Paper ID #50464

Optimizing Question Bank Size to Influence ”Gaming the System” Behavior

Nicole Qian, Bucknell University

Nicole Qian is an undergraduate student at Bucknell University double-majoring in computer science and mathematical economics and with additional interests in philosophy, religion, psychology, linguistics, and women and gender studies.

Sing Chun Lee, Bucknell University

Sing Chun is a Computer Science faculty member at Bucknell University. While pursuing his PhD in geometry processing and augmented reality, he discovered a passion for computer science education. His current work explores how mixed reality (XR) can support CS learning, guided by multimedia learning theory and cognitive load theory. He is also committed to promoting undergraduate research and to making advanced computational ideas more accessible to undergraduates. Outside of academia, he enjoys immersing himself in MR experiences and discussing the future of XR, especially its potential to transform education, with fellow enthusiasts.

Lea Wittie, Bucknell University

Lea Wittie is an Associate Professor and co-chair in the department of Computer Science in the Engineering College at Bucknell University.

Rachel Landsman, Bucknell University

Rachel Landsman is an Assistant Professor of Economics at Bucknell University. Her research focuses on experimental methods and the role of beliefs in social decision-making. Her work examines how institutional design and incentives shape judgment and decision-making, and she applies this broader research agenda to topics such as gender, including stereotype formation, gender differences in competition, and the attribution of blame to women leaders. She teaches behavioral economics, experimental economics, and microeconomic theory.

Optimizing Question Bank Size to Influence “Gaming the System” Behavior

Nicole Qian
nq001@bucknell.edu
Computer Science
Bucknell University

Lea Wittie
lwittie@bucknell.edu
Computer Science
Bucknell University

Rachel Landsman
rbl013@bucknell.edu
Economics
Bucknell University

Sing Chun Lee
singchun.lee@bucknell.edu
Computer Science
Bucknell University

Abstract

In a Teach-for-Mastery (TfM) course, instructors often use question banks to generate multiple random quizzes to let students practice the same concept until mastery. While honest learners effectively master the concepts this way, other students can exploit the no-penalty structure to “game the system” without learning. In this paper, we collected six semesters of TfM classroom data (6,625 quiz-attempt records from 106 students), built a behavior model, and studied how the size of the question bank influences this “gaming the system” behavior.

We borrow insights from game theory, including Nash equilibrium and sequential instructor-student interactions, to model how course policies shape student strategies. We classify students into “gamers” and “non-gamers” using clustering on retake intensity and attempt duration, and we estimate a production function that links student time-on-task and instructor preparation effort to achievement under a policy variable $\pi = 1/N_q$. We then vary N_q and solve for the resulting Nash equilibrium effort separately by topic, using the estimated model and type-specific utility functions.

Our equilibrium results suggest that increasing N_q has the largest impact when question banks are small, with diminishing marginal changes once banks reach moderate sizes. This provides practical guidance for instructors on where question-writing effort is most valuable: topics with very small or narrowly varying banks are the strongest candidates for expansion. In our case study, Topic 5 (fixed at $N_q = 4$) is the clearest redesign candidate; the model suggests expanding to approximately $N_q \approx 10$ to 20 captures most of the incentive shift toward higher effort, whereas topics already operating with larger banks show limited gains from further expansion. In summary, we address a common challenge in TfM: how to mitigate “gaming the system” behavior while still encouraging learning, and provide a general framework for evaluating and selecting TfM policy settings using routinely collected learning management system data.

1 Introduction

Keeping students focused on genuine learning can be challenging. Educators often struggle to maintain student focus, and as more distractions enter students' learning environments, there is a growing pull away from genuine knowledge acquisition. Students are constantly finding new ways to get better grades while doing less work, and as technologies like artificial intelligence improve, the return on actually learning material seems to decrease. Instructors may find their jobs increasingly difficult as they deal with these uncooperative students.

There has been work analyzing various models of student-instructor interaction, inspired by economic theory. Many such models provide a definition of what student and instructor "effort" are. Following economic norms, there is also a recurring theme of utility functions, with varying levels of complexity. These studies often apply economic concepts to quantify the interactions of students and instructors in the context of economics courses. However, these factors can change in a different setting, such as the classroom for a computer science (CS) course.

Students within the same major are unlikely to behave similarly, so it is an even further stretch to assume that the behavior of economics students will be generalizable to CS students. Instructors must adapt to their students' demographics, and we can take a mixed-methods approach combining economics and CS concepts to propose a new course structure that maximizes learning in the CS classroom. In particular, a Teach-for-Mastery (TfM) course, which offers a unique setting beyond standard graded homework and tests, will be the primary focus of our work.

TfM is a classroom structure that prioritizes understanding the material over grades. One way to implement this is to give students unlimited attempts on assignments; they may take a quiz, for example, as many times as needed until they receive the desired grade. This takes the pressure off of getting a certain grade in one try and instead aims to give students the space to really understand the material. While this sounds like a positive approach, one issue that arises is "gaming the system" behavior. This is when a student repeatedly takes the test and answers questions based solely on memorizing previous questions. This strategy relies on cycling through the quiz many times until enough of the questions are recognized, and the student is able to provide the correct answers through memorization. "Gamers", or students who exhibit "gaming the system" behavior, are unwilling to thoroughly learn the material, and TfM gives them the leeway to take advantage of a system that is meant for people who genuinely want to understand what they are learning. So, **how can we design retake-friendly quizzes, particularly the size and structure of the question bank that they draw from, to discourage gaming while keeping instructor workload manageable?**

One solution is to create a question bank large enough that gamers do not see repeated questions, thereby making their strategy no longer worthwhile. While this discourages gaming, it can also require the instructor to put in drastically more effort if they have to come up with too many questions. In order to find a balance between these issues, we analyze the interactions between students, both gamers and those who genuinely want to learn (non-gamers), and instructors in a CS classroom. We can do this using utility models from economic theory and data collected over multiple semesters of a real CS course that follows a TfM structure.

