Provision of Online Public Goods: Evidence From a Student Discussion Board

Hammad Shaikh*
University of Stavanger

March 5, 2026

Abstract

How does a policy designer maintain an actively participating online community in the presence of free-rider incentives? This paper studies the question in the context of voluntary online student discussion boards – a prominent feature of distance education used to support learning at scale. I first collect rich survey and administrative data covering nearly 1200 undergraduate students in a foundational programming course at a large public university. The data allow me to observe the number of unique posts read and written, and characterize students' altruistic attitudes. I then conduct two randomized informational interventions, successfully nudging students to sign-up and then contribute further to the discussion board. The first intervention informs students about availability of the voluntary discussion board, and the next intervention is a writing exercise that helps them internalize the spillover value of peer discussion. I find that having access to and participating in the discussion board significantly improves students' learning outcomes. To measure the extent of free-riding present, I develop and estimate a behavioural model of peer discussion board provision. The estimated model allows me to simulate participation in the discussion board when all students cooperate to maximize aggregate learning. I find that students in the control group provide 21% less contributions than the simulated social optimal allocation.

*I would like to thank Robert McMillan, Aloysius Siow, and Román Andrés Zárate for their guidance and support. Thanks also to Marc-Antoine Chatelain, Choek In Fok, Elaine Guo, Alexander Hempel, Faisal Ibrahim, Daniel Indacochea, Catherine Michaud Leclerc, Swapnika Rachapalli, Baxter Robinson, Jessica Wagner, Matt Walshe, Joseph Jay Williams, Farhan Yahya, and Jiaqi Zou. The study also benefited from comments by participants in the Graduate students in Economics of Education Zoom seminar, University of Toronto's Public Economics Discussion Group, and International Research Society for Public Management conference. Sameul Khan, Sam Maldanado, and Qi Yin Zheng provided valuable data support. Anonymous course instructors helped me navigate the online learning platforms. The experiments and data collection are approved by the Research Ethics Board of the anonymized institution of this study. The study is registered at the AEA RCT Registry (ID 0008787). Funding from the Ontario Graduate Scholarship for this project is greatly appreciated. All remaining errors are my own. Contact: Department of Economics, University of Stavanger Business School, Kjell Arholms gate 35, 4021 Stavanger. Please send comments to hammad.shaikh@uis.no.

1 Introduction

The emergence of several vast online communities over the past two decades is a striking feature of modernity. Prominent among these are Wikipedia, reddit, GitHub, Stack Overflow, and Quora. Such online communities have obvious public good aspects in that they are generally available to everyone with access to the internet (they are non-excludable), and participation in the community does not limit the usage of others (they are non-rivalrous). As the participants in these online environments are typically both users and creators of the content provided by the online platform, the quantity and quality of that content depends on the voluntary efforts exerted by community members. Yet maintaining a high level of contributions in practice is challenging due to members having clear incentives to free-ride on other users' content creation efforts.[1] This prompts the natural question: How does a policy designer maintain an actively participating online community in the presence of free-ridership incentives?

My study addresses this question in the context of online student discussion boards – an increasingly common adjunct to traditional university course instruction, and also a fundamental component of so-called Massive Open Online Courses (MOOCs). I employ rich student-level administrative and survey data together with two online field experiments to investigate the effectiveness of 'nudging' strategies in promoting an active peer learning environment. Online learning technologies are often viewed as being scalable cost-saving alternatives to the traditional in-person classroom setting (Deming et al., 2015). It is difficult, however, for instructors to provide individual support to a vast body of students, often having varying class and work schedules, a challenge we have yet to solve (Bowen, 2012). In that respect, enhancing the provision of student discussion forums is a potential appealing way of supporting student learning while providing education *at scale* to a massive amount of students.

The specific setting for my analysis is a large online introductory programming course offered each 12-week semester at a research-intensive Canadian university. The course follows the pedagogy of a 'flipped-classroom,' where students are required to learn content independently by watching

[1] We see potential evidence of free-ridership, for example, on Stack Overflow, an online question-and-answer platform for various coding languages. According to Xu et al. (2020), 90 percent of the content is attributable to 10 percent of active users, which suggests that many users only browse existing posts. Similarly on Coursera, a large provider of massive open online courses, Cohen et al. (2019) find that majority of the discussion board contributions come from 20% of the students.

videos and doing practice problems posted on an online platform, prior to the in-person lecture, which focuses mostly on problem solving. Given that students in the course learn most of the material on their own online, the course employs an online student discussion board to further support students. The discussion board facilitates learning by allowing students to help each other while the instructor can moderate the discussion asynchronously. In this setting, students must register for the discussion board prior to being able to consume and contribute content. That is, students face both extensive margin (i.e., sign-up) and intensive margin (i.e., writing content and reading posts) decisions. Student participation in the discussion board is completely voluntary.

I collected data from the Winter 2020 academic term involving nearly 1200 students who consented to participate in the research. The administrative data includes discussion board registration status, student contributions, and the number of unique posts read. These data enable me to monitor students participation activity on the online discussion board precisely. Students can 'consume' the discussion board privately by browsing and reading posts, or they can contribute by asking unique questions, answering unresolved problems, and commenting on the posts of their peers. I find that students enrolled in the discussion board on average write 3.5 posts, even though students on average read around 120 unique posts. I supplement the administrative data with survey data, collecting information relating to how students utilize the discussion board and further eliciting their behavioural characteristics, such as their altruistic attitudes and attentiveness. I find that 65% of the students report to using the online discussion forum primarily to learn course concepts by browsing posts written by their peers.[2] Further, more attentive students are more likely to sign-up to the discussion board, and more altruistic students contribute more. The evidence clearly suggests that free-ridership behaviour is present in this environment given the natural presence of less altruistic students in a large classroom.

Several incentive design mechanisms have been proposed to resolve under-contribution in canonical public good settings (e.g. mechanisms proposed in Groves and Ledyard (1977) or Falkinger (1996)). These are, however, infeasible in many online environments where monetary transfers are not applicable, or there is no central authority to enforce transfers. Combatting free-ridership in online environments thus tends to involve using mechanisms that do not rely on transfers but

[2] A question in the baseline survey prompted students to rank the following categories to represent their discussion board usage (top = most representative, bottom = least representative): ask questions, answer peer questions, learn course concepts by browsing other peers' posts, and make comments on peers' posts.

rather on aggregate donations exceeding some threshold.[3] Such mechanisms are appropriate in public online settings where contributions themselves are monetary, for example, donations to a online charity. However, they clearly cannot be incorporated in a discussion board setting where contributions (i.e., written content) cannot be refunded if the threshold for provision is not achieved.

Instead of directly changing student incentives to participate on the online discussion board through transfers, I consider a nudging approach that is widely employed in the behavioural economics and education literature.[4] The interventions considered in this study can be characterized as "targeted informational prompts," as they present students with information that prompts them to take a specific action. I deploy two randomized nudges, separately targeting the extensive sign-up margin, and the intensive participation margin. At the beginning of the course, students receive the first nudge, an online sign-up activity that encourages them to register for the discussion board. The sign-up activity uses screenshots and an animated GIF to provide clear, step-by-step instructions to register for the discussion board, and informs students about its functionality. The second nudge, deployed later in the course, is an online writing activity that leverages quotes from a previous cohort of students to encourage students to be active contributors rather than only consumers (i.e., free-riders). The reflective writing prompts in this intervention are designed to help students internalize the spillover benefits from engaging in peer discussion.

On the extensive margin, I find that the sign-up activity is successful in encouraging registration early on in the course, increasing student utilization of the online discussion board by approximately 4.5 weeks. I then use the random assignment to the sign-up activity as an instrument for the number of weeks students are registered: an additional five weeks utilizing the discussion board increases homework and final exam achievement by 0.14 SD and 0.08 SD, respectively. I also use the survey data to investigate the mechanisms underlying this intervention, finding that the sign-up activity is highly effective in informing inattentive students about the online peer discussion platform, that is, students who are less aware of the discussion board and who are unfamiliar with its functionality at the outset of the semester.

[3] The Provision Point Mechanism (PPM) Rondeau et al. (1999) or Conditional Contribution Mechanism (CCM) Oechssler et al. (2022) are based on aggregate community donations exceeding some threshold, and are returned to the initial donor otherwise. In the CCM setting, participants can make investments dependent on the contribution of others. For example, a participant can choose to contribute \$10 only if at least two other members also contribute \$10, otherwise they contribute nothing and are refunded their initial \$10 investment.

[4] For example, see Thaler (2018), Harackiewicz and Priniski (2018), and Damgaard and Nielsen (2018).

On the intensive margin, I find that the reflective writing exercise increased the number of posts written by around 26%, and the number of unique posts read by 11%. Although the intervention was designed to increase contributions, consumption of content naturally increases as students are discouraged from posting duplicate questions, and answering involves reading a peer's question. Information regarding students' time usage on the discussion board from the survey data enable me to construct a measure of the hours students spend utilizing the discussion board.[5] Using the writing activity as an instrument for the time spent using the discussion board, I find that spending an extra hour of participation increases students' final exam grade by 0.037 SD.

To further understand the mechanisms underlying free-riding behaviour and the variation in contribution activity observed across student types, I develop a model for the provision of an online public good. The model features a single instructor (i.e., principal) and multiple students who vary in their altruism and baseline human capital (i.e., they are heterogenous agents); where a student's altruism governs the extent spillover learning benefits from own contribution are internalized, and the human capital dictates the quality of contribution. In this model, the instructor's objective is to maximize the aggregate learning of their students from participating in the discussion board, whereas the students make participation decision to maximize only their own learning net of effort costs. Solving the instructor and student problems results in the following key implications: 1) students will under-contribute relative to the social optimal if there are positive learning externalities from content creation efforts that they do not internalize, and 2) altruism and human capital are complimentary characteristics for inducing contributions (e.g. pure altruistic and 4.0 GPA students provides the most content).

After validating the model by finding empirical evidence to support the model implications, I use a combination of survey and administrative data to estimate the model using an iterated maximum likelihood estimation routine. The marginal benefit and marginal cost of contribution parameters are credibly identified using the field experiments and the survey data.[6] I then employ the model to infer student contribution behaviour under the following counterfactual regimes: 1) when all students internalize the spillover benefits of their contributions, and 2) students receive bonus credit for writing high quality posts which are endorsed by the instructor. I find that student contributions

[5] Section 4 details how the time students spend in the discussion board is measured using the data.

[6] Identification argument for the structural parameters is described in Section 8.7.

increases by 21% when all students cooperate to maximize aggregate learning of students in the discussion board. Additionally, the average final exam grade increases by around 0.04σ under the cooperative allocation. To mitigate such free-riding behaviour, I carry out another counterfactual where students are given 0.5% bonus credit for making a contribution that is publicly endorsed by the instructor. Under this incentive scheme, I find that low-altruism students, yet high human-capital, significantly write more posts, increasing overall contributions by 18%. That is, providing incentives for contributing can help combat free-riding incentives and increase the quality of the discussion board, especially through encouraging high human capital students who otherwise would contribute less.

Overall, the findings in this paper contribute to our understanding of how to encourage engagement and mitigate free-riding in online learning communities. I provide the first causal evidence showing that utilizing an online peer discussion board helps support learning at scale. The experimental variation in student engagement allows me to estimate and credibly identify a behavioural model of online public goods provision, helping shed light on the mechanisms underlying free-riding behaviour. While the analysis in this paper focuses on student participation activity on a online discussion board, the proposed approach is quiet general and can be applied in other online public good settings such as on GitHub or Stack Overflow, helping understand how members in a given online community formulate their participation decision.

The rest of the paper is organized as follows. The next section places my analysis in the context of the related literature. Section 3 provides information about the sample and describes the online discussion board environment. Section 4 outlines sources of data collection and also presents descriptive statistics. The experimental design and the key features of the interventions are discussed in Section 5, and corresponding results are presented in Section 6. A model for online public good provision in the context of student discussion boards and corresponding theoretical implications are set out in Section 7. I estimate the proposed model and discuss identification in Section 8, Section 9 presents counterfactual analyses using the estimated model, and Section 10 concludes.

2 Literature Review

This paper builds on several prior literatures. These include research on the provision of online public goods, the structural behavioural economics literature, and papers that assess the efficacy of digital education technology. In this section, I discuss the related literature in turn and highlight the ways in which my paper contributes to each of them.

First, my paper relates to the literature on online public goods that focuses on identifying the determinants that motivate users to contribute content. For example, Edelman (2012) uses observational data of around 40000 question and answers from Google answers, a Q & A website where the answerers are paid, to study the key determinants of answerers pay. The study finds that longer answers with more urls are associated with higher pay per minute. Although there are online communities where contributions are incentivized through monetary reward, most of the economic literature and my paper focuses on a setting where contributions are voluntary. For instance, Chen et al. (2010) evaluate the affect of presenting social information to users on their subsequent voluntary contribution of movie ratings. The authors conduct a field experiment with 398 users on MovieLens and find heterogenous affects from presenting the median contribution of other users on the platform: those below the median rate more movies, whereas those above the median reduce movies rated. A prominent strand of the online public goods literature studies motivations of users to contribute to Wikipedia. In particular, Gallus (2017) assess the effectiveness of providing pure symbolic awards (non-monetary) in alleviating the problem of declining retention rates of new wikipedia editors. The author conducts a field experiments involving 4007 editors and finds that retention rates for the group of editors that were given a symbolic award are 20% higher in the following month than the control group. Chen et al. (2018) focuses on what motivates experts (e.g., professors) to contribute Wikipedia. Through a field experiment that involves emailing professors, the study finds that the prospective for the expert to be cited on highly viewed articles motivates them to contribute. A common theme across this literature is that it focuses on using messaging systems to deploy various information to users to promote contributions. In contrast, my paper investigates student motivations for contributing to an online peer discussion board through having students participate in various interactive online activities. This design enables me to study how students interpret and react to the information they are presented, analysis that is limited in emailing interventions as data on engagement with emails is typically unobserved.

Second, my paper is related to the literature that integrates behavioural insights into estimable models of decision making to design public policies. This literature often leverages experimental data to credibly identifying the structural parameters. For example, Allcott and Taubinsky (2015) conduct an experiment involving store customers to measure the extent to which consumers are inattentive to the energy savings of LED light bulbs. The authors elicit willingness to pay for various light bulbs through an iPad survey and randomize a portion of the participants to receive information about the corresponding energy costs. They find modest evidence of inattention that would justify a subsidy of 3 dollars per LED light bulb. Taubinsky and Rees-Jones (2018) conduct an online shopping experiment in which consumers face various tax rates, and find substantial evidence of inattention to taxation. Augenblick (2018) measures the present bias of study participants who are paid to complete unpleasant tasks by asking them repeatedly on different days the number of blurry greek letter transcriptions they would like to complete on a given date. By randomizing wages and the number of days after the decision the tasks are carried, the author finds some evidence of present-bias exhibited by the study participants. A large fraction of the literature that identifies structural behavioural parameters uses experiments that are conducted in laboratory settings (Card et al., 2011). Stringent experimental configurations are typically required to credibly identify parameters such as the extend to which someone is altruistic. I add to the literature by using survey data to elicit a student's altruism and incorporating that into a estimable model of human-capital adjusted contributions.

