chance that a survey respondent indicates a peer that *does not* share any group with him to be in a group of a subsequent course. The estimate of 3.8 pp in column 3 shows that interacting in at least in one group doubles that chance. Further, conditional on interacting at least once, the additional effect of meeting in more than one group is 6.7 pp, which means again a huge increase of 87% in the chance of reporting a peer compared to 7.7% rate of those meeting once. Column 4 shows that peers' characteristics are not driving peer report either.

Each measure of peer report – conditional or unconditional – and results associated to each one suggest

5.2 Effects on performance

Table 5 shows estimates for equations 3, 4 and 5 and, for the latter, also reports results separately for the subgroups of low and high-ability students. The first column shows that with the most basic specification there is no evidence of peer effects on performance. In this exercise, the coefficient on *High peers* is close to zero and not statistically significant. It would give an estimate of the average effect of adding one high-ability peer in replacement of a low-ability one since group size and the number of other peers are constant.

Table 5: Peer effects on performance

Dependent variable: Standardized exam					
	All Students			Only High	Only Low
	(1)	(2)	(3)	(4)	(5)
High Peers	0.001	-0.032	-0.029	0.001	-0.052
	(0.017)	(0.025)	(0.026)	(0.031)	(0.045)
Frequent Low Peers (%) × Low Peers		-0.012	0.001	-0.013	-0.002
		(0.023)	(0.026)	(0.037)	(0.044)
Frequent High Peers (%) × High Peers		0.051*	0.045*	0.052	0.035
		(0.023)	(0.022)	(0.033)	(0.035)
Low Peers in other groups			-0.007	-0.011	0.004
			(0.017)	(0.024)	(0.020)
High Peers in other groups			0.013	0.023+	0.005
			(0.010)	(0.013)	(0.017)
Observations	791	791	791	427	364
Students	135	135	135	135	135
Groups	67	67	67	67	67

Notes: Standard errors clustered by student and groups in parentheses. $^{+}$ $p < 0.1$, * $p < 0.05$

Following the discussion about the problems of this basic peer effects specification, column 2 adds interactions of the number of low and high-ability peers with the fraction of these peers that appear in some other group. Since we did not find any effect by solely changing group composition, this allows us to look at what happens following a change in the potential interaction of peers.

It is simpler to interpret the result with an example. Suppose in a given group we will keep constant the total number of low or high-ability groups. Also consider that, initially, none of the low-ability peers of a given student in this group appears in some other group. The coefficient on *Frequent Low Peers (%) × Low Peers* shows what would be the average effect on the student's performance if we replace all these *nonfrequent* low-ability peers with other low-ability peers interacting with the student also in some other group. That is, given the number of low-ability peers, the effect of having 100% of them as frequent peers compared to having zero. Results show that the estimate for this case is statistically zero. However, applying the same reasoning to high-ability peers, the estimate of 0.051 now indicates that there would be a positive and significant effect. On average, half of a student's high-ability peers in a group fall into the category of frequent peers. If we would replace all non-frequent high-ability peers with high-ability frequent peers, the average effect on performance would be $0.5 \times 0.051 = 2.5\%$ of a standard deviation.

Yet, for two students in the same group, it is more likely that the fraction of frequent high-ability peers in the group is greater for the one who meets a larger number of high-ability peers in all other groups. If high-ability peers positively impacted performance, then the estimate of 0.051 would be upward biased by this feedback from interaction in other groups. This is why we include in column 3 the total exposition to low and high-ability peers outside the group as controls in the regressions (equation 5). The coefficient on *Frequent High Peers (%)* $\times$ *High Peers* falls to 0.045, which means that, by the same above argument, replacing all non-frequent high-ability peers with other frequent high-ability peers would have an average effect of 2.2%.

Finally, we separate the analysis by groups of ability in columns 4 and 5. The effects of high-ability frequent peers are no longer statistically significant, and this might be a result of losing statistical power since regressions now use smaller samples. However, point estimates are still positive and suggest that the effect would be stronger for high-ability students (+5.2% of an sd) than for low-ability ones (+3.5% of an sd). Besides, the estimate on *High peers* suggests that for high-ability students, changing the number of high-ability peers would not have any effect as long as the share of frequent high-ability peers remains constant. The same does not apply to low-ability students: there would be a negative effect of having more high-ability peers (-5.2% of an sd), but this could be at least partially offset if some of them were frequent peers. We explore this in more detail in the next exercise.