2 Related Work

Teach-for-Mastery (TfM) is a teaching style derived from the "mastery learning" strategy developed by Benjamin Bloom in 1968. Mastery learning aims to mimic the level of attention and direction a student would get from one-on-one tutoring, but in a way that is applicable to everyone in the classroom[1]. There are multiple defining steps in mastery learning: first, an education phase that allows students to learn the material. Then, there are formative assessments, which provide a system for instructors to give feedback and identify students who may need more support, and in which topic areas. Both students and instructors can determine whether a student has met a defined standard for mastery. Typically, this standard is represented as a percentage, commonly set at around 80-90$\%$. Students who have achieved mastery during the first assessment can then continue on with new material or enhance their understanding of the current material. Those who have not yet achieved mastery are assigned remedial exercises to supplement their knowledge. Once they have completed these tasks, students can retake the assessments to test for mastery, though the assessments are generally structured somewhat differently than the first iteration[2]. It is important to note that TfM includes two kinds of assessments: summative and formative. Summative assessments are comprehensive tests taken at the end of the course, while formative assessments are given regularly throughout the course to gauge mastery of material as learning occurs. Formative assessments are more essential to the idea behind TfM,[1], and we will focus more on them in this paper.

There has been significant support for TfM strategies[2], but they also come with risks. As technology becomes more integrated into the learning environment, and generations grow up with the ability to adapt to new technology, the use of online assessment software could pose a challenge. Students have been shown to find new ways of cheating and committing academic dishonesty when able to hide behind a screen[3]. While some behaviors are not as obviously wrong as others, they may still slow or negate learning. One example of this is "wheel spinning," a term that describes behavior in which students learn without understanding the material and make no effort to change their situation. They look like they may be progressing, even though they are not. Another behavior, which will be the primary focus of this paper, is "gaming the system." This describes a behavior in which students recognize the technical features of the assessment technology and abuse them to pass, without actually learning the material. While Li *et al.*[4] distinguish between "harmful gaming" and "non-harmful gaming," they acknowledge that both remain gaming and negatively impact education. In a setting like ours, where students can retake online quizzes as often as they need, there is a clear opening for this gaming behavior to fester. Therefore, we will later propose a framework to address these issues.

To properly digest our framework, it is important to understand some key economic concepts. The first is utility functions. Utility functions provide a numerical value that represents the satisfaction or utility a consumer derives from a good or service. One of the most common and simplest utility functions is a Cobb-Douglas function, which has the general form of $Ax_1^{\alpha_1}...x_n^{\alpha_n}$, where A is a positive constant, and the exponents are non-zero constants. Another common utility function is the Constant Elasticity of Substitution (CES) function, which has the general form $A(\sum_{i=1}^{n} a_i^{\rho} x_i^{\rho})^{\frac{\gamma}{\rho}}$, where α, γ, ρ, and A are nonzero constants, and α and A are positive. The key feature of the CES function is that its parameters have a constant elasticity of substitution within their domains[5].

The next important concept is the Bayesian Nash equilibrium (NE). A Bayesian game is a game with incomplete information. This means that all the players in the game are allowed to hold information that no other player is aware of. For example, if we model the classroom setting as a Bayesian game in which the students and the instructor are the players, the instructor may have information about assessments that the students are unaware of. On the other hand, students may have test-taking strategies that the instructor is unaware of. This Bayesian game will serve as the setting for our search for an NE. An NE is a moment in the game in which neither player has an incentive to change their strategy. Both players are at the point where they have the most amount of utility that they believe they can get, so neither of them is willing to make a different decision. When an NE is reached in a Bayesian game, we get a Bayesian Nash equilibrium[6].

To address the issue of "gaming the system," we will draw on economic concepts to model a TfM classroom setting as a game. Correa and Gruver[7] provide a starting point on how we might set this up. They introduce the achievement production function, which is influenced by both student and instructor efforts. Students and instructors both still have their own utility functions, however, that model their unique constraints and preferences. Correa and Gruver provide ideas on what constraints should be used for utility and production functions and suggest the use of both Cobb-Douglas and CES functions for modeling purposes. They also provide preliminary calculations on finding the equilibrium between students and instructors.

Part of the instructor's utility comes from how much they can motivate students to learn: the more content students master, the happier the instructor is. Therefore, the instructor has an incentive to convince students to put effort into learning. Bonesrønning[8] provides some insight into related economics in education literature here. As mentioned previously, achievement can be measured by both the instructor and the student. Students, though, can also have *perceived achievement*, and the instructor can adjust how students' perceived achievement changes. If an instructor is known to have an easier grading system, students are likely to spend less effort studying because they assume they can achieve more without trying as hard as they would under a harder grading system. While it might seem to follow that the grading should be as harsh as possible, Bonesrønning indicates that there is a point at which students will give up and return to not putting in any effort at all if they perceive the difficulty to be too high and no longer worth the effort.

There needs to be a balance in how difficult instructors make assessments seem, but there is not much previous work on which grading practices are most optimal[8]. In the context of a TfM course that uses retakeable quizzes as formative assessments, the number of potential questions can help the instructor adjust difficulty levels. Furthermore, the number of questions in the question bank should ensure that students routinely encounter new ones, even when retaking the quiz multiple times. At the same time, this should also discourage students from partaking in "gaming the system" behavior. We will explore these ideas more in-depth in the paper.

3 Vignette: Strategic Behavior in Teach-for-Mastery Course Design

In a Teach-for-Mastery (TfM) course, instructors frequently rely on large question banks to generate randomized quizzes that promote repeated practice and conceptual understanding. For example, Wittie[9] describes implementing TfM in an undergraduate programming languages course by replacing traditional exams with short, one-topic "take-when-ready" assessments that draw a random question from a pool of many variants (roughly 30 per topic), with built-in delays between re-attempts to discourage memorization-based retakes. A sufficiently large and diverse bank is generally thought to discourage shallow memorization and reduce "gaming the system" behavior. More importantly, from a strategic perspective, choices about the size, composition, and distribution of difficulty in the question bank constitute a policy setting that directly shapes students' study strategies.