Third, my paper adds to the literature assessing the efficacy of educational tools by evaluating the value-added of an online peer discussion board.[7] The studies that directly focus on assessing the efficacy of student discussion forums are scarce and largely provide only correlational evidence. Cheng et al. (2011) evaluate the effectiveness of students participating in an online voluntary discussion forum in a large undergraduate introductory psychology course. Through conducting a descriptive multiple regression analysis, the authors find that students who post on the online forums achieved around a 0.05 SD higher grade on their midterm and final exam. AlJeraisy et al. (2015) assess the impact of discussion board availability on the extensive margin by comparing outcomes across two sections of students enrolled in a small Business administration course, with

[7] Escueta et al. (2017) contains an survey of the literature on education technology.

only one section using a peer discussion board. The authors find that students in the cohort with the discussion board achieved a higher course grade and also self-reported an higher level of satisfaction. Since students can select into the section of their choice, the estimate for the value-added of the discussion board may be biased.[8] In contrast to this work, this paper leverages experimentally induced variation in the propensity to register for the discussion board to evaluate the achievement gains from discussion board utilization.

3 Institutional Background

Course Structure The course assumes no prior programming knowledge and teaches programming fundamentals using Python. It is offered every semester and typically enrols around 1,000-1,500 students in the Fall and Winter terms, and around 200-400 students in the Summer term. Although the course is offered at the first-year level in the Computer Science (CS1) department, it consists of CS-majors and non-majors alike and is not exclusive to first-year students; many students who enrol have no programming experience.

The course structure follows a 'flipped classroom' pedagogy. At the beginning of each week, students watch instructional videos and complete a set of multiple-choice recall problems for course credit. During in-person lectures, they focus further on practising problems and revise concepts related to the videos they watched. After the lectures and at the end of the week the students complete a set of for-credit online multiple choice questions and open-ended Python coding problems. The coursework also includes a midterm, and a final exam. In addition, students can obtain course credit by participating in two research surveys that are deployed during the start and end of the course.[9] The graded homework assessments count for 25% of the students' overall course grade. The midterm counts for 30% of the students' course grade, and the final exam count for the remaining 45%. The course progresses through new course content on a weekly basis and runs for approximately 12 weeks, after which the students write a final exam.

The Online Peer Discussion Board. The course offers an interactive online course discussion board where students can discuss course material. All students in the course are able to sign up

[8] The authors use propensity score matching to control for selection into the two sections. Although, the remaining matched sample has less than 100 students, reducing power in identifying achievement differences across sections.

[9] The research surveys typically account for 1-2 percent of their course grade.

for the online discussion board, where they can participate by asking questions, help their peers by writing answers, and engage in discussion with peers by commenting on existing questions and/or answers. Students who do not sign up cannot read posts or contribute content in the discussion board. Posts on the discussion board are organized per week as new content is introduced weekly. Appendix A.3 describes the user interface of the discussion board in more detail. Although encouraged, participation in the discussion board is completely voluntary, and students are not awarded any additional course credit for registering for it or creating content. The peer discussion board is monitored by the instructors of the course. Each post is eligible to receive upvotes for good quality by students and instructors. When instructors upvote a students' post, they receive a "endorsed by instructor [full name]" label.

The Online Learning Platform The weekly homework modules are hosted on an open-source online learning environment created by the computer science department of this institution. The environment is an interactive online platform that allows education providers to bundle video instruction together with multiple choice and open-ended programming problems. Distinct content on the online homework platform is separated by weeks, and each week students are assigned to watch videos and complete follow-up problems. This online learning platform is extensively studied in Shaikh (2020A, 2020B).

Learning Management System. Canvas is the learning management system (LMS) employed by the course involved in this study. All students registered in the course have access to Canvas as it is used by the university. It is used to set up and organize a digital learning environment. In my setting, Canvas is used by instructors to post announcements, manage course deadlines, and release student grades. As discussed in detail in Shaikh (2020A), instructors can also send messages to students through Canvas.

4 Data and Descriptive Statistics

This study uses a combination of rich student-level administrative and survey data to characterize student participation activity in the online discussion board for different types of students. All data are gathered from the introductory programming course (introduced above), offered during the Winter 2020 academic semester in a large research-intensive Canadian university. The data are collected and merged together from the following sources: online surveys, a learning management system, an online homework platform, and the online peer discussion board. The writing activity (as mentioned above) uses survey data gathered in Fall 2020. Pilot data from Summer 2020 was also collected to finalize the interventions considered in this study.

Student Survey Data. All students in the course are requested to complete three research surveys [10]. The baseline survey collects information on students' demographics, and elicits behavioural characteristics such as their altruistic attitudes. The next survey elicits student motivations to participate in the discussion board and their utilization of learning resources. The final survey gathers data on interaction with peers aside from the discussion board, and elicits student feedback on the interventions they received. Each survey takes around 20 - 25 minutes to fill out, and is voluntary, although students are given around 1% course credit for completing each survey. The response rate is around 92% for the baseline survey and 87% for the end-line survey.

The baseline survey contains a consent form, which asks students to participate in the study by allowing their data to be used for the purposes of academic analysis and research. In addition to the baseline survey, students are also given an opportunity to consent to be a participant of the research study on the online homework environment. Overall the consent rate is around 84%. The sample of total consenting students who completed the course consists of 1184 student. My data consists of 1184 consenting students who completed the course in Winter 2020.

Student Achievement Data. Student achievement data are collected from the weekly online homework, the midterm, and the final exam. The availability of this high frequency achievement data allows me to assess student learning throughout the course. The primary measure of learning is the students' grade on the cumulative final exam, which is standardized to be mean 0 and standard deviation 1.

[10] All surveys were deployed online using the Qualtrics survey platform. The questionnaires are available upon request.

4.1 Student Discussion Board Activity Data

Student discussion board registration status is collected at the weekly level. Thus, I observe the number of weeks a student is registered for the discussion board. Additionally, I observe time-stamped data on all contributions (question, answers, or comments), and the number of unique posts viewed by students each week. For each contribution, I can also observe the number of instructor and peer upvotes. Overall, I observe discussion board registration, contributions, and 'consumption' decisions.

Online Discussion Board Study Time. Although the administrative data set includes the number of posts written and read at each stage of the learning process, the time spent on these activities is not observed. To fill this gap, the final survey asks students the minutes they spend on average writing and reading a post. The administrative data on engagement, and corresponding student-level survey data on average time use, is used together to measure the minutes spent on the discussion board at each learning stage.[11]

4.2 Student Behavioural Characteristics

The baseline survey elicits students' altruistic attitudes through a series of questions. Each question is measured on a 7-point Likert's scale, and a student's response can vary between strongly disagree (i.e., a response value of 1) to strongly agree (i.e., a response value of 7). For example, they are asked whether "I see myself as someone who is willing to volunteer to take notes for students with disability". The accessibility services at the university considered in this study allows students to have a peer share notes with them. Clearly, students who volunteer to take notes and share them with a disabled peer are likely to be more altruistic. The additional three questions that are included in altruism index are presented in Appendix B.1. To construct a continuous index of altruism, replies to all relevant questions are aggregated together so that they are increasing in altruistic attitudes.

[11]For example, suppose a student views 11 posts and writes 5 questions. If in the survey data this student reports to taking 3 minutes to view a post, and 6 minutes to write a question, then their estimated time spent on the discuss board is $11 \times 3 + 5 \times 6 = 63$ minutes.

4.3 Summary Statistics

Table 4 presents a rich set of summary statistics related to student demographics, characteristics, behavioural information, homework participation, and discussion board activity. Although computer science graduates primarily composed of males Baer and DeOrio (2020), there is no significant gender disparity in my sample as 49% of the students are female. Around 77% of students live outside of campus residence and need to commute to attend their classes. The course is offered as a first-year course, but is not exclusive to first-year students, as around 45% of students are beyond their first year. Consistent with only 52% of students being domestic Canadians, only around 31% of students speak English at home.

Panel B of Table 1 shows that 87% of students do not have any programming experience prior to taking the course. From those that registered, the average student contributes 3-4 posts and views around 122 posts.

4.4 Self-reported Discussion Board Usage Statistics

The second survey collects detailed information about how students who registered for the discussion board use the online peer learning environment. Figure 6 shows the frequency distribution of how often students visit the discussion board for those that signed up for the discussion board. Around 68% of the students report to visiting the discussion board at least once per week and less than 10% report to almost never using it after registration. Therefore majority of the students somewhat actively view the discussion board on a weekly basis.

To learn about how students interact with the discussion board, the second survey asked students the extent to which they used the discussion board for asking their own questions, answering peer questions, writing comments on peer posts, and just browsing posts and learning by reading their peers posts. Figure 7 illustrates distribution of student responses to how they use the discussion board. The results suggest that students primarily use the discussion for learning course material by browsing and reading content contributed by others. Furthermore, the most popular type of contribution is asking questions, whereas answering peer questions is the least frequent. This self-reported data is consistent with the realized discussion board participation data shown in Table 4 as the views to contribution ratio is around 35 distinct post views per contribution. This suggest potential free-ridership behaviour is present in this setting.

5 Experimental Design and Nudges

This section describes the rationale behind the interventions that were deployed and outlines the experimental design for allocating students to treatment.

5.1 Design of experiments

The sample frame eligible to receive the nudges consists of all students who consented to participate in research during the Winter 2020 academic term. As discussed in the previous section, the data collection results in a sample of 1186 study participants. The study followed a double-blind protocol for implementing the randomized interventions. That is, students were not informed of their treatment status but were aware that a study was being conducted for the purposes of improving course design. The course instructors were aware of the interventions that were being deployed but were not informed about the students' treatment status. I performed the randomizations on an anonymized dataset, and I was not part of the instructional team.

5.2 Description of Nudges

The two interventions considered in this study are both online activities that are embedded within a incentivized research survey. The intervention are designed to target students to take an specific action and are composed of the following elements: 1) present students with information related to the targeted action e.g., instructions to sign-up for discussion board), 2) ask students to write reflections based on information presented (e.g., reflect on benefits to utilizing the discussion board) , and 3) prompt them to a take specific action (e.g., register to the discussion board). The design of the nudges is inspired by insights from psychology and behavioural economics research (Damgaard and Nielsen, 2018).

Discussion Board Sign-up Activity. The sign-up activity is designed to promote discussion board registration. The activity is composed of the following elements: 1) it presents a link to the discussion board sign up page, 2) it uses screen shots to illustrate key steps for sign up, 3) it summarizes all steps into an animated GIF, and 4) students are disclosed information about the discussion board. The online activity consists of two pages: the first page contains the instructions, and the second includes information about the discussion board. The informational page discusses the functionality of the discussion board, and also discloses the proportion of students' questions that have been answered by either a peer or the instructor. The activity had an interactive component as students were prompted to reflect on the information they were presented (see Appendix C.1 for the self-reflection questions).

For all the students who completed the baseline survey and did not register for the discussion board within the first week of the course, half of them received the sign-up activity. Such a balanced assignment will maximize power if the variance in potential outcome distributions across control and treatment is identical (Tabord-Meehan, 2018). Additionally, students in the treatment group are randomly assigned to receive either sign-up instructions only (i.e., page 1), discussion board information (i.e., page 2), or both. Table 5 shows the assignment of students to control and treatment. The factorial design helps assess which elements of the sign-up activity are most effective for nudging students to utilize the discussion board. The primary analysis combines the three conditions into a single treatment group that is compared with the control group.

Discussion Board Contributions Writing Activity. In the context of education, value-affirmation interventions involve students engaging in writing activities which prompt students to relate course materials to their own personal values. In these type of interventions students are often initially prompted to select values from a given list and write about why they are meaningful to them. The list students are shown can include values such as having a sense of humour, humbleness, independence, creativity and athletic ability. After students have identified values important to them, they are then prompted to relate those values to preparing for various course assessments such as assignments and exams.

Based on the value-affirmation literature, I design a writing activity which involves students reading quotes written by students from a prior cohort and writing about how content in the quotes relates to their own personal values.[12] The quotes expressed students perspective on the importance of contributing to the discussion board by asking questions, engaging in discussion

[12]The quotes were collected through surveys conducted in the Fall 2019 term and are used to design the value-affirmations module in Winter 2020.

with peers through comments, and answering other peers' questions. The activity is designed to help students reflect on the value of being an active contributor for creating an effective learning community, helping students internalize the spillover benefits to contributing. Table 6 shows the assignment of students to control and treatment.

5.3 Statistical Validity of Experiments

I now discuss the statistical validity of the experimental design by showing the following: 1) pre-treatment characteristics are balanced across the control and treatment group, 2) there is no differential attrition by treatment status, and 3) results are robust to spillovers.

Independence of treatment assignment The aim of the experiments is to identify Average Treatment Effect (ATE) effects of interest. The ATE is identified as students are randomly to the control and treatment group across both interventions and that it is mandatory for treatment group students to participate in the activity presented.[13] I investigate the validity of the random assignment by testing whether the pre-treatment student demographics and characteristics are balanced across the control and treatment group. I do so by standardizing each pre-treatment control and regressing these on the treatment status. Figure 8 shows a well-balanced control and treatment group for the sign-up activity. Similarly, Figure 9 shows that students who are assigned to the contributions writing activity are statistically identical in their demographics and characteristics at baseline.

Student Attrition. Student attrition is natural in my setting as students who initially enrolled and consented to participate in the study can choose to drop out from the course afterwards. In my sample of 1388 students who initially agreed to participate in the study, around 88% of them completed the course. Table 7 reports the differences in the proportion of attriters in the control and treatment groups across both interventions studied in this paper. The analysis suggests that none of the interventions caused students to dropout of the course directly as all treatment coefficients are close to 0 and the corresponding p-values are larger than 0.1.

[13] There are a few students try to avoid participating in answering questions by writing nonsensical responses. These observations are dropped from the analysis.

Well-defined Treatment Assignment. The validity of the experiment requires the following conditions to hold: 1) treatment level is unique so that potential outcomes are well defined and 2) treatment applied to one unit does not effect the outcomes of other units. The intensity of the treatment is homogenous across the treatment group as all students within a treatment group receive the same online activity. Therefore the potential outcomes corresponding to the experimental conditions are well defined. Next let us discuss the possibility of spillovers.

Since students interact with each other during regular lecture time to work towards solving problems, it is possible that students in the treatment group who received some nudge will interact with the control group who did not receive the online activity. Assuming the nudge increases an outcome of interest (e.g. more discussion board contributions), that reasonably can result in positive spill overs to the control group through information sharing (e.g. treatment group students sharing lessons learned from online activity) or peer effects (control group student following behaviour of treatment group student). Such positive spillover effects will result in downward biased effect sizes. I use the final survey data to argue that spill over resulting from students in the treatment group discussing the contents of the intervention to students in the control group is negligible in the next section.

6 Theoretical Framework

This section presents a conceptualization of an online learning environment where educational content is voluntarily provided by the student themselves. This framework formalizes the under-provision problem that arises and quantify the extent of free-riding behaviour. For simplicity, I model only the intensive margin contributions activity.

6.1 Environment Setup

Consider N students who have access to the online peer discussion board. They can participate by writing content (w) in the form of questions (w^q) and answers (w^a), where $w_i = w_i^q + w_i^a \geq 0$. Then, $W = \sum_{i=1}^{N} w_i$ represent all content written to the online forum, and similarly define $W_{-i} = \sum_{j \neq i} w_j$.

Students can vary in their human capital (h) and altruism (a).[14] We normalize the lowest types

[14] Allowing students to be heterogenous in their altruism and human capital is motivated by the data as they are key determinants of discussion board participation.

as $h_L = 1$ (i.e., GPA of 1 out of 4) and $a_L = 0$ (i.e., purely egoistic), then define higher types as $h = \gamma_h h_L, a = a_L + \gamma_a$, where $\gamma_h \in (1,4]$ and $\gamma_a \in (0,1]$.

