Identifying a Cumulative Learning Technology: Evidence from Online Learning

Hammad Shaikh*
University of Stavanger

November 30, 2025

Abstract

An inherent feature of learning in most disciplines, especially STEM, is its cumulative structure, which makes developing advanced skills challenging. This paper credibly estimates a dynamic learning technology of student effort in a foundational university course with a cumulative structure. Doing so is incredibly difficult because effort inputs are typically unobserved and are dynamic endogenous choices. To address this, I use rich panel data on nearly 3,700 undergraduates in an online introductory programming course, which precisely tracks study time throughout the course. Then I carry out a field experiment which generates period-by-period exogenous variation in effort allocation, enabling me to identify dynamic interactions across effort inputs in the learning technology. I find evidence of dynamic learning complementarities as the marginal benefit to studying in each learning period is increasing in prior knowledge accumulated. I then develop and estimate a multi-stage behavioral model of effort supply, using the experiment to identify benefit and cost parameters at each stage. The model informs policy simulations of grading schemes, showing that decreasing assignment weights across the course—rather than equal weights—encourages effective effort allocation and mitigates procrastination by boosting effort early on when foundational skills are acquired. The findings have implications for effective learning strategies and approaches to course design in cumulative learning environments.

*Thanks to Victor Aguirregabiria, Isaiah Andrews, Carolina Arteaga, Stephen Ayerst, Gabriel Carroll, Marc-Antoine Chatelain, Cheok In Fok, Jiaying Gu, Elaine Guo, Alexander Hempel, Tanjim Hossain, Faisal Ibrahim, Daniel Indacochea, Raji Jayaraman, Catherine Michaud Leclerc, Étienne Makdissi, Aloysius Siow, Todd Stinebrickner, David Price, Swapnika Rachapalli, Baxter Robinson, Marc-Antoine Schmidt, Fernando Saltiel, Eduardo Souza-Rodrigues, Eva Vivalt, Jessica Wagner, Matt Walshe, Joseph Jay Williams, Farhan Yahya, Román Andrés Zárate, Ruizhi Zhu, Jiaqi Zou, and participants in the University of Toronto's Public Economics Discussion Group, Empirical Microeconomics Seminar, Canadian Economic Association Conference, Norwegian Economic Association Conference, Leuven Economics of Education Conference, Workshop on Education Economics and Policy, University of Stavanger, Lahore University of Management, and the University of Aarhus, for their comments. Sameul Khan, Sam Maldanado, and Qi Yin Zheng provided valuable data support, and 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 in this study. The study is registered at the AEA RCT Registry (ID 0008787). A previous version of this paper was circulated under the title "Improving Online Learning Through Course Design: A Microeconomic Approach". Financial support from an Ontario Graduate Scholarship 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

Higher education has expanded rapidly over the past decade as technical and analytical skills are fundamental to the development of any modern technological society (Hanushek and Woessmann, 2015). The number of students conferring STEM degrees in the US, for example, has increased from around 472,000 in 2009 to around 733,000 in 2018 (US National Centre for Education Statistics). In particular, the demand for Computer Science has grown more rapidly than many other fields (Light, 2024). Policy makers continue to call for a dramatic increase in the supply of STEM majors, and substantial government funding has been allocated towards promoting computer science in particular.[1]

A defining feature of many foundational university courses is their cumulative learning structure, especially in STEM.[2] Mastery of advanced material requires first acquiring proficiency in earlier, more basic concepts, and insufficient effort at early stages can cause students to permanently fall behind.[3] Because most degree programs require such foundational prerequisites, the cumulative structure extends beyond individual courses to entire majors. This creates challenges for both students, who must adapt their study habits and sustain continuous effort to keep pace with the cumulative structure of the material, and instructors, who must design courses that motivate students to invest effort early and ensure that fundamental concepts are effectively taught. Students also vary widely in their prior preparation and organizational skills, making large introductory classes difficult to design effectively. Given these challenges, many students struggle to complete majors with a strong cumulative structure such as in STEM, and drop out during their first two years of university.[4] Consequently, understanding the dynamics of the learning process is especially important in these foundational courses, as they serve as mandatory gateways to advanced study across a wide range of disciplines.

This paper is the first to credibly estimate the cumulative learning technology which maps dynamic effort inputs into learning outcomes in a foundational university course. The specified technology is cumulative with multiple learning periods and allows for past knowledge to persist

[1]The US Department of Education spends around $400 – $500 million on promoting STEM education annually. Around 25% of this funding is typically allocated towards computer science education.

[2]It is worth noting that many courses outside of STEM also have a cumulative structure. For example, in a language learning course, students will first learn the alphabet, common words, and then learn to speak simple sentences.

[3]Introduction to programming, for example, has a standard cumulative structure where students learn numerical operations, functions, and then algorithms.

[4]Stinebrickner and Stinebrickner (2012) and Stinebrickner and Stinebrickner (2014) find that drop-out decisions can be largely explained by students learning about their academic abilities and being overoptimistic about succeeding in a science program at college entry.

into the future. The learning technology also incorporates dynamic learning complementarities as the productivity of effort (rate of learning with study time) in a learning period can depend on previous effort exertion. Estimating the learning production function helps inform effective learning strategies and approaches to course design.

Identifying such a cumulative learning process is incredibly challenging for several reasons. First, repeated learning measures to observe the knowledge accumulation process are not typically available. Second, precise students' effort inputs such as study time allocation are typically unobserved. Although data on student performance on various cognitive assessments are readily available, administrative data on students' precise study time allocation is nearly nonexistent. Third, the effort inputs across the cumulative learning process are dynamic endogenous choice variables. Then period-by-period exogenous variation in effort is required to identify the dynamic learning complementarities, even though just the availability of a single exogenous shock in effort is a rare occurrence.

I develop a unique approach to address all these challenges and credibly estimate the cumulative learning process in the higher education context for the first time. I first gather unique administrative data from a prominent online STEM course which enables me to precisely measures effort allocation and corresponding knowledge accumulation throughout the entire course. I then carry out a field experiment which generates exogenous variation in effort allocation, credibly identifying each element of the cumulative learning process. The estimates of the learning technology help inform effective learning strategies for students in foundational university courses. Lastly, I use rich survey data to study heterogeneity in effort allocation across different types 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 uses an open-source online learning platform where students learn content on their own by watching videos and doing practice problems posted on a weekly basis.[5] In addition to weekly low-stakes homework assessments, the course also includes two high-stakes assessments: a midterm and a final exam. Given that students in the course learn most of the material through self-study, the course also employs a voluntary online student discussion board to further support students. The discussion board facilitates learning by allowing students to interact with each other, discussing the course material and collaborating on assignments in an instructor-moderated online environment.

I collect data on nearly 3,700 students who consented to participate in the research.[6] The rich

[5]The online course features around 130 videos and 400 practice problems.

[6]The student consent rate is around 87% and over 40,000 student-week observations are available.

administrative data include time-stamped student interactions with the online learning environments throughout the entire semester. These data enable me to measure total online study time at each stage of the learning process precisely. Evidence from the administrative data suggests a lack of online participation activity by a non-trivial proportion of students. For example, each week around $15 - 20\%$ of students spend no time whatsoever doing the homework. I supplement the administrative data with survey data (as mentioned), collecting demographic information from students and further eliciting their behavioural characteristics, such as their attentiveness and forward-looking (versus myopic) perspectives. I find that more attentive students have a higher propensity to start the low-stakes homework assignments and have higher awareness of course resources. I also find that students who are more myopic tend to be less likely to do optional (ungraded) homework problems that are available throughout the course. The survey evidence also reveals substantial heterogeneity in effort allocation according to a student's behavioural 'type.'

To identify the cumulative learning process, I conduct a field experiment to generate exogenous variation in students' effort allocation. The interventions considered in this study can be characterized as 'targeted informational reminders nudges,' as they prompt a student to take a specific action, provide simple instructions for doing so, and lastly, remind the student to complete the task. The nudge is a homework reminder message which is deployed repeatedly across several weeks throughout the course and is aimed at promoting further participation in online homework. The reminder message informs students of the upcoming homework deadline, prompts them to set aside time in their schedule to work on the homework, and provides a direct link to the homework. I find that receiving an additional reminder message, on average, induces students to spend an extra 23 minutes on the corresponding homework assessment. The reminder messages are most useful for inattentive students, who are less likely to be aware of upcoming homework deadlines.

The deployment of randomized homework reminders throughout the course provides an opportunity to estimate the cumulative learning technology as a function of students' study time allocation. Reflective of the actual course structure, the learning process is estimated separately across three learning stages: basic, intermediate, and advanced. The benefit of effort in a learning stage depends on the cumulative technology, which has two endogenous variables – the total study time in the current learning stage and the knowledge accumulated in the previous learning stage. I construct instruments for both endogenous inputs using the number of randomly assigned homework reminders at each learning stage, allowing me to identify the parameters of the cumulative learning technology. I find a positive marginal return to effort at each stage of the learning process and document evidence of dynamic effort complementarities. For example, in the advanced learning stage, I find an additional hour of online study time increases final exam grades by 0.11 SD, and this marginal benefit increases by 0.07 SD for every 1 SD increase in intermediate stage knowledge.

To further explore the student effort allocation process, I develop a multi-stage estimable model of student effort supply. The model features a single instructor (i.e., the principal) and multiple students (i.e., the agents), the latter exhibiting heterogeneity in their baseline knowledge, English language proficiency, and whether they are forward-looking or not. Myopic students set effort at each learning stage independently without internalizing that the productivity of studying in the future is increasing in current knowledge accumulation, while forward-looking students internalize the cumulative benefits when allocating their study time earlier in the course. The instructor's objective is to maximize the learning of a representative student net of effort costs, whereas students exert effort throughout the course to maximize their expected course grade. Students with higher baseline knowledge are more productive (i.e., have higher marginal returns to effort) in the basic learning stage. Whereas students with a lower English proficiency have a higher cost to exerting effort. The solution to the students' problem indicates that myopic students misallocate their effort and underinvest in low-stakes homework assessments covering foundational programming skills. Moreover, the solution to the instructor's problem demonstrates that, when designing an optimal grading scheme for a course in the presence of a cumulative learning technology in which dynamic complementarities are strong, a myopic (forward-looking) student's learning is best served by assessments whose weights decrease (increase) throughout the course.

Next, I use the administrative dataset including students' precise study time allocation to estimate the model (also showing empirical evidence that the model implications just outlined are consistent with the data). The benefit of effort in a learning stage depends on the dynamic learning technology that is estimated using the field experiment. I then estimate the cost function using maximum likelihood estimation. Here I assume the cost function is convex and linearly separable across the learning stages. Considering each stage independently, the convexity parameter of the cost function can be identified. To see how, note that students randomly nudged to work harder earlier in the course will be more productive learning more advanced skills at a subsequent learning stage. That is, exogenous shifters in the marginal benefit of effort are available while holding the marginal cost of effort constant (i.e., conditioning on English proficiency) to identify the cost function.[7] Overall the field experiment produces 75 distinct sequences of effort shocks across the learning stages (basic, intermediate, and advanced), allowing identification of the marginal benefit and cost of effort structural parameters.

[7] See Section 8.7 for details underlying the identification argument for the effort cost function.

The estimated model enables me to project student effort allocation and corresponding learning outcomes as a function of the grading scheme implemented by the instructor. I calibrate the proportion of myopic students 68% using the survey data and find that the estimated optimal assignment grading weights are gradually decreasing as the course progresses. Relative to the grading scheme used in the existing course with equal assignment weights, I find that implementing the optimal grading scheme is predicted to increase final exam performance for myopic students by 0.11 SD. The achievement gain arises from myopic students effectively allocating their study time by investing more effort at earlier learning stages of the course under the optimal weights, thereby obtaining proficiency in the basic concepts that serve as foundational building blocks for rest of the course. Next, I simulate the optimal homework weights while varying the share of myopic students in the course. I find that incentives should be front-loaded when majority of students are myopic, middle-loaded when a modest share of students are myopic, and end-loaded when the vast majority of students are forward-looking.

While the model is estimated using student effort supply data in an online course with a given course structure, the proposed framework is quite general and can accommodate other courses. The model informs our understanding of how different types of students allocate their effort across the assessments under a given grading scheme, helping illuminate how to design the incentives in courses with heterogenous students.

Overall, the findings in this paper contribute to our understanding of how the effort allocation of students can be made more efficient through appropriately designing the grade scheme, following the systematic approach I develop. The experimental variation in student engagement allows me to estimate and credibly identify a behavioural model of effort allocation which sheds new light on the optimal design of course grading schemes. For a foundational course with many myopic students, the simulations indicate that instructors should assign more weight to assessments given earlier in the course. Doing so will lead myopic students to appropriately front-load their effort (as noted). As most university courses have multiple assessments, the instructor can distribute the grading incentives unevenly, as informed by the model simulations.

This paper builds on several prior literatures. These include research that estimates cumulative education production functions and those that uses field experiments to evaluate the efficacy of various educational interventions, among others. First, my paper contributes to a body of research estimating the marginal learning returns to student effort, an essential input to the education production function. Stinebrickner and Stinebrickner (2008) estimates the returns to effort using self-reported diary data on time use from college students, and also collect information on whether their roommate owns video games. The authors use roommate ownership of videos games as an instrument for study time and find that increasing study time by 1 hour a day increases the first semester GPA by around 0.3 - 0.4. Ersoy (2021) uses administrative data from a popular language learning platform, Duolingo, to measure causal learning returns to effort. In the study, students learning Spanish are randomly assigned to complete a different number of lessons. The author finds that spending around a hour completing 9 lessons results in increasing achievement by 0.057 SD on tests that are external to the online learning environment. My paper contributes to this literature by using administrative data from a large prominent STEM course together with a field experiment to measure the causal learning returns to study time and assignment completion.

Second, my paper also relates to the literature estimating cumulative education production functions. Such cumulative technology maps present and historical inputs to current learning outcomes. Todd and Wolpin (2007) estimate a cumulative production function for children as a function of child ability and history of family inputs. The authors find that lagged family inputs are significant predictors of cognitive achievement. Consistent with a cumulative technology, Aizer and Cunha (2012) find larger IQ gains from preschool enrolment for children with higher stocks of early human capital. Gilraine (2016) uses year-to-year variation in school accountability to identify dynamic complementarities in school inputs. The author finds a 0.18 SD increase in test scores for students who are in schools that were subject to school accountability in two consecutive periods relative to those subject to accountability only in the previous period. In a recent paper related to my study, Cotton et al. (2026) uses online study-time tracking and randomized monetary incentives to credibly estimate a structural model of learning for 5th and 6th-grade students. Unlike my study, where students acquire skills through natural online learning over a semester, the authors offer a web-based math platform for 2 weeks as an additional learning opportunity for students. They find that low productivity–the rate at which an hour of study converts into a durable skill–is a strong predictor of academic struggles. Appendix A.1 describes the education production function literature in more depth. Although the education production function literature includes parental and school inputs, time varying students inputs are scarce. Consequently, policies designed to increase educational attainment by targeting either parental or school inputs may implicitly assume that students' effort is held constant. Identifying the cumulative education production as a function of student effort is important for designing dynamic policies to encourage students to learn effectively. I contribute to this literature by estimating the cumulative learning technology as a function of students' study time allocation, focusing on the important setting of learning programming—a fundamental STEM skill. It uses exogenous variation in student study time at each learning stage to identify dynamic complementarities in student effort inputs.

Finally, the field experiments in my paper contribute to a large body of recent work investigating the efficacy of behavioural nudges in promoting desirable academic behaviours in higher education.[8] Smith et al. (2018) conduct a field experiment to evaluate the efficacy of a personalized email message which informed students how their assignment grade will influence their final grade, based on their current grade in the course. The authors find that students who received the message achieved a 4 percentage point higher grade on the assignment. Oreopoulos et al. (2018) investigate the effectiveness of a planning module, which involved a group of randomly selected students building a weekly calendar and receiving follow-up reminders from an upper-year coach. The authors find that the online planning module marginally increased self-reported weekly study time, but the increase in weekly study time did not result in an increase in academic performance outcomes. Clark et al. (2020) conduct an experiment to test whether college students who set goals exert more effort and achieve improved learning outcomes. The authors find setting task-based goals increased task completion and subsequent course performance. However, setting performance-based goals did not result in a significant increase in learning. Appendix A.2 includes a more extensive list of related papers exploring student effort choices. I add to this literature by using administrative data on student effort and showing that targeted informational reminders can improve achievement by nudging inattentive students to participate further in learning activities.

The rest of the paper is organized as follows. The next section provides information about the sample and describes the online learning environments. Section 3 outlines sources of data collection and also presents descriptive statistics. The experimental design and the key features of the interventions are discussed in Section 4, and corresponding results are presented in Section 5. A model of dynamic effort supply is discussed in Section 6, and the cumulative learning process is estimated 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 Institutional Background

The setting for the study is a first-year undergraduate online introductory programming course offered at a large research-intensive public university in Canada. This section describes the course structure and the platforms used to facilitate student learning in the course.

[8] Kizilcec et al. (2020) and Harackiewicz and Priniski (2018) discuss a variety of behavioural interventions in the literature that are focused on improving academic outcomes in higher education.

Cumulative Course Structure. The course assumes no prior programming knowledge and teaches programming fundamentals using Python (see Appendix B.1 for the course outline). 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 content can be naturally partitioned into three stages: basic, intermediate, and advanced. Weeks 1 - 4 cover the foundational concepts of programming, such as variable declaration and loops. Then, weeks 5 - 8 cover intermediate concepts such as nested loops and dictionaries. Finally, building on the basic and intermediate learning stages, the course concludes by covering advanced concepts, such as algorithms and objected oriented programming. Although the content in week 1 requires no prior programming experience, the topics covered in all others weeks are cumulative, as they build on concepts covered in past weeks.[9]

The coursework consists of low-stakes weekly homework assessments and also higher stakes assessments, which include a midterm and a final exam. The homework in the first two weeks is optional (i.e., they are ungraded) to allow students to practice interacting with the online learning environment.[10] The midterm and final exams are typically written in weeks 5 and 13, respectively. Additionally, students can obtain course credit by participating in two research surveys that are deployed at the start and end of the course.[11] 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 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. Appendix B.2 further illustrates the user interface

[9] For example, learning nested lists and nested loops in week 6 requires an understanding of basics of loops and lists, covered in weeks 4 and 5, respectively.

[10] Students can enrol into the course up until the second week. Making the homework in the first two weeks optional also reduces the logistical burden on instructors, as otherwise students who enrol late may demand alternative make-up assessments, or request a grading scheme adjustment.