From an economic game perspective, the instructor first specifies the question bank and the quiz randomization mechanism. Students then observe and infer these policy settings as they complete the quizzes and adapt their behavior accordingly: deciding how broadly to study, how much effort to allocate to different topics, and how to prioritize this course relative to others. This act-react sequence naturally aligns with a sequential, leader–follower structure, where the instructor's TfM course policy is the leader's actions, and students' learning strategies are the followers' best responses.

Consider, for example, a topic whose bank contains only a few predictable questions. Students may rationally adopt a shallow memorization strategy, *e.g.*, focusing on recalling specific examples or past mistakes rather than developing conceptual mastery. Conversely, a sufficiently large and diverse bank may encourage students to deeply study and master the concept through learning a variety of examples. In either case, students' collective learning behavior settles into a **Nash equilibrium** determined by the instructor's policy.

At equilibrium, given their inferred policy settings, no student is interested in changing their study plan. From this perspective, question bank design is not merely administrative work; it essentially shapes the equilibrium landscape of students' learning behaviors in a TfM course. Therefore, modeling instructor–student interactions as an economic game allows instructors to use equilibrium analysis to inform policy settings. By anticipating how question bank design influences students' equilibrium learning behaviors, instructors can design a TfM course that more effectively aligns students' learning behaviors with the intended outcome of conceptual mastery. ***Broadly speaking, our proposed approach aims to inform an instructor of whether students' equilibrium learning behaviors settle into the "gaming the system" behavior, for a given course design, and provide insights into which question banks are worth expanding their sizes.***

4 Mathematical Framework

Having motivated the instructor-student interaction as a strategic setting, we now turn to a general modeling framework for TfM courses. Our goal in this section is to formalize the key elements of TfM question bank design and students' learning behaviors in a way that is applicable across a broad range of TfM course designs. To this end, we adapt Correa and Gruver's extended economic theory of education[7] to the context of TfM courses with randomized quizzes. We introduce the notation, specify the game's structure, and characterize the resulting equilibrium learning behavior. After presenting the framework, we illustrate its use in later sections through a case study that demonstrates how instructors can apply equilibrium analysis to assess whether a particular TfM question bank design likely promotes mastery and determine the optimal question bank size.

4.1 Achievement Production Function

Let π be a TfM course policy set by the instructor. We follow Correa and Gruver's theory to model the student's achievement as an outcome jointly produced by the instructor's effort (E_I) and the student's effort (E_{S_i}). We represent student i's achievement (A_i) by a production function (ν):

$$A_i = \nu(E_{S_i}, E_I|\pi), \tag{1}$$

where the input E_I is the instructor's effort to set up the policy π and prepare the TfM course according to π, and E_{S_i} is the student i's effort spent on the course responding to π. The output A_i is the student i's achievement produced by E_I and E_{S_i} under π.

This production function ν satisfies standard assumptions: (i) There is no negative achievement. (ii) There is no achievement if the student puts no effort. (iii) The achievement increases in both the instructor's and the students' effort. And (iv) The achievement increase is decreasing as effort increases. That is, extra effort helps, but the more effort you have already put in, the smaller the additional benefit. Formally, these define the following constraints for ν:

$$\nu(E_{S_i}, E_I|\pi) \geq 0 \text{ , } \nu(0, E_I|\pi) = 0 \text{ , } \nu'(\cdot|\pi) > 0 \text{ and } \nu''(\cdot|\pi) < 0. \tag{2}$$

These assumptions reflect a simple educational intuition: a well-prepared course and a motivated learner both raise achievement, but neither can increase learning indefinitely. In a TfM context, this production view helps us ask: How much additional learning can we expect when the policy (*e.g.*, expands the question bank) changes, and how do different students respond to such changes?

4.2 Effort Allocation and the Time Constraint

Each student $i = 1, \ldots, N$ and the instructor have a limited time to spend on the course. Following the economic model, we assume that the effort E_{S_i} and E_I are measured in time, and that each student and instructor has a capped total time available for the course, denoted by T_{S_i} and T_I. The time that they do not spend on the course is called the leisure time, denoted by L_{S_i} and L_I. Together, they form the time constraints on effort allocation: $E_{S_i} + L_{S_i} = T_{S_i}$ and $E_I + L_I = T_I$. This simple constraint captures a key reality in educational settings: **every additional hour spent studying or preparing the course is an hour taken away from something else**. These time constraints are the foundation for understanding how TfM policies influence the opportunity cost of learning and teaching.

4.3 Student Types and Incomplete Information

Students differ in how they weigh mastery learning against time spent. We model this through a **student type** τ_i, which influences how much they value achievement relative to leisure. The instructor does not know the individual types of each student but does have a reasonable sense of the overall student population, represented by the distribution $D(\tau)$. This creates a **Bayesian game**:

- The instructor designs the policy π by anticipating the distribution of behaviors from different student types $D(\tau)$.
- Each student observes/infers the policy π and chooses effort based on their personal type τ_i.

By modeling it as a Bayesian game, the framework captures behavioral heterogeneity in TfM courses, from mastery learners to students inclined toward "gaming the system" behaviors.

4.4 Utility Functions

A key challenge in understanding student behavior in TfM courses is that students do not respond to π in the same way. Some students are motivated to achieve deep mastery, others seek to minimize the time they spend, and many balance these goals depending on workload, confidence, or competing commitments. Likewise, instructors must balance their desire for high student achievement with the practical limits of their own time and resources.