6.2 Contributing to the Discussion Board

Students can learn from each other through participating in the discussion board. Human capital enters the model through a learning production function to account for the quality of contributions: $L_i(w_i, w_{-i}; h_i, h_{-i}) = L_i(w_i h_i + \sum_{j \neq i} w_j h_j)$, where a positive learning externality is incurred through the aggregate human capital-adjusted contributions of other students, $W^h_{-i} = \sum_{j \neq i} w_j h_j$.[15] I assume the learning production is concavely increasing ($\frac{dL_i}{dw_i} > 0$ and $\frac{d^2 L_i}{dw_i^2} < 0$). The student's altruism governs the extent to which the spillover benefit their contribution has on other students is internalized. That is, a student with $a_i > 0$ gains the following benefit from participation:

$$\underbrace{(1-a_i)L_i}_{\text{egotism-weighted own learning}} + \underbrace{a_i(\sum_{j \neq i} L_j + L_i)}_{\text{altruism-weighted total learning}} = \underbrace{L_i}_{\text{own learning}} + \underbrace{a_i \sum_{j \neq i} L_j}_{\text{altruism-weighted other learning}} .$$

Each student has a private cost to writing content that is represented by $C_i(w_i)$. I assume the cost function is convex in contributions ($\frac{dC_i}{dw_i} > 0$ and $\frac{d^2 C_i}{dw_i^2} \geq 0$). The surplus of student i from contribution w_i posts to the discussion board is therefore:

$$S_i(w_i, w_{-i}; h_i, a_i, h_{-i}) = L_i(w_i h_i + \sum_{j \neq i} w_j h_j) + a_i \sum_{j \neq i} L_j(w_i h_i + \sum_{j \neq i} w_j h_j) - C_i(w_i).$$

Taking the contributions of others as given, a student will select their contribution amount to maximize their own surplus.

Equilibrium Concept. Given N students that have registered for the discussion board, the Nash Equilibrium (NE) of this simultaneous provision of public goods game is collection of contributions $w^* = (w_1^*, w_2^* \ldots, w_N^*)$ such that no student has incentives to deviate. That is, no student i can improve their surplus by deviating from w_i^* given the remaining $N-1$ students contribute according to $w_{-i}^* : S_i(w_i^*, w_{-i}^*) \geq S_i(w_i, w_{-i}^*)$ for all $w_i, i \in \{1, 2, \ldots, N\}$. The equilibrium is achieved through

[15] In the canonical model of voluntary public goods contribution, total provision is defined as the unweighted total sum of all agents' monetary contributions (e.g., charitable donations). In contrast, the online context considered here involves non-monetary contributions, and so a quality element is incorporated to define the total value of the public good as $W^h = \sum_{i=1}^n f_i(w_i) = \sum_{i=1}^N h_i w_i$.

solving the following set of N first order conditions simultaneously:

$$\underbrace{\frac{dS_i}{dw_i} = 0}_{\text{first order condition}} \iff \frac{d}{dw_i}\left[L_i(w_ih_i + \sum_{j\neq i} w_jh_j) + a_i \sum_{j\neq i} L_j(w_ih_i + \sum_{j\neq i} w_jh_j) - C_i(w_i)\right] = 0$$

$$\iff \underbrace{h_i\frac{dL_i}{dw_i}}_{\text{own marginal benefit}} + \underbrace{a_i}_{\text{altruism weight}} \times \underbrace{h_i\sum_{j\neq i}\frac{dL_j}{dw_i}}_{\text{marginal spillover benefit}} = \underbrace{\frac{dC_i}{dw_i}}_{\text{own marginal cost}},$$

for $i = 1, \ldots, N$. Assuming a Cournot-Nash assumption and taking the contributions of other students as given, the FOC can characterize a student's best response function. Let $w_i^R(w_{-i}; a_i, h_i, h_{-i})$ represent the reaction function which satisfies the FOC for any given amount of contributions by others, w_{-i}, own type h_i, a_i, and human capital of remaining students h_{-i}. Then at the Nash Equilibrium w^*, all best response functions intersect, and so each students contribution w_i^* is a best response to the actions of other students w_{-i}^*.

The Existence of a Nash Equilibrium. In this framework, we can show the Nash equilibrium exists under some reasonable assumptions. In particular, assume students select their contributions from a closed interval: $w_i \in [0, w_i^{max}]$, where w_i^{max} is the most contributions a student can make given their study time budget. Furthermore, assume the learning production function, $L_i(w_ih_i, W_{-i}^h)$, is continuous, differentiable, and concavely increasing in own contribution w_i for any given aggregate human capital-adjusted contribution of other students W_{-i}^h. I also assume the cost function $C_i(w_i)$ is continuous, differentiable, and convexly increasing in own contribution w_i. Then for any given W_{-i}^h, we can show the surplus function, $S(w_ih_i, W_{-i}^h)$ is concave in own contribution w_i, as it is simply a sum of concave functions:

$$S_i(w_ih_i, W_{-i}^h) = \underbrace{L_i(w_ih_i + \sum_{j\neq i} w_jh_j)}_{\text{concave}} + \underbrace{a_i}_{>0} \underbrace{\sum_{j\neq i} L_j(w_ih_i + \sum_{j\neq i} w_jh_j)}_{\text{concave}} + \underbrace{(-C_i(w_i))}_{\text{concave}}.$$

This property is desirable as now each student has a well defined best response for any given human capital-adjusted contributions made by other students as illustrated in Figure 1.

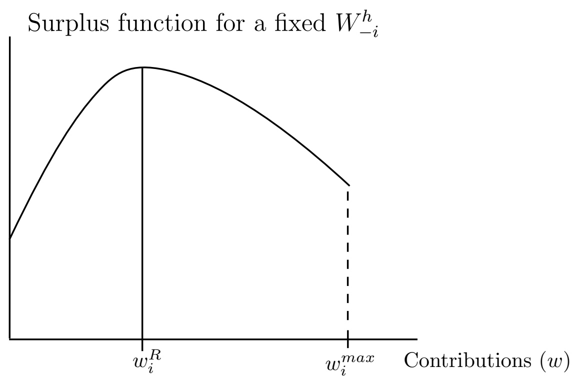

Figure 1: Illustrates concave surplus function on a closed interval

It can be shown that when the surplus function is concave on a closed and bounded set (i.e., compact), the then at least one Nash Equilibrium will be present (Osborne and Rubinstein, 1994). This is because under these conditions, the function that defines the collection of best response functions $(w_1^R, w_2^R, \ldots, w_N^R)$ has a fixed point.[16]

Uniqueness of Nash Equilibrium. Leveraging insights from Cornes and Hartley (2007) about the aggregative structure of public goods game, I can show the Nash equilibrium of the model is unique. The aggregative structure implies the reaction function for each student depends only on their own type, and the human capital-adjusted contribution of others: $w_i^R(\sum_{j\neq i} h_j w_j; a_i, h_i)$, for $i = 1, \ldots, N$. Furthermore, the reaction function strictly decreases in the human capital-adjusted contribution of others, i.e., $\frac{dw_i^R}{d(\sum_{j\neq i} h_j w_j)} < 0$. To show this, note that the reaction function w_i^R satisfies student $i's$ FOC for any given W^h_{-i}:

$$h_i L_i'(w_i^R, W^h_{-i}) + a_i h_i \sum_{j\neq i} L_j'(w_i^R, W^h_{-i}) - C_i'(w_i^R) = 0.$$

Implicitly differentiating the above equation with respect to W^h_{-i} results in the following:

[16] Let $X = \{w \in R^N : w_i \in [0, w_i^{max}]$ for i $=1, 2, \ldots, N\}$ represent the collection of possible allocations. Then define a function $B : X \to X$ which includes a mapping for each student that represents their reaction function. Under the concavity of the surplus function, this will have a fixed point, $B(w^*) = w^*$, which is the Nash-equilibrium vector of student contributions.

$$\frac{dw_i^R}{dW_{-i}^h} = \frac{\overbrace{-h_i(L_i'' + a_i \sum_{j \neq i} L_j'')}^{>0}}{\underbrace{h_i^2 L_i''}_{<0} + \underbrace{a_i h_i^2 (\sum_{j \neq i} L_j'')}_{<0} + \underbrace{(-C_i'')}_{<0}} < 0,$$

since by assumption $a_i \geq 0, h_i > 0, L_i'' < 0$ (concave learning production), and $C_i'' > 0$ (convex costs). That is, the reaction function is downward sloping with respect to W_{-i}^h, and its steepness depends on the student's type. At the equilibrium, the remaining students will also best respond to w_i resulting in $(W_{-i}^h)^R(w_i) = \sum_{j \neq i} h_j w_j^R(w_i)$. This aggregate reaction function is strictly decreasing in w_i since we can show by a similar argument that each component of the summation $\frac{dw_j^R}{dw_i} < 0, j \neq i$. It can be shown that for $N > 2$, the aggregate reaction function is steeper then $i's$ reaction function, i.e., $|((W_{-i}^h)^R)'(w_i)| > |(w_i^R)'(W_{-i}^h)|$. Then, as illustrated in Figure 2, there will exist a unique solution $(w_i^*, (W_{-i}^h)^*)$.[17] Since this pairing must be unique for each student i, the Nash equilibrium $(w_1^*, w_2^*, \ldots, w_N^*)$ is unique.

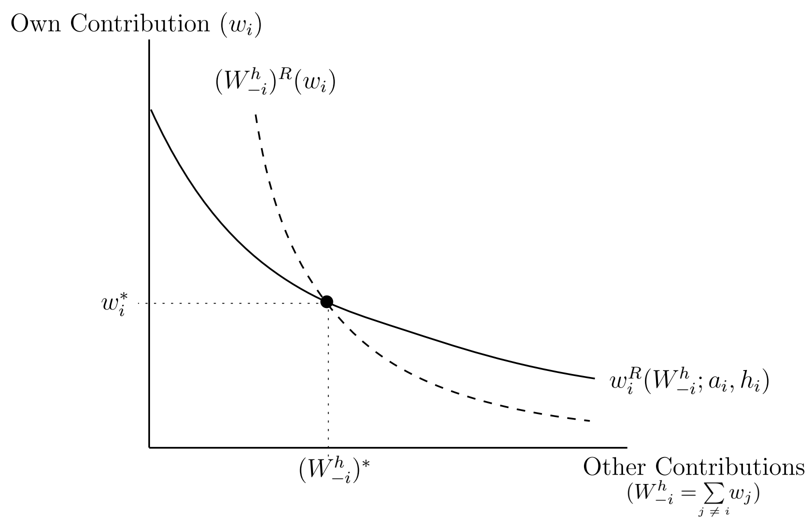

Figure 2: Unique of the equilibrium illustrated as the intersection of a student's reaction curve and the aggregate reaction of the remaining students

[17] Infinite number of solutions are possible if all students are of the same type, but this is not practical in this setting. The possibility of no solution is also ruled out as we previously showed at least one equilibrium exists.

6.3 Implications of the Voluntary Contributions Model.

This model can be used to describe how free-ridership incentives and composition of student types affects overall learning for those registered in the discussion board. I define a utilitarian benchmark for desirable participation activity as students collectively cooperating to maximize total student learning net of the total participation time:

$$\max_{w_1 \geq 0, \ldots, w_N \geq 0} \{\sum_{i=1}^{N} L_i(w_i, w_{-i}; h_i, h_{-i}) - \sum_i C_i(w_i)\}.$$[18]

Let $w^{*,c} = (w_1^{*,c}, w_2^{*,c}, \ldots, w_N^{*,c})$ represent the resulting cooperative allocation where students completely internalize the learning externalities their contributions have on others. Relative to the Nash equilibrium $w^* = (w_1^*, w_2^*, \ldots, w_N^*)$, we can show that each student at least increases their contribution under the cooperative allocation.[19] Then, the severity of free-ridership for a student is naturally defined as extent to which the Nash contribution deviates from the cooperative amount: $\tau_i^{FR} = \frac{w_i^{*,c} - w_i^*}{w_i^{*,c}} I(w_i^{*,c} > 0)$.

Total learning that results from the peer discussion board depends on the type of students; altruism influences the quantity of contribution, and human capital governs the quality. It can be shown that human capital and altruism are complementary in the sense that overall student learning in the decentralized environment is increased, as all students become higher human capital and more altruistic. However, it is possible for altruistic students to crowd-out the contributions of higher human capital students to reduce overall learning.

The following Proposition summarizes the consequences of reducing free-riding behavior.

Proposition 1 *The Nash equilibrium is not Pareto efficient when students do not internalize the externalities of their contributions. The efficient allocation involves no free-riding. However, reducing the extent of free-riding through recruiting altruistic, but low human capital students can increase provision at the cost of reducing overall learning.*

Proposition 1 implies that eliminating free-riding incentives through promoting further cooperation by each student will result in higher average learning per student. However, policies that aim to

[18]This can also be interpreted as the planners problem, which in this context would be the course instructor.

[19]Students contribute until their marginal benefit of doing so exceeds their marginal cost. Under the cooperative objective, marginal benefit is at least (strictly if not a pure altruist) increased for each student while marginal costs are unchanged, and hence the amount of contributions at least increases.

increase participation activity by recruiting altruistic, but low human capital students can harm overall learning, as the average quality of contribution will decrease.

Inefficiency of Decentralized Allocation. The decentralized Nash allocation w^* results from each student maximizing their own surplus: $\max_{w_i \geq 0} S_i(w_i, w_{-i}; h_i, a_i, h_{-j})$. This is not Pareto-efficient as we can show that collective allocation $w^{*,c}$ is Pareto dominant. Firstly, notice since the learning production is concave, students provide more when cooperating relative to the decentralized solution as illustrated in Figure 3. Therefore, total provision of the public good increases under the collective allocation. We can show individual student surplus increases across w^* and $w^{*,c}$ as follows:

$$\Delta S_i = \frac{dL_i}{dw_i}\Delta W + a_i \Delta W \sum_{j \neq i} \frac{dL_j}{dw_i} - \Delta w \frac{dC_i}{dw_i} = \Delta W(\frac{dL_i}{dw_i} + a_i \sum_{j \neq i} \frac{dL_j}{dw_i}) - \Delta w \frac{dC_i}{dw_i} = \frac{dC_i}{dw_i}(\Delta W - \Delta w) > 0.$$

Intuitively, $w^{*,c}$ results in a higher surplus than w^* as the spill-over learning benefit resulting from the increased total provision is superior to the increase in own costs.

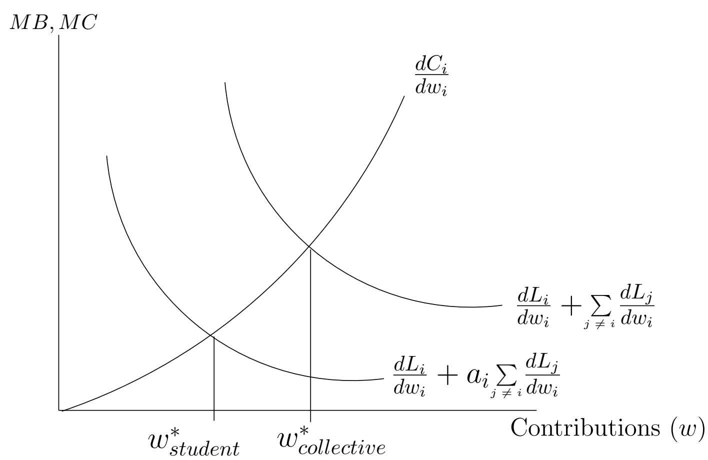

Figure 3: Comparing contributions under decentralized and cooperative settings

No-free riding when externalities internalized: Clearly, there is no free-riding (i.e., $\tau_i^{FR} = 0$) when externalities from contributions are completely internalized as the Nash equilibrium equals the cooperative equilibrium when all students are pure altruists (i.e., $a_i = 1$).