5.3 Heterogeneous effects of group change

The median group in our sample is composed of six high-ability and six low-ability students. Suppose we want to evaluate the average effect of replacing one low-ability student with a high-ability student. That is, we would like to compare the expected outcome in a new group with seven high-ability and five low-ability students to the expected outcome with the previous group formation.

The previous discussion suggests that to perform this analysis we should consider both the heterogeneity along students' ability levels and whether the replacement of students changes the fraction of frequent peers a given student will have in the group. In this exercise, we will consider two polar cases in terms of the potential interaction induced by such change: (i) Less interaction: the addition of a nonfrequent high-ability peer in replacement of a frequent low-ability one; and (ii) More interaction: the addition of a frequent high-ability peer in replacement of a nonfrequent low-ability. Based on estimates from equation 5, the next figures report point estimates and confidence intervals for the appropriate linear combination of coefficients representing the effect in each case for the desired outcome.

5.3.1 Effects on performance

Figure 2 shows that low-ability students would face an average decrease of 5% of a standard deviation in performance if the change reduced their potential interaction within the group and considering a 90% confidence interval we reject the hypothesis of no effect. However, if the same group composition was done by increasing their potential peer interaction, the point estimate falls to -0.01, but intervals at usual confidence levels now include zero. The importance of peer interaction is also evident when the exercise considers the subset of high-ability students.

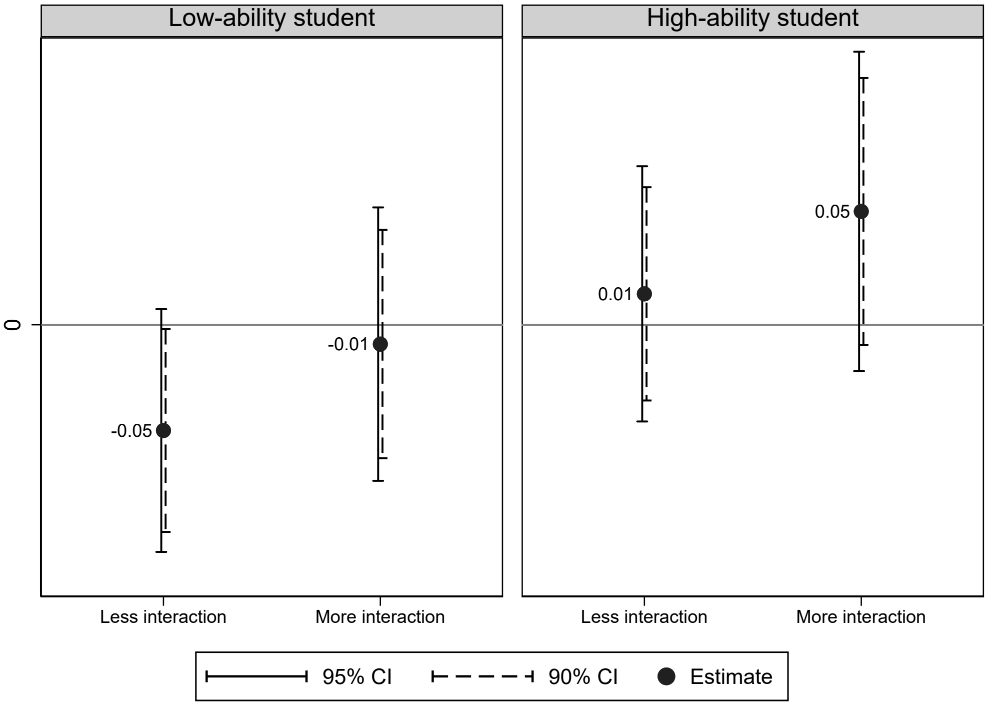

Figure 2: Effects on performance – *Less interaction* means the addition of a nonfrequent high-ability peer in replacement of a frequent low-ability one; *more interaction* means the addition of a frequent high-ability peer in replacement of a nonfrequent low-ability.