To systematically and analytically capture these trade-offs, we use **utility functions**, a standard tool in economics. A utility function summarizes an agent's preferences on how much they value learning relative to leisure, and formalizes the decisions they make under limited time. This allows the model to account for heterogeneous student motivations, to represent the instructor's objectives in designing a TfM course, and to predict how different policy settings will influence effort and learning behaviors. We denote the utility functions as follows:

$$U_{S_i}(A_i, L_{S_i}|\tau_i) \text{ and } U_I(A_1, A_2, \ldots, A_N, L_I), \tag{3}$$

where U_{S_i} is the student i's utility function, which captures the trade-off between their achievement A_i and their remaining leisure time L_{S_i}, under the student i's type τ_i, reflecting how much they personally care about mastery versus minimal efforts; and U_I is the instructor's utility function, which captures the trade-off among all students' achievement $A_i, i = 1, \ldots, N$ and the instructor's remaining leisure time L_I. Similar to the achievement production function, the utility function satisfies the same four constraints as in Eq. 2:

$$U(\cdot,\cdot) \geq 0 \,,\, U(0,\cdot) = 0 \,,\, U'(\cdot,\cdot) > 0 \text{ and } U''(\cdot,\cdot) < 0, \tag{4}$$

which means utility should (i) be non-negative, (ii) be zero if achievement is zero, (iii) increase with achievement and leisure time, and (iv) increase less with additional achievement and leisure time. This formulation captures behaviors educators readily recognize: a student may study intensively until the benefit of studying longer becomes too small, or an instructor may continue to refine the course until additional preparation yields little additional learning.

4.5 Bayesian Nash Equilibrium

In a TfM course, both students and the instructor make decisions about how much effort to invest, but they do so in response to the course policies that structure the learning environment. Once the instructor designs a policy (such as the question bank size, number of quiz attempts, or feedback frequency), students observe/infer that policy and decide how much time to devote to the course. Their decisions depend on how much they value mastery, how much total time they have available, and how effective they expect additional effort to be under the current policies. Mathematically, each student chooses the level of effort that maximizes their own utility:

$$E^*_{S_i}(\pi, \tau_i) = \arg\max_{E_{S_i}} U_{S_i}(A_i, L_{S_i}|\tau_i), \tag{5}$$

which means: **given the instructor's policy π, each student chooses the study effort that works best for them**.

Education Interpretation of Bayesian Nash Equilibrium

A Bayesian Nash equilibrium is reached when:

- No student would benefit from adjusting their study effort, given the policy and their types.
- The instructor would not benefit from modifying policies, given students' expected responses.

In other words, equilibrium describes the steady-state behavior of a TfM course: ***students settle into predictable study patterns, and the instructor's policy is optimal given those patterns.*** Therefore, analyzing this equilibrium allows instructors to evaluate whether a TfM policy promotes mastery learning or inadvertently creates incentives for gaming.

5 Applying the Framework to a Teaching for Mastery (TfM) Course

The mathematical framework presented above provides a general framework for reasoning about how TfM course policies influence students' learning behaviors. We now apply this framework to a specific TfM course by focusing on question bank design as the primary instructional policy. This application illustrates how equilibrium analysis can help instructors evaluate whether their course design is encouraging the desired behavior.

5.1 Course Context and Policy Specification

To demonstrate how the proposed framework operates in practice, we apply it to an undergraduate TfM course that uses randomized quizzes generated from question banks. In this setting, the instructor selects a policy π that determines key features of the assessment environment, including the size of the question bank, the diversity of problem types, and the randomization rules used to generate quizzes. These design choices influence not only the instructor's preparation effort but also how students allocate their time across practice, review, and other learning activities.

In this course, student effort is measured by quiz attempts and time spent on practice problems, while instructor effort is reflected in the development and maintenance of the question banks. Achievement is operationalized using quiz performance and pass rates. By specifying how these course elements map onto the framework's variables, we can use equilibrium analysis to evaluate whether different question bank designs encourage mastery learning or unintentially allow for "gaming the system" behavior.

5.2 Data Collection

The data used in this study consist of routinely collected course records generated through normal instructional practice. IRB approval was obtained prior the study (IRB# 2526-061), and consent forms were sent to students via emails. Our analysis draws on archival Moodle quiz records from a TfM course spanning six academic semesters (Spring 2022, Spring 2023, Spring 2024, Fall 2024, Spring 2025, and Fall 2025) with a total of 233 students enrolled. Out of 233 students, 106 provided consent, 3 declined, and 122 did not respond. We exported the quiz-attempt logs of the students who consented from Moodle and assigned anonymized identifiers before analysis.

The question bank design comprises 36 quizzes over ten topics. Across six semesters, the exported Moodle logs yield a total of 6,625 quiz-attempt records, each corresponding to one attempt by an anonymized student on a single quiz. Question banks vary substantially in size, from small banks containing as few as two to five questions to larger banks containing up to 30 questions. Each quiz attempt consists of multiple questions randomly drawn from the relevant question bank. Consequently, the number of possible quiz forms (*i.e.*, unique question combinations) ranges from 2 to 20,160, depending on how many questions are used in a quiz.

Because quizzes function as alternate assessments of the same underlying topic, we aggregate at the topic level for analysis. Some quizzes assess overlapping concepts across multiple topics; in these cases, we include the same quiz in each relevant topic aggregate. This structural variation yields a wide range of question bank designs that support structural estimation. Tab. 1 summarizes variation in question bank size and the corresponding number of attempt records.

Practically speaking, this heterogeneity implies that different designs present different incentives for mastery learning vs. "gaming the system" behavior. This variation allows us to examine how question bank size is associated with students' learning behaviors in the TfM course - specifically, how retake patterns, time spent on quizzes, and scores relate to mastery learning and "gaming the system" behavior.

	Variation	SP 22	SP 23	SP 24	FA 24	SP 25	FA 25	Total
Topic 1	3-30	146	450	405	117	338	249	1705
Topic 2	1-30	102	380	319	86	250	178	1315
Topic 3	3-6	34	68	72	22	51	31	278
Topic 4	2-30	25	143	120	31	126	73	518
Topic 5	4	15	30	43	17	15	25	145
Topic 6	2-7	31	81	78	16	65	48	319
Topic 7	3-14	71	180	143	38	113	84	629
Topic 8	1-30	103	308	279	75	218	158	1141
Topic 9	4-7	13	40	64	13	48	36	214
Topic 10	5-24	12	64	98	29	89	69	361
# of Students	-	8	17	12	26	25	19	107[1]

Table 1: Variation in question bank size and the corresponding number of attempt records by topic and semester. The last column shows the total number of records included in the analysis, and the last row shows the total number of students who provided consent.