[11] Both surveys were deployed online using the Qualtrics survey platform. Students earn a 1% bonus credit for each research survey they complete.

of this learning environment. The online learning platform for the course contains around 133 instructional videos and 401 problems, assigned across 12 weeks through homework assessments.

The Online Peer Discussion Board. The course offers an interactive online course discussion board where students can discuss course material. Students 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. Posts on the discussion board are organized per week as new content is introduced weekly. Appendix B.3 describes the user interface of the discussion board in more detail. Students must register on the discussion board before they can read or write their own posts. Although encouraged, participation in the discussion board is completely voluntary, and students are not awarded any additional course credit for participation.

Learning Management System. Canvas is the learning management system (LMS) employed by the course involved in this study. 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. Students also have a message inbox on Canvas that is separate from their institutional email. Instructors can send student messages through Canvas, and such messages are received in students' Canvas messages inbox and also automatically forwarded to students' institutional email.

3 Data and Descriptive Statistics

This study uses a combination of rich student-level administrative and survey data to characterize online learning participation behaviour for different types of students. All data are gathered from the introductory programming course (introduced above), offered during the Winter 2020, Fall 2020, Winter 2021, and Summer 2021 academic semesters at 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. Pilot data were also gathered from the Summer 2019 and Fall 2019 cohorts and served to finalize the design of the primary data collection that is the focus of this section.[12] The timeline of the complete data collection exercise is shown in Figure 12.

[12]Pilot data collection involved having 30-minute recorded interviews with several students, conducting online surveys using various software, and prototyping various interventions.

Student Surveys. The baseline survey collects information about students' demographics and elicits information about their behavioral characteristics; the final survey gathers data on various course inputs, interactions with peers, and elicits student feedback about different components of the course. 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 91% for the baseline survey and 86% for the end-line survey.[13] 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 87%. The sample of total consenting students who completed the course consists of 3,686 students.

Student Activity on the Learning Management System. I collect student activity data from student interaction reports captured on Canvas. This includes the total number of announcement views, aggregate page views, and a daily list of all students enrolled in the course. The list of students enrolled in the course is retrieved daily to track attrition of students from the sample over the study period.[14]

Student Achievement. Student achievement data are collected from the weekly online homework, the midterm, and the final exam. This high-frequency achievement data allows me to assess student learning throughout the course. The primary measure of student learning is their cumulative final exam grade, which I standardized to have a mean of 0 and standard deviation of 1.[15]

Student Discussion Board Activity Data. Students' 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 also observe time-stamped data on all contributions (question, answers, or comments), and the number of unique posts viewed by students each week. Overall, I observe discussion board registration, contributions, and 'consumption' decisions.

[13]The research surveys are announced through the learning management system, and students who did not complete the survey two days before the deadline received a reminder to do so.

[14]The data are retrieved using the Canvas Application Program Interface.

[15]The final exam is a 3-hour comprehensive assessment that evaluates overall understanding of introductory programming in Python.

3.1 Student Study Time Data

Numerous types of student interactions with the online platforms are observed in the administrative data, recorded to the nearest second. That is, the online platforms serve as a monitoring device in terms of students' learning activities. For example, observed interactions include the times when students log in or out, play or pause an instructional video, submit problem solutions, and write in the discussion board. The availability of such rich time-stamped interaction-level data enables me to construct a precise measure of online study time at each stage of the learning process. The study time measure includes minutes spent watching instructional videos, working on homework problems, and reading and writing posts on the discussion board. I will now outline the construction of study time for the online homework, and for the online peer discussion board separately.

Online Homework Study Time. I couple the students' time-stamped online interactions together with a basic clustering algorithm to identify periods of learning activity at each stage of the course. The time-stamped interaction data are used to measure the minutes spent watching instructional videos and doing homework problems. The procedure is built around the empirical observation that students tend to study in approximately 30-minute blocks throughout the week (e.g., Tuesday from 6 - 6:30 pm). Each block of homework activity begins with students interacting with the online learning platform for at least 5 minutes, and concludes after 5 minutes of inactivity. Video watching time is computed based on when students play or pause the instructional videos. Students' time spent attempting homework problems is measured using information when students submit problems and click to view the next problem. Then, online homework study time is constructed by aggregating all blocks of learning activity for each stage of the course.

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 average time they spend on average writing and reading a post in minutes (see Appendix C.4 for survey questions). The administrative data on student engagement, and corresponding student-level survey data on average time use are used together to measure the minutes spent on the discussion board at each learning stage.[16]

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

Distribution of Online Study Time. Total online study time at each learning stage aggregates minutes spent on the online learning environment together with the minutes spent on the online discussion board. Figure 1 below shows the distribution of online study time at each learning stage of the course. This figure shows that, on average, student study time increases as the course progresses. Students exerting more effort at later stages of the course is consistent with the grading incentives, as the homework in the first two weeks is optional, the midterm falls in the intermediate stage, and the final exam is at the end of the advanced stage.

Figure 1: Distribution of Study Time by Learning Stage

(a) Basic Stage (b) Intermediate Stage (c) Advanced Stage

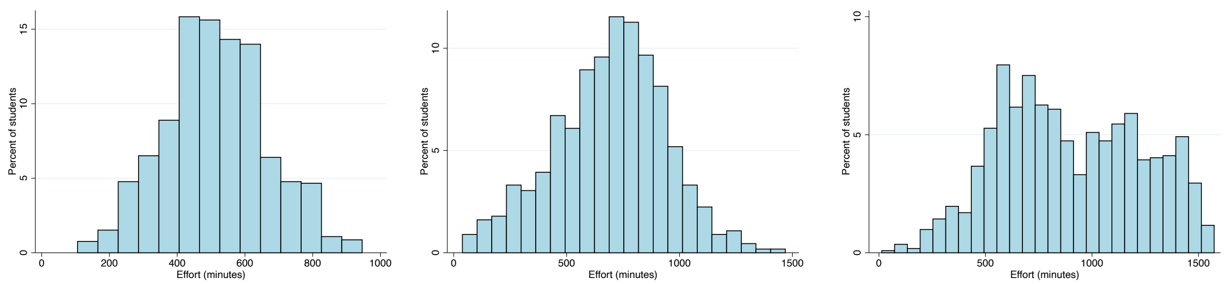

Notes: The figure presents the distribution of total online study time for each stage of the course: basic, intermediate, and advanced. All histograms use a bin width of 60-minutes.

3.2 Student Attentiveness

The baseline survey elicits student attentiveness through a series of questions. Each question is measured on a 7-point Likert 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). To measure attentiveness, one question students are asked whether "I tend to read all the instructor announcements for this course." All questions relating to students' attentiveness are included in Appendix C.2. To construct a continuous index of the behavioral responses, replies to all relevant questions are aggregated together so that they are increasing in attentiveness. The distribution of student attentiveness is shown in the following figure:

The apparent left skew of the attentiveness distribution suggests that most students self-report themselves to be attentive.

Forward-looking vs. Myopic Perspective The baseline survey elicits student forward-looking perspectives through a series of questions. Each question is measured on a 7-point Likert 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, one question they are asked is "I consider myself

Figure 2: Distribution of Attentiveness

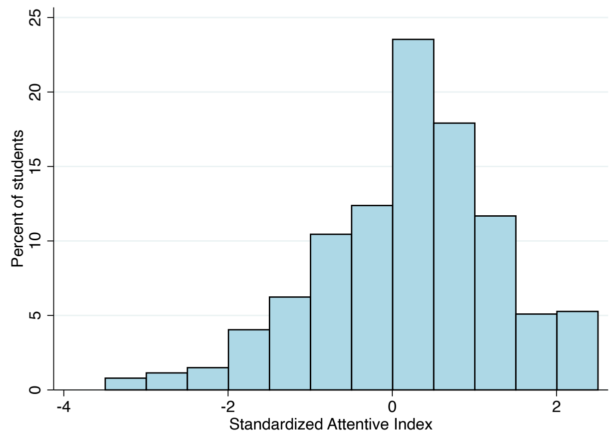

Notes: The figure presents the distribution of the standardized attentiveness index. All histograms use a bin width of 0.5 SD.

to be a forward-looking person who has clear plans about the future." To construct a continuous index of the forward-looking index, replies to all relevant questions are aggregated together so that they are increasing in the attribute of interest. See Appendix C.2 for the questions that formulate the forward-looking index. The distribution of student forward-looking views are shown in the following figure:

Figure 3: Distribution of Forward-looking Perspective

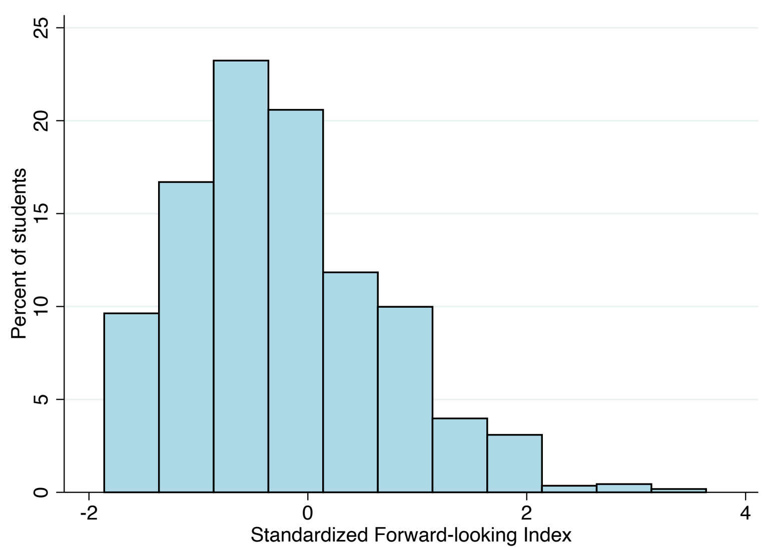

The right skew exhibited by the distribution of students' forward-looking perspectives suggests that a large fraction of students are myopic.

I dichotomize the forward-looking behavioral measures for expositional simplicity in later sections. The questions pertaining to students' forward-looking perspectives have their response values increasing with students' forward-looking views. I then characterize students as being forward-looking (versus myopic) if they responded with at least a value of 5 (out of 7) to each question.

3.3 Summary Statistics

Table 1 presents a rich set of summary statistics related to student demographics, characteristics, behavioral information, homework participation activity, and overall study time. Although computer science graduates are primarily male (Baer and DeOrio, 2020), there is no significant gender disparity in my sample as 49% of the students are female. Panel A shows the course is offered as a first-year course, but is not exclusive to first-year students, as around 38% of students are beyond their first year. Additionally, around 28% of students are pursuing non-STEM majors. Consistent with only 53% of students being domestic Canadians, only around 29% of students speak English at home. Appendix C.1 contains the survey questions used for gathering the student demographics and other characteristics.

Panel B of Table 1 shows that 87% of students do not have any programming experience prior to taking the course. Panel C shows that around 76% of students are attentive, and 32% of students are forward-looking. Panel D indicates that around 16% of students do not attempt the low-stakes homework each week. On average, students spend around 25 minutes watching videos each week, and 2 hours working on each homework assignment. Students who registered for the discussion board spend around 28 minutes each week on making posts or viewing content.

4 Experimental Design and Description of Interventions

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

4.1 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, Fall 2020, Winter 2021, and Summer 2021 academic terms. As discussed in the previous section, the data collection results in a sample of 3686 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. Prototyping interventions during the pilot data collection in Fall 2020 informed the design of the interventions presented in this section.

4.2 Description of the Intervention

The interventions considered in this study can be categorized as 'targeted informational reminders' as their design includes the following elements: 1) they prompt students to take a specific action, 2) they provide information on how to clearly execute the action, and 3) they serve as a reminder for the specified task. The design of the nudge is inspired by insights from psychology and behavioural economics research (Damgaard and Nielsen, 2018). In particular, the intervention is designed to nudge inattentive students who may have a tendency to forget homework deadlines.

Homework Reminder Messages. The reminder messages were aimed at promoting students to further participate in their weekly low-stakes homework. Reminders are only sent for the graded homework assessments after week 2.[17] The homework reminder is composed of the following three elements: 1) reminding students of the upcoming homework deadline, 2) prompting them to set aside time in their schedule to next make progress on the homework, and 3) including a direct link to the homework assessment. Appendix D.1 shows the template of the homework reminder message. The reminder messages were sent within 48 hours after the homework assignment was released and are deployed using the learning management system (i.e., Canvas). Students would receive the reminder both in their Canvas and institutional email inbox.

For students who had not completed the homework before the deployment of the reminder message, half of them are randomly assigned to receive a homework reminder.[18] The reminder messages were sent throughout the course and were re-randomized with each deployment. Consequently, the number of total homework reminders a student receives approximately follows a binomial distribution with 10 trials and a 0.5 probability of success. Figure 13 illustrates the assignment of students to the number of homework reminders.[19] Across the three learning stages (as noted earlier), stu-

[17] The courses instructors would make important announcements in the first two weeks to get students started with the course. Consequently, the reminders were not sent during this week to avoid crowding out the instructors' announcements.

[18] Each week, only around $5\% - 10\%$ of students completed the homework within 48 hours of release.

[19] Alternatively, students could have been uniformly randomized to receive between 0 to 10 reminders at the start of the course. This design was not implemented due to caveats that were discovered while piloting the reminder messages in the Summer 2019 and Fall 2019 cohorts. Since a small portion of students complete the homework soon after its release, these students should not receive a reminder. Additionally, it was important to check with the

dents can receive one of 75 distinct sequences of reminders, calculated as $3 \times 5 \times 5$. For example, a student could receive no reminders in the basic stage, 2 reminders in the intermediate stage, and finally 3 reminders at the advanced stage.

4.3 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 Intent to Treat (ITT) effects of interest.[20] The ITT is identified as students are randomly assigned to a control of treatment group each week. I investigate the validity of the random assignment by testing whether the pre-treatment student demographics and characteristics are balanced across the experimental conditions. I do so by standardizing each pre-treatment control and regressing these on the number of reminders received. Figure 14 shows that students who are assigned to receive an extra homework reminder 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 4091 students who initially agreed to participate in the study, around 90% of them completed the course. Table 2 examines whether the number of reminders received impacts the propensity to drop out. The analysis suggests that the reminder messages did not cause 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.

Well-defined Treatment Assignment. For the treatment allocation to be well-defined, the following two assumptions must hold true: 1) the treatment level is unique so that potential outcomes are well defined, and 2) the treatment applied to one student does not affect learning outcomes of other students. The intensity of the homework reminders is homogenous across the treatment groups as all students in the treatment condition receive the same reminder. Therefore, the potential outcomes corresponding to the experimental conditions are well defined.

instructor each week that the reminder message would not crowd-out any important announcement that they may of wanted to make.

[20] I cannot directly observe whether the students opened or read the reminder.

Next, I discuss the possibility of spillover effects across students. Since students can interact with each other on the discussion board and work towards solving problems, it is possible that students in the 'treatment' group who received the reminder will interact with the 'control' group who did not receive any reminder. Assuming the reminder increases an outcome of interest (e.g. more participation on homework problems), that can result in positive spillovers to the control group through information sharing (e.g., answering questions of control group students) or peer effects (e.g., control group student mimicking behaviour of treatment group student). Such positive spillover effects will result in downward biased effect sizes.[21]

Although the experimental design does not guard against such spillovers in this setting, the online nature of the course mitigates standard in-person student interactions that would typically be present. Additionally, I am able to leverage certain features of the data collection for robustness analysis. The baseline survey collected data on whether students are in a study group, the number of other students in the course they study with, and how frequently they meet. The final survey also directly asked students whether they discussed information shown in the reminder messages with other students. I use this survey data to discuss the robustness of my primary results to potential spillovers in the next section. Appendix C.5 includes the survey questions about student peer interactions.

5 Experimental Results

This section discusses the results from the field experiment described in Section 4, and outlines the corresponding empirical methodology. For simplicity of exposition and interpretation, the analysis is carried out by aggregating the data so that the parameters of interest are estimated by a cross-sectional regression with cohort fixed effects.

5.1 The Effect of Homework Reminders on Homework Participation

To measure the effect of receiving reminder messages on students' homework participation, I estimate the following specification:

$$D_{ic} = \delta_0 + \delta_1 RemindersFreq_{ic} + \pi_c + X_{ic}'\Delta + \epsilon_{ic},$$

where D_{ic} is either the number of homework assessments completed or the total hours a student

[21] Upwards biased estimates due to information spillover are possible, but unlikely. For example, suppose students in the control group discover that their peer in the treatment group received a reminder. Then the control group student may feel discouraged and exert less effort because they are not being supported by the instructional team.

spends studying; π_c is cohort fixed effects; $RemindersFreq_{ic}$ is total number of homework reminders a student receives. Control variables X_{ic} include student demographics and other pre-treatment characteristics listed in Panels A and B of Table 1.

Table 3 presents the results from estimating the above specification. On average, receiving 5 additional reminders induces students to complete an extra homework assessment. Additionally, the estimates show that receiving an extra reminder message increases the time spent on homework by 23 minutes. Since students spend around 2.4 hours each week on the online homework platform, receiving a homework reminder increases corresponding homework study time by around 16%. This effect size is also statistically significant at the 1% significance level, with an F-statistic exceeding 100. Consequently, the frequency of reminders received will provide a strong first stage for inducing exogenous variation in study time.[22]

Figure 15 illustrates the average number of homeworks completed by the number of reminder messages received. Clearly, receiving more reminders encourages students to complete more homework. The figure also suggests that the marginal increase in homework completion is decreasing with more reminders, although the apparent diminishing returns to homework reminders are not statistically significant.[23]

Mechanisms Underlying the Homework Reminders. Next, I examine whether the reminder messages are more helpful for less attentive students. The final survey asked whether the students found the reminder emails to be helpful in keeping on track with the homework assessments. Figure 16 illustrates the relationship between finding the reminders useful and a student's attentiveness. The significant negative linear association suggest that less attentive students are more likely to be helped by the homework reminders. The evidence suggests that the homework reminders are most effective at encouraging inattentive students to exert more effort.

5.2 The Effect of Homework Participation on Learning

Since students choose their level of homework participation, associating homework participation with learning outcomes will likely result in biased estimates due to omitted variable bias. For example, students who have a higher innate programming ability will obtain better grades on course assessments, while exerting less effort than students with lower innate programming ability. As a result, the returns to homework participation will be downward-biased through this unobserved

[22]The F-statistic exceeds the threshold of 104.6 stated in Lee et al. (2022) for assessing a strong first stage.

[23]Regressing homework completion on the reminder frequency and the square of the reminder frequency results in a negative, but statistically insignificant, coefficient on the quadratic term.

programming ability channel. To circumvent such issues of endogeneity, I use random assignment to the number of homework reminders received as an instrument for homework participation.