Moving from the change that decreases interaction to the one increasing it, the estimate goes from 1% to 5% of a standard deviation increase, and zero is barely within the 90% confidence interval in the latter case.

5.3.2 Effects on network position

If the previous results are indeed explained by differences in potential interaction induced by the group change, we should see effects on *actual* interaction. To test this hypothesis we will use the subset of students that responded to the survey and assume that if there is a link between two students, then both students should report each other in the survey.[3]

Results from columns 3 and 4 of table 4 already show that increased interaction between pairs of students makes it more likely the unilateral report the desire of having a peer in subsequent groups. Now, for each pair of students in each group, we compute an indicator of whether there is a match in such a report. That is, the variable is equal to one if students forming a pair reported each other in the survey. Finally, the dependent variable to be used in equation 5 is an indicator of whether a student has any match in the group.

In groups composed of six high-ability and six low-ability students, there is a 23% chance that a given student has match in peer reporting. Thus, the statistically significant estimates of 0.06 and 0.05 in figure 3 shows that changing group composition with an increase in the potential interaction means roughly a 22% increase in the chance of having a match in the

[3]Here we use the unconditional report as we believe this is a stronger measure of the interaction between two students.

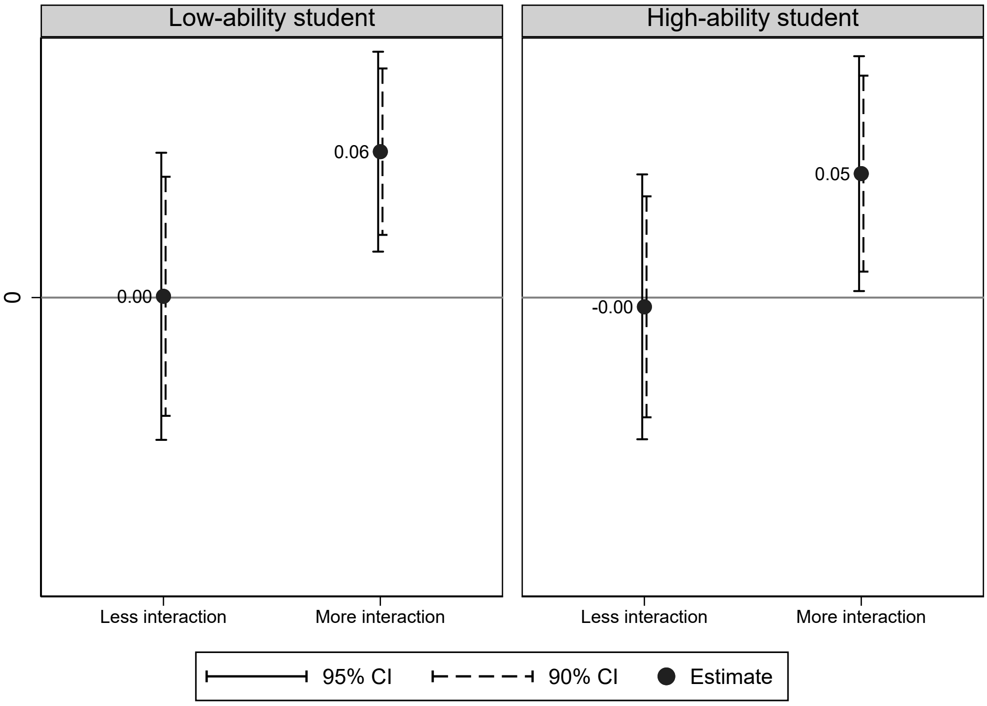

Figure 3: Effects on network position – *Less interaction* means the addition of a non-frequent high-ability peer in replacement of a frequent low-ability one; *more interaction* means the addition of a frequent high-ability peer in replacement of a nonfrequent low-ability.

group. We interpret this as an evidence that increased *actual* interaction is what explain the attenuation of a negative peer effects for low-ability students and the positive peer effects for high-ability ones.