5.3 Choices of Achievement Production Function and Utility Functions

To use the framework for equilibrium analysis, we need to model the achievement production function ν (Eq. 1) and utility function U (Eq. 3). We first used a correlation matrix to examine relationships among the collected data (such as time spent, number of attempts, and scores), but we did not find any strong linear relationship. We then used our educator's judgment and defined the student's effort as the total recorded time a student spent. Let G_i be the number of attempts student i made and $\bar{t}_i$ be the average time the student spent per attempt. Then the student's effort is $E_{s,i} = G_i\bar{t}_i$. We also assume the instructor's effort to prepare N_q questions follows an exponential model $E_I = e^{\kappa_0+\kappa_1 N_q}$, capturing compounding marginal effort as the question banks expand. *i.e.*, the more questions in the bank, the harder to prepare a new question for the same topic. In this model, we can interpret κ_0 as the fixed effort required to prepare the quiz, and κ_1 as the linear factor for preparing N_q questions in the bank. Furthermore, we assume that production is affected by the policy $\pi = 1/N_q$ (*i.e.*, only the number of questions in the bank in our case). Modeling the policy input as $1/N_q$ ensures the predicted achievement is decreasing in bank size (for $\beta > 0$), consistent with the idea that larger banks dilute exposure to any particular question and increase

[1]One student retook the course. In our analysis, we assigned the same anonymized identifier to this student, so all quiz records of this student in two semesters were treated as a single individual's behavior records.

effective task difficulty. Using the classic Cobb-Douglas function[10] from economics, we model the achievement function as:

$$A_i = \nu(E_{S_i}, E_I | \pi) = E_{s,i}^{\alpha} \pi^{\beta} E_I^{1-\alpha-\beta} \tag{6}$$

Because the question bank size is our policy variable, we also impose a policy monotonicity condition: holding effort fixed, larger banks weakly reduce achievement,

$$\frac{\partial \nu(E_{S_i}, E_I | \pi)}{\partial \pi} \leq 0. \tag{7}$$

This captures that larger banks reduce question repetition and raise effective difficulty, limiting memorization-based strategies.

To model student preferences over performance and effort, we adopt a Constant Elasticity of Substitution (CES) utility function[11]. The CES function provides a flexible yet parsimonious way to represent how students trade off achievement against leisure time, while allowing the degree of substitutability between the two components to be governed by a single curvature parameter ρ. Specifically, we assume

$$U_{s,i}(A_i, L_{s,i}) = \Big(\omega_i \, A_i^{\rho} + (1 - \omega_i) \, L_{s,i}^{\rho}\Big)^{1/\rho}, \tag{8}$$

where the weight $\omega_i \in (0, 1)$ governs the relative importance of achievement versus leisure, and ρ controls the elasticity of substitution between A_i and $L_{s,i}$ (with smaller ρ implying stronger complementarity and larger ρ implying greater substitutability).

5.4 Identifying Student Types

For our purposes, we classify students into two behavioral types: "gamers" and "non-gamers". Gamers exhibit "gaming the system" behaviors by rapidly retaking quizzes to recognize previously seen questions and increase scores through memorization rather than conceptual learning. Non-gamers, in contrast, primarily use retakes as intended, *i.e.*, to practice, receive feedback, and improve mastery.

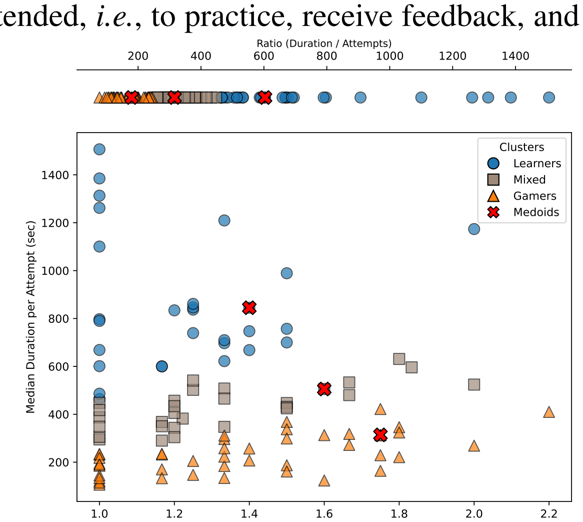

Because Moodle logs do not directly reveal students' intent, we operationalize these types using behavioral proxies derived from quiz-attempt records. For each student, we compute the average number of attempts per quiz and the median duration per attempt. We use the average number of attempts to capture overall retake intensity across quizzes, and the median duration to reduce sensitivity to extreme outliers (*e.g.*, leaving a browser tab open for a while and immediately finishing a quiz). We cluster students using k-medoids on the ratio of these two indicators[2]. To allow for a potential "mixed" behavioral pattern, we set $k = 3$ and interpret the resulting clusters as "gamer", "learner",

[2]We use k-medoids rather than k-means for its robustness to non-Gaussian distributions and outliers.

and “mixed”. For subsequent analysis, we adopt a conservative labeling approach and merge the mixed cluster with learners, yielding a binary partition into “gamer” and “non-gamer”. This choice reduces the risk of over-classifying students as gamers. As shown in the above inset, an example of our clustering results, we identified three behaviorally distinct clusters: a low-attempt/high-duration group (“learner” highlighted in yellow), a high-attempt/low-duration group (“gamer” highlighted in teal), and a mixed group (highlighted in purple).