Crowd-out due to altruism: We can show that student $i's$ contribution strategically decreases in the contribution of others, i.e., $\frac{dw_i^R}{dw_j} < 0, j \neq i$. We can derive this from the FOC of student i as follows:

$$\frac{dw_i^R}{dw_j} = \frac{\overbrace{-h_i h_j (l_i'' + a_i l_j'')}^{>0}}{\underbrace{l_i''(h_i)^2}_{<0} + \underbrace{a_i h_i^2 (\sum_{j \neq i} l_j'')}_{<0} + \underbrace{(-c_i'')}_{<0}} < 0.$$

Furthermore, we can also show that more altruistic students react by providing more:

$$\frac{dw_i^R}{da_i} = \frac{\overbrace{-(\sum_{j \neq i} l_j') h_i}^{<0}}{\underbrace{l_i''(h_i)^2}_{<0} + \underbrace{a_i h_i^2 (\sum_{j \neq i} l_j'')}_{<0} + \underbrace{(-c_i'')}_{<0}} > 0.$$

Since students strategically substitute their contribution efforts for the contributions of others, altruistic behaviour crowds-out the contributions from other students:

$$\frac{dw_j^R}{da_i} = \underbrace{\frac{dw_j^R}{dw_i}}_{<0} \underbrace{\frac{dw_i^R}{da_i}}_{>0} < 0.$$

Learning Consequences Due to Altruistic Crowd-Out. Let total learning be represented by $L = \sum_{i=1}^{n} L_i(w_i, W_{-i}; a_i, h_i, h_{-i})$. Then we can represent the change in total learning with a student becoming more altruistic as:

$$\frac{dL}{da_i} = \underbrace{\underbrace{\frac{dL_i}{dw_i}\frac{dw_i}{da_i}}_{\Delta \text{ own contribution}>0} + \underbrace{\frac{dL_i}{dW_{-i}} \sum_{j \neq i} h_j \frac{dw_j}{da_i}}_{\Delta \text{ other spillover } <0}}_{\Delta \text{ own learning } \lesseqgtr 0} + \underbrace{\underbrace{\sum_{j \neq i} \frac{dL_j}{dw_j}\frac{dw_j}{da_i}}_{\Delta \text{ other contribution}<0} + \underbrace{\frac{dL_j}{dW_{-j}}[\frac{dw_i}{da_i} h_i + \sum_{k \neq i,j} \frac{dw_k}{da_i} h_k]}_{\Delta\text{spillover } \lesseqgtr 0}}_{\Delta \text{ other learning } \lesseqgtr 0} \lesseqgtr 0.$$

Intuitively, an altruistic student can crowd-out contributions of higher human capital students, which can decrease overall learning.

6.4 Incorporating the Writing Activity Nudge

Now we integrate the self-reflecting writing activity nudge discussed in the previous section into this framework. Let $z_i \in \{0, 1\}$ be an indicator variable representing whether student i is in the control group ($z_i = 0$) or is in the treatment group ($z_i = 1$).

The self-reflection writing exercise was aimed to reduce free-ridership and promote active participation. The nudge involved reading and reflecting on real quotes from a previous cohort of students on the importance of peer discussion for improving the overall learning experience for all registered students (as mentioned). This nudge can further the extent to which externalities are internalized by changing the weight placed on the learning of the other students: $\min\{a_i(1+\lambda), 1\}$. Incorporating the nudge, the surplus of student i now:

$$S_i = L_i + \min\{a_i(1 + z_i\lambda), 1\} \sum_{j \neq i} L_j - C_i,$$

where z_i. Proposition 2 formulates hypothesis for how the writing activity nudge will influence contribution activity.

Proposition 2

Having students reflect on spillover benefits of their contributions increases participation for students who not completely altruistic (i.e., $a_i < 1$).

The Writing Activity Increases the Marginal Benefit to Contributing The self-reflection nudge increases the extent to which the externality is internalized by λ. Consequently, as illustrated in Figure 4, the marginal benefit curve shifts outwards. Given the marginal time cost of participation is not affected by the self-reflection nudge, the contributions in equilibrium increase.

7 Empirical Results

This section discusses the results from the fields experiments described in Section 5, and outlines the corresponding empirical methodology.

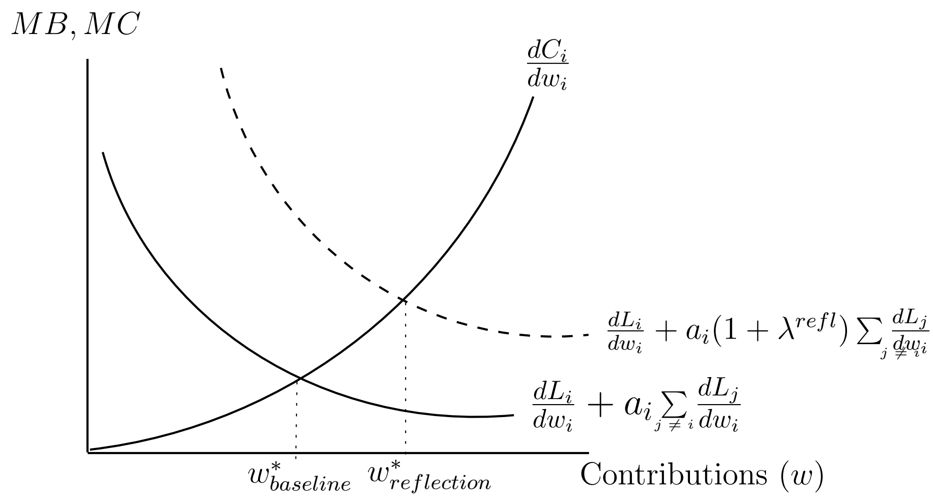

Figure 4: Illustrates how self-reflection nudge affects contributions

7.1 The Effect of Sign-up Activity on Discussion Board Accessibility

To measure the effect of the sign-up activity on discussion board accessibility, I estimate the following specification:

$$Y_i = \alpha_0 + \alpha_1 SignupActivity_i + X_i'\gamma + \epsilon_i,$$

where Y_i denotes either an indicator for signing up for the discussion board, or the number of weeks a student utilizes the online peer forum; $SignupActivity_i$ is an indicator denoting whether a student receives the sign-up activity nudge. Control variables in X_i include student demographics and pre-treatment characteristics listed in Panels A and B of Table 4.

Figure 10 illustrates the effect of being assigned the sign-up activity nudge on discussion board utilization. Showing clear registration instructions and/or providing information to students about the discussion board increased utilization by around 4.5 weeks relative to the control group. The magnitude is large when considering that none of the students who were eligible for treatment had registered for the discussion board at the start of the course, and the intervention more than doubled the duration of utilization (i.e. from around 3 weeks to 7.5 weeks).[20] This effect size is also statistically significant at the 1% significance level with an F-statistic exceeding 100. Table

[20] Among students that sign up for the discussion board, only 3% are inactive. The vast majority of students registered to the discussion board at least view a few posts each week – behavior that can be tracked by the researcher.

8 shows that the sign-up activity increases discussion board registration by end of the course by 17 percentage points. Additionally, the treatment effects are also stable following the inclusion of pre-treatment control variables and cohort fixed effects. Looking ahead, assignment to the sign-up activity nudge will provide a strong first stage for inducing exogenous variation in discussion board utilization.

Mechanisms Underlying the Sign-up Activity. I investigate the mechanisms underlying the sign-up activity by examining how the nudge affects the extent to which different types of students are informed about the discussion board. I investigate whether the sign-up activity informs inattentive students about the online discussion board. To do so, I use student responses from a question that was embedded at the end of the baseline survey (i.e., post sign-up activity) which elicited whether students are well informed about the existence and functionality of the discussion board. Figure 11 displays the propensity to be well informed about the discussion board as a function of students' attentiveness by treatment status. Clearly, the nudge was successful in closing the information gap between inattentive and attentive students.

Next, I investigate whether the component of the sign-up activity that informs students about the functionality of the discussion board induces utilization more than when only registration instructions are provided. Figure 12 shows that receiving only the functionality information increases discussion board utilization by an extra week relative to receiving only registration instructions. (The increase in discussion board utilization for students who receive both the instructions and information is not statistically different from receiving the functionality information only.)

7.2 Effect of Reflective Writing Activity on Contributions

To measure the effect of the writing activity on student discussion board participation, I estimate the following specifications:

$$\begin{cases} log(\widetilde{Write}) = \beta_0 + \beta_1 WritingActivity_i + X_i'\alpha + \epsilon_i \\ log(Read_i) = \beta_0 + \beta_1 WritingActivity_i + X_i'\alpha + \epsilon_i \end{cases},$$

where $WritingActivity_i$ is an indicator for being assigned the writing activity, and $\widetilde{Write} =$

$Write_i + (Write_i^2 + 1)^{1/2}$ is the inverse hyperbolic sine transformation for contributions. This transformation is recommended in Bellemare and Wichman (2020) for data that is right skewed with a non-trivial portion of 0 value observations. Table 9 presents the results from estimating the above model. Columns 3 and 4 of the table show that the reflective writing exercise increased the number of posts written by around 26% and reading by 11%. Since reading naturally precedes writing in a discussion board setting, reading activity increases with contributions.

Figure 13 examines the value affirmations treatment effect more closely by illustrating distribution of log contributions across the control and treatment group. The p-value of 0.0083 resulting from the Kolmogorov–Smirnov equality of distributions test suggests the intervention significantly changed the distribution of contributions for the treatment group. That is, the reflective writing exercise mainly induced further content creation for those near the bottom left of the distribution who have modest contribution activity.

7.3 Effect of Reflective Writing Activity on Quality of Contributions

Next I evaluate whether the writing activity had any impacts on the quality of contributions made by estimating the following specification:

$$UpVotes_i = \beta_0 + \beta_1 WritingActivity_i + X_i'\alpha + \epsilon_i,$$

where $UpVotes_i$ is the total number of up votes received by student i on their contributions made to the discussion board.

7.4 Learning Value-Added from Discussion Board Utilization

Identifying the causal effect of discussion board utilization on final exam achievement is challenging, given students are likely to self-select into registering for the online discussion board based on their expected learning benefit. For example, attentive and motivated students are going to be more likely to utilize the discussion board and also likely to perform better on the final exam, resulting in upwards bias. To circumvent such endogeneity issues, I use the sign-up activity as an instrument for discussion board utilization on the extensive margin, and the writing activity as an instrument on the intensive margin.. These are valid instrumental variable as they are randomly assigned to students (i.e., it is independent), do not directly affect the exam grade (i.e., is excludable), and directly affects discussion board utilization (i.e., is relevant).

Extensive Margin Value Added. To evaluate the learning gains from discussion board utilization on the extensive margin, I estimate the following 2SLS model:

$$\begin{cases} A_i = \theta_0 + \theta_1 WeeksRegistered_i + X_i'\Theta + \epsilon_i \\ WeeksRegistered_i = \tau_0 + \tau_1 SignupActivity_i + X_i'T + \epsilon_i \end{cases},$$

where A_i denotes achievement outcome such as the final exam grade or the average homework grade; $WeeksRegistered_i$ is the number of weeks a student is registered for the discussion board.

Table 10 presents the results for the effect of discussion board accessibility on student achievement. The results indicate that an extra 5 weeks of discussion board accessibility increases mean homework and final exam achievement by 0.14 SD and 0.07 SD, respectively. The different effect sizes across the homework and final exam are consistent with the course rule that allows students to discuss the homework problems on the online peer discussion board, but that forbids students from discussing questions from their online exam with one another. These magnitudes are large, equivalent to increasing the course grade by half a letter grade (e.g., B+ to A-). A back-of-the-envelope calculation predicts that 31% of students in the control group would have received half a letter higher course grade had they also been assigned to the sign-up activity.

Intensive Margin Value Added. To evaluate the learning gains from discussion board utilization on the intensive margin, I estimate the following 2SLS model:

$$\begin{cases} A_i = \gamma_0 + \gamma_1 DiscussionMins_i + X_i'\Gamma + \epsilon_i \\ DiscussionMins_i = \pi_0 + \pi_1 WritingActivity_i + X_i'\Pi + \epsilon_i \end{cases},$$

where A_i denotes achievement outcome such as the final exam grade or the average homework grade; $DiscussionMins_i$ is the number of minutes students spend reading and writing posts on the discussion board (constructed as discussed in Section 4).

Table 11 presents the results for the effect of the time spent on the discussion board on student achievement. The results indicate that spending an extra hour participating in the discussing board increases mean homework and final exam achievement by 0.062 SD and 0.037 SD respectively. Consistent with the extensive margin results, participating in the discussion board has a bigger effect on homework achievement than the final exam.

7.5 Robustness to Spillover Effects

I present empirical evidence to support the view that the results presented in this section are not severely affected by spillover effects. Only around 13% of students in the final survey attested to discussing contents of either the sign-up activity or the writing activity with their peers. Then, the information spillovers from students in the treatment group to the control group are likely small. Additionally, around 19% of students in the course are involved in a study group where they study course content with their peers. I investigate whether the effect sizes for either the sign-up or writing activities vary according to whether students are in a study group. The results for this analysis are presented in Table 12. I find that the sign-up activity is less effective for students in a study group, although the relevant estimates are only marginally statistically significant. In contrast, the effect of the writing activity does not vary according to whether a student is in a study group.

8 Estimating a Model of Online Public Goods Provision

In this section, I describe the estimation of the model introduced in Section 6. The model takes advantage of the administrative data on contributions activity and baseline human capital, as well as survey data on students' altruistic attitudes.

8.1 Validating Theoretical Model Implications

The model introduced in Section 6 predicts that altruistic students contribute more to the discussion board, conditioning on their human capital which governs the marginal private learning benefit from contributing. To test whether this implications holds in the data, Figure 5 shows the relationship between the number of contributions made and a student's altruism after conditioning on baseline human capital.

The figure makes clear that on average, contribution activity increases with students' altruistic

Figure 5: Mean Contributions by Students' Altruism

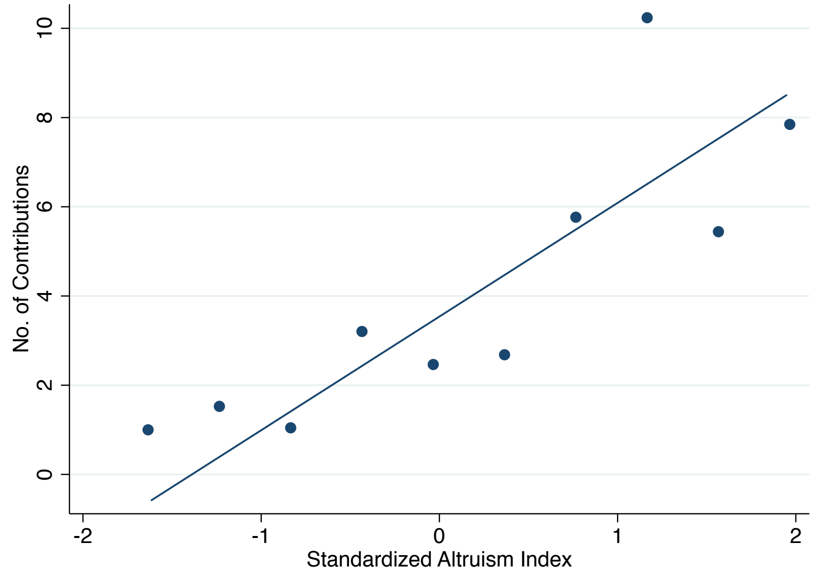

Notes: The figure presents a binned scatter plot showing the relationship between contribution activity and students' altruistic attitudes. The student altruism index is constructed using a series of survey questions (see Appendix B.1).

attitudes. This is consistent with the models prediction that altruistic students internalize the spill over benefit of their contribution to the learning of their peers.