I argue that email reminders are a valid instrument for homework participation as they are randomly assigned to students (i.e., independent), do not directly affect learning outcomes (i.e., are excludable), and promote students to complete homework successfully (i.e., they are relevant).[24]

Relevance. A theoretical model presented in Ericson (2017) shows that reminders can be helpful in task completion when individuals have a limited memory. I find empirical evidence consistent with the models implication as Figure 16 shows that the reminder messages are perceived to be most beneficial by more inattentive students. Additionally, Table 3 shows that on average the reminder messages are successful in significantly increase effect exertion (as previous discussed).

Exclusion. Exclusion is violated if receiving homework reminders affects learning through channels aside from homework participation. This is plausible if the reminder message induces students to also participate further, for example, on the online discussion board. Receiving reminders over time may also help students build better time management and organization skills, enabling them to learn more by more efficiently using a given amount of study time. As the reminder message directly targets the homework assessment, it is unlikely to affect participation on the discussion board. Study habits and organizations ability will be well established prior to enrolling in the course. It is unlikely a low intensity nudge such as receiving a few reminder messages will have persistent effects on long term study habits.

Independence. Since students are randomly assigned to whether they receive a homework reminder for each homework assessment, then the number of reminders they receive is independent of any observed or unobserved determinants of learning. Concerns around finite sample imbalances across the type students who receive many reminders versus few reminders are alleviated by the large sample of students involved in this study. Evidence to support the independence assumption is presented in 14 (as previously discussed).

I now employ the frequency of reminders received as an instrument to estimate the causal effects of homework participation on homework performance using the following 2SLS model:

[24] I assume a homogeneous treatment effects framework for simplicity. Under a heterogeneous effects framework, monotonicity of the instrument and the absence of defiers can also be reasonably argued.

$$\begin{cases} ExamGrade_{ic} = \lambda_0 + \lambda_1 D_{ic} + \pi_c + X'_{ic}\Pi + \epsilon_{ic}, \\ D_{ic} = \phi_0 + \phi_1 RemindersFreq_{ic} + \pi_c + X'_{ic}\Gamma + \epsilon_{ic} \end{cases}$$

where $ExamGrade_{ic}$ denotes the final exam grade. Table 4 presents the 2SLS results. The estimates show that completing one extra homework increases the final exam grade by around 0.18 SD. Additionally, an extra hour spent studying through doing online homework increases final exam grade by 0.09 SD. These estimates are statistically significant at the 1% level. The large effects reflect the fact that the homework is the primary source of learning the course material in this online course.

5.3 Robustness to Spillover Effects

I now present two pieces of evidence supporting the view that the main results presented in this section are not severely affected by spillover effects from treated to control students. First, only around 9% of students in the final survey attested to discussing contents of the reminder messages with their peers at least once. Therefore, information spillovers from the treatment to the control group would be expected to be small. Second, around 17% of the students in the course are in study groups, where they meet at least once a month and discuss course material. I investigate whether the treatment effects for the reminder messages vary according to whether students are in a study group at baseline. The analysis is presented in Table 5. The results suggest that the efficacy of the reminder messages does not vary according to whether a student is in a study group.

6 Theoretical Framework

In this section, a learning environment with a cumulative structure is conceptualized, where students acquire knowledge by exerting effort over multiple learning stages. The framework formalizes an effective learning strategy in courses with a cumulative structure.

6.1 The Environment

Consider N students in a course, who allocate total study time or 'effort' (e) across three learning stages $t \in \{basic, int, adv\}$. Then, let L_i^t denote the amount of learning for student i during stage t. Students can vary in their baseline human capital (h). All students are assumed to be forward-looking who internalize the cumulative learning process when allocating effort (an assumption that we will relax later).

6.2 The Student Effort Choice Problem

Students allocate their effort to maximize their total knowledge net of effort costs as follows:

$$\max_{(e_i^t)_t} L_i(e_i^{basic}, e_i^{int}, e_i^{adv}; h_i) - C(e_i^{basic}, e_i^{int}, e_i^{adv}), \tag{1}$$

where $C(\cdot)$ is an convex function that is increasing in effort, representing the cost of effort exertion. The learning technology is concave in each effort input, increasing in effort and baseline human capital. For simplicity of the exposition, we assume total learning can be separated as follows:

$$L_i(e_i^{basic}, e_i^{int}, e_i^{adv}; h_i) = L_i^{basic}(e_i^{basic}, h_i) + L_i^{int}(e_i^{int}, L_i^{basic}) + L_i^{adv}(e_i^{adv}, L_i^{int}).$$

That is, the amount of learning L_i^t in a given period depends on present effort e_i^t, and previous knowledge L_i^{t-1}.[25] Similarly, I also assume the cost of effort is separable across learning stages:

$$C(e_i^{basic}, e_i^{int}, e_i^{adv}) = C(e_i^{basic}) + C(e_i^{int}) + C(e_i^{adv}),$$

imposing that student fatigue does not spillover across learning periods, a reasonable assumptions if student 'burn out' from exerting too much effort is not a major concern in the learning setting. For expositional simplicity, I also assume the cost function is identical across learning periods.

Students exert effort until the marginal benefit of effort exceeds the marginal cost of effort. The first order condition that characterizes effort exerted at the basic stage is:

$$\frac{dL_i}{de_i^{basic}} = \frac{dL_i^{basic}}{de_i^{basic}}\left[1 + \frac{dL_i^{int}}{dL_i^{basic}}\left(1 + \frac{dL_i^{adv}}{dL_i^{int}}\right)\right] = \frac{dC}{de_i^{basic}}.$$

Assuming a cumulative learning structure with previous knowledge persisting into future learning periods, then $\frac{dL_i^{int}}{dL_i^{basic}} > 0$ and $\frac{dL_i^{adv}}{dL_i^{int}} > 0$. That is, exerting effort in the basic stage has spillover learning benefits in the intermediate and advanced learning stages. Intermediate learning stage effort is determined by the following equation:

$$\frac{dL_i}{de_i^{int}} = \frac{dL_i^{int}}{de_i^{int}} \times \left[1 + \frac{dL_i^{adv}}{dL_i^{int}}\right] = \frac{dC}{de_i^{int}}.$$

Finally, effort in the advance learning stage is characterized by:

$$\frac{dL_i}{de_i^{adv}} = \frac{dL_i^{adv}}{de_i^{adv}} = \frac{dC}{de_i^{adv}}.$$

[25] Given that course structure is assumed to be cumulative, L_i^{t-1} is used as a sufficient statistic for all prior knowledge accumulation. Prior knowledge at the basic stage is denoted by L_i^{-1} and is the baseline knowledge h_i.

As the advanced stage is the last learning period, there is no future spillover learning benefit from exerting effort in the advanced stage. I show that under reasonable assumptions (see Appendix E), the effective learning strategy in a course with a cumulative structure is to front-load effort allocation so that $e_i^{basic,*} > e_i^{int,*} > e_i^{adv,*}$. The optimal effort exerted decreases across the learning stages as illustrated by the following figure:

Figure 4: Effective Effort Allocation When Learning is Cumulative

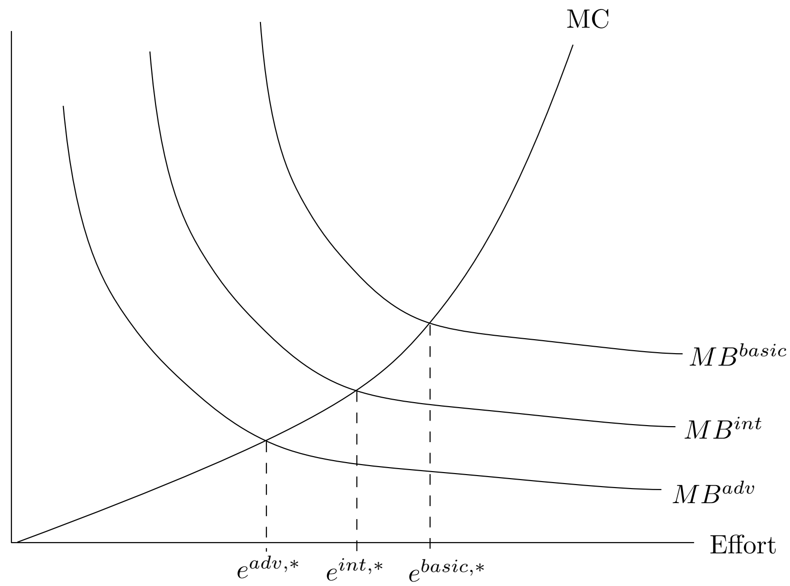

Notes: The figure illustrates the optimal effort allocation across the basic, intermediate, and advanced learning stage as the intersection of the respective marginal benefit and marginal cost of effort curves. The marginal benefit of effort shifts downwards across learning stages under a reasnoable cumulative learning structure.

6.3 Stylized Example

To intuitively illustrate the implications of the model, consider a course with a cumulative structure and two learning stages. Suppose that students learn basic concepts in the first half of the course and advanced concepts in the remaining half. That is, $t \in \{basic, adv\}$.

Parameterization of the Learning Technology and Cost Function. Let the following simple learning technologies represent the cumulative learning process:

$$L_i^{basic} = \alpha_0 + \alpha_1 e_i^{basic} + \alpha_2 h_i,$$

$$L_i^{adv} = \beta_0 + \beta_1 e_i^{adv} + \beta_2 L_i^{basic} + \beta_3 e_i^{adv} \times L_i^{basic}.$$

A positive marginal benefit of effort at both learning stages implies that $\alpha_1 > 0$ and $\beta_1 > 0$.

Since the advanced learning stage is cumulative, then clearly $\beta_2 > 0$. Finally, assuming effort exertion in the basic stage increases the productivity of advanced stage effort (i.e., there are dynamic complementarities in effort), then $\beta_3 > 0$.

The cost of effort is assumed to be linearly separable and represented by a quadratic cost function:

$$c(e_i^t) = \frac{(e_i^t)^2}{2} \text{ for } t \in \{basic, adv\}.$$

The Students' Optimal Effort Choice. The effective learning allocation acrosss the basic and advanced stage to maximize learning net of effort costs is given by:

$$e_i^{basic,*} = \frac{\alpha_1 \left[(1+\beta_2)(1-\beta_3\alpha_1)(1+\beta_3\alpha_1) + \beta_3(\beta_1 + \alpha_0\beta_3 + \alpha_1\beta_3(1+\beta_2) + \alpha_2 h_i)\right]}{(1-\beta_3\alpha_1)(1+\beta_3\alpha_1)},$$

$$e_i^{adv,*} = \frac{\beta_1 + \beta_3\left[\alpha_0 + \alpha_2 h_i + \alpha_1^2(1+\beta_2)\right]}{(1-\beta_3\alpha_1)(1+\beta_3\alpha_1)}.$$

Note that if there are no dynamic learning complimentarities (i.e., $\beta_3 = 0$), then the optimal effort allocation simplifies to $e_i^{basic,*} = \alpha_1(1+\beta_2)$, and $e_i^{adv,*} = \beta_1$. Additionally, suppose the course structure was not cumulative, and completely distinct content was covered in both learning stages (i.e., $\beta_2 = 0$). Then effort allocation at each stage would be the corresponding marginal effort benefit: $e_i^{basic,*} = \alpha_1$ and $e_i^{adv,*} = \beta_1$.

Simulating Optimal Effort Choice. Next, I will simulate an example effort allocation using the stylized model. Assuming learning performance is measured on a scale from 0 to 100 and effort is measured is hours of study time, I will consider the following reasonable parameter values: $\alpha_0 = 5, \alpha_1 = 3, \alpha_2 = 0.1, \beta_0 = 0, \beta_1 = 2, \beta_2 = 0.8, \beta_3 = 0.2$, and $h_i = 70$. Then the optimal effort allocation is 15.7 hours in the basic stage, and 11.9 hours in the advanced stage. That is, the student studies around 28 hours in this course, and slightly front-loads there effort allocation to the basic stage.

Figure 5 illustrates the fraction of achievable learning achieved as a function of basic stage effort when 28 hours in total area available for studying. As learning is cumulative in this setup, it is important to sufficiently exert effort in both learning periods, rather than focusing entirely on a single learning period.

Figure 5: Learning and Effort Allocaton

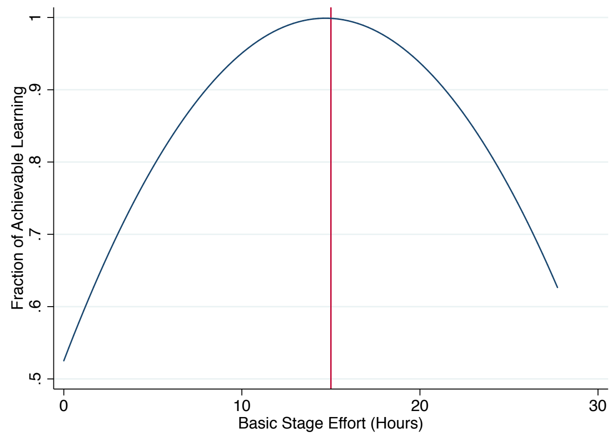

Notes: The figure shows learning as a function of basic stage effort with total study time fixed. The vertical line denotes optimal effort allocation to maximize learning.

6.4 The Instructor's Optimal Grading Scheme.

The model is flexible and can be naturally modified to reflect the actual course structure and introduce grading weights for the course assessments and student heterogeneity (see Appendix E). As the assessment in the advanced stage is cumulative, I assume the instructor wants students to allocate effort to maximize their advanced stage grade net of effort cost.[26] The optimal weights to incentivize the students to effectively allocate their study time varies according to whether the student is myopic or forward-looking. The following figure summarizes the model implications, simulating the proportion of total effort that is exerted in the basic stage as a function of the advanced grading weight (e.g., 70% advanced and 30% basic). I chose reasonable parameter values are used to represent a cumulative course structure.

The figure shows that myopic students' relative basic effort supply decreases sharply as less weight is assigned to the basic stage. In contrast, forward-looking students always allocate effort to the basic stage even if it does not count towards any course credit. The instructor and forward-looking student have the same objective function when all weight is placed on the advanced stage. Therefore, assigning 100% weight to the advanced stage is the optimal grading scheme for forward-looking students. However, for the myopic students, the instructor needs to assign appropriate weight to the basic stage to incentivize students to adequately exert effort during that stage. If the course structure is sufficiently cumulative, then the instructor should assign more weight to

[26]If the course lacks a cumulative structure and topics are independent, the instructor can simply assign grading weights proportional to each topic's importance. However, when the course is cumulative, students' performance on the final exam is best reflective of their learning.

Figure 6: Proportion of Total Effort Expended in the Basic Stage by Grading Scheme

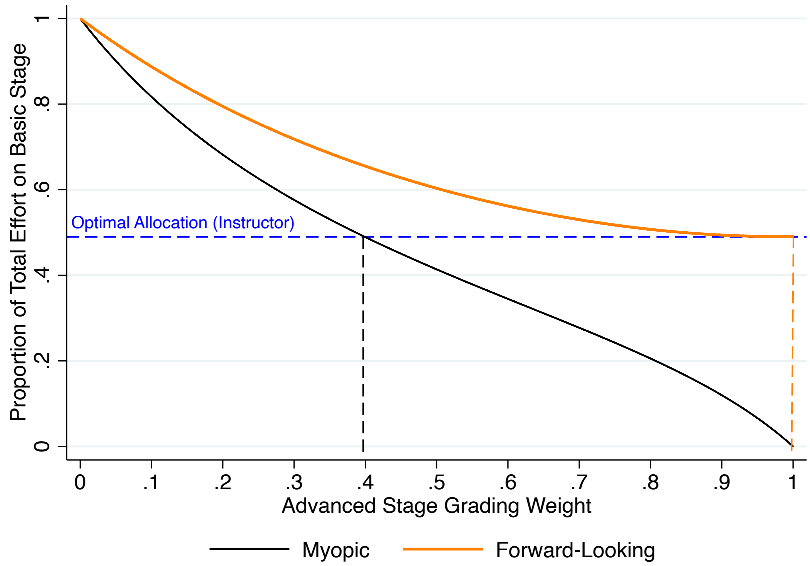

Notes: The figure presents the proportion of total effort allocated to the basic stage by myopic (black) and forward-looking (orange) students for a given advanced stage grading weight. The horizontal line denotes the optimal effort allocation from the instructors perspective (blue). The intersection points denote the optimal grading weights by student type.

the basic stage than the advanced stage (e.g., 60% to basic and 40% to advanced as illustrated). The discrepancy in the steepness of the effort allocation functions in Figure 6 shows that the effort allocation for myopic students are much more influenced by the grading scheme design than those of the forward-looking students. Consequently, the model implies that the dynamic allocation of grading weights is most important when a large share of students in the course are believed to be myopic.

7 Estimating a Cumulative Education Production Function

In this section, I describe the estimation of a multi-stage education production function. The parameterization of the production function is informed by the actual structure of the introductory programming course under consideration, noting that students learn across three distinct learning stages: basic (e.g., lists and loops), intermediate (e.g., nested lists and nested loops), and advanced (e.g., algorithms). Further, the estimation takes advantage of the unique data in this setting, the administrative data allowing me to observe in a precise way both the total online study time spent on each learning stage and the corresponding learning associated with each stage.

7.1 Specifying the Learning Technology

The technology maps effort inputs into contemporaneous learning for a given stage of the learning process. 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: Effort inputs across learning activities are substitutes within a period

To reasonably constrain the number of unknown parameters in the technology, I assume that effort investments in different online learning activities such as watching videos, doing problems, and time spent on the discussion board are perfect substitutes. That is, I only consider the aggregate observed online learning effort across these various activities in my analysis.[27]

Assumption 2: The learning technology is linear and additive in inputs

First, I assume the learning technology is linear and additive in effort and prior knowledge. The linear structure allows me to identify the marginal benefit of effort using instrumental variable estimation. Consistent with the cumulative nature of programming, I also assume the learning technology is cumulative.

Assumption 3: The learning technology is cumulative

Second, I assume the technology is cumulative, allowing learning beyond the basic stage to build upon previously attained knowledge. For example, I allow learning in the intermediate stage to be increasing in the knowledge accumulated in the basic stage. The cumulative technology reflects the cumulative course structure as programming topics build on each other.