5.5 Parameter Estimations

Achievement Production Function

We estimated the parameters of the achievement function via structural estimation, fitting the model to observed quiz attempt behavior using likelihood-based methods[12]. To connect the achievement function to observed performance, we use a binary mastery indicator $Y_{i,q} \in \{0, 1\}$ (pass/fail) and assume

$$P(Y_{i,q} = 1|\nu) = \frac{1}{1 + e^{-(\theta_0 + \theta_1 \log \nu)}} := P(\nu), \tag{9}$$

so that higher predicted achievement corresponds to a higher probability of passing. We estimate the parameters $\alpha, \beta, \kappa_0, \kappa_1, \theta_0, \theta_1$ (from Eq. 6 and 9) using the SciPy package[3] by maximizing the corresponding log-likelihood of Eq. 9 subject to the necessary constraints (Eq. 2 and 7).

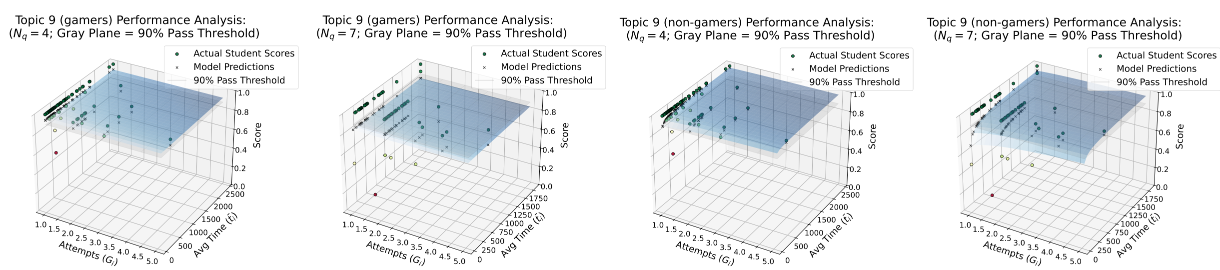

Figure 1: Each figure overlays actual student outcomes (circles) and model-predicted probabilities (×) across the effort dimensions, with the translucent plane indicating the 90% pass threshold. Figures compare question bank sizes ($N_q = 4$ and $N_q = 7$) and student types (gamers vs. non-gamers). Overall, higher predicted achievement corresponds to higher observed pass rates, and predicted pass probabilities are systematically higher for observed passes than fails.

Fig. 1 provides a visual assessment of model fit on Topic 9. The fitted probabilities are broadly aligned with the observed mastery outcomes: attempts with higher predicted achievement tend to exhibit higher observed pass rates, and the distribution of predicted pass probabilities is shifted upward for observed passes relative to fails. While the model is intentionally parsimonious, the estimated specification remains consistent with the production function constraints in Eq. 2. The Cobb-Douglas function ensures nonnegative achievement and implies diminishing marginal returns to effort, and the estimated exponents yield a monotone relationship between achievement and both students’ and instructor’s effort. In addition, the policy term $\pi = 1/N_q$ enters with a positive exponent, so predicted achievement decreases with question bank size (holding effort fixed). These patterns suggest the fitted model is suitable for the equilibrium analyses that follow.

[3]Estimation is implemented in Python using SciPy[13].

Student Utility Functions

When choosing a value for ω_i and ρ (Eq. 8), we consulted previous literature on the economics of education. There were not many existing works that suggested actual numerical values for the elasticity of substitution of student achievement for leisure. What we did find was a work from Foster and Gehrke[14] that compared investment into human capital versus investment into schooling. Their results found an elasticity of substitution (σ) of 0.290, which corresponds to a ρ of -2.45, related by $\sigma = \frac{1}{1-\rho}$. This was the closest proxy we could find, so we used it in our analysis.

ω_i represents the weight that is assigned to the academic achievement of student i, so we know that this value should be higher for non-gamers and lower for gamers, by the nature of these types. The weights placed on neutral goods, which leisure and achievement would be for non-gamers, correspond to a value of 0.5 for ω_i[15]. We did not find existing literature that could serve as a proxy for gamer weight, so we followed economic benchmark standards to set ω_i to 0.25. Note that the choice of these values must satisfy the utility constraints (Eq. 4).

6 Results

6.1 Nash Equilibrium Analysis

After specifying the utility functions, we analyze Nash equilibrium behavior by solving Eq. 5 under varying question bank sizes N_q from 1 to 200. Across all ten topics, the equilibrium predictions exhibit the same qualitative pattern: the optimal time-on-task ($G_i\bar{t}_i$) increases with question bank size (N_q) for both gamers and learners, rises most sharply when N_q is small, and then gradually flattens as N_q becomes larger. By approximately $N_q = 40$ to 60, the curves are nearly unchanged, indicating that further expansion is unlikely to generate meaningful additional increases in $G_i\bar{t}_i$ at equilibrium. For the topics in our case study, this range therefore provides a practical upper bound for question bank design.

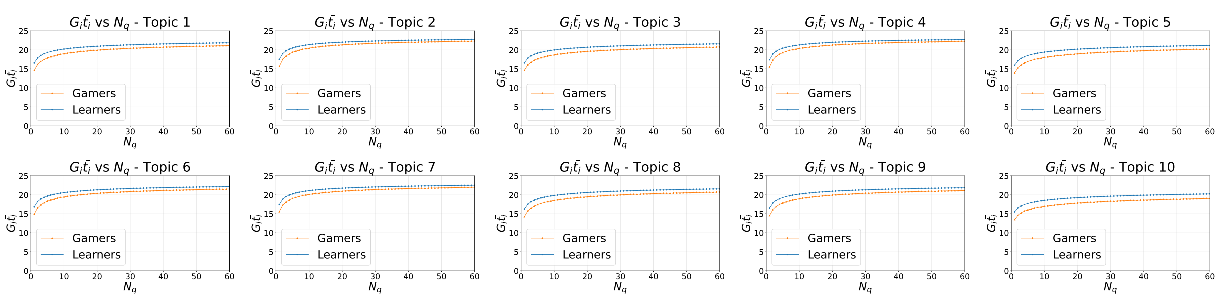

Figure 2: Nash equilibrium predictions of optimal time-on-task ($G_i\bar{t}_i$) as a function of question bank size (N_q) for each topic. For each panel, the curves report the equilibrium solution of Eq. 5 over a range of N_q, shown separately for the two student types (learners vs. gamers) using the estimated topic-specific parameters. Note: learners here include the mixed group.

