

Thesis Abstract

Peer Effects in MBA Program (Job Market Paper)

In the paper on Peer Effects in MBA program I document the existence and magnitude of peer effects within small groups in a Masters of Business Administration program at a leading Canadian university.

A novel and unique dataset I acquired allows me to directly assess the potential spillover effects among the students. In the program of interest students are assigned by the administrators to groups of six and the assignment is based on a number of observable characteristics. To supplement the data, I conducted a survey among four cohorts of MBA students. I asked them about their study habits, their typical roles on their teams, evaluated their personality characteristics, and gathered information about their social networks. I analyze this data using the linear regression model accounting for cohort/section fixed effects and potential correlation of shocks within each team.

First, I confirm the existence of peer effects among MBA students. In particular, I find that high fraction of peers with Science/Engineering degrees has a negative effect on grades in Finance course. It also appears that higher fraction of top GMAT score peers has a somewhat negative effect, while higher fractions of bottom admission GPA peers may have a positive effect on grades. Second, more interestingly, I was able to show that the peer effects are heterogeneous in their nature. Students are not affected by their peers equally: the most striking result is that grouping students by admission GPA score may be beneficial to their academic outcomes. I find that low ability (in terms of admission GPA) peers have a positive effect on low ability students, and high ability peers have a stronger and more significant positive effect on high ability teammates. Third, using the survey data, I find that high GMAT students tend to score lower on the “Agreeableness” characteristic. And Science/Engineering students overwhelmingly prefer to “play devil’s advocate” during the meetings, getting involved in arguments with their colleagues. Taking these two findings together, I posit that the driving force behind the peer effects may be the psychological atmosphere in the group: students in teams with more agreeable peers, who do not engage in needless arguments may be better off academically. Similarly, lower ability students may feel more comfortable contributing to the discussion and asking questions if they are studying with similar ability peers, making their study process more efficient and improving their grades.

My final finding addresses one of the potential issues with peer effect research. Often, researchers use a cumulative GPA or another average grade as an outcome of interest. Using the grades from two different courses I show that peer effects may be course specific. This means that if we look at an average grade or a cumulative GPA as an outcome, we may be missing important peer effects.

Peer Effects in Small Teams: Testing Team Allocation Rules

Many MBA programs in North America divide their incoming class into small learning teams – with the goal of providing students a group of peers to study with, for group projects and other group activities a program may host. The assignment to groups is often done based on some basic, common sense rules that administration deems appropriate or necessary. These rules are aimed to even out the teams in terms of educational backgrounds, gender, immigration status and ability. However, there is little

published research that looks at the efficiency of these rules, considers how students might affect each other, or tests different team allocation guidelines.

In this paper, I use the results from my paper “Peer Effects in MBA Program” (Foltin 2017) to run a group creation simulation and test a variety of group assignment rules and their effect on students’ grades. I incorporate the peer effects I found in the previous paper, as well as account for heterogeneity of these peer effects by including a number of cross terms.

I test the random allocation of students across the teams, allocation by admission GPA and GMAT (both homogeneous and mixed), by gender, by immigration status and by undergraduate degree. I also see whether making minor adjustment to existing group composition would make a significant difference in grades.

I find that grouping students by admission GPA results in highest average grade. A number of other rules also produce slightly higher average grade than the original allocation. I check how different types of students are affected by every group assignment rule to ensure that the increased average is not driven by one group benefitting at the expense of the others. These findings could be used to improve the existing team allocation guidelines within MBA program or in any other academic settings.

On Methodology of Social Network Data Collection: Comparison of Two Common Methods

(With Tim Conley)

There are two traditional ways of collecting social network data: asking participants to pick names from a given roster (recognition) or asking them to name their friends (recall). There is no existing research in economics that compares and contrasts the results of these two methods.

Using unique data from a survey where we asked MBA students about their connections within their class both ways, we can directly compare the resulting networks. In addition, we have information about one of the underlying social networks of the class – students' administratively-assigned study group compositions – which helps us make comparisons between the reported and actual connections between students.

In the “recall” question, we asked students to write down up to seven names of their friends in the program, not limiting them with any definitions of friendship. In the “recognition” question, we asked them to pick names from a list, answering three questions: who do they chatted with over the course of the last two weeks, who do they go to for a school related help or advice and who do they socialize with outside of school hours. This resulted in three different, overlapping networks. We compare the resulting networks to draw conclusions about how different are the results of recognition and recall questions.

We use the actual team network to comment on the possible false negatives in the networks. There is a very high probability that students assigned to the same team have talked to each other over the course of the last two weeks. Looking at the overlap of team network and reported network allows us to comment on how many peers students forget to mention in their recollection and recognition questions.

Our findings will be informative to the future social network data collection and it will provide some basis for choosing one way of gathering friendships data over another.