Assumption 4: The learning technology includes dynamic complementarities in effort

Third, I allow for the productivity of study time in the present stage to depend on the knowledge accumulated in the previous stage. That is, dynamic interactions across effort inputs may be present in the production function. Putting all three assumptions together, the learning technology in the basic stage is as follows:

$$L_{ic}^{basic} = \alpha_0 + \alpha_1 e_{ic}^{basic} + \alpha_2 h_{ic} + \delta_c + \alpha_3 e_{ic}^{basic} \times h_{ic} + \epsilon_{ic}^{basic}, \tag{2}$$

where L_{ic}^{basic} is the basic stage homework performance; e_{ic}^{basic} is the total online study time at the basic learning stage; h_{ic} denotes baseline programming experience; δ_c is cohort fixed effects, and ϵ_{ic}^{basic} is a mean 0 stochastic error term. In equation 2, $\alpha_1 > 0$ implies a positive marginal benefit

[27] Alternatively I could use a CES–style production function to aggregate effort across different activities and flexibly allow for complimentary in–effort inputs. However, such a model will require many more types of exogenous effort shocks (e.g., for watching videos, doing problems, and contributing to the discussion board) than are available in the data for identification.

of effort. The extent to which prior programming knowledge persists to the basic stage is captured by $\alpha_2 > 0$ (Assumption 2). For $\alpha_3 > 0$, the marginal learning gains from basic effort exertion are increasing in baseline knowledge (Assumption 3). The learning technology at the intermediate and advanced learning stages are analogously defined as:

$$L_{ic}^{int} = \beta_0 + \beta_1 e_{ic}^{int} + \beta_2 L_{ic}^{basic} + \delta_c + \beta_3 e_{ic}^{int} \times L_{ic}^{basic} + \epsilon_{ic}^{int}, \tag{3}$$

$$L_{ic}^{adv} = \lambda_0 + \lambda_1 e_{ic}^{adv} + \lambda_2 L_{ic}^{int} + \delta_c + \lambda_3 e_{ic}^{adv} \times L_{ic}^{int} + \epsilon_{ic}^{adv}, \tag{4}$$

where L_{ic}^{int} is a sufficient statistic for previously accumulated knowledge in equation 4, reflecting the cumulative course structure. Given a positive marginal benefit of effort at each learning stage, dynamic complementarities in effort across the stages implies $\beta_3 > 0$ and $\lambda_3 > 0$. As the midterm is based on the basic stage material, the technology mapping basic stage effort to midterm performance is defined analogously to equation 2. Similarly the technology that maps effort inputs to the cumulative final exam performance is defined analogously to equation 4. The parametrization of the dynamic learning technology is flexible, yet tractable. I discuss possible extensions of the learning technology in Appendix F.

Identification of the Learning Technology. Identifying the cumulative technology requires exogenous variation in student effort, learning stage by learning stage. The marginal benefit parameters are identified using the exogenous variation in online learning participation within a student across the learning stages induced by the randomly assigned homework reminders throughout the course. Consistent with the cumulative course structure, the learning technology at each stage of the learning process is a function of present period total study time and previously accumulated knowledge. Although the field experiment allows for the possibility to construct many instruments, I construct two simple instruments for the two endogenous variables. In particular, I instrument for both endogenous variables by using the number of randomly assigned homework reminders a student receives at each learning stage. Therefore the repeated homework reminders identify marginal benefit parameters.

7.2 Estimation of the linear cumulative technology.

The marginal benefit of effort parameters are estimated using 2SLS by using the number of randomly assigned reminders a student receives at each learning stage to construct the relevant instruments (as noted). For the basic learning technology, the number of reminders received at the basic stage is used to instrument for total basic stage study time. The intermediate learning technology has two endogenous variables: the intermediate stage effort and basic stage knowledge. I use the number of reminders received at the basic and intermediate stages separately as instruments to estimate the intermediate learning technology. Following Angrist (2006), I estimate the learning technology with two endogenous variable as follows:

$$
\begin{cases}
L_{ic}^{int} = \beta_0 + \beta_1 e_{ic}^{int} + \beta_2 L_{ic}^{basic} + \delta_c + \beta_3 e_{ic}^{int} \times L_{ic}^{basic} + \epsilon_{ic}^{int}, \\
e_{ic}^{int} = \rho_0 + \rho_1 RemindersFreqBasic_{ic} + \rho_2 RemindersFreqInt_{ic} + \delta_c + \epsilon_{ic}^{int} \\
L_{ic}^{basic} = \gamma_0 + \gamma_1 RemindersFreqBasic_{ic} + \gamma_2 RemindersFreqInt_{ic} + \delta_c + \epsilon_{ic}^{basic}
\end{cases}
$$

where $RemindersFreqBasic_{ic}$ is the number of reminders received in the basic stage, and $RemindersFreqInt_{ic}$ is the number of reminders received in the intermediate stage. As the reminders are effective in inducing effort exertion at each learning stage, we expect $\rho_2 > 0$ and $\gamma_1 > 0$. The advanced learning technology is estimated analogously.

7.3 Discussion of Parameter Estimates

The marginal benefit parameter estimates are shown in Table 6. The estimates show a positive marginal benefit of effort at each learning stage. For example, at the basic stage, an additional hour spent studying increases students basic stage homework grade by around 0.14 SD for the average student. The estimates also show that exerting effort on the low-stakes homework assessments results in higher midterm and final exam performance. For example, in the advanced learning stage, I find an additional hour of online study time increases final exam grades by 0.11 SD, and this marginal benefit increases by 0.07 SD for every 1 SD increase in intermediate stage homework performance. The evidence is consistent with the online homework assessments being a primary source of learning in this course.

The estimates also indicate evidence of dynamic interactions in effort inputs across the learning stages. Complementarities in present effort exertion and previous knowledge are present in both the intermediate and advanced learning stages as β_3 and λ_3 are both statistically significant from 0. For example, the marginal benefit of studying for an hour in the advanced learning stage is increasing by 0.08 SD for each every 1 SD increase in intermediate homework performance. The results are consistent with most students having no prior programming experience, and also reflect the cumulative learning structure of programming. That is, students are accumulating programming knowledge over time through their effort exertion across varying programming topics that naturally build on each other.

Overall, the results suggest that students learning effectively should appropriately front-load their effort allocation, thereby becoming proficient in foundational skills that serve as the building blocks for rest of the course. Figure 7 illustrates the proportion of total effort, on average, allocated over the three learning stages and the corresponding grading weight.

Figure 7: Average Effort Allocation and Grading Scheme Across Learning Periods

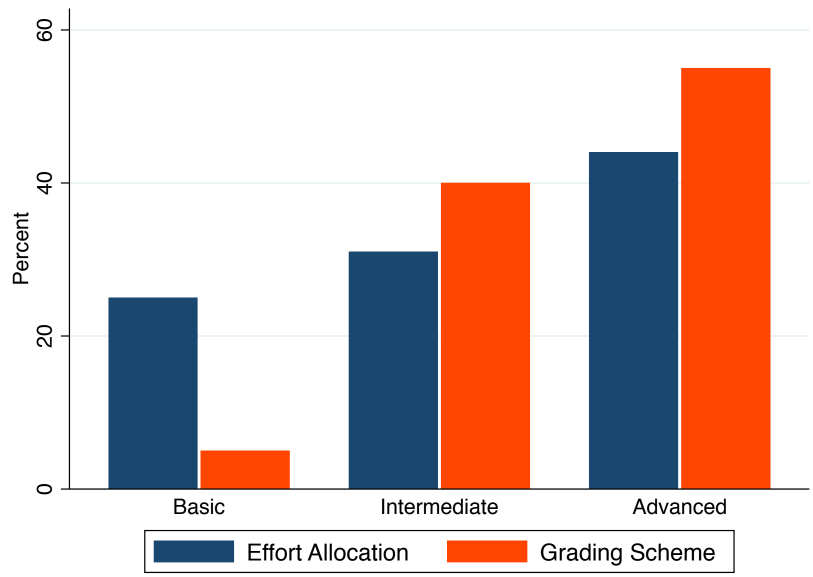

Notes: The figure presents the mean proportion of total effort allocated across the basic (homework only), intermediate (homework + midterm), and advanced learning stages (homework + exam). The corresponding grading weights are also presented.

Figure 7 shows that contrary to the effective learning strategy implied by the results, the average student's efforts are loaded towards the later learning stages. Such an effort allocation is consistent with the course grading scheme which increases as the course progress, as is typical in most courses.

8 Estimating a Model of Cumulative Learning

In this section, I describe the estimation of the model introduced in Section 6. The model is informed by the actual structure of the introductory programming course under consideration, noting that students learn across three distinct learning stages: basic (e.g., loops), intermediate (e.g., nested loops), and advanced (e.g., algorithms). Further, the model takes advantage of the unique data in this setting, the administrative data allowing me to observe in a precise way both the total online study time spent on each learning stage and the corresponding learning associated with each stage. Students' English proficiency, and forward-looking status are inferred from the survey data (see Appendix C.2 and C.3 for details).

8.1 Validating Theoretical Model Implications

The model presented above predicts that myopic students invest less effort early on in the course than forward-looking students. To test whether this implication holds in the data, Figure 8 shows the average proportion of total study time allocated to each learning stage by forward-looking students and myopic students.

Figure 8: Allocation of Effort Across Learning Stages

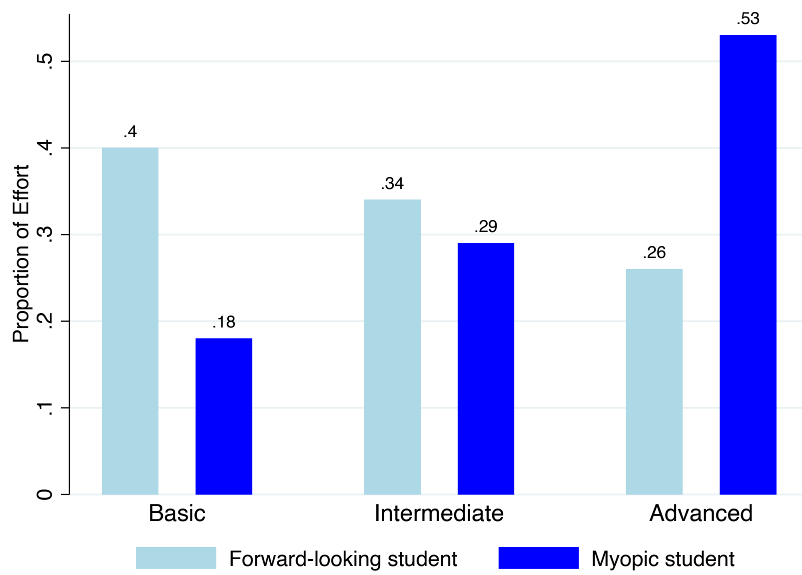

Notes: The figure presents the mean proportion of total effort allocated across the basic, intermediate, and advanced learning stages. Within each learning stage, the proportion of effort allocated is show the following order: (1) forward-looking student, and (2) myopic student.

The figure makes clear that on average, myopic students' effort increases as the course progresses, whereas forward-looking students front-load their study time allocation. The myopic students' effort allocation is consistent with the actual course grading scheme, which is increasing as the course progresses, as the initial two online assignments are ungraded, followed by the graded midterm in the middle, and the final exam with the most weight at the end. This evidence is consistent with the model prediction that myopic students exert less effort learning foundational skills, as they do not internalize the cumulative benefits of exerting effort early on, given the course structure, when choosing effort.

8.2 Estimating the Model

To rationalize the multi-stage effort-setting process by myopic and forward-looking students under a given grading scheme, I now estimate the three-learning stage model discussed earlier. Since students in the model choose effort optimally in order to balance their expected benefit against the costs, I estimate the cumulative learning technology and the implied learning benefits as well as a convex cost to effort exertion. The estimated model will serve as a foundation for conducting counterfactual experiments that inform the optimal grading scheme design, as anticipated in the Introduction. This approach is analogous to the one followed in Macartney et al. (2021). Those authors estimate a model of teacher effort as a function of accountability incentives under the No Child Left Behind Act of 2001, and use the estimated model to conduct counterfactual analyses of alternative incentive schemes.

8.3 Specifying the Cost Function

I impose some minimal structure on the cost function for tractability, making the following assumptions.

Assumption 1: The cost function is linearly separable across learning stages

The assumption that the cost function is linearly separable across the learning stages amounts to:

$$C(e_i^{basic}, e_i^{int}, e_i^{adv}; E_i) = C^{basic}(e_i^{basic}; E_i) + C^{int}(e_i^{int}; E_i) + C^{adv}(e_i^{adv}; E_i).$$

The assumption implies that exerting effort in a given learning stage does not affect the marginal cost of studying in any future period. A violation to this would be if students 'burn out' from exerting too much effort early in the course, impairing performance subsequently. As most students in the data allocate the majority of their study time in the advanced stage (as shown in Figure 8), the separability assumption is reasonable.

Assumption 2: The cost function is convex

The cost function for a given stage is specified as a power function:

$$C_i^t(e_i^t; E_i) = \kappa_t \exp(-\gamma_t E_i) \frac{(e_i^t)^{1+\gamma_t}}{1+\gamma_t} \text{ for } t \in \{basic, int, adv\},$$

where κ_t is the steepness, and $\gamma_t > 0$ represents the convexity. Intuitively, a convex cost of effort reflects the tendency for students to become fatigued the longer time they spend studying. The costs can vary at the student-level as those with a higher English proficiency will have a lower cost from exerting a unit of effort.

8.4 Estimating the Marginal Benefit of Effort Parameters

The marginal benefit of effort parameters estimates are taken from the estimated cumulative technology in the previous section.

8.5 Estimating Marginal Cost of Effort Parameters

After estimating the cumulative learning process, I estimate the marginal cost of effort parameters (scaling and convexity) $\Theta^t = (\kappa_t, \gamma_t)$ using maximum likelihood estimation (MLE) at each learning stage t. I construct the likelihood (l) using an 'implementations error' approach (Bernheim et al., 2019). That is, I assume students implement the optimal effort choice with error:

$$\underbrace{e_i^t}_{\text{Observed effort}} - \underbrace{e_i^{t,*}(\Theta^t)}_{\text{Optimal effort from model}} \sim \underbrace{N(0, \sigma_t^2)}_{\text{Deviation from optimal distribution}}, t \in \{basic, int, adv\}.$$

Then the resulting log-likelihood function is:

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

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

1. Start with an initial value of $\tilde{\Theta}^t$.
2. Compute $e_i^*(\tilde{\Theta}^t)$ for students $i = 1, \ldots, n$.
3. Use $\tilde{\Theta}^t, e_i^*(\tilde{\Theta}^t)$, and e_i for each student to compute likelihood $l(\tilde{\Theta}^t)$.
4. Update $\tilde{\Theta}^t$ to $\tilde{\Theta}^{t\prime}$ using Newton's method to take the next step to improve model fit.
5. Iterate through steps 2-4 until convergence: $|\tilde{\Theta}^t - \tilde{\Theta}^{t\prime}| < 10^{-6}$.

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

$$\hat{\Theta}_{MLE}^t = \arg\max_{\Theta^t} l(\Theta^t).$$

Maximum likelihood estimation is carried out in three learning stages since the cost function is linearly separable.[28] The marginal cost parameter estimates are shown in Table 7. The increasing estimates for the convexity parameter across learning stages are consistent with the course becoming progressively more difficult, increasing the rate at which fatigue from studying is accumulated.

8.6 Goodness of Fit

To evaluate the model fit, I start by comparing the distribution of observed effort to the effort distribution predicted by the model. Figure 17 shows clearly that the model fits the data well, as the mean and variance of effort distribution predicted by the model at each learning stage are closely aligned with the observed data. Additionally, I examine the association between observed effort and the corresponding effort implied by the model. Figure 18 shows that there is a strongly linear relationship between the model implied effort and observed effort at each learning stage. That is, students who exert effort well beyond the average are also predicted by the model to be exerting higher amounts of effort.

8.7 Identification of Model Parameters

Next I discuss the features of the data and sources of exogenous variation that drive the values of the parameter estimates, following best practices for structural research outlined in Andrews et al. (2020). I use rich micro-data and the field experiment to identify the marginal benefit and marginal cost of effort parameters at each learning stage.

Identification of Marginal Benefit Parameters. Identifying the cumulative technology requires randomization in student effort, learning stage by learning stage. The marginal benefit parameters are identified using the exogenous variation in online learning participation within a student across the learning stages induced by the randomly assigned homework reminders throughout the course. Consistent with the cumulative course structure, the learning technology at each stage of the learning process is a function of present period total study time and previously accumulated knowledge. I can instrument for both endogenous variables by using the number of randomly assigned homework reminders a student receives at each learning stage. Therefore the repeated homework reminders identify marginal benefit parameters.

[28] I also repeat the iterative MLE for several initial values and find the resulting estimates are fairly stable, suggesting a global optimum is achieved.

Identification of Marginal Cost Parameters: The marginal cost parameter vector is $\Theta_t = (\gamma_t, \kappa_t)$ where t indexes the relevant learning stage. While these parameters are estimated jointly using MLE, I will discuss moments in the data induced through the homework reminders that identify the cost of effort function at each stage. The key cost parameter is γ_t, which represents the convexity of the cost function. Since the cost function is assumed to be linearly separable, identification can be considered at each stage separately. Through log-linearizing the first-order conditions of the students' problem, it can be shown that γ_t is proportional to the inverse effort elasticity with respect to the marginal learning benefit after conditioning on English proficiency. Identification of the convexity parameter therefore requires exogenous variation in effort across students within each learning stage. Since the reminders were deployed at each learning stage, they served to exogenously shock study time at each learning stage. Consequently, I can use random assignment to receiving reminders to identify the convexity parameter of the cost function at each stage.

Intuitively, consider two distinct myopic students i and j, who have the same English proficiency, i.e., $E_i = E_j$. Then both students have the same marginal cost to exerting effort at each stage. Suppose, however, that $L_i^{basic} > L_j^{basic}$, as student i further completes the homework at the basic stage due to randomly receiving a basic stage homework reminder. That is, the homework reminder served as an exogenous shifter of the marginal benefit of effort in the intermediate stage. Then differencing the log-linearized first order condition across the two students within the intermediate stage results in the following equation:

$$log\left(\frac{e_i^{int,*}}{e_j^{int,*}}\right) = \frac{1}{\gamma_{int}} log\left(\frac{\beta_1 + \beta_3 L_i^{basic}}{\beta_1 + \beta_3 L_j^{basic}}\right).$$

Then, the above equation makes clear that γ_{int} can be uniquely recovered from the data given that exogenous variation in basic stage knowledge between students with identical English proficiency is available. Given linear separability of the cost function, the identification argument is analogous for γ_{basic} and γ_{adv}.[29] Finally, the scaling parameter κ_t at each learning stage is identified from the corresponding effort-setting first order condition as it is a function of identified quantities.