Fig. 2 presents the equilibrium predictions of optimal time-on-task ($G_i\bar{t}_i$) as a function of question bank size (N_q) for each topic. As expected from the model, the learner curve lies above the gamer curve in every topic. The gap is modest, and the two curves are generally close to parallel. Thus, increasing N_q raises the predicted equilibrium effort for both groups, while preserving their relative ordering: learners continue to exhibit higher time-on-task than gamers at comparable bank sizes.

The more important topic-level differences in Fig. 2 are not the directions of the curves, but their levels and the points at which they begin to flatten. Topics 2, 4, and 7 show relatively high equilibrium $G_i\bar{t}_i$ even at low N_q. Topics 1, 6, and 9 occupy an intermediate position but follow the same saturating pattern. Topics 3, 5, 8, and especially 10 lie at lower levels overall, indicating that these topics require a larger bank size to reach the same equilibrium effort achieved more quickly in topics 2, 4, and 7. However, a lower curve does not by itself imply that a topic should be expanded. Whether expansion is needed depends on where the currently observed bank sizes in Tab. 1 fall along the corresponding equilibrium curve.

Interpreting Fig. 2 together with Tab. 1 allows us to identify where additional question-writing is likely to matter most. Topic 5 is the clearest candidate for expansion. Its observed bank size is fixed at $N_q = 4$, and Fig. 2 shows that this value lies in the steep portion of the Topic 5 equilibrium curve, where even a modest increase in N_q would produce a noticeable increase in equilibrium effort. Topics 3, 6, and 9 lead to the same conclusion. Their observed bank-size ranges ($N_q = 3$–6, 2–7, and 4–7, respectively) also lie mainly in the steep part of the curves, suggesting that increasing these banks to around $N_q = 10$ would move them toward a higher equilibrium effort.

By contrast, topics whose observed bank sizes already extend well into the moderate range are lower-priority candidates for further expansion. In particular, Topics 1, 2, 4, and 8 already reach $N_q = 30$, where the corresponding equilibrium curves have largely flattened. Topic 10 also extends to $N_q = 24$; although its curve lies relatively low, part of its observed range is already beyond the steepest region of the curve. Thus, for these topics, adding more questions may still increase equilibrium effort, but the predicted marginal gain is smaller than for topics whose current bank sizes remain concentrated at low N_q.

Overall, the equilibrium analysis provides both a prioritization rule and a practical target for redesign. ***Topics whose current bank sizes still lie in the steep portion of their equilibrium curves should be expanded first***, because they offer the largest predicted behavioral gains. Under the current course design, Topic 5 is the strongest candidate for expansion, followed by Topics 3, 6, and 9. For these topics, expanding the bank to at least around $N_q = 10$ appears justified by the model. In contrast, topics already reaching moderate or large bank sizes appear closer to saturation, so further expansion is likely to yield only limited additional benefit.

6.2 Sensitivity Analysis

Because the equilibrium results depend on the utility student parameters chosen in Eq. 8, we conduct a sensitivity analysis to examine whether the equilibrium analysis in Fig. 2 is robust to alternative parameter values. We examined the parameter sensitivity of each topic, and they all exhibited a similar pattern. As an example, Fig. 3 illustrates this analysis for Topic 7 by varying ω_i and ρ (σ) separately for gamers and learners. We focus on these two parameters because they encode the behavioral assumptions introduced in Sec. 5.3: ω_i governs how much weight a student places on achievement relative to leisure, whereas ρ (σ) governs the curvature of that trade-off through the elasticity of substitution. In our baseline specification, we set $\rho = -2.45$ ($\omega = 0.29$) using the closest proxy available from prior literature, and we allow ω_i to differ by student type, with a higher value for learners than for gamers, reflecting the assumption that learners place relatively more weight on achievement. Fig. 3, therefore, serves as a check on whether our equilibrium results are driven by these particular parameter choices.

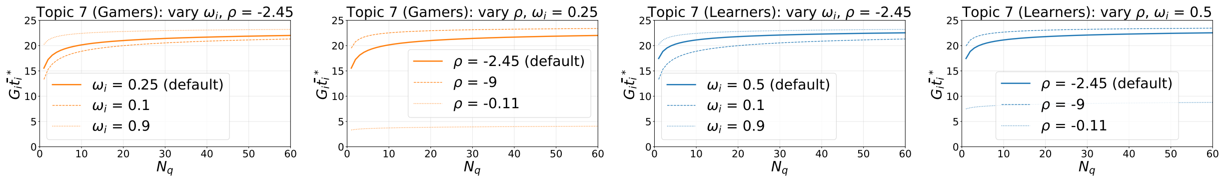

Figure 3: Sensitivity analysis for Topic 7. The panels show Nash equilibrium predictions of optimal time-on-task ($G_i\bar{t}_i$) as a function of question bank size (N_q) under alternative utility parameters. The left and third panels vary ω_i while holding ρ fixed at its default value, and the second and fourth panels vary ρ while holding ω_i fixed at its default value, shown separately for gamers and learners.

Fig. 3 shows that varying ω_i mainly shifts the equilibrium curves vertically while preserving their overall shape. For both gamers and learners, a higher ω_i leads to a higher equilibrium $G_i\bar{t}_i$, as expected, because students place a greater weight on achievement relative to leisure. However, the qualitative pattern remains unchanged: equilibrium effort still rises quickly when N_q is small and then gradually flattens as N_q increases. This is important for interpretation. Changing ω_i alters the predicted effort level and the gap between gamers and learners, but it does not materially change the analysis of the effect of question bank expansion. In other words, ω_i mainly controls *how much* effort each type is willing to supply, rather than *whether* larger banks increase effort.