6 Conclusion

Using data of students in higher education in an *active learning* environment we showed that increased opportunity for peers to meet is an important way to strengthen peer interaction so that positive spillovers on performance arise.

Actually, increase interaction with high-ability peers attenuate negative peer effects from high-ability peers on low-ability students and explain positive peer effects on high-ability students. In short, estimates relying solely on the variation of the share of low and high-ability peers could not reflect this important heterogeneity in the potential and actual interaction between the two groups.

A back-of-the-envelope calculation assuming that the maximization of the frequency of interaction with high-ability has a homogeneous effect for high-ability students indicates a 9.9% increase in average performance (not standardized). It is important to understand if this is a persistent effect since we analyze students in their first exams at university. Finally, a word of caution is needed since our results still come from reduced-form estimates, and – as evidence has shown – it may fail to uncover important endogenous peer responses.

References

Booij, A. S., Leuven, E., and Oosterbeek, H. (2017). Ability peer effects in university: Evidence from a randomized experiment. *The review of economic studies*, 84(2):547–578.

Brady, R. R., Insler, M. A., and Rahman, A. S. (2017). Bad company: Understanding negative peer effects in college achievement. *European Economic Review*, 98:144–168.

Bramoullé, Y., Djebbari, H., and Fortin, B. (2009). Identification of peer effects through social networks. *Journal of econometrics*, 150(1):41–55.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2011). Robust inference with multiway clustering. *Journal of Business & Economic Statistics*, 29(2):238–249.

Carrell, S. E., Fullerton, R. L., and West, J. E. (2009). Does your cohort matter? measuring peer effects in college achievement. *Journal of Labor Economics*, 27(3):439–464.

Carrell, S. E., Sacerdote, B. I., and West, J. E. (2013). From natural variation to optimal policy? the importance of endogenous peer group formation. *Econometrica*, 81(3):855–882.

Epple, D. and Romano, R. E. (2011). Peer effects in education: A survey of the theory and evidence. In *Handbook of social economics*, volume 1, pages 1053–1163. Elsevier.

Ferman, B. (2019). A simple way to assess inference methods. *arXiv preprint arXiv:1912.08772*.

Garlick, R. (2018). Academic peer effects with different group assignment policies: Residential tracking versus random assignment. *American Economic Journal: Applied Economics*, 10(3):345–69.

Goldsmith-Pinkham, P. and Imbens, G. W. (2013). Social networks and the identification of peer effects. *Journal of Business & Economic Statistics*, 31(3):253–264.

Hanushek, E. A., Kain, J. F., Markman, J. M., and Rivkin, S. G. (2003). Does peer ability affect student achievement? *Journal of applied econometrics*, 18(5):527–544.

Hoxby, C. (2000). Peer effects in the classroom: Learning from gender and race variation. Technical report, National Bureau of Economic Research.

Lee, L.-F. (2007). Identification and estimation of econometric models with group interactions, contextual factors and fixed effects. *Journal of Econometrics*, 140(2):333–374.

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

Sacerdote, B. (2001). Peer effects with random assignment: Results for dartmouth roommates. *The Quarterly journal of economics*, 116(2):681–704.

Sacerdote, B. (2011). Peer effects in education: How might they work, how big are they and how much do we know thus far? In *Handbook of the Economics of Education*, volume 3, pages 249–277. Elsevier.

Zimmerman, D. J. (2003). Peer effects in academic outcomes: Evidence from a natural experiment. *Review of Economics and statistics*, 85(1):9–23.

A Appendix

A.1 Survey responses

Table 6 shows the output of a logistic regression where the dependent variable is an indicator of survey response. Women were more liklely to answer the questionnaire, but responding is not related to student's ability.

Table 6: Probability of responding to the survey

	Answered
Math Ability Score	0.12
	(0.25)
Writing Ability Score	-0.06
	(0.20)
Woman	1.57^{**}
	(0.64)
Low-ability student	-0.51
	(0.62)
Mid-ability student	-0.70
	(0.58)
Observations	135

Robust standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$