In summary, as the cost of effort function is characterized by two unknown parameters, at least two moments of the data at each stage are required for identification. Such moments can be constructed using students average effort allocation using the field experiment at each stage, conditioning on students' English proficiency.

[29] I treat baseline knowledge h_i as exogenous within the model.

9 Counterfactual Experiments

This section conducts a set of counterfactual experiments that would be difficult to carry out in the field at a large scale. It does so using the estimated marginal benefit and cost of effort parameters from the behavioural model of student effort supply. First, I simulate the grading scheme that maximizes learning for the existing distribution of myopic and forward-looking students. I compare the simulated final exam grade to the learning acquired under the existing course structure. Aside from total effort exertion, the simulations show that the specific allocation of effort across the stages of the course is also relevant, as learning in programming is cumulative with strong dynamic complementarities. After establishing the optimal grading scheme for a course where all students are either myopic or forward-looking, I simulate the optimal grading scheme for any share of myopic students.

Identifying the Student Effort Response to Different Grading Schemes. To study how student effort allocation changes with the design of the grading incentives, the ideal experiment would randomize a large number of students to a variety of different grading schemes. Although this experimental variation in incentives is unavailable in practice, the model leverages within-student variation in grading weights along with the observed differential responses to incentives by myopic and forward-looking students to identify how changing the grading weights influences the effort-setting process. All marginal benefit and cost of effort parameters in the previous section are estimated using unweighted learning measures and are held constant in the counterfactuals.

9.1 Optimal Assignment Weights

Setup of Optimal Assignment Weights Counterfactual. It should be noted that many universities require instructors to place more weight on the final exam than the midterm. Assigning a large fraction of the weight to the midterm may decrease course completion as students who perform unexpectedly poorly on the midterm through receiving a negative shock (e.g., they become sick on midterm day) may drop out early. I consider a realistic counterfactual exercise where the weights on the midterm and final exam remain fixed, and the assignment weights can vary. Aligned with the actual course grading scheme, I set the midterm weight to be 30% and the final exam weight, 45%. Since myopic and forward-looking students respond differentially to grading incentives, the instructor also internalizes the proportion of myopic students in the course when setting the grading scheme; using survey data, I infer that around 68% of the students are myopic.[30]

[30] I explore the sensitivity of the main simulations to changes in this proportion below.

The grading scheme is also catered to the representative student, with average English language proficiency and average prior programming experience. Aside from the marginal benefit and cost parameters, the remaining fixed parameters of the simulation setup are as follows:

$$w_{midterm} = 0.30, w_{final} = 0.45, h_i = \bar{h}, E_i = \bar{E}, p_m = 0.68.$$

Then, the instructor's problem is to set assignment weights to incentivize the representative student to allocate effort to maximize their grade in the advanced learning stage, captured by their performance on the cumulative final exam:

$$\max_{w_{basic}, w_{int}, w_{adv}} p_m[L_i^{final}(e_i^{adv,*}; e_i^{int,*}, e_i^{basic,*}, \bar{h}, f_i = 0) - \sum_t C(e_i^{t,*}; \bar{E}, f_i = 0)] -$$
$$(1 - p_m)[L_i^{final}(e_i^{adv,*}; e_i^{int,*}, e_i^{basic,*}, \bar{h}, f_i = 1) - \sum_t C(e_i^{t,*}; \bar{E}, f_i = 1)],$$
$$\text{subject to } w_{basic} + w_{int} + w_{adv} = 0.25,$$

where the cumulative learning technology and cost functions are as estimated in the previous section. Since myopic students' effort-choice is considered independently at each stage, their effort at each learning stage is a function of the stage-level assignment weight and prior knowledge in the basic and intermediate stage. That is, $e_i^t(f_i = 0) = e_i^t(w_t; f_i = 0, L_i^{t-1})$ for $t \in \{basic, int\}$. Given the large weight on the final exam, myopic students do internalize that doing homework in the advanced stage helps prepare for the final exam, i.e., $e_i^{adv}(f_i = 0) = e_i^{adv}(w_{adv}, w_{final}; f_i = 0, L_i^{int})$. As the forward-looking students internalize the cumulative course structure, their effort at each stage is a function of current and future assessment weights, and prior knowledge:

$$e_i^{basic,*}(f_i = 1) = e_i^{basic,*}(w_{basic}, w_{int}, w_{adv}, w_{mid}, w_{final}; f_i = 1, h_i),$$

$$e_i^{int,*}(f_i = 1) = e_i^{int,*}(w_{int}, w_{adv}, w_{final}; f_i = 1, L_i^{basic}),$$
$$e_i^{adv,*}(f_i = 1) = e_i^{adv,*}(w_{adv}, w_{final}; f_i = 1, L_i^{int}).$$

(See Appendix E for a fuller discussion of how myopic and forward-looking students choose effort by balancing the benefits and costs of effort exertion.) As the course includes a fixed number of instructional videos and assignment problems at each learning stage, there is a natural upper bound on effort exertion for the average student.[31] In the counterfactual simulations, the students can

[31] It is possible for a student to watch videos repeatedly, however only a few students watch all the videos more

only exert effort up to a reasonable amount of study time $\bar{e}^t$, where the upper bound is inferred from the administrative data. That is, the model implied effort at each learning stage is $\max\{e_i^{t,*}, \bar{e}^t\}$ for $t \in \{basic, int, adv\}$. Then the optimal assignment weights are determined through the following iterative routine:

1. Start with an initial value of $w_{assign} = (w_{basic}, w_{int}, w_{adv})$ and set $p_m = 0.68$.
2. Compute effort $e_i^{t,*}(w_{assign})$ for myopic ($f_i = 0$) and forward-looking ($f_i = 1$) students for each stage $t \in \{basic, int, adv\}$.
3. Use effort allocations to compute learning $L_i^{final}(e_i^{adv,*}; e_i^{int,*}, e_i^{basic,*}, \bar{h})$ for myopic ($f_i = 0$) and forward-looking ($f_i = 1$) students.
4. Use effort allocations to compute total cost of effort $\sum_t C(e_i^{t,*})$ for myopic and forward-looking students.
5. Evaluate instructor's objective function using knowledge accumulated in the advanced stage and the total effort cost as computed in steps 3 and 4, respectively.
6. Update w_{assign} to w'_{assign} using Newton's method to take the next step.
7. Iterate through steps 2-6 until convergence.

The above procedure is repeated for numerous initial values of w_{assign}. The simulated assignment weights that maximize learning of a representative student net of effort costs are then selected.

Simulation of Optimal Assignment Weights. Figure 9 compares the optimal assignment weights with the existing course grading scheme. In contrast to the existing grading scheme, the optimal counterfactual grading scheme weights are decreasing across the learning stages: $w_{basic} = 0.11, w_{int} = 0.09, w_{adv} = 0.05$. The optimal grading scheme is expected to increase myopic students' final exam grade by 2.5 pp (0.11 SD) while having no significant impact on the forward-looking students. Under the optimal grading scheme, myopic students are incentivized to front-load their effort allocation, while mitigating distortions introduced to the optimal effort allocation of forward-looking students. The following figure shows the proportion of effort exerted by the representative myopic and forward-looking students under the optimal grading scheme.

Figure 10a clearly shows that under the simulated grading scheme a myopic student exerts a larger fraction of their total effort in the basic stage relative to the actual course grading scheme.

than once.

Figure 9: Optimal Assignment Weights

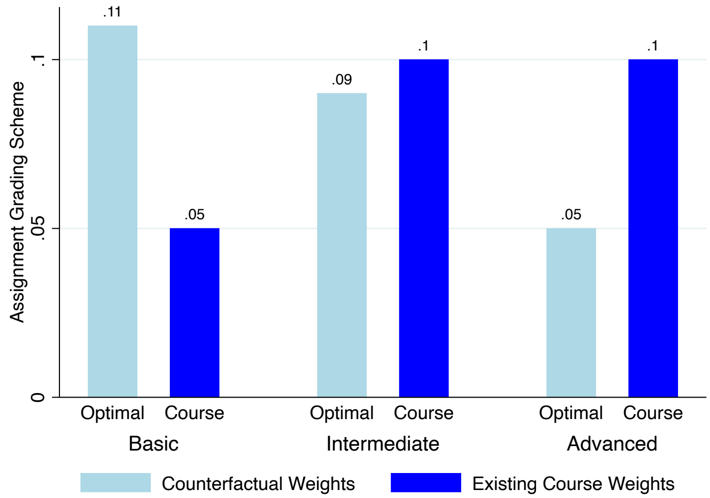

Notes: The figure presents the optimal assignment weights together with the actual course grading scheme. Within each learning stage, the assignment weights are shown in the following order: (1) optimal simulated weights, and (2) the actual course assignment weights.

Figure 10: Allocation of Effort Across Learning Stages by Grading Scheme

(a) Myopic Student

(b) Forward-looking Student

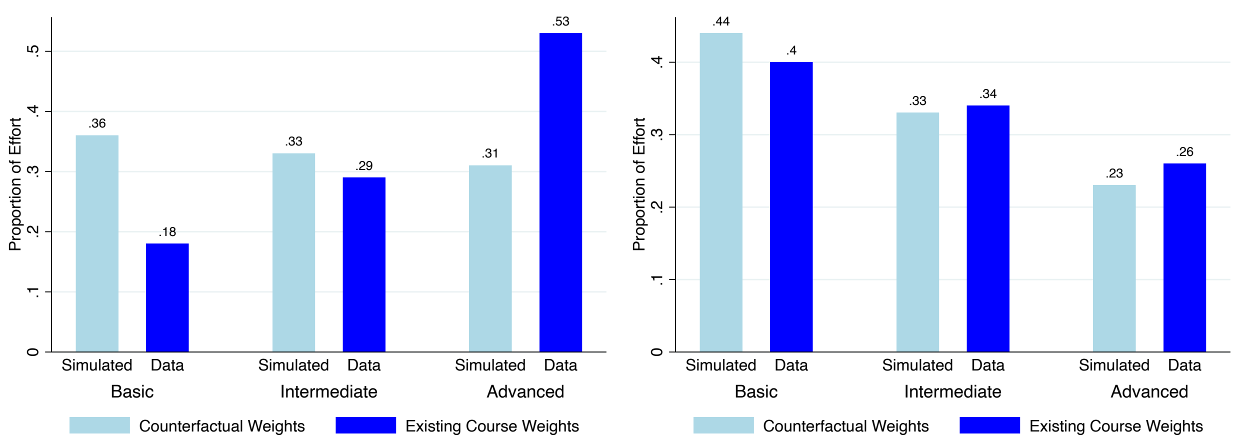

Notes: The figure presents the relative distribution of total study time allocation across the learning stages. Effort allocation under the optimal grading scheme and the actual course grading weights are presented separately for myopic (left) and forward-looking (right) students.

Myopic students effort supply can be easily influenced in the basic stage by assigning it more weight. In contrast, a forward-looking student's effort supply in the basic stage is less malleable to changes in the assignment grading weights. This is because, the forward-looking student internalizes the cumulative course structure and the increased productivity of study time in later stages resulting from exerting effort in the basic stage. Figure 10b shows forward-looking students only slightly increase their effort in earlier stages under the front-loaded counterfactual grading scheme. Comparing Figures 10a and 10b also shows that forward-looking students in the actual course front-load their effort to a greater degree than myopic students under the simulated optimal grading scheme. The evidence thus suggests that student learning can be further improved if both myopic and forward-looking students are incentivized to allocate their type-specific effort most efficiently through receiving personalized grading schemes.

Optimal Assignment Weights by Proportion of Myopic Students. I will now generate the optimal grading scheme for any distribution of myopic and forward-looking students. To do so, I simulate the optimal assignment weights as the proportion of myopic students increases from 0 to 1. The following figure shows how the optimal grading weights transition from being loaded on the advanced stage, intermediate stage, and then the basic stage as the share of myopic students increases.

Figure 11: Optimal Assignment Weights by Share of Myopic Students

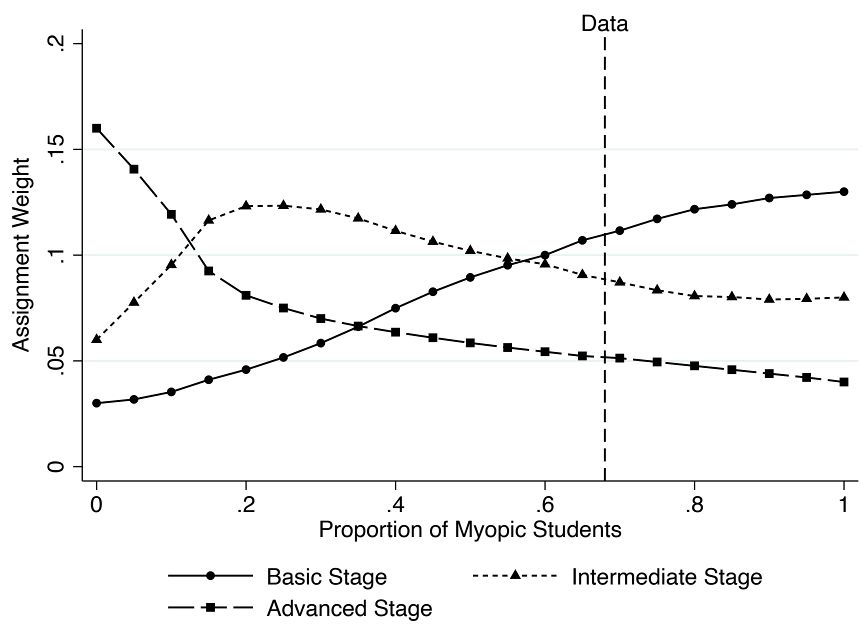

Notes: The figure presents the optimal assignment weight for each learning stage by 0.05 increments in the proportion of myopic students. The intersection of the assignment weights with the dashed vertical line indicates the optimal grading scheme in the actual course with a 68% share of myopic students.

Figure 11 shows that as the share of myopic students in a course increases from 0 to 20%, there is a rapid decline in the weight assigned to the advanced stage and a gradual increase in the basic assignment weight. The corresponding steep increase in the intermediate assignment weight suggests that the advanced stage weight is substituted more on the intermediate stage than the basic stage. That is, for a course with a modest 20% share of myopic students, the optimal assignment weights are loaded towards the middle of the course. Assigning the majority of the weight to the middle of the course caters to both myopic and forward-looking students. The myopic students are incentivized to exert effort prior to entering the advanced stage, while mitigating losses in learning to forward-looking students by deviating the incentives more on the intermediate rather than the basic stage. Next, as the share of myopic students exceeds 55% as observed in the actual course, the assignment weights are loaded most to the basic stage. As a result, front-loading the assignment weights is optimal when a modest majority of students are myopic.

10 Conclusion

The share of students pursuing education in STEM has been increasing rapidly over the past decade, especiallly in Computer Science. Although STEM graduates are a vital input to any modern society, learning advanced skills can be challenging given their cumulative nature. Identifying the cumulative learning process is beneficial for understanding effective learning strategies and approaches to design in foundational university courses. In light of the benefits, this paper presented an approach to credibly estimating the learning process while circumventing the typical issues of measurement error and endogenity of effort inputs.

The empirical approach I developed involves several components. I employed rich administrative data from a large pre-existing foundational online STEM course that has a cumulative structure. I then conducted randomized interventions that were successful in nudging students to spend more time learning and complete more online assignments throughout the course. The administrative dataset, including precise measures of student study time allocation, together with the field experiment were then used to estimate a cumulative education production function. Exogenous experimental variation arising from the field experiment served to credibly identify the marginal benefit each stage of the cumulative learning process built into the STEM course.

The findings presented help to inform the design of large foundational courses that apply to both online and traditional in-person setups. First, completing low-stakes online assignments throughout the course is essential for student learning: spending an extra hour on online assignments increases final exam grades by 0.09 SD (noting that online homework is the key means of learning in the course). Second, I find evidence of strong dynamic learning complementarities in the cumulative education production function. That is, the productivity of effort learning advanced skills is increasing in prior knowledge. Contrary to this evidence, however, I find that students effort allocation increases as the course progresses. Then, given a cumulative course structure, an instructor can consider setting the assignment grading weights so that students exert sufficient effort early on learning the fundamentals.

The counterfactual analysis reveals that courses with a cumulative course structure and the presence of many myopic students, an instructor can usefully adjust the assignment grading weights. That is, in contrast to the actual course grading scheme with homework incentives increasing across the three learning stages (basic, intermediate, and advanced), I find that the simulated assignment weights that maximize learning are decreasing as the course progresses. The optimal weights serve to encourage myopic students to adequately invest effort to learn foundational skills early in the course, increasing their final exam grade by 0.11 SD. Additional simulations suggest that more generally, in a course with a strong cumulative structure, incentives should be front-loaded when majority of students are myopic, middle-loaded when a modest share of students are myopic, and end-loaded when the vast majority of students are forward-looking.

References

Aizer, Anna and Flavio Cunha (2012) "The production of human capital: Endowments, investments and fertility," Working Paper 18429, National Bureau of Economic Research, http://www.nber.org/papers/w18429.

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

Angrist, Joshua D (2006) "Instrumental variables methods in experimental criminological research: what, why and how," *Journal of Experimental Criminology*, 2 (1), 23–44.

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.

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

Clark, Damon, David Gill, Victoria Prowse, and Mark Rush (2020) "Using goals to motivate college students: Theory and evidence from field experiments," *Review of Economics and Statistics*, 102 (4), 648–663.

Cotton, Christopher S, Brent R Hickman, John A List, Joseph Price, and Sutanuka Roy (2026) "Why don't struggling students do their homework? disentangling motivation and study productivity as drivers of human capital formation," *Journal of Political Economy*, 134 (1), 000–000.

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

Ericson, Keith Marzilli (2017) "On the interaction of memory and procrastination: Implications for reminders, deadlines, and empirical estimation," *Journal of the European Economic Association*, 15 (3), 692–719.

Ersoy, Fulya (2021) "Returns to effort: experimental evidence from an online language platform," *Experimental Economics*, 24 (3), 1047–1073.

Gilraine, Michael (2016) "School accountability and the dynamics of human capital formation," http://tinyurl.com/Gilraine-JMP.

Hanushek, Eric A and Ludger Woessmann (2015) *The knowledge capital of nations: Education and the economics of growth*: MIT press.

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, Rene 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.

Lee, David S, Justin McCrary, Marcelo J Moreira, and Jack Porter (2022) "Valid t-ratio Inference for IV," *American Economic Review*, 112 (10), 3260–90.

Light, Jacob (2024) "Student demand and the supply of college courses," *Available at SSRN 4856488*.