8.2 Estimating the Model

To rationalize the free-ridership present in online public goods, I now estimate the model discussion in Section 6. In the model, students decide their contribution activity by balancing their benefits against the costs. Estimating the model will serve as foundation to conducting counterfactual analysis for measuring the extent of free-ridership present and policies for mitigating free-riding incentives.

8.3 Specifying the Learning Technology

The total benefit of contribution is composed of the private learning benefit and a public learning benefit for other students in the classroom. The learning technology for student i maps the human capital adjusted contributions in the discussion board to a learning benefit. While the true technology is unknown, I impose minimal structure on the learning technology to serve as a first-order approximation, using the following assumptions:

Assumption 1: The learning technology is concave

First, I assume the learning technology is concave. That is, there are diminishing learning returns from additional content being provided in the discussion board. This is reflective of the fact that in discussion board settings that the topics that students are commonly struggling with will be discussed, and more nice questions are likely to be posted and discussed later.

Assumption 2: Students benefit from the contributions of their peers

Second, I assume that students utilizing the discussion board will derive learning benefits from the posts written by their peers. Since the vast major of students in the course are new programmers, it is reasonable to assume they will gain some learning benefit from discussion about course content and homework problems that is contributed to the discussion board by other students.

Assumption 3: Some degree of substitutability between own contribution and peers' contribution is present

Third, I assume that own contributions and the contributions made by peers are substitutable to an extent. For example, if student i posts a question about a homework problem many students are struggling with, this will be substitute to an extent for other similar questions other students may have.

Considering the above assumptions, I specify the learning technology using a CES learning production function as follows:

$$L_i(w_i, \sum_{j \neq i} w_j; h_i, h_{-i}) = \left(\rho_1 (h_i w_i)^{\rho_2} + (1 - \rho_1)(\sum_{j \neq i} h_j w_j)^{\rho_2} \right)^{\frac{1}{\rho_2}},$$

where ρ_1 represents the share of learning benefit due to own contribution, and ρ_2 is the elasticity of substitution between own contributions and the contributions of peers.

8.4 Specifying the Cost Function

I impose a linearity assumption on the cost function for tractability.

Assumption 4: The cost function is the average time spent writing

The cost function is simply modelled as the time students spend on the discussion board:

$$C_i(w_i) = t_{i,w} w_i,$$

where $t_{i,w}$ is the average time it takes student i to write a question as informed by survey data. That is, the marginal cost of contribution is observed and there are no unknown cost parameters.

8.5 Estimating the Marginal Benefit of Contribution Parameters

The model parameters $\rho = (\rho_1, \rho_2)$ are estimated using data covering students' discussion board participation and their type for those who registered to the discussion board. The descriptive statistics for this data are shown below:

Variable	Mean	SD	Source
No. of total contributions (w)	3.55	14.09	Administrative
Prior cGPA (h)	3.21	0.711	Administrative
Altruism index (a)	0.61	0.17	Survey
Minutes required to write post (t^w)	5.62	2.75	Survey
Total available hours for studying (T)	71.81	42.01	Survey
No. of Students		694	

Denote the observed contributions as w_i and the model implied contribution as w_i^*. I estimate the marginal benefit of effort parameters $\rho = (\rho_1, \rho_2)$ using an iterative maximum likelihood estimation (MLE) routine. I can construct the log-likelihood using an 'implementations error' approach (Bernheim et al., 2019). That is, I assume students implement the model implied optimal contribution amount with error:

$$\underbrace{w_i}_{\text{Observed contributions}} - \underbrace{w_i^*(\rho)}_{\text{Optimal contributions from model}} \sim \underbrace{N(0, \sigma^2)}_{\text{Deviation from optimal distribution}}.$$

Then the resulting log-likelihood function is:

$$l(\rho; (w_i)_i) = -n \log(2\pi) + \frac{n}{2} \log(\sigma_\epsilon^2) - \frac{1}{2\sigma_\epsilon^2} \sum_{i=1}^{N} (w_i - w_i^*(\rho))^2.$$

The marginal benefit parameter estimates are then uncovered numerically by carrying out the following steps, iteratively maximizing the likelihood function:

1. Start with an initial value of $\tilde{\rho}$.

2. Compute $w_i^*(\tilde{\rho})$ by solving a system of FOCs for students $i = 1, \ldots, n$

3. Use $\tilde{\rho}, w_i^*(\tilde{\rho})$, and w_i for each student to compute likelihood $l(\tilde{\rho})$.

4. Update $\tilde{\rho}$ to $\tilde{\rho}'$ using Newton's method to take the next step.

5. Iterate through steps 2-4 until convergence: $|\tilde{\rho} - \tilde{\rho}'| < 10^{-6}$.

The estimation routine results in parameter estimates that maximize the likelihood function:

$$\widehat{\rho}_{MLE} = \arg\max_{\rho} l(\rho).$$

The marginal benefit parameters estimates are presented in the following table.

Table 1: Benefit of Writing Parameter Estimates (MLE)

Parameter	Estimate (SE)
$\widehat{\rho_1}$ (own contribution share of learning production)	0.0931** (0.0437)
$\widehat{\rho_2}$ (elasticity of substitution)	0.5273*** (0.1478)

Notes: Benefit parameters estimated using maximum likelihood estimation. Standard errors appear in parentheses and are calculated using the outer-product of gradients method. Significance levels are represented by $^{*}p < 0.1,^{**}p < 0.05,^{***}p < 0.01$.

8.6 Goodness of Fit

To evaluate whether the model fits the data, I start by comparing the distribution of the observed contributions to the contributions distribution predicted by the model.[21] Figure 14 shows clearly that the model fits the data well, as the mean and variance of contributions distribution implied by the model is closely aligned with the observed writing activity distribution. Additionally, I examine the association between observed contributions and the corresponding contributions implied by the model. Figure 15 shows that there is a strong linear relationship between the model implied writing activity and the actual writing activity. That is, on average, the model correct predicts the level of student writing activity.

8.7 Identification of Model Parameters

Following the best practices for structural research outlined in Andrews et al. (2020), I now discus the features of the data and sources of exogenous variation that drive the values of the param-

[21] The model implied contributions can be any real positive number given the continuous nature of the estimation routine. I round it to the nearest integer for investigating the goodness of fit.

eter estimates. I use the administrative discussion board participation data and the randomized reflective writing activity to identify the marginal benefit parameters.

Since the mapping between learning and contributing in the discussion board is endogenous, identifying the learning technology requires exogenous variation in writing activity. As the specified learning technology as two unknown parameters, then at least two moments of the data are required for identification. Such moments can be constructed using students' contribution activity under the control and treatment groups of the writing activity, holding constant baseline human capital. For purposes of the identification argument, if we consider the spillover benefits to be a component of the cost function, then the writing activity reduces cost to contributing while holding fixed the private marginal benefit of doing so.

9 Counterfactual Analysis

This section conducts two counterfactual experiments that would be impractical to implement in the field at a large scale. It does so using the estimated marginal benefit and cost of contribution parameters from the behavioural model of online public good provision. First, I am simulate the human capital adjusted contributions made to the discussion board when all students cooperate and internalize the spillover benefits of their contributions. Next, I simulate the effects of providing 0.5% bonus credit for making high quality contributions. For both counterfactual exercises, I then compare the simulated final exam grade under the counterfactual regime to the final exam grade achieved in the actual course. Both simulations together are informative of the extent of free-ridership present and the efficacy of providing extrinsic incentives to combat free-riding.

9.1 Measuring the Extent of Free-Riding

To measure the extent of free-riding present in online peer discussion boards, I will first simulate students' contributions assuming they are all purely altruistic. That is, I set $a_i = 1$ for all students in the counterfactual. Then every student internalizes the positive spill over learning benefits from contributing and contributes according to the following system of equations:

Table 2: Contributions and Learning Under Cooperative Allocation

Statistic	Data	Simulation	Difference
Average contributions	3.62	4.38	0.69
Average final exam grade	81.74	82.83	1.09

Table 3: Contributions and Learning Under Incentivized Allocation

Statistic	Data	Simulation	Difference
Average contributions	3.62	4.27	0.65
Average final exam grade	81.74	82.81	1.07

$$\underbrace{h_i \frac{dL_i}{dw_i}}_{\text{own marginal benefit}} + \underbrace{h_i \sum_{j \neq i} \frac{dL_j}{dw_i}}_{\text{marginal spillover benefit}} = \underbrace{\frac{dC_i}{dw_i}}_{\text{own marginal cost}},$$

for $i = 1, \ldots, N$. The following table compares the simulated cooperative allocation to the observed data.

The results from Column 3 indicate that the average student contributions increase by 21% and the average final exam graded increases by around 1.1 percentage points. That is, substantial free-ridership is present in the online discussion board and results in less learning for students utilizing the discussion board.

9.2 Providing Bonus Credit for High Quality Contributions

In the setting considered, contributions can receive public endorsements from the instructor. That propensity to receive a instructor endorsement for a contribution is 21%. That is, around 1 in 5 contributions are endorsed by the instructor as being of high quality. Let $p(h_i)$ be a non-parametric function that represents the propensity of having a contribution endorsed as a function of a students baseline human capital. In the data, it is seen that students with higher human capital have a higher propensity to receive endorsements. Suppose students are given 0.5% bonus credit for each endorsement they receive from the instructor. Then, on average, a student's marginal benefit increases by $0.5 \times p(h_i) \times w_i$. The following table compares simulated allocation where students receive 0.5% bonus credit for instructor endorsement to the observed data.

The results from Column 3 indicate that the average student contributions increase by 18%.

Further investigation shows not present here shows that the rise in contribution activity is mostly from higher human capital students, as they are more capable of making high quality contributions. The final exam grade increases by almost 1.1 percentage points when students are incentivized for good quality contributions.

10 Conclusion

The number of people participating in online learning communities has been expanding over the last two decades, and dramatically so in the two last years in response to the COVID-19 pandemic. Although these communities can provide free valuable information that is easily accessible to anyone with internet access, their success is reliant on active contributors who are voluntarily creating content. The presence of free-riding incentives can reduce the overall learning value-added of the online community relative to when all users internalize the spill over benefits of their contributions. In light of these prevalent challenges, this paper presented a microeconomic approach to measure free-riding and investigate policies that encourage cooperation in the context of online student discussion boards.

I presented three main findings that help to inform the value-added of student discussion boards and policies to mitigate the obvious free-riding incentives. First, the discussion board serves as an effective tool for supporting student learning: an additional five-weeks of utilization increases final exam grades by 0.08 SD. Second, there is substantial free-riding as the overall contributions to the discussion board would increase by 21% under the scenario in which all students cooperate and internalize the spill over learning benefits to their contributions. Third, I found evidence for policies that can successfully encourage contributions and mitigate free-riding. In particular, having students reflect on the spill over benefits of their contributions through a written activity, or providing sufficient bonus credit for making good quality posts both increase contributions to the discussion board. Interventions that target specifically low altruism and high human capital students will result in the largest achievement gains for the learning community.

While this paper considers student participation in an online peer discussion board, there are good reasons to think that the findings can be generalized to other online learning communities. The paper's findings contribute to our understanding of how to support the learning of students at scale, and how to keep students engaged in a voluntarily learning environment. Heterogeneity of participants will naturally arise in any large online learning environment. Then given the presence of less altruistic users, free-riding is to be expected on massive online communities where participation is voluntary. Most online environments are capable of showing users prompts to reflect on the value of being an active contributor, or publicly endorsing the contributions of their active users. Then the free-riding mitigation policies suggested in my paper can also be deployed in many other online public good settings.

References

AlJeraisy, Mashael Nasser, Heba Mohammad, Ayham Fayyoumi, and Waleed Alrashideh (2015) "Web 2.0 in education: The impact of discussion board on student performance and satisfaction.," *Turkish Online Journal of Educational Technology-TOJET*, 14 (2), 247–258.

Allcott, Hunt and Dmitry Taubinsky (2015) "Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market," *American Economic Review*, 105 (8), 2501–38.

Andrews, Isaiah, Matthew Gentzkow, and Jesse M Shapiro (2020) "Transparency in structural research," *Journal of Business & Economic Statistics*, 38 (4), 711–722.

Augenblick, Ned (2018) "Short-Term Time Discounting of Unpleasant Tasks," working paper, http://faculty.haas.berkeley.edu/ned/Augenblick_ShortTermDiscounting.pdf.

Baer, Amy and Andrew DeOrio (2020) "A longitudinal view of gender balance in a large computer science program," in *Proceedings of the 51st ACM Technical Symposium on Computer Science Education*, 23–29.

Bellemare, Marc F and Casey J Wichman (2020) "Elasticities and the inverse hyperbolic sine transformation," *Oxford Bulletin of Economics and Statistics*, 82 (1), 50–61.

Bernheim, B Douglas, Stefano DellaVigna, and David Laibson (2019) *Handbook of Behavioral Economics-Foundations and Applications 2*: Elsevier.

Bowen, José Antonio (2012) *Teaching naked: How moving technology out of your college classroom will improve student learning*: John Wiley & Sons.

Card, David, Stefano DellaVigna, and Ulrike Malmendier (2011) "The role of theory in field experiments," *Journal of Economic Perspectives*, 25 (3), 39–62.

Chen, Yan, Rosta Farzan, Robert Kraut, Iman YeckehZaare, and Ark Fangzhou Zhang (2018) "Motivating Contributions to Public Information Goods: A Personalized Field Experiment on Wikipedia."

Chen, Yan, F Maxwell Harper, Joseph Konstan, and Sherry Xin Li (2010) "Social comparisons and contributions to online communities: A field experiment on movielens," *American Economic Review*, 100 (4), 1358–98.

Cheng, Cho Kin, Dwayne E Paré, Lisa-Marie Collimore, and Steve Joordens (2011) "Assessing the effectiveness of a voluntary online discussion forum on improving students? course performance," *Computers & Education*, 56 (1), 253–261.

Cohen, Anat, Udi Shimony, Rafi Nachmias, and Tal Soffer (2019) "Active learners' characterization in MOOC forums and their generated knowledge," *British Journal of Educational Technology*, 50 (1), 177–198.

Cornes, Richard and Roger Hartley (2007) "Aggregative public good games," *Journal of Public Economic Theory*, 9 (2), 201–219.

Damgaard, Mette Trier and Helena Skyt Nielsen (2018) "Nudging in education," *Economics of Education Review*, 64, 313–342.

Deming, David J, Claudia Goldin, Lawrence F Katz, and Noam Yuchtman (2015) "Can online learning bend the higher education cost curve?" *American Economic Review*, 105 (5), 496–501.

Edelman, Benjamin (2012) "Earnings and ratings at Google Answers," *Economic Inquiry*, 50 (2), 309–320.

Escueta, Maya, Vincent Quan, Andre Joshua Nickow, and Philip Oreopoulos (2017) "Education technology: An evidence-based review."

Falkinger, Josef (1996) "Efficient private provision of public goods by rewarding deviations from average," *Journal of Public Economics*, 62 (3), 413–422.

Gallus, Jana (2017) "Fostering public good contributions with symbolic awards: A large-scale natural field experiment at Wikipedia," *Management Science*, 63 (12), 3999–4015.

Groves, Theodore and John Ledyard (1977) "Optimal allocation of public goods: A solution to the" free rider" problem," *Econometrica: Journal of the Econometric Society*, 783–809.

Harackiewicz, Judith M and Stacy J Priniski (2018) "Improving student outcomes in higher education: The science of targeted intervention," *Annual review of psychology*, 69, 409–435.

Kizilcec, René F, Justin Reich, Michael Yeomans et al. (2020) "Scaling up behavioral science interventions in online education," *Proceedings of the National Academy of Sciences*, 117 (26), 14900–14905.

Oechssler, Joerg, Andreas Reischmann, and Andis Sofianos (2022) "The conditional contribution mechanism for repeated public goods–The general case," *Journal of Economic Theory*, 203, 105488.