Macartney, Hugh, Robert McMillan, and Uros Petronijevic (2021) "A Quantitative Framework for Analyzing the Distributional Effects of Incentive Schemes," Working Paper 28816, National Bureau of Economic Research, https://www.nber.org/papers/w28816.

Oreopoulos, Philip, Richard W Patterson, Uros Petronijevic, and Nolan G Pope (2018) "When studying and nudging don't go as planned: Unsuccessful attempts to help traditional and online college students," Working Paper 25036, National Bureau of Economic Research, https://www.nber.org/papers/w25036.

Smith, Ben O, Dustin R White, Patricia C Kuzyk, and James E Tierney (2018) "Improved grade outcomes with an e-mailed grade nudge," *The Journal of Economic Education*, 49 (1), 1–7.

Stinebrickner, Ralph and Todd R Stinebrickner (2008) "The causal effect of studying on academic performance," *The BE Journal of Economic Analysis & Policy*, 8 (1).

——— (2014) "A major in science? Initial beliefs and final outcomes for college major and dropout," *Review of Economic Studies*, 81 (1), 426–472.

Stinebrickner, Todd and Ralph Stinebrickner (2012) "Learning about academic ability and the college dropout decision," *Journal of Labor Economics*, 30 (4), 707–748.

Todd, Petra E and Kenneth I Wolpin (2007) "The production of cognitive achievement in children: Home, school, and racial test score gaps," *Journal of Human capital*, 1 (1), 91–136.

Tables

Table 1: Student Level Summary Statistics

Variable	Mean	Std. Dev.	Min.	Max.	N
Panel A: Demographics					
I(Female)	0.491	0.499	0	1	3686
I(First year of university)	0.618	0.485	0	1	3686
I(Domestic student)	0.531	0.499	0	1	3686
I(Speaks English at home)	0.286	0.452	0	1	3686
I(First generation university)	0.173	0.378	0	1	3686
I(Mother at least college graduate)	0.668	0.459	0	1	3686
I(Father at least college graduate)	0.700	0.458	0	1	3686
Panel B: Characteristics					
I(Has some prior programming experience)	0.134	0.341	0	1	3686
I(Course required for major)	0.736	0.441	0	1	3686
I(Pursuing STEM major)	0.717	0.448	0	1	3686
Panel C: Behavioural Characteristics					
I(Student is attentive)	0.762	0.426	0	1	3686
I(Student is forward-looking)	0.315	0.465	0	1	3686
Panel D: Online Homework Participation					
I(Started weekly homework)	0.843	0.367	0	1	3686
I(Completed weekly homework)	0.671	0.471	0	1	3686
Weekly unique minutes of videos watched	24.52	9.211	0	38.35	3686
Weekly minutes spent doing problems	122.49	63.107	0	434.13	3686
Panel E: Discussion Board Participation					
I(Registered for course discussion board)	0.791	0.407	0	1	3686
No. of total contributions	3.57	14.087	0	237	2911
No. of unique posts viewed	121.88	149.28	0	1022	2911
Weekly minutes spent on discussion board	28.61	14.51	0	187.66	2911

Notes: Table presents descriptive statistics related to student demographic and characteristics, discussion board participation and online homework activity. Statistics shown in Panel A and B are formulated using self-reported student responses on the baseline survey (see Appendix C.1). Panel C uses survey data to characterize students as attentive (see Appendix C.2). Panel D statistics are formulated from the administrative data of the online homework platform (C.4). Finally, the statistics shown in Panel E are computed using data gathered from the discussion board.

Table 2: Student Attrition and Treatment Allocation

	(1) I(Dropped course)	(2) I(Dropped course)
No. of reminder messages received	0.0116	-0.0061
	(0.0898)	(0.0192)
Controls	No	Yes
No. of Students	4091	4091
R-squared	0.0015	0.13

Notes: Table shows differential attrition rate by intensity of the treatment condition. Controls include pre-treatment student demographics and characteristics included in Panels A and B of Table 1, and cohort fixed effects. Significance levels are represented by $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$.

Table 3: Homework Reminders and Online Homework Participation

	(1) Homework Completion[a]	(2) Homework Completion	(3) Study Time[b]	(4) Study Time
No. of reminders received	0.184***	0.187***	23.83***	22.71***
	(0.0126)	(0.0128)	(2.1663)	(2.0783)
Controls	No	Yes	No	Yes
F-stat for treatment	213.15	217.61	121.31	120.24
Adjusted R-square	0.125	0.358	0.133	0.326
No. of Students	3686	3686	3686	3686

Notes: [a]Homework completion is defined as the student attempting all problems with a positive score. [b]Total minutes spent watching videos and working on homework problems. Students can receive at most 10 reminder messages. Controls include pre-treatment student demographic and characteristics included in Panels A and B of Table 1, and cohort fixed effects. Significance levels are represented by $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$.

Table 4: Student Online Learning Participation and Final Exam Grade (2SLS)

	(1) Exam Performance[a]	(2) Exam Performance	(3) Exam Performance	(4) Exam Performance
Homework Completion[b]	0.181***	0.176***		
	(0.0540)	(0.0516)		
Total Study Time (Hours)[c]			0.084***	0.091***
			(0.0312)	(0.0322)
Controls	No	Yes	No	Yes
Adjusted R-square	0.092	0.262	0.083	0.241
No. of Students	3686	3686	3686	3686

Notes: [a]Standardized final exam grade. [b]Homework completion is defined as the student attempting all problems with a positive score. [c]Time spent watching videos and working on homework problems. Total number of reminders received is used as an instrument for homework participation. Controls include pre-treatment student demographic and characteristics included in Panels A and B of Table 1, and cohort fixed effects. Significance levels are represented by $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$.

Table 5: Efficacy of Reminders by Study Group Involvement

	(1) Homework Completed[b]	(2) Homework Completed
I(Study group) $\times$ No. of reminders	0.061	0.043
	(0.0537)	(0.0317)
Controls	No	Yes
Adjusted R-square	0.142	0.371
No. of Students	3686	3686

Notes: Indicator for whether a student is in a study group and the number of reminders received 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. Significance levels are represented by $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$.

Table 6: Benefit of Effort Parameter Estimates (2SLS)

Parameter	Estimate (SE)
Panel A: Basic Learning Stage Performance	
$\widehat{\alpha_1}$ (basic minutes study)	0.00241*** (0.00061)
$\widehat{\alpha_2}$ (baseline knowledge)	0.114** (0.0553)
$\widehat{\alpha_3}$ (basic effort $\times$ baseline knowledge)	0.00126* (0.00072)
Panel B: Midterm Performance	
$\widehat{\theta_1}$ (basic minutes study)	0.00221*** (0.00059)
$\widehat{\theta_2}$ (baseline knowledge)	0.093* (0.0492)
$\widehat{\theta_3}$ (basic effort $\times$ baseline knowledge)	0.00214 (0.00358)
Panel C: Intermediate Learning Stage Performance	
$\widehat{\beta_1}$ (intermediate minutes study)	0.00212*** (0.00052)
$\widehat{\beta_2}$ (basic knowledge)	0.166*** (0.0442)
$\widehat{\beta_3}$ (int. effort $\times$ basic knowledge)	0.00111** (0.00054)
Panel D: Advanced Learning Stage Performance	
$\widehat{\lambda_1}$ (intermediate minutes study)	0.00178** (0.00087)
$\widehat{\lambda_2}$ (intermediate knowledge)	0.183*** (0.0441)
$\widehat{\lambda_3}$ (adv. effort $\times$ int. knowledge)	0.00137** (0.00069)
Panel E: Final Exam Performance	
$\widehat{\pi_1}$ (adv. minutes study)	0.00189** (0.00095)
$\widehat{\pi_2}$ (intermediate knowledge)	0.155*** (0.0418)
$\widehat{\pi_3}$ (adv. effort $\times$ int. knowledge)	0.00123** (0.00061)
No. of students	3686

Notes: All assessment performances are standardized. Baseline knowledge is a standardized measure that aggregates prior programming knowledge and cGPA. The number of reminders received at each stage are the instrumental variables for study time and prior stage knowledge. Significance levels are represented by $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$.

Table 7: Cost of Effort Parameter Estimates (MLE)

Parameter	Estimate (SE)
Panel A: Basic Learning Stage	
$\widehat{1+\gamma_1}$ (Convexity)	2.1367***
	(0.5014)
$\widehat{\kappa_1}$ (Steepness)	0.00038***
	(0.00014)
Panel B: Intermediate Learning Stage	
$\widehat{1+\gamma_2}$	2.3531***
	(0.4424)
$\widehat{\kappa_2}$	0.00038***
	(0.00013)
Panel C: Advanced Learning Stage	
$\widehat{1+\gamma_3}$	2.5121***
	(0.5430)
$\widehat{\kappa_3}$	0.00039***
	(0.00014)
No. of students	3686

Notes: Performance within each stage is standardized. Cost 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 $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$.

Figures

Figure 12: Timeline of Data Collection

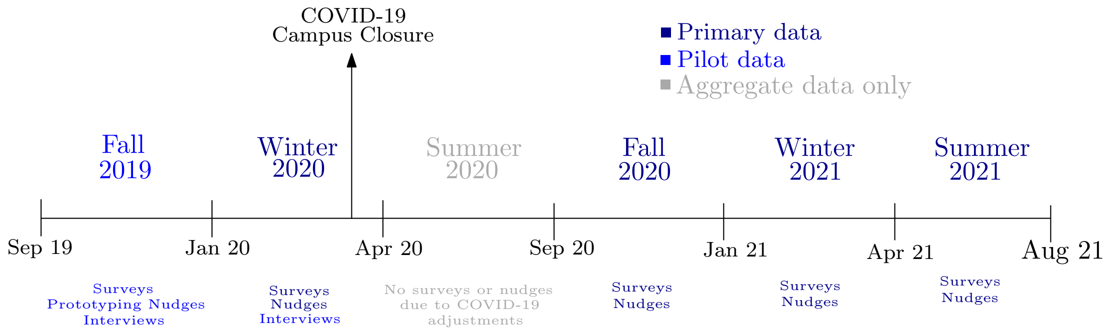

Notes: The figure illustrates the timeline of the data collection. The primary data collection is gathered from the Winter 2020, Fall 2020, Winter 2021, and Summer 2021 cohorts. Pilot data is gathered from the Fall 2019 cohort and involved conducting interviews with students and instructors, surveying students, and prototyping interventions.

Figure 13: Distribution of Homework Reminders Received

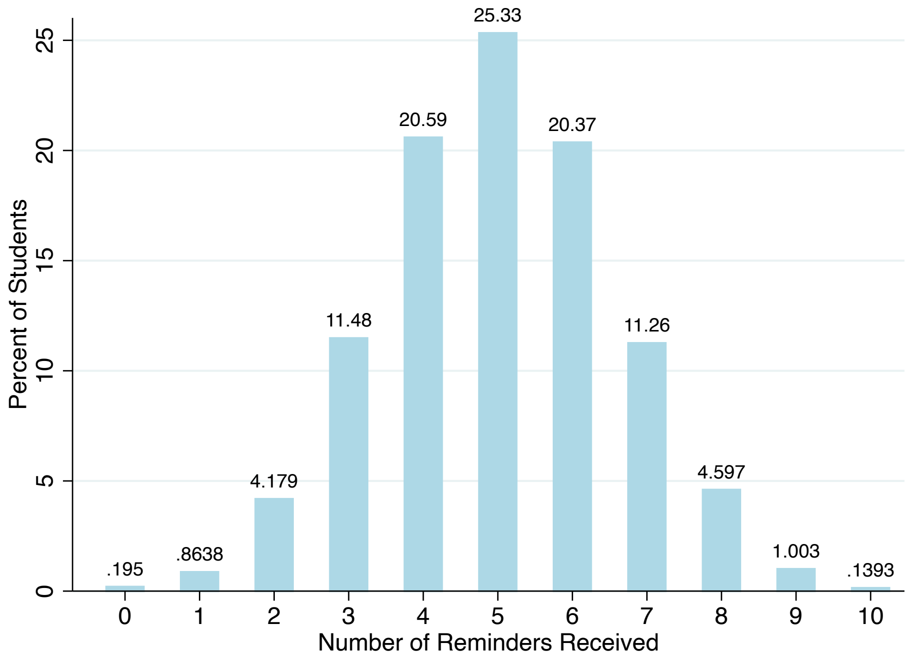

Notes: The figure shows the distribution of reminders received by students. For each of the 10 graded homework assessments, half of the students are randomly selected to receive a reminder. Then each student is eligible to receive between 0 and 10 homework reminders in total.

Figure 14: Student Demographic and Characteristics Balance Check for Homework Reminders

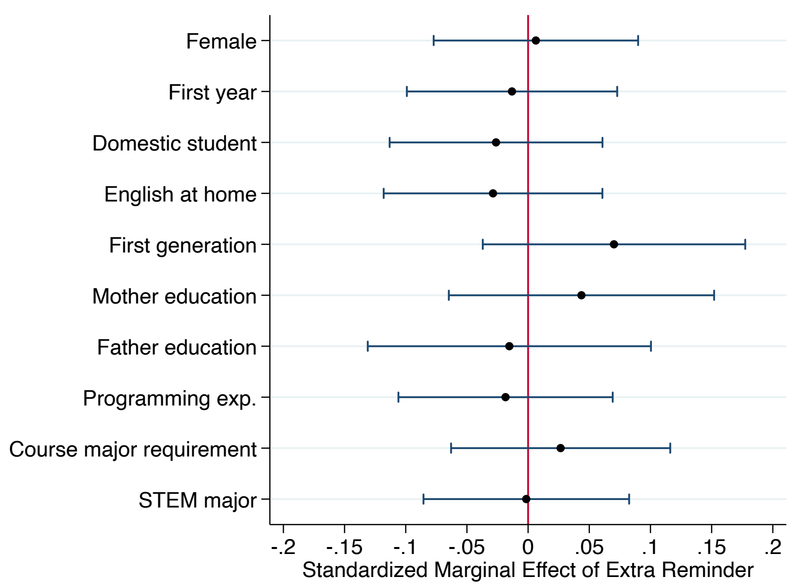

Notes: The estimates displayed are computed by regressing each standardized variable presented on the vertical axis on the number of reminders received. The error bars represent 95% confidence intervals.

Figure 15: Homework Completed and Reminder Messages

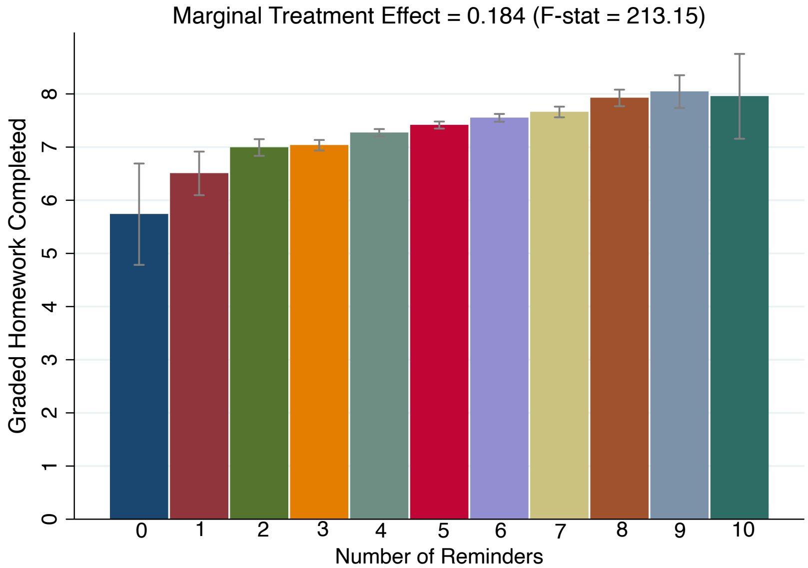

Notes: The figure shows the average number of online homework assessments completed by the number of reminders received. There are 10 graded homework assessments. The error bars represent the 95% confidence intervals.

Figure 16: Reminder Messages Attentiveness Mechanism

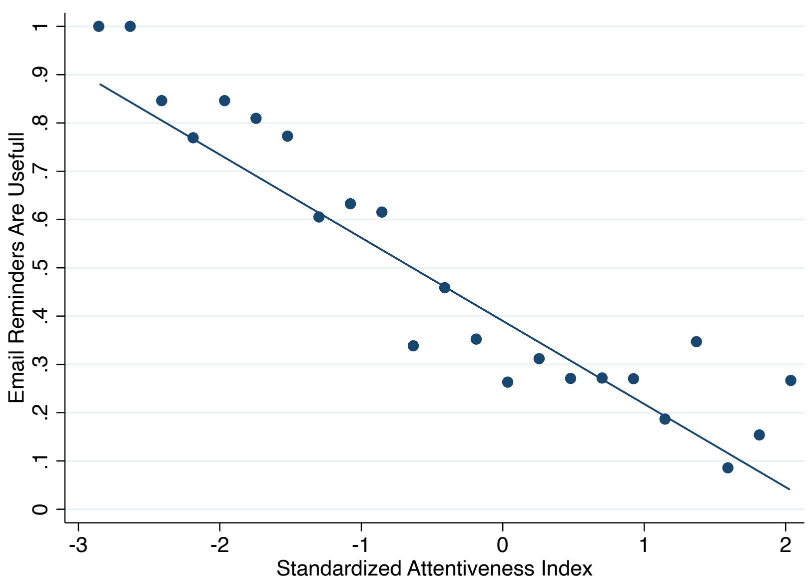

Notes: The figure presents a binned scatter plot showing the relationship between finding the reminders helpful for keeping on track with homework and student attentiveness. Whether a student finds reminders to be useful is inferred from the survey data. The student attentiveness index is constructed using a series of survey questions (see Appendix ??).

Figure 17: Distribution of Model Implied Study Time and Observed Study Time

(a) Basic Stage

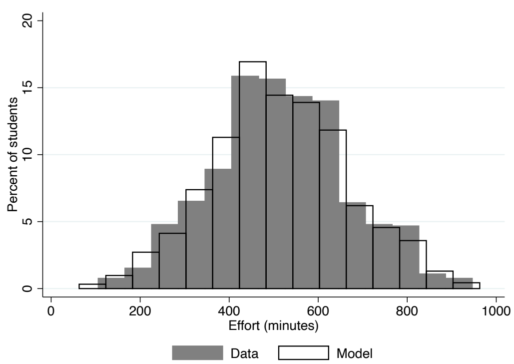

(b) Intermediate Stage

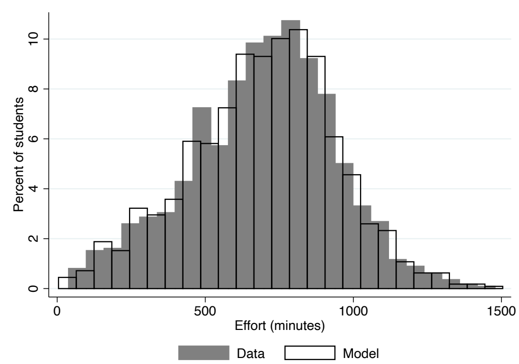

(c) Advanced Stage

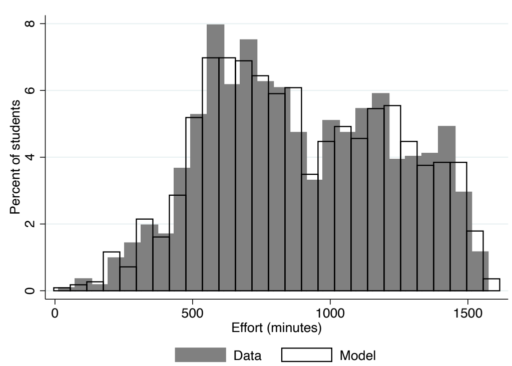

Notes: The figure presents the distribution of observed study time (grey) overlaid together with the distribution of model implied study time (white) at each learning stage.

Figure 18: Model Implied Study Time and Observed Study Time

(a) Basic Stage

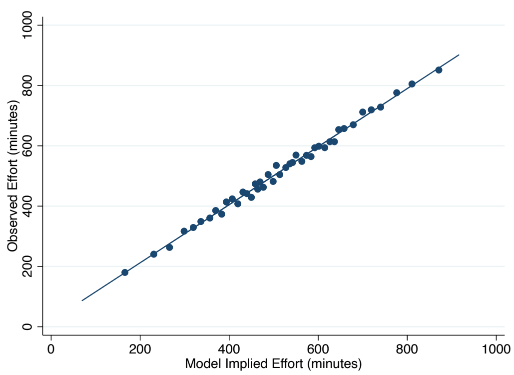

(b) Intermediate Stage

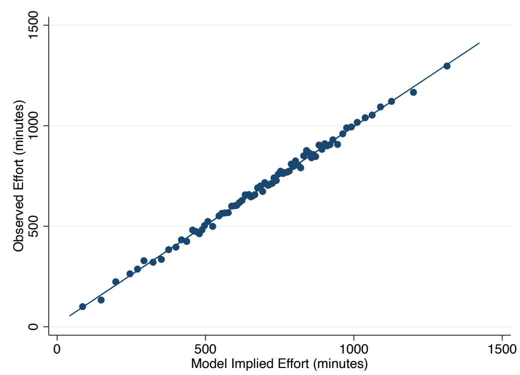

(c) Advanced Stage

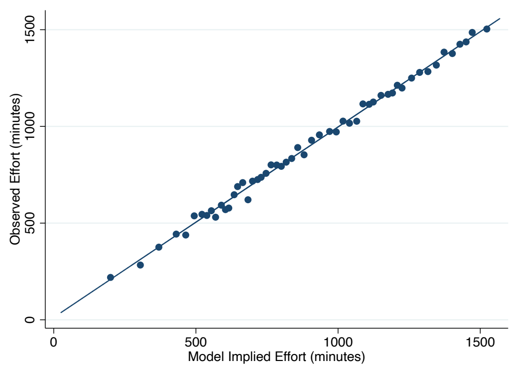

Notes: The figure presents a binned scatter plot showing the relationship between observed study time (vertical axis) and model implied study time (horizontal axis).

A Appendix: Summary of Related Literature

A.1 Related Education Production Function Literature

Table 8: Research Exploring Education Production and Dynamic Complementarities

Title	Authors (Date, Publisher)	Data	Research Design	Main Results
The Production of Cognitive Achievement in Children: Home, School, and Racial Test Score Gaps	Todd and Wolpin (2007, JHC).	N = 7700 individuals who are aged 14-21 in the NLSY79-CS.	Estimate cumulative production function using the mothers ability, child ability, and history of family and school inputs.	Lagged home inputs are significant predictors of present achievement. Overall, estimates suggest the learning process is cumulative.
The Technology of Skill Formation	Cunha and Heckman (2007, AER).	NA (Conceptual Framework)	Develop model of human capital accumulation which features dynamic complementarities in parental inputs.	Model suggests it is important to invest during early childhood stage (e.g. pre-school), more so than later stages (e.g. tuition reduction programs).
The Production of Human Capital: Endowments, Investments, and Fertility	Aizer and Cunha (2012, NBER WP).	N = 30,039 children from 1963 - 1970 whose mothers were involved in the National Collaborative Perinatal Project (NCPP).	Use introduction of Head Start in 1996 as instrument for investment.	Consistent with dynamic complementarities, authors find larger IQ gains from preschool for children with the highest stock of early human capital.
School Accountability and the Dynamics of Human Capital Formation	Gilraine (2018, Working Paper).	N = 3,310 school-year observations from public schools in North Carolina.	Leverages year-to-year variation in school accountability resulting from whether there are at least forty students belonging to a specific demographic group.	Author finds a 0.18σ increase in test scores for students who are in schools that were subject to school accountability in two consecutive periods relative to those in schools subject to accountability only in the previous period.
Does EdTech Substitute for Traditional Learning? Experimental Estimates of Educational Production Function	Bettinger et al. (2020, NBER WP).	N = 6253 grade 3 students in Russia. Teachers had access to computer assisted learning software to help students learn math and language by solving assigned problems.	Students randomized to 1) no computer assisted learning (control), 2) 45-minute computer assisted learning, 3) 90-minute computer assisted learning. Time spent learning using the software was a direct substitute for traditional learning.	Education production function is concave in computer assisted learning. Estimates suggest a hybrid of computer assisted learning and traditional learning is optimal.

A.2 Related Student Effort Literature

Table 9: Research on Exploring Student Effort

Title	Authors (Date, Publisher)	Data	Research Design	Main Results
The Effect of Time Spent Online on Student Achievement in Online Economics and Finance Courses	Calafiore and Damianov (2011, JEE).	N = 438 students enrolled in online Economics and Finance courses during the Spring and Fall 2008 in large public university in south Texas.	Multiple and logistic regression analysis using prior cGPA, age, gender, and major as a control variables. Use sessions logs from Blackboard to track time usage.	Even after conditioning on prior cGPA, time spent on course activities is a significant predictor of performance and earning a better letter grade in the course.
"Making it count": incentives, student effort and performance	Chevalier, Dolton, and Luhrmann (2018, JRS).	N = 424 introductory economic students across two cohorts enrolled at a large college of the University of London. Students are followed across 20 weeks.	Variation in incentives across weeks of either 1) additional study material conditional on quiz participation, 2) 20 GBP book voucher for best quiz performance, or 3) quiz grade counts towards course grade.	Additional study material for participation and book vouchers are ineffective in increasing quiz participation. Grade incentives significantly increases quiz participation and also results in improved exam grades.
Financial Incentives and Educational Investments: The Impact of Performance-Based Scholarships on Student Time Use	Lisa Barrow and Cecilia Elena Rouse (2018, EFP).	N = 5160 high school seniors in California.	Students randomized to performance based (obtain a C average) postsecondary scholarships of $1000 − $4000.	Financial incentives induce more time usage on educational activities and allocate less time on work and leisure.
What sets college thrivers and divers apart? A contrast in study habits, attitudes, and mental health	Beattie et al. (2019, EL)	N = 3849 students enrolled in introductory economics in 2017 at University of Toronto.	Compare student characteristics and habits across thrivers and divers.	Thrivers study around 15 hours per week, seven more hours per week than divers (8 hours per week).
When Study and Nudge Don't Go As Planned: Unsuccessful Attempts to Help College Students	Oreopoulos et al. (2018, NBER WP).	N = 9503 students from University of Toronto (N = 3438) and Western Governors University (N= 6065) in the 2017-18 academic year.	Students randomly assigned to 1) personality test (control) or 2) planning module (build weekly calendar + assigned coach).	Despite marginal increase in study time for those in treatment group, null effects on course grades and retention.
Using Goals to Motivate College Students: Theory and Evidence from Field Experiments	Clark et al. (2020, ReStat)	N = 2004 students for task-based experiment, and N = 1967 for performance based experiment. First year introductory course.	Students randomly assigned to control or goals treatment. Fall 2013 cohort for performance-based goals, and Fall 2014 for task-based goals.	Task-based goals increased task completion and resulted in significant performance gains. Although, performance-based goals are not as effective.

A.3 Related Structural Behavioural Economics Literature

Table 10: Research Estimating Relevant Behavioural Parameters

Title	Authors (Date, Publisher)	Data	Research Design	Main Results
Short-Term Time Discounting of Unpleasant Tasks	Augenblick (2018, WP)	N = 79 subjects who make several decisions overtime resulting in 8875 observations.	Ask repeatedly within a week how many unpleasant real effort tasks they want to do. Task to transcribe blurry greek letters. Randomized wages and number of days from decision task is to be carried out.	Choose 42 tasks start time one week away, and drops to 40 (one day away), 38 (1 hour), and 36 (imminent) as deadline approaches. Estimate discount factor $\delta = 0.85$ when task 1-week away, and some present bias as $\beta \in [0.9, 0.95]$.
What Motivates Effort? Evidence and Expert Forecasts	Dellavigna and Pope (2018, ReStud)	Around N = 9861 Mturkers. Additionally had N = 213 experts make forecast for each treatment condition.	Mturkers randomly allocated to one of 18 treatment conditions ($N \approx 550$ each). Task was to press A-B consecutively on keyboard fast as possible within 10 minutes.	Psychological treatments more effective than baseline, but less impactful than monetary incentives. No evidence of present bias ($\beta = 1$) and altruism ($a = 0$).
Salience and Taxation: Theory and Evidence.	Chetty, Looney and, Kroft (2009, AER)	Focused on cosmetics, hair care accessories, and deodorant. Around 750 distinct products. Intervention was for three weeks.	Treatment group was tax-inclusive prices for the chosen toiletries within a store. Control group 1 are similar products like toothpaste within same aisle and store. Control group 2 are toiletries sold in other stores.	Estimate a degree of intention parameter of $\theta = 0.75$, where $\theta = 1$ is full awareness of the taxation. Although, the result is imprecise due to the quasi-experimental design that relies on week-to-week varition.
Evaluating Behaviourally Motivated Policy: Experimental Evidence from the Lightbulb Market.	Allcott and Taubinsky (2015, AER)	N = 1087 survey participants who were customers at the store.	Participants of iPad survey randomized into information treatment. Treatment group included energy costs for CFLs vs. incandescents. Control group did not get any information on energy costs.	Evidence for small inattention to energy savings with $\theta = 0.05$. Justifies a subsidy of 3 dollar per LED bulb due to inattention.
Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment.	Taubinsky and Rees-Jones (2018, ReStud)	N = 2998 online consumers purchase common household products.	Research participants given 20 dollars to buy one of the randomly chosen products. Randomize into 1) no tax, 2) standard sales tax, 3) triple the sales tax.	Find substantial inattention to taxes and vast heterogeneity in attentiveness. Individuals react to non-salient taxes as if they were 25% their size (i.e., $\theta = 0.25$). Less inattention of $\theta = 0.5$ when taxes are tripled.

A.4 Related Course Design Literature

Table 11: Research Exploring Course Design and Implications of Homework Participation

Title	Authors (Date, Publisher)	Data	Research Design	Main Results
Procrastination, Deadlines, and Performance: Self-Control by Precommitment	Ariely and Wertenbroch (2002, PS).	N = 169 students in a semester (14 weeks) course at MIT.	Students randomly assigned into 1) equally spaced deadlines, 2) end deadlines, and 3) free-choice (set own deadlines).	Performance in equally spaced deadlines dominates self-imposed deadline. However, self-imposed deadlines enhanced performances more than maximally delayed deadlines.
Requiring a Math Skills Unit: Results from a Randomize Experiment	Pozo and Stull (2006, AER:P&P).	N = 273 students in principles of macroeconomics from the spring of 2004 from Western Michigan University.	Randomly assign students to one of two sections with the treatment group requiring the completion of a math unit as a part of the final grade.	Requiring a graded math unit increases participation in homework and raise overall course performance by 2 percentage points, implying an increased letter grade for 26 percent of the class.
Experimental Evidence on the Effect of Grading Incentives on Student Learning in Spain	Artes and Rahona (2013, JEE)	N = 289 students in Public Finance offered in fall of 2019 at University of Madrid.	Students enrolled in morning and afternoon version of class alphabetically by their last name. Experimenters randomized 2 out of 4 problem sets to be graded in each section. Exam questions could be linked to a problem set.	Graded problem sets increased final exam score by 8 percentage points. Students with lower baseline knowledge benefit most from graded homework.
The Role of Homework in Student Learning Outcomes: Evidence from a Field Experiment	Grodner and Rupp (2013, JEE); Citations = 81.	N = 423 students enrolled in microeconomics in the spring of 2008 in a mid-sized university in North Carolina.	Students randomized to two grading schemes 1) 10% homework and four 22.5% tests, and 2) four 25% term tests.	The homework-required group had 5-6% higher test averages; 10-14% higher for those who failed the first test. Students in the homework-required group are also 6 p.p. more likely to complete the course.
The Impact of Assignments and Quizzes on Exam Grades: A Difference-in-Difference Approach	Latif and Miles (2020, JSE)	N = 124 students enrolled in an introductory statistics course in a small business school in Canada.	Authors use data over time across 3 course sections and employ differences-in-differences approach to leverage introduction of quizzes and assignments in only some of the sections. Control section does not have either quiz nor assignment.	Introduction of homework assessment significantly improves midterm grade. Although, no significant results found with the introduction of quizzes.

B Appendix: Institutional Details

B.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.

B.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 1: Outline for Learning Numerical Operations

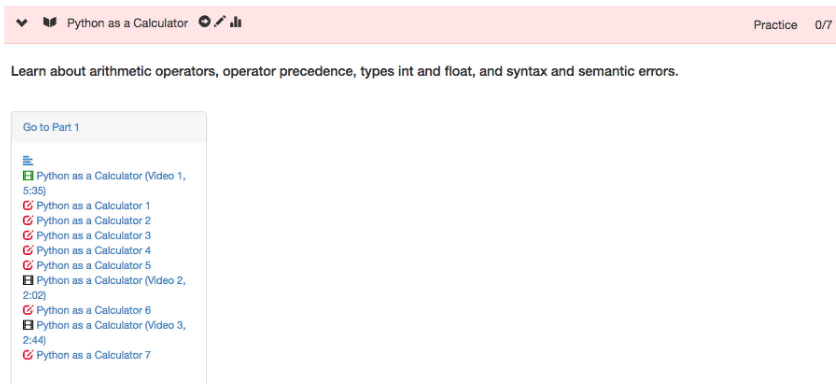

Students begin the course by watching a video about numerical operations in Python:

Figure 2: Video on Numerical Operations

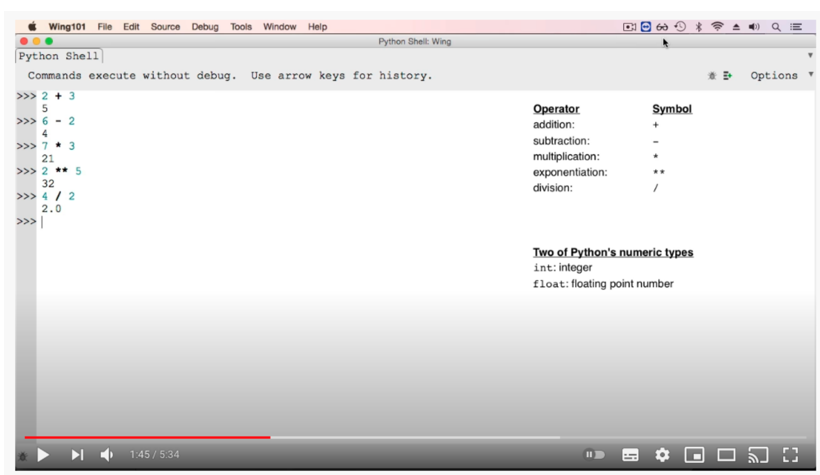

The following figure shows an example of a follow-up coding problem:

Figure 3: Sample Coding Exercise

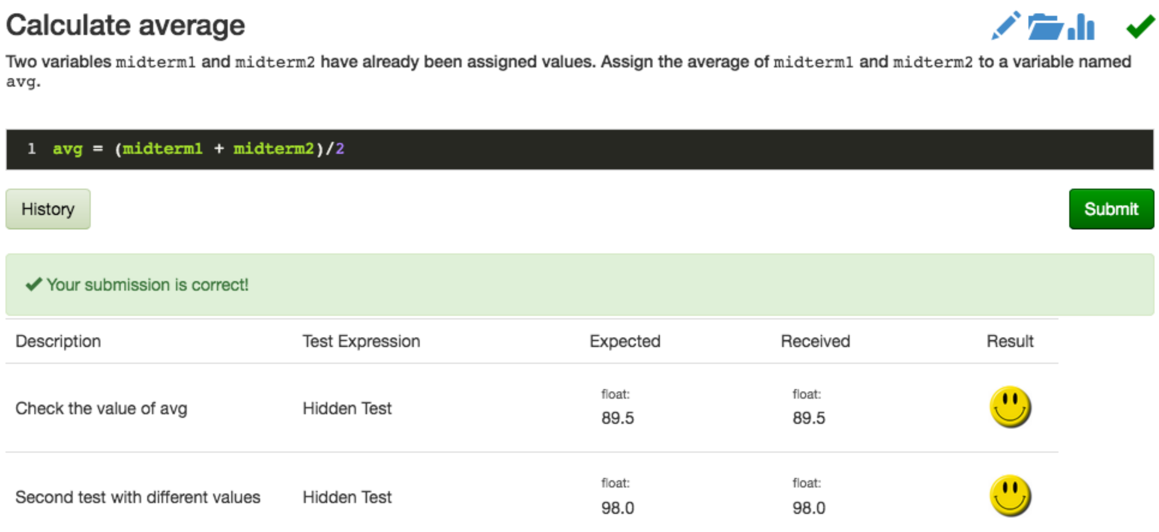

B.3 Online Peer Discussion Board

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 12: 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.

C Appendix: Data

C.1 Measuring Student Demographics and Characteristics

Student demographics and other characteristics shown in Panels A and B of Table 1 are used as pre-treatment controls for most regression specifications. Several of the student controls are constructed using the following questions on the baseline survey. Aside from the options listed below, students could also opt-out from answering the question by selection "Prefer not to answer".