Rondeau, Daniel, William D Schulze, and Gregory L Poe (1999) "Voluntary revelation of the demand for public goods using a provision point mechanism," *Journal of public Economics*, 72 (3), 455–470.

Tabord-Meehan, Max (2018) "Stratification trees for adaptive randomization in randomized controlled trials," *arXiv preprint arXiv:1806.05127*.

Taubinsky, Dmitry and Alex Rees-Jones (2018) "Attention variation and welfare: theory and evidence from a tax salience experiment," *The Review of Economic Studies*, 85 (4), 2462–2496.

Thaler, Richard H (2018) "From cashews to nudges: The evolution of behavioral economics," *American Economic Review*, 108 (6), 1265–87.

Xu, Lei, Tingting Nian, and Luis Cabral (2020) "What makes geeks tick? a study of stack overflow careers," *Management Science*, 66 (2), 587–604.

Tables

Table 4: Student Level Summary Statistics

Variable	Mean	Std. Dev.	Min.	Max.	N
Panel A: Demographics					
I(Female)	0.492	0.499	0	1	1186
I(Lives outside campus)	0.771	0.420	0	1	1186
I(First year of university)	0.558	0.496	0	1	1186
I(Speaks english at home)	0.252	0.434	0	1	1186
I(Mother at least college graduate)	0.635	0.481	0	1	1186
I(Father at least college graduate)	0.728	0.444	0	1	1186
Panel B: Other Characteristics					
Prior programming experience (9-points scale)	2.689	1.675	1	9	1186
Hours studied per week for this course	4.511	2.53	0	30	1186
I(Has undergraduate academic mentor)	0.301	0.458	0	1	1186
I(Course required for major)	0.612	0.487	0	1	1186
I(Pursuing STEM major)	0.674	0.469	0	1	1186
No. other courses taken with discussion board	2.359	1.571	0	9	1186
Panel C: Discussion Board Participation					
I(Registered for course discussion board)	0.585	0.493	0	1	1186
No. of total contributions	3.556	14.087	0	237	694
No. of questions posted	1.399	3.107	0	41	694
No. of answers posted	1.081	9.144	0	161	694
No. of instructor endorsed answers	0.531	5.094	0	115	694
No. of posts viewed	121.882	149.279	0	1022	694
No. of days online	20.569	21.16	0	130	694

Notes: Table presents descriptive statistics related to study student demographic, characteristics, and discussion board participation and activity outcomes. Statistics shown in Panel A and B are formulated using self-reported student responses on the baseline survey. Finally the statistics shown in Panel C are computed using data gathered from the discussion board.

Table 5: Assignment of Students to Sign-Up Activity Control and Treatment Groups

Assignment Group	Number of Students	Percent of Students
Control	525	49.8
Instructions only	176	16.7
Information only	175	16.6
Both instructions and information	179	16.9

Table 6: Assignment of Students to Writing Activity Control and Treatment Groups

Assignment Group	Number of Students	Percent of Students
Control	298	48.8
Treatment	313	51.1

Table 7: Student Attrition and Treatment Allocation

	(1) I(Dropped course)	(2) I(Dropped course)	(3) I(Dropped course)	(4) I(Dropped course)
I(Assignment to sign-up activity)	0.0133		-0.023	
	(0.0427)		(0.0161)	
I(Assignment to writing activity)		0.0148		0.0052
		(0.0578)		(0.0116)
Controls	No	No	Yes	Yes
No. of Students	1102	758	1102	758
R-squared	0.0017	0.0014	0.14	0.12

Notes: Table shows differences in attrition rate across the control and treatment group for both interventions considered in this study. Controls include pre-treatment student demographics and characteristics included in Panels A and B of Table 1. Only students who did not register for the discussion board prior to sign-up activity were eligible to receive the treatment. Only students who signed up for the discussion board are eligible to receive the writing activity. Significance levels are represented by $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$.

Table 8: Effect of Sign-up Activity on Discussion Board Utilization

	(1) I(Signed Up)[a]	(2) I(Signed Up)	(3) Weeks Registered[b]	(4) Weeks Registered
I(Receives sign-up activity)	0.193***	0.186***	4.684***	4.329***
	(0.0284)	(0.0294)	(0.2174)	(0.2284)
Control Mean	0.48	0.48	3.152	3.152
Controls	No	Yes	No	Yes
Adjusted R-square	0.137	0.239	0.257	0.348
F-stat for treatment	33.09	32.15	210.13	211.34
No. of Students	1055	1055	1055	1055

Notes: [a]The outcome variable is whether students signed up for the course discussion board by end of the course. [b]Outcome is the number of weeks the student is registered for the discussion board. Student demographics and other characteristics include variables present in Panel A and B in Table 1 respectively. Only students who had not registered for the discussion board prior to the baseline survey were eligible for treatment. Significance levels are represented by $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$.

c

Table 9: Effect of Writing Activity on Reading and Writing

	(1) Writing[a]	(2) Log(Reading)[b]	(3) Writing[b]	(4) Log(Reading)
I(Receives writing activity)	0.2657***	0.1183***	0.253***	0.1274***
	(0.0398)	(0.0483)	(0.1374)	(0.0484)
Controls	No	Yes	No	Yes
Adjusted R-square	0.084	0.293	0.248	0.318
No. of Students	611	611	611	611

Notes: [a]The outcome variable is the hyperbolic since inverse of the number of contributions made.. [b]Outcome is the logarithm of the number of unique posts read. Student demographics and other characteristics include variables present in Panel A and B in Table 1 respectively. Only students who had registered for the discussion board eligible for treatment. Significance levels are represented by $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$.

Table 10: Effect of Discussion Board Utilization on Learning Outcomes (2SLS)

	(1) Exam Performance[a]	(2) Exam Performance	(3) Homework Performance[b]	(4) Homework Performance
No. of weeks registered	0.014**	0.016**	0.026***	0.027***
	(0.0064)	(0.0067)	(0.0087)	(0.0078)
Controls	No	Yes	No	Yes
Adjusted R-square	0.083	0.195	0.117	0.382
No. of Students	1055	1055	1055	1055

Notes: [a]Standardized final exam grade. [b]Standardized average homework performance all homeworks. Assignment to the sign-up activity is used as an instrument for the number of weeks registered. Controls include pre-treatment student demographic and characteristics included in Panels A and B of Table 1. Significance levels are represented by $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$.

Table 11: Student Online Learning Participation and Final Exam Grade (2SLS)

	(1) Homework Performance[a]	(2) Homework Performance	(3) Exam Performance	(4) Exam Performance
Discussion Board Study Time (Hours)[c]	0.067**	0.067**	0.032*	0.037**
	(0.0321)	(0.0311)	(0.0171)	(0.0173)
Controls	No	Yes	No	Yes
Adjusted R-square	0.087	0.191	0.081	0.221
No. of Students	611	611	611	611

Notes: [a]Standardized average homework grade across all homeworks. [b]Standardized final exam grade [c]Time spent reading and writing posts on the discussion board. Assignment to the writing activity is used as an instrument for time spent on the discussion board. Controls include pre-treatment student demographic and characteristics included in Panels A and B of Table 1. Significance levels are represented by $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$.

Table 12: Efficacy of Nudges and Study Group Involvement

	(1) I(Signed Up)	(2) I(Signed Up)	(3) No. of Contributions	(4) No. of Contributions
I(Study group) × I(Receives sign-up activity)	-0.0212*	-0.0193*		
	(0.0123)	(0.0104)		
I(Study group) × I(Receives writing activity)			0.021	0.027
			(0.0335)	(0.0257)
Controls	No	Yes	No	Yes
Adjusted R-square	0.124	0.274	0.082	0.196
No. of Students	1055	1055	611	611

Notes: Indicator for whether a student is in a study group and treatment status are also included in the estimation. Student demographics and other characteristics include variables present in Panel A and B in Table 1 respectively. Only students who had not registered for the discussion board prior to the baseline survey were eligible for sign-up activity. Only students who signed up to the discussion board are eligible to receive the writing activity. Significance levels are represented by $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$.

Figure 6: Discussion Board Visit Frequency

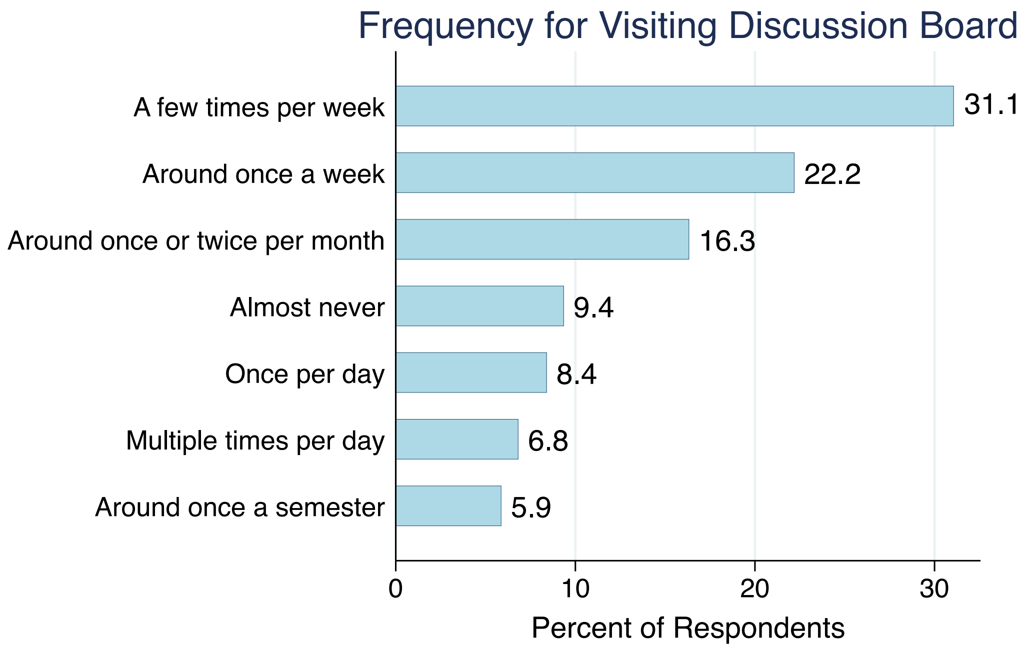

Notes: The figure shows the distribution of student responses to a survey question that asks them about the extent to which they visit the discussion board on a weekly basis.

Figures

Figure 7: Student self-reported discussion board usage

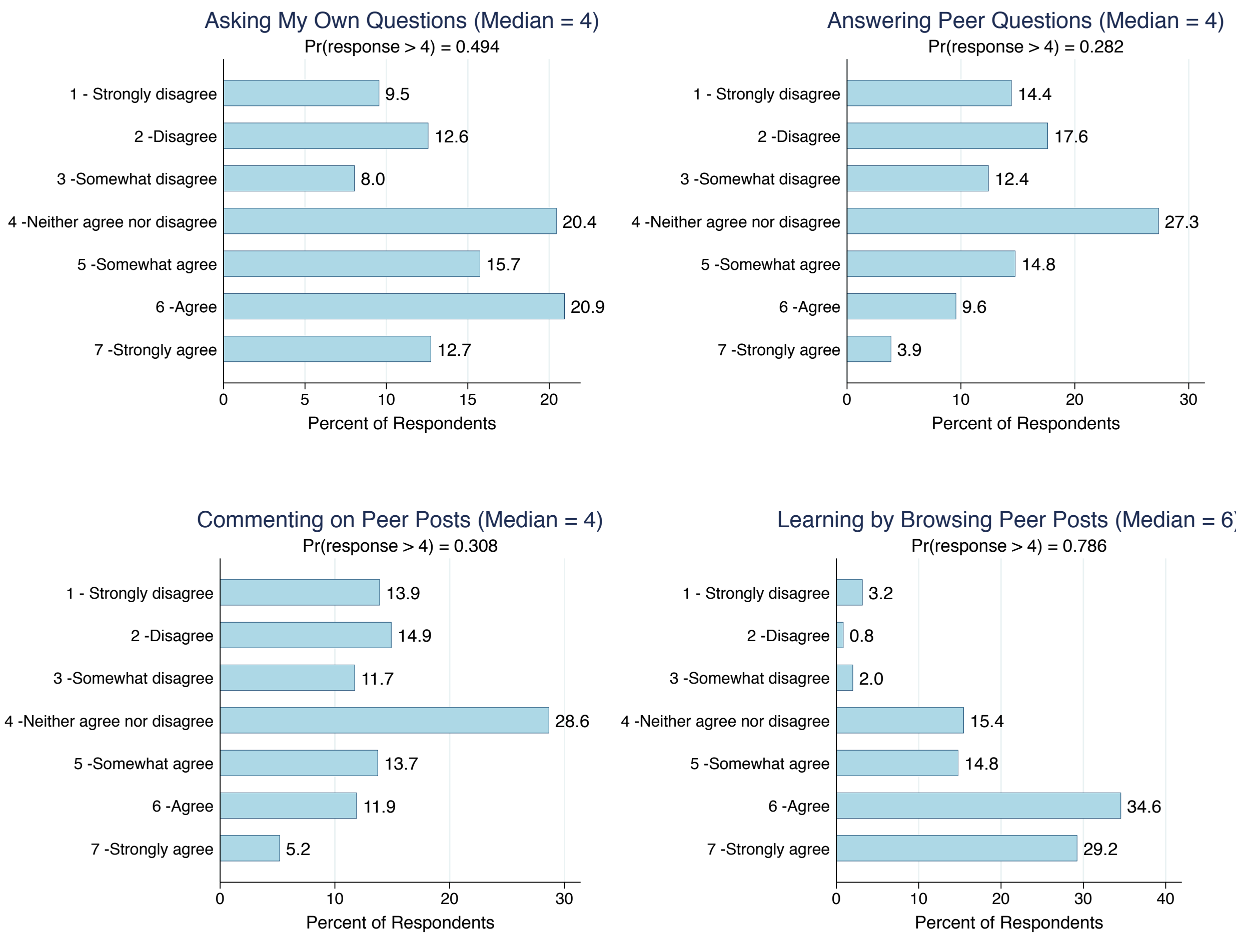

Notes: The figure shows the distribution of student responses to a set of Likert scale survey questions that asks them about how they utilize the discussion board.

Figure 8: Student Demographic and Characteristics Balance Check for Sign-up Activity

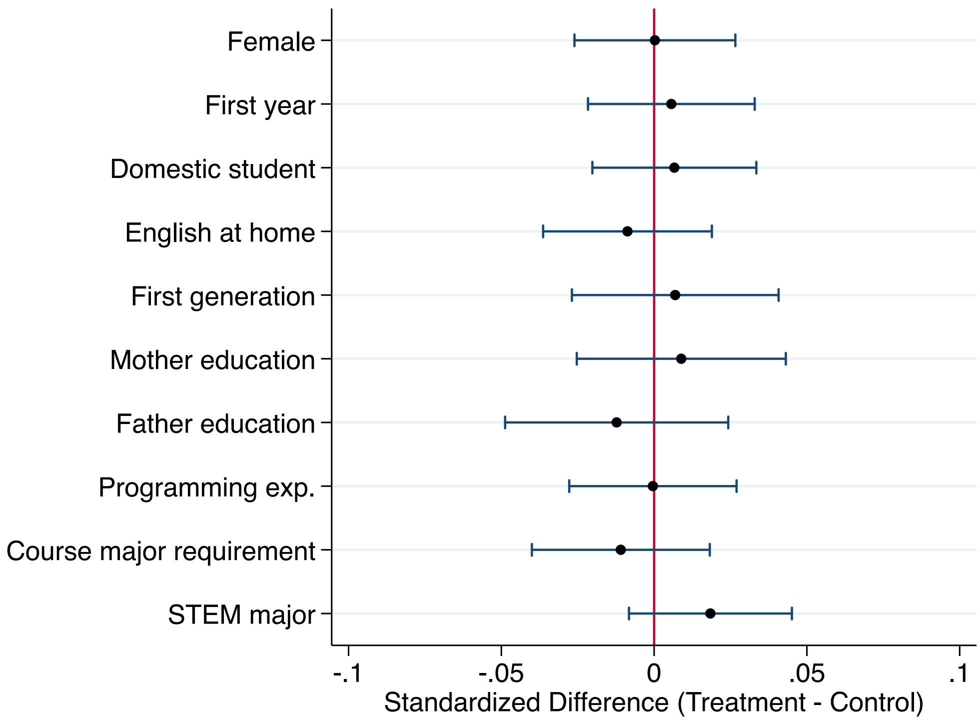

Notes: The figure reports the mean differences in the baseline student demographics and characteristics across the control and treatment groups for the sign-up activity. The estimates displayed are computed by regressing each standardized variable presented on the vertical axis on an indicator for receiving the sign-up activity. The error bars represent 95% confidence intervals..