- What is your gender identification?
  - Male; Female; Other
- Are you the first one in your immediate family to attend university?
  - Yes; No
- What is your mother's highest level of education?
  - Did not finish high school; high school graduate; some college; college graduate; graduate degree (e.g., masters or doctorate)
- What is your father's highest level of education?
  - Did not finish high school; high school graduate; some college; college graduate; graduate degree (e.g., masters or doctorate)
- Is English your native/first language?
  - Yes; No
- What language do you speak at home? *(open response)*
- How would you describe your prior experience with programming?
  - I have never programmed before; I have written a few lines of code; I have written basic programs before; I have extensive experience programming
- Which of these is closest to your (intended) program of study?
  - Computer Science; Commerce; Humanities; Life Sciences; Physical and Mathematical Sciences; Social Sciences; Other

C.2 Measuring attentiveness and forward-looking perspective

The baseline survey elicits a student's attentiveness and forward-looking perspective 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 attentiveness.

Table 13: Survey Components for Forward-looking Perspective

Survey question (7-point scale)	Relationship with Forward-looking
I consider myself to be a forward-looking person who has clear plans about the future	Increasing
I tend to think about how working hard on the present homework will make doing future homework easier	Increasing
I tend to think about how working hard on homework each week will help me do better on the exam	Increasing
When I have multiple deadlines, I think ahead and plan how to split my time before I begin working	Increasing
I have a good sense of my expected career trajectory after completing my current degree	Increasing

Table 14: Survey Components for Attentiveness

Survey question (7-point scale)	Relationship with attentiveness
I tend to read all the instructor announcements for this course each week	Increasing
I have read the course syllabus in detail	Increasing
I know how to access office hours	Increasing
I know when office hours are held	Increasing
I tend to forget about my assignment deadlines	Decreasing

I characterize students as being forward-looking if they responded with at least 5 to each question in Table 14. Similarly, Students are classified as attentive if they responded with at least 5 to the first three questions in Table 14, and at most 3 to the last question.

C.3 Measuring English proficiency

The tables below include the 7-point Likert scale survey questions used to measure a student's English proficiency. Responses are aggregated so the resulting index variable is increasing in English proficiency.

Table 15: Survey Components for English Proficiency

Survey question	Relationship with English proficiency
I am comfortable writing in English	Increasing
I am comfortable speaking in English	Increasing
I have difficulty reading and listening in English	Decreasing
I have difficulty learning by watching instructional videos in English	Decreasing

C.4 Measuring Study Time

I use the time-stamped online interactions to construct a measure of total study time for each learning stage. Students primarily spend their time on the online homework platform. Additionally, students participate in the online peer discussion board by writing and reading posts.

C.4.1 Study time on online homework

The administrative data includes time-stamps for when students log-in, log-out, click to play/pause videos, submit a solution to a problem, and various other interactions with the platform. I develop a simple algorithm that uses the time-stamped data to measure the number of minutes of videos watched (v) and minutes spent doing homework problems (h). The algorithm leverages the fact that students tend to study in around 30-minute blocks throughout the week. The blocks of study time are identified to the nearest 5-minute of inactivity and aggregated together. Then, for each learning stage t, the time spent on the online homework is:

$$e_{i,t}^{H} = v_{i,t} + h_{i,t}.$$

C.4.2 Study time on peer discussion board

Although the administrative data includes the number of posts written (w) and unique posts read (r), the time spent on these activities is not included. To fill this gap, the final survey asks students how many minutes on average they spend writing (m^w) and reading a post (m^r). The survey questions eliciting time costs to discussion board participation are as follows:

- Roughly how much time (in minutes) do you believe it takes you to write an average quality discussion board post (i.e. make new question or answer peer question)?
  - Minutes it takes to write a question [numeric response]
  - Minutes it takes to write a answer [numeric response]
- In a hypothetical scenario, suppose you were given 10 minutes to browse the discussion board and read posts (question or answers). How many posts do you think you could read in detail in that time period?
  - Number of questions carefully read in 10 minutes [numeric response]
  - Number of answers carefully read in 10 minutes [numeric response]

Then, for each learning stage t, the time spent on the discussion board is:

$$e_{i,t}^{D} = m_i^w w_{i,t} + m_i^w r_{i,t}.$$

C.4.3 Total study time

Time spent across the online homework and discussion boards aggregated at each learning stage to construct study time:

$$e_{i,t} = e_{i,t}^{H} + e_{i,t}^{D}$$

C.5 Survey questions eliciting student peer interactions

The surveys include the following questions to measure the extent to which students interact with other peers in the course.

- Are you in a study group for [CourseCode]?
  - I am in a study group officially recognized by [institution name]
  - I am in another study group with students from this course
  - No
- Around how many students in the course do you study with per week? [Numerical Entry]
- Around how many hours per week do you study with other students in this course? [Numerical Entry]
- I discussed the contents of the homework reminder messages with other students in the course [Likert Scale]

D Appendix: Nudges

The reminder messages are designed using various behavioural insights such as implementation intentions, utility value, and self-reflection. Kizilcec et al. (2020), Harackiewicz and Priniski (2018) and Damgaard and Nielsen (2018) provides excellent reviews on the behavioural nudging literature in education.

D.1 Homework reminder messages

The homework reminders are sent through the learning management system. Students receive the reminder in their personal university email inbox and a notification of the message on the learning management system. The template for the homework reminder is as follows.

Hi [Student Name],
The homework is due by [Deadline]. Please take a moment to think about the following prompts:

When will you next work on this week's homework? Can you set aside time on your schedule to progress on the homework?

Some students find it valuable to just open up the online homework system and spend a minute on a problem. Here is the link to the homework: [Link to Homework]

[Course Code] Learning Support Team

E Theory Appendix

E.1 Front-loading student effort

As course material typically becomes more challenging as the course progresses, the period specific marginal benefit of effort decreases across learning stages: $\frac{dL_i^{basic}}{de_i^{basic}} > \frac{dL_i^{int}}{de_i^{int}} > \frac{dL_i^{adv}}{de_i^{adv}}$. The persistence of learning from one period to the next will likely be less than 100%: $\frac{dL_i^{int}}{dL_i^{basic}} < 1$ and $\frac{dL_i^{adv}}{dL_i^{int}} < 1$. Suppose that the persistence of knowledge from basic to intermediate exceeds the persistence from intermediate to advanced (i.e., $\frac{dL_i^{int}}{dL_i^{basic}} > \frac{dL_i^{adv}}{dL_i^{int}}$) then marginal total learning benefit of effort in the basic stage exceeds the intermediate stage (i.e., $\frac{dL_i}{de_i^{basic}} > \frac{dL_i}{de_i^{int}}$). Although such an assumption may be strong, $\frac{dL_i}{de_i^{basic}} > \frac{dL_i}{de_i^{int}}$ holds as long as $\frac{dL_i^{int}}{dL_i^{basic}}$ is sufficiently large. Under these assumptions, the marginal total learning benefit of effort is decreasing as the course progresses

$$\frac{dL_i}{de_i^{basic}} > \frac{dL_i}{de_i^{int}} > \frac{dL_i}{de_i^{adv}}.$$

E.2 Extending model to include grading weights

Consider N students in a course, who allocate total study time or 'effort' (e) across three learning stages $t \in \{basic, int, adv\}$. Then, let L_i^t denote the amount of learning for student i during stage t. Students can vary in their baseline human capital (h), English proficiency (E), and whether they are forward-looking (f). Human capital h_i and English proficiency E_i are standardized to have mean 0 and standard deviation 1. The indicator f_i takes on a value of 0 to denote myopic students, and a value of 1 to represent forward-looking students.[32] Let p_m be the proportion of myopic students in the course. The instructor sets the grading weight w_t for each learning period.

The timeline of the model is as follows. First, the instructor specifies the grading weights $(w_t)_t$. Then, given the grading scheme, students allocate their effort (i.e., study time) across the course $(e_i^t)_t$. For forward-looking students, I solve this model by backwards-induction. Thus, I first begin therefore begin by discussing the student effort choice problem, and then outline the instructor's problem.

E.3 The Student Effort Choice Problem

Forward-looking students internalize the cumulative learning process and allocate their effort to maximize their course grade net of effort costs:

[32] Heterogeneity in attentiveness is not included in the model as descriptive analysis not presented here shows inattentive students effort choices are much less influenced by the grading scheme relative to myopic students.

$$\max_{(e_i^t)_t} \sum_t w_t L_i^t(e_i^t; L_i^{t-1}) - C(e_i^{basic}, e_i^{int}, e_i^{adv}; E_i), \tag{5}$$

where $C(\cdot)$ is a convex function representing the cost of effort exertion. The learning technology at each stage of the learning process is cumulative. As a result, the amount of learning L_i^t in a given period depends on present effort e_i^t, and previous knowledge L_i^{t-1}.[33] I assume the learning production functional is concave, increasing in effort and baseline human capital, whereas the cost function is convex and increasing in effort. Additionally, cost of effort decreases with English proficiency.[34]

In contrast, the myopic students do not internalize the cumulative course structure, and thus focus on each learning stage separately. As a result, they allocate effort to maximize their grade in each learning stage net of effort costs:

$$\max_{e_i^t} w_t L_i^t(e_i^t; L_i^{t-1}) - C(e_i^t; E_i) \text{ for each } t \in \{basic, int, adv\}. \tag{6}$$

Let $e_i^{t,*}$ denote the optimal effort of student i in stage t resulting from solving their effort-choice problem.

Incorporating Additional Assessments. The effort allocation model is flexible and can be adapted to incorporate beyond three assessments. In the actual course, students are assigned online assignments at each learning stage, a midterm based on material covered in the basic stage, and a cumulative final exam. Then the forward-looking students allocate effort to maximize:

$$\max_{(e_i^t)_t} \sum_t w_t L_i^t(e_i^t; L_i^{t-1}) + w_{mid} L_i^{mid}(e_i^{basic}; h_i) + w_{final} L_i^{final}(e_i^{adv}; L_i^{int}) - C(e_i^{basic}, e_i^{int}, e_i^{adv}; E_i). \tag{7}$$

Myopic students' effort choice problem in the basic and intermediate stage is as represented in equation 6. For the advanced stage, however, the myopic students internalize that studying during the advanced stage benefits both their assignments and final exam performance.[35] The final exam has the highest weight of all assessments and the university announces the exam schedule at the advanced stage. Therefore, even the myopic students internalize the importance of exerting effort during the advanced stage and do so by maximizing:

[33] Given that course structure is assumed to be cumulative, L_i^{t-1} is used as a sufficient statistic for all prior knowledge accumulation. Prior knowledge at the basic stage is denoted by L_i^{-1} and is the baseline knowledge h_i.

[34] Students who are proficient in English are going to have an easier time understanding the contents of the videos and interpreting the homework problems.

[35] Students in this setting do not study for the midterm and final exam distinctly from the assignments. The assignments also serve as the main source of preparation for the high-stakes assessments.

$$\max_{e_i^{adv}} w_{adv} L_i^{adv}(e_i^{adv}; L_i^{int}) + w_{final} L_i^{final}(e_i^{adv}; L_i^{int}) - C(e_i^{adv}; E_i). \tag{8}$$

E.4 Instructor's Grading Scheme Design Problem

Suppose the primary evaluation in the advanced learning stage is a comprehensive and cumulative final exam. Then, the instructor's goal is to have the representative student allocate their effort throughout the course to maximize their grade on the final exam net of effort costs. Since p_m proportion of the students in the course are myopic, the instructor weights learning and effort costs across representative myopic and forward-looking students as follows:

$$p_m[L_i^{final}(e_i^{adv}; L_i^{int}, \bar{h}, f_i = 0) - C(e_i^{basic}, e_i^{int}, e_i^{adv}; \bar{E}, f_i = 0)] - \\ (1 - p_m)[L_i^{final}(e_i^{adv}; L_i^{int}, \bar{h}, f_i = 1) - C(e_i^{basic}, e_i^{int}, e_i^{adv}; \bar{E}, f_i = 1)], \tag{9}$$

where $\bar{h}$ is the average baseline human capital, and $\bar{E}$ is the average English proficiency.[36] The instructor uses the final exam grade as the primary outcome measure, as it is cumulative and best represents the totality of the course material relative to other assessments in earlier stages. To induce students to efficiently allocate their effort, the instructor sets the grading scheme $(w_t)_t$ to maximize the above objective.

E.5 Stylized Example

To intuitively illustrate the implications of the model, consider a course with a cumulative structure and two learning stages. Suppose that students learn basic concepts in the first half of the course and advanced concepts in the remaining half. That is, $t \in \{basic, adv\}$. Additionally, students only vary according to whether they are forward-looking ($f_i = 1$) or myopic ($f_i = 0$).

Parameterization of the Learning Technology and Cost Function. Let the following simple learning technologies represent the cumulative learning process:

$$L_i^{basic} = \alpha_1 e_i^{basic},$$

$$L_i^{adv} = \beta_1 e_i^{adv} + \beta_2 L_i^{basic} + \beta_3 e_i^{adv} \times L_i^{basic}.$$

[36]The average baseline knowledge and English proficiency are similar across myopic and forward-looking students in the data. Then, in the model, I assume student characteristics are independent of students' forward-looking mentality.

A positive marginal benefit of effort at both learning stages implies that $\alpha_1 > 0$ and $\beta_1 > 0$. Since the advanced learning stage is cumulative, then clearly $\beta_2 > 0$. Finally, assuming effort exertion in the basic stage increases the productivity of advanced stage effort (i.e., there are dynamic complementarities in effort), then $\beta_3 > 0$.

The cost of effort is assumed to be linearly separable and represented by a quadratic cost function:

$$c(e_i^t) = \kappa \frac{(e_i^t)^2}{2} \text{ for } t \in \{basic, adv\},$$

where κ represents the steepness of the cost function.

The Students' Optimal Effort Choice. Let us assume there is no bonus credit, and so $w_{basic} + w_{adv} = 1$. The forward-looking student sets effort by maximizing her expected course grade net of costs using backwards induction, beginning at the advanced stage. The resulting effort allocation across the basic and advanced learning stages are:

$$e_i^{basic,*}(f_i = 1) = \frac{\alpha_1[\kappa(1 - w_{adv}(1 - \beta_2)) + \beta_1\beta_3 w_{adv}^2]}{(\kappa - w_{adv}\alpha_1\beta_3)(\kappa + w_{adv}\alpha_1\beta_3)},$$

$$e_i^{adv,*}(f_i = 1) = \frac{w_{adv}[\beta_3\alpha_1^2[\kappa(1 - w_{adv}(1 - \beta_2)) + \beta_1\beta_3 w_{adv}^2] + \beta_1(\kappa - w_{adv}\alpha_1\beta_3)(\kappa + w_{adv}\alpha_1\beta_3)]}{\kappa(\kappa - w_{adv}\alpha_1\beta_3)(\kappa + w_{adv}\alpha_1\beta_3)}.$$

In contrast, the myopic students do not internalize the cumulative learning process, and supply effort as follows:

$$e_i^{basic,*}(f_i = 0) = \frac{\alpha_1(1 - w_{adv})}{\kappa},$$

$$e_i^{adv,*}(f_i = 0) = \frac{w_{adv}[\beta_3\alpha_1^2(1 - w_{basic}) + \kappa\beta_1]}{\kappa^2}.$$

Clearly, myopic students exert less effort than forward-looking students in the basic learning stage as long as the assessment in the advanced stage is cumulative (i.e., $\beta_2 > 0$), and there are dynamic complementarities in effort inputs (i.e., $\beta_3 > 0$). Otherwise, if the course covers distinct and unrelated topics (i.e., $\beta_2 = \beta_3 = 0$) then both myopic and forward-looking students allocate effort identically:

$$e_i^{basic,*}(\beta_2 = \beta_3 = 0) = \frac{\alpha_1(1 - w_{adv})}{\kappa}, \text{ and } e_i^{adv,*}(\beta_2 = \beta_3 = 0) = \frac{\beta_1 w_{adv}}{\kappa}.$$

F Appendix: Parametrization of the Learning Technology

Discussing Other Parameterization of the Learning Technology. The main specification of the cumulative learning process is represented by equations 2, 3, and 4. These are designed to be easily interpretable and minimalistic representation of the cumulative STEM learning process. I now discuss more complex representations of the cumulative learning process that may better fit the data.[37] For example, to allow for diminishing marginal returns to effort exertion, I consider the following specification in the basic stage:

$$L_i^{basic} = \alpha_0 + \alpha_1 e_i^{basic} + \alpha_2 (e_i^{basic})^2 + \alpha_3 h_i + \alpha_4 e_i^{basic} \times h_i + \epsilon_i^{basic}, \tag{10}$$

with $\alpha_2 < 0$ representing diminishing returns to effort in the basic stage. Adapting similar specifications to the intermediate and advanced stage will also allow for diminishing marginal returns to effort, but with a more complex characterization of concavity due to the dynamic complementarities.

Instead of using the prior stage grade as a sufficient statistics for previous knowledge, cumulative learning can be modelled entirely using historical effort inputs. For example, for learning in the intermediate stage could be represented by:

$$L_i^{int} = \beta_0 + \beta_1 e_i^{int} + \beta_2 e_i^{basic} + \beta_3 h_i + \beta_4 e_i^{basic} \times e_i^{int} + \beta_5 e_i^{int} \times h_i + \epsilon_i^{int}. \tag{11}$$

Now we have two separate dynamic interaction parameters, β_4 and β_5.

Instead of imposing linearity, the accumulation of knowledge can be represented by a CES production function as follows:

$$L_i^{int} = \lambda(\theta_1 (e_i^{int})^{\sigma} + \theta_2 (L_i^{basic})^{\sigma})^{\frac{1}{\sigma}} + \epsilon_i^{int}, \tag{12}$$

where $\sigma > 0$ governs the elasticity of substitution between present effort and previous knowledge. Assuming $\theta_1 > 0$ and $\theta_2 > 0$, dynamic complimentarities will be reflected by $\sigma < 1$. The CES production function can be linearized as a translog production function for estimation:

$$\begin{aligned} ln(L_i^{int}) = {} & ln(\lambda) + \theta_1 ln(e_i^{int}) + \theta_2 ln(L_i^{basic}) + \theta_{11} ln^2(e_i^{int}) + \\ & \theta_{22} ln^2(e_i^{int}) + \theta_{22} ln^2(L_i^{basic}) + \theta_{12} ln(e_i^{int}) ln(L_i^{basic}) + \epsilon_i. \end{aligned} \tag{13}$$

[37] I omit the cohort index for simplicity of the exposition.

When the coefficients associated with the quadratic term are 0, note that equation 3 and equation 13 have an analogous structure. That is, the education production function being estimated by equations 2, 3, and 4 are sufficiently flexible to capture features of a CES production function.