Figure 9: Student Demographic and Characteristics Balance Check for Writing Activity

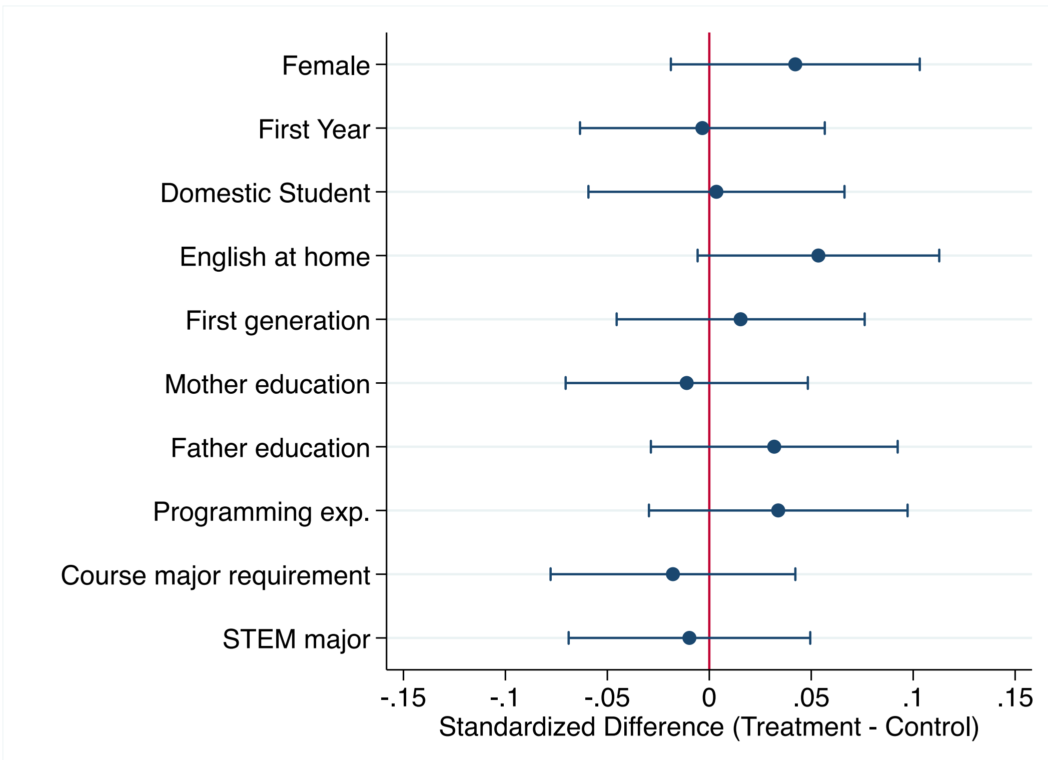

Notes: The figure reports the mean differences in the baseline student demographics and characteristics across the control and treatment groups for the writing activity. The estimates displayed are computed by regressing each standardized variable presented on the vertical axis on an indicator for receiving the writing activity. The error bars represent 95% confidence intervals..

Figure 10: Sign-up Activity and Discussion Board Utilization

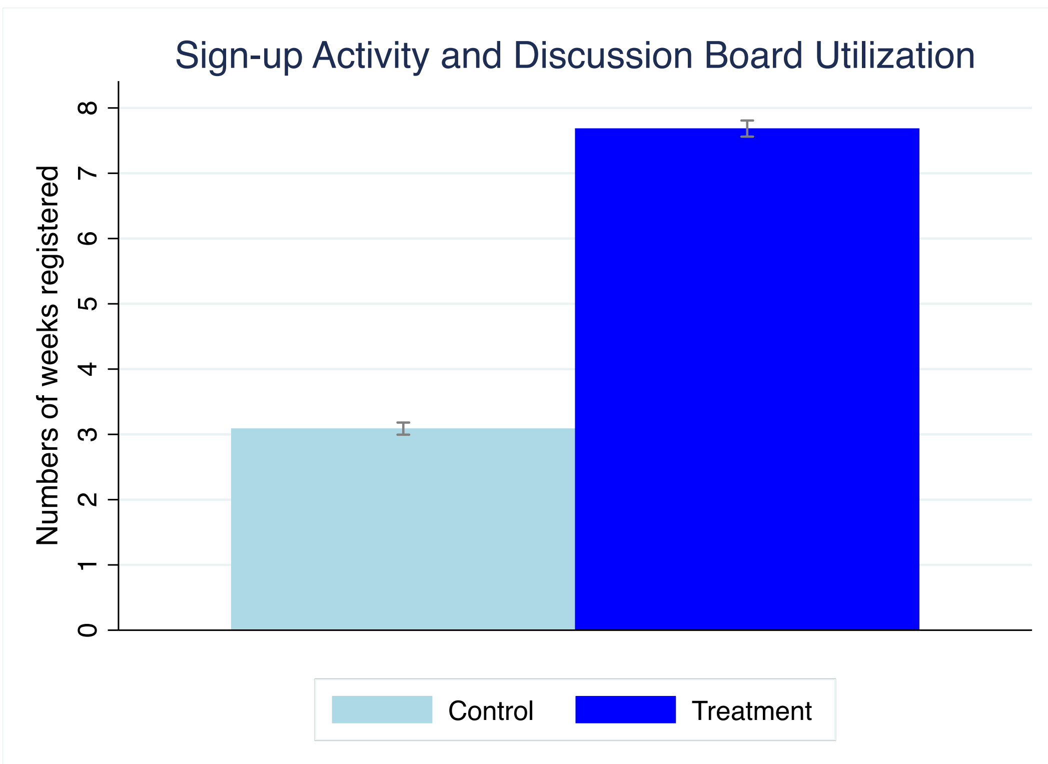

Notes: The figure shows the average number of weeks students are registered on the discussion board across the control (left) and treatment (right) group. The maximum number of weeks a student can be registered is 12-weeks. The treatment group consists of the students who were randomly assigned to receive the sign-up activity near the start of the course. The error bars represent the 95% confidence intervals.

Figure 11: Sign-up Activity and Being Informed About Discussion Board

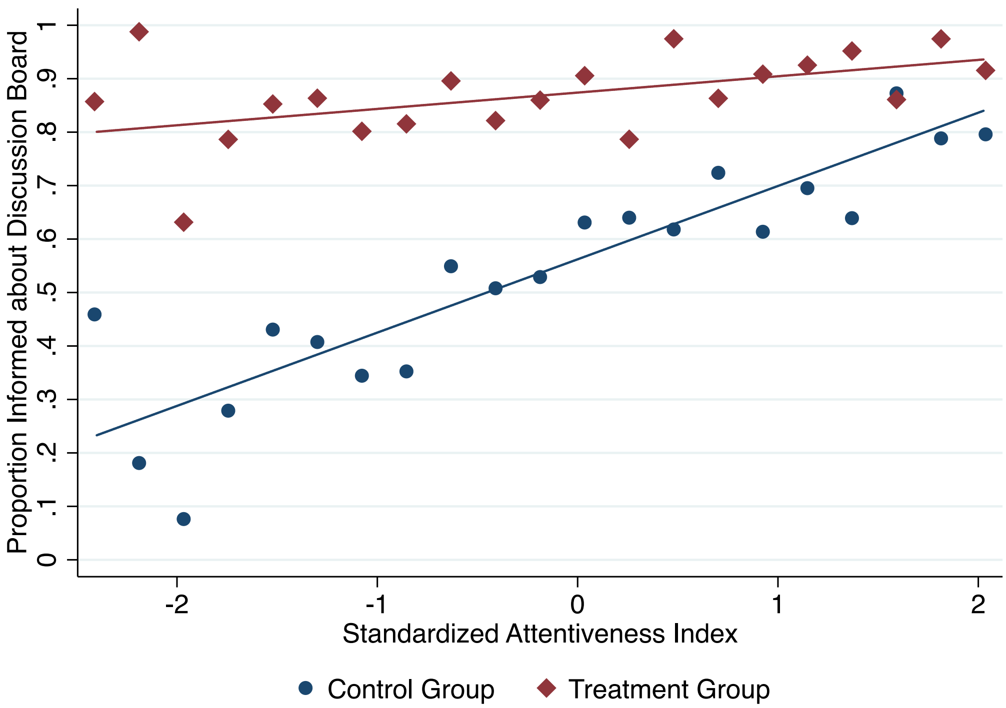

Notes: The figure presents a binned scatter plot showing the relationship between being informed about the discussion board and students' attentiveness by their sign-up activity treatment status. Students in the treatment group were randomly assigned to receive the sign-up activity near the start of the course. Whether a student is informed about the discussion board is inferred through a question on the baseline survey. The student attentiveness index is constructed using a series of survey questions.

Figure 12: Variations of the Sign-up Activity and Discussion Board Utilization

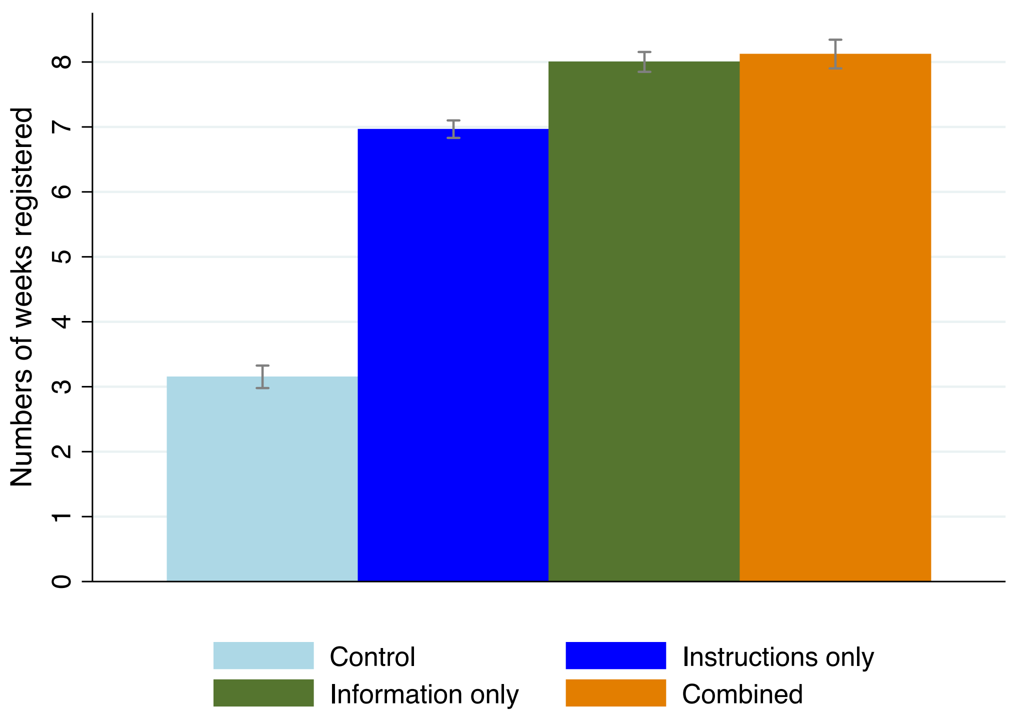

Notes: The figure shows the average number of weeks students are registered on the discussion board across the control and the variations of the sign-up activity. The maximum number of weeks a student can be registered is 12-weeks. The mean estimates are presented in the following order: (1) control, (2) instructions only, (3) information only, and (4) both instructions and information. The error bars represent the 95% confidence intervals.

Figure 13: Effect of Writing Activity on Contributions

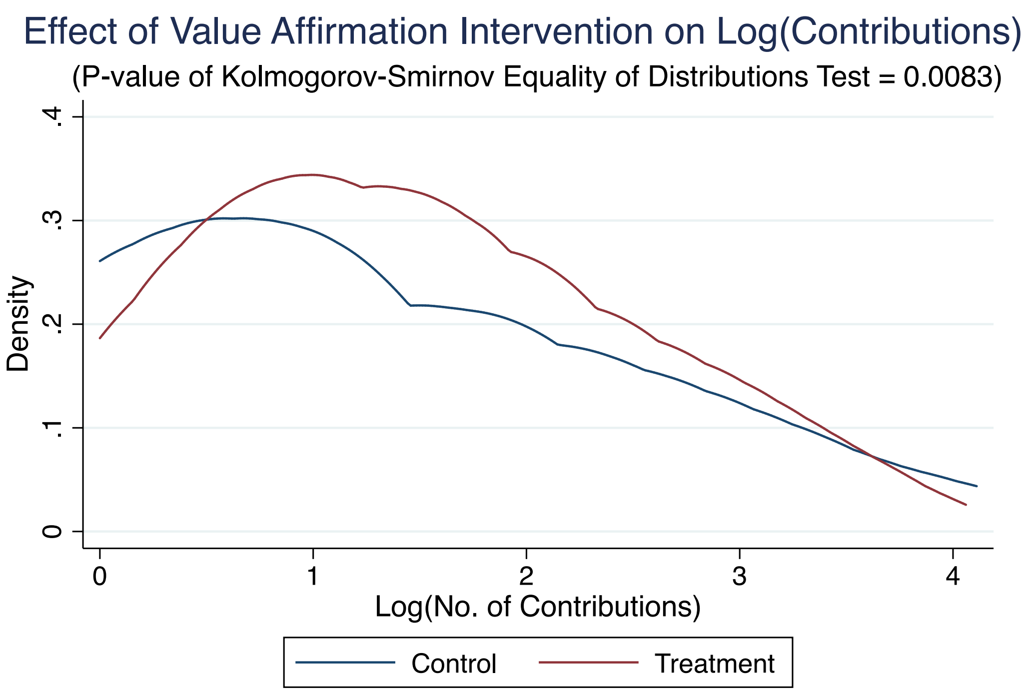

Notes: The figure presents the distribution of contributions made by treatment status of the writing activity. The hyperbolic sine inverse transformation is applied to the number of contributions as many students do not contribute.

Figure 14: Distribution of Model Implied Contributions and Observed Contributions

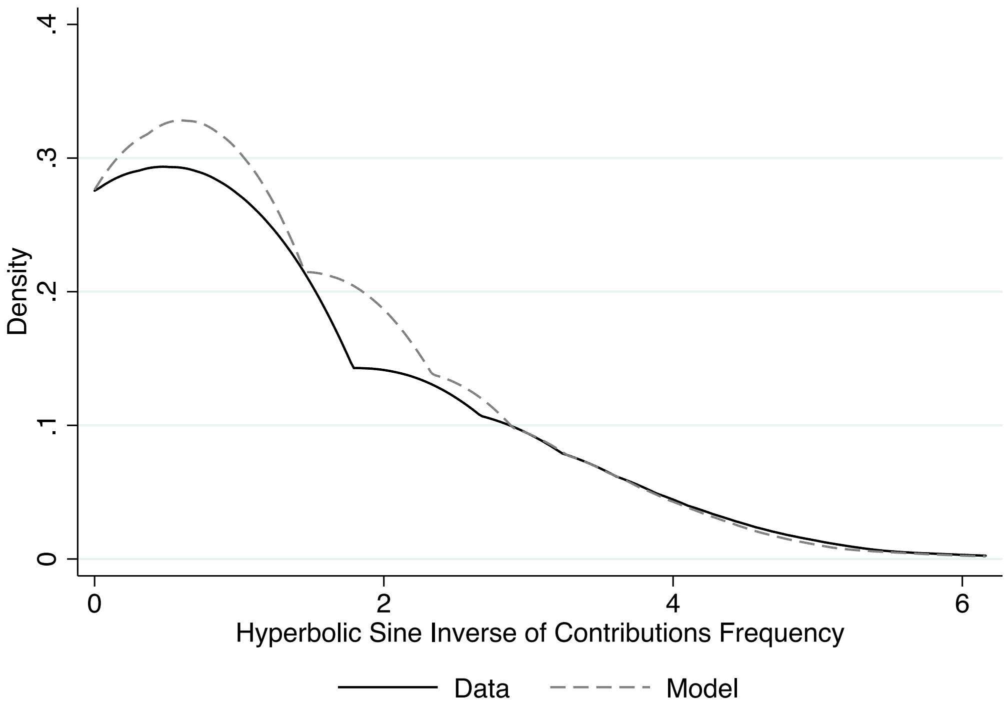

Notes: The figure presents the distribution of observed writing activity (solid) overlaid together with the distribution of model implied writing activity (dashed).

Figure 15: Model Implied Writing Activity and Observed Writing Activity

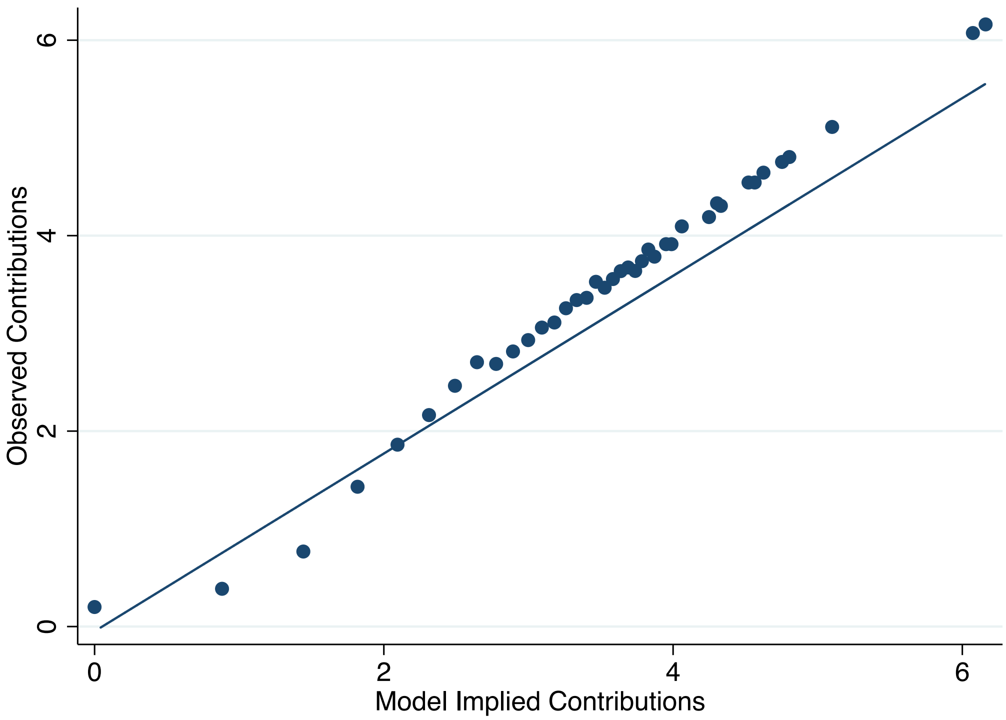

Notes: The figure presents a binned scatter plot showing the relationship between observed contributions (vertical axis) and model implied contributions (horizontal axis).

A Appendix: Institutional Details

A.1 Course Outline

The course is taught over 12 weeks. Learning the principles of programming can be broken down into the following three stages: 1) basic concepts (e.g., variables and loops), 2) intermediate concepts (e.g., nested loops and parallel lists), and 3) advanced higher order concepts (e.g., algorithms and object oriented programming). That is, the course have a cumulative structure where topics build on each other. The following table includes the syllabus for the foundation programming course.

Week	Topics Coverage
1	Numerical operations, variable assignment, and common coding errors
2	Defining functions and string variables
3	Conditional statements (if, elif, and else) and boolean variables
4	Loops (for and while)
5	Properties of lists (e.g., aliasing and mutability)
6	Nested lists and nested loops
7	Tuples, dictionary, and parallel lists
8	Palindromes classification algorithm and more about lists, tuples, and dictionaries
9	Good programming practices for testing and debugging code (e.g., unit tests)
10	Search and sorting algorithms (e.g., binary search and bubble sort)
11	Writing classes and methods
12	More object oriented programming (classes and methods)

The course employs two online learning platforms: an online homework environment and an online peer discussion board.

A.2 Online Homework Environment

Each week students receive an online homework module where students watch videos and then subsequently solve homework problems. Students login to the platform, and are given an outline for the videos they should watch and are presented with the follow-up coding problems. The online learning platform hosts a total of 133 videos (7.1 hours) and 401 follow-up homework problems. All homework problems are graded through an automatic artificial intelligent system. The following table presents summary statistics for the weekly content available on the platform.

Variable	Mean	SD
No. of videos assigned per week	11.1	4.4
Minutes of video lectures assigned per week	35.4	14.402
No. of questions assigned per week	33.3	13.614
Proportion of coding questions per week	0.22	0.121
No. of weeks	12	

Students provided an outline for how to learn a topic:

Figure 16: Outline for Learning Numerical Operations

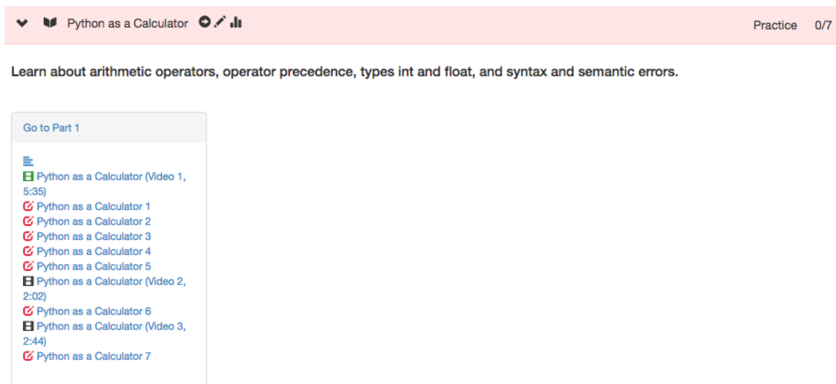

Students begin the course by watching a video about numerical operations in Python:

Figure 17: Video on Numerical Operations

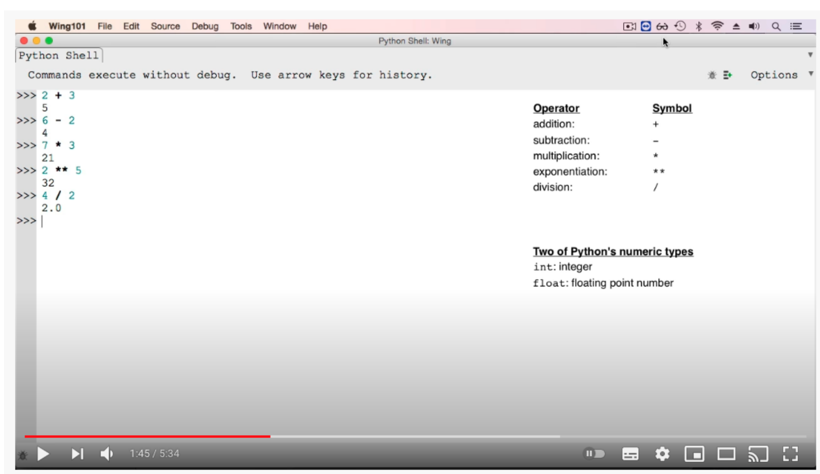

The following figure shows an example of a follow-up coding problem:

Figure 18: Video on Numerical Operations

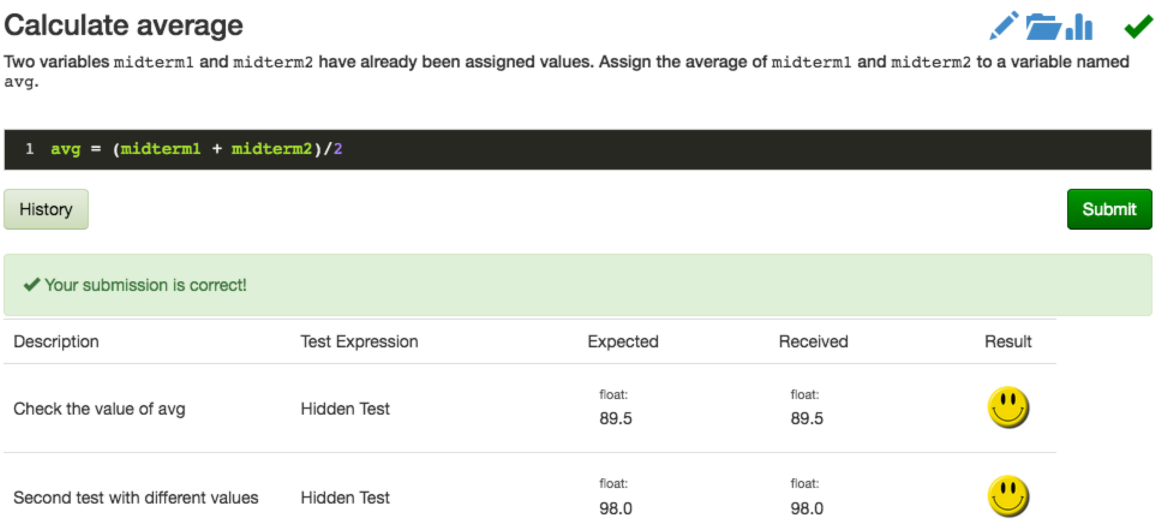

A.3 Online Peer Discussion Board

Students can use the online peer discussion board to get help with course material through asking questions. The questions are answered by peers, and answers can be validated by TAs or instructors. Students can also comment on either questions and answers. Comments can be used to further clarify the question, or give ideas on how to start solving the problems. The following table shows an example of student interactions on the discussion board.

Table 13: Example of Student Peer Interaction on Discussion Board

Interaction Type	Response
Question	How do we write a new line in a file using python?
Answer	Similar to how you would create a new line in a print function: file = open("somefile.txt", "w") file.write('\n') file.close() I hope that helps.
Comment	The code in the answer works, but note that opening a file in write mode will delete the contents of the file. Use append mode if you want to add to the file.

B Appendix: Data

B.1 Measuring altruism

The baseline survey elicits a student's altruistic through a series of questions. Responses to the relevant questions are aggregated so that they are increasing in the attribute of interest. The tables below include the 7-point Likert scale survey questions used to measure altruism.

Table 14: Survey Components for Altruism

Survey question (7-point scale)	Relationship with Forward-looking
I see myself as someone who is helpful and unselfish in sharing notes with peers	Increasing
I see myself as someone who is careful not to disturb the learning of others during lectures	Increasing
I see myself as someone who feels little concern about struggling peers	decreasing
I see myself as someone who is willing to volunteer to take notes for students with a disability	Increasing

B.2 Survey questions to study information spillovers

- I discussed the discussion board information received during the baseline survey with other students in the course [Likert Scale]
- I discussed the writing exercise that involved reflecting on quotes from a prior cohort of students on how they utilized the discussion board with other students in the course [Likert Scale]

C Appendix: Nudges

The nudges are designed using various behavioural insights, for example, the presence of inattentive students for designing the sign-up activity. Kizilcec et al. (2020), Harackiewicz and Priniski (2018), and Damgaard and Nielsen (2018) provides excellent reviews on the behavioural nudging literature in education.

C.1 Sign-up activity

The sign-up activity is designed to help students internalize the information provided through self-reflection. Students were given the following questions as a part of the activity.

- (Open Response) Based on the instructions presented, approximately how long does it take to sign up for the discussion board?
- (Disagree-Agree) Are you aware that 87% of questions asked have received a response?
- (Disagree-Agree) I am aware students can use the discussion board to learn through ...
  - asking their own questions
  - answering questions of other students
  - engaging in discussion with peers by commenting on posts
- (Open Response) How could you potentially benefit from the discussion board?
- (Multiple Choice) Would it be worthwhile for you to sign-up to the discussion board? Here is the link to sign-up page: [Link]
  - Yes, I have just registered for the discussion board
  - Yes, I will sign-up to the discussion board this week *[Follow-up prompt to schedule day and time]*
  - No, I will not sign-up to the discussion board *[Follow-up prompt to request reasoning for not using discussion board]*

C.2 Contributions writing activity

The writing activity leverages survey data from the Fall 2019 cohort on how students utilize the discussion board. The writing activity is outlined as follows:

1. Students are presented with a pair of (free-rider, contributor) real quotes.

    - "Here are some quotes from a couple of Fall 2019 students on how they primarily used the discussion board."
    - Student A: "I mostly use it to browse through existing questions and answers to see if anyone else has the same questions to me or not."
    - Student B: "I primarily use Piazza to answer my peers' questions as a way of reinforcing my understanding of course material. I find it helpful since if I'm constantly applying course knowledge it will be less likely for me to forget about the concepts."

2. Students classified which quote their own behaviour resembled.

    - "From student A or B above, which student does your own discussion board usage in the course thus far most resemble?"
        - Similar to only student A
        - Similar to only student B
        - Similar to both student A and B
        - NOT similar to either student A or B
    - "Please take a moment and share us your perspective on how student A and B above differ in their discussion board usage" (Open Response)

3. Students reflected on which behaviour is optimal for learning.

    - "To what extent do you agree with the following statements:"
    - "Student [A or B] effectively uses the discussion board to **benefit their own learning**" (Disagree - Agree)
    - "Student [A or B] effectively uses the discussion board to **benefit the learning of other classmates**" (Disagree - Agree)

- “Suppose you were to use Piazza primarily like Student A or Student B. What type of skills do you think you would gain under both scenerios:”
  - “Skills gained from using the discussion board like student [A or B]” (Open Response)
- “Between student A and B, whose behavior is better suited to advance their own learning while also benefiting their peers?”
  - Student A
  - Student B
  - Student’s A behaviour is equally effective as student B

4. Steps 1-3 for (negative,positive) perspective on peer discussion. The following pair of real quotes students shown to students are:

   - “Here are some quotes from a couple of Fall 2019 students on their perspectives about engaging in peer discussions.”

   - Student C: “I don’t think there are questions and discussions on there that I need help with so I don’t use it for discussing concepts with peers.”

   - Student D: “The discussion board helps me to connect with my fellow peers. Additionally, It is useful to see what other students are having problems with, and if I can help them, or how others can help them. When students combine their knowledge to create an answer, it usually covers different perspectives and different approaches you might not have thought about on your own.”

5. Students prompted to write about how being an active contributor is important for building an effective online learning community.