By contrast, the equilibrium predictions are more sensitive to the choice of ρ (σ). When ρ (σ) is varied from the default value of -9 (0.1) to -2.45 (0.29), the resulting curves remain qualitatively similar, with only moderate changes in the equilibrium level. However, when ρ (σ) is increased to -0.11 (0.9), the equilibrium effort drops substantially and the response to N_q becomes much weaker for both gamers and learners. Thus, compared with ω_i, the curvature parameter ρ has a stronger influence on the quantitative magnitude of the equilibrium predictions. This is also why, in our setting, we keep ρ (σ) fixed across student types. Conceptually, ρ (σ) governs the general form of the trade-off between achievement and leisure in the CES utility function, whereas ω_i captures type-specific preferences within that common structure. Allowing ρ (σ) to differ across gamers and learners would make it more difficult to interpret whether observed differences arise from distinct preferences over achievement or from fundamentally different substitution structures.

Overall, Fig. 3 suggests that the main qualitative implication of our model is robust: increasing question bank size raises equilibrium student effort most strongly when N_q is small, with diminishing returns as N_q becomes larger. At the same time, the sensitivity analysis clarifies how the utility parameters affect interpretation. The parameter ω_i changes the vertical position of the equilibrium curves and therefore influences how strongly we distinguish gamers from learners in equilibrium. A higher ω_i for gamers, for example, would make their behavior appear closer to that of learners, whereas a lower ω_i for learners would compress the gap between the two groups. By contrast, ρ (σ) affects the overall strength of the equilibrium response itself. Therefore, our conclusions about which topics are likely under-expanded remain qualitatively stable, but the precise quantitative targets should be interpreted with greater caution when they depend strongly on ρ (σ).

7 Education Implications and Conclusion

This study presents a practical, data-based approach for improving TfM question bank design using the activity records that learning platforms already collect. We (i) group students based on how they use retakes (how often they retry, how long they spend, and how their scores change), (ii) fit a model that links student effort and instructor design choices to achievement, and (iii) use the model to predict how students would respond if the question bank size changed. The main takeaway for education is straightforward: question bank size is not just a technical setting. It shapes which strategies "work" for students, whether spending time learning is rewarded, and whether repeated attempts with minimal effort can still lead to a pass.

A key design result is that the equilibrium effect of question bank expansion depends on where a topic's current bank size lies on its equilibrium curve. Fig. 2 shows that equilibrium effort rises quickly when N_q is small and then gradually flattens as N_q becomes larger, whereas Tab. 1 shows that topics enter this curve at different points. This means that the instructional value of writing more questions is topic-specific. In our case study, Topic 5 is the clearest priority for expansion because its observed bank size is fixed at $N_q = 4$, which lies in the steeply changing part of the equilibrium curve. Topics 3, 6, and 9 are also strong candidates for expansion because their observed bank-size ranges remain similarly small. By contrast, Topics 1, 2, 4, and 8 already reach $N_q = 30$, where the equilibrium curves are much flatter.

From a classroom implementation perspective, the framework supports an instructor-friendly workflow each semester: (i) export quiz-attempt logs, (ii) aggregate topic-level behavioral summaries, (iii) identify topics where short-duration retakes are associated with passing (a potential signature of gaming), and (iv) use the framework to evaluate counterfactual question bank designs such as whether the question bank sizes are currently lying on the steeper parts of the equilibrium curves.

In conclusion, we demonstrate that routinely collected Moodle records can support structural modeling and equilibrium analysis that connect question bank design choices to predicted student behavior. Beyond descriptive analytics, the proposed approach provides prescriptive guidance by identifying where additional questions are likely to have the greatest impact on equilibrium effort, and by offering a principled way to select a cost-effective target for N_q. Broadly speaking, a 10-20 question bank size seems to be the best design across all topics. Importantly, although we focus on question bank size in this study, the framework is general and can be applied to other assessment policy settings: for example, retake limits, scoring rules (highest score vs. average), time limits, feedback timing, and mastery thresholds, by modifying the policy variable and solving the corresponding equilibrium. Future work will extend this framework through prospective studies that experimentally vary question bank size and retake policies, and by developing practical data-collection guidelines (*e.g.*, consistent quiz settings, timing granularity, and item tagging) that improve identifiability and enable broader adoption. Our longer-term goal is a set of evidence-based course design principles for TfM policy that are scalable for instructors and robust against "gaming the system" behavior.

References

[1] S. V. DeWeese and J. J. Randolph, "Effective Use of Correctives in Mastery Learning," Tech. Rep., Feb. 2011, eRIC Number: ED523991. [Online]. Available: https://eric.ed.gov/?id=ED523991

[2] M. Winget and A. M. Persky, “A Practical Review of Mastery Learning,” *American Journal of Pharmaceutical Education*, vol. 86, no. 10, p. ajpe8906, Dec. 2022. DOI:10.5688/ajpe8906. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0002945923007386

[3] M. Garg and A. Goel, “A systematic literature review on online assessment security: Current challenges and integrity strategies,” *Computers & Security*, vol. 113, p. 102544, Feb. 2022. DOI:10.1016/j.cose.2021.102544. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0167404821003680

[4] J. Li, L. Fang, Y. Liu, J. Xie, and X. Wang, “Ineffective Learning Behaviors and Their Psychological Mechanisms among Adolescents in Online Learning: A Narrative Review,” *Behavioral Sciences*, vol. 14, no. 6, p. 477, Jun. 2024. DOI:10.3390/bs14060477. [Online]. Available: https://www.mdpi.com/2076-328X/14/6/477

[5] X. Wang and Y. Fu, “Some Characterizations of the Cobb-Douglas and CES Production Functions in Microeconomics,” *Abstract and Applied Analysis*, vol. 2013, pp. 1–6, Dec 2013. DOI:10.1155/2013/761832. [Online]. Available: http://www.hindawi.com/journals/aaa/2013/761832/

[6] T. C.-C. Tan and S. R. Da Costa Werlang, “The Bayesian foundations of solution concepts of games,” *Journal of Economic Theory*, vol. 45, no. 2, pp. 370–391, Aug. 1988. DOI:10.1016/0022-0531(88)90276-1. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/0022053188902761

[7] H. Correa and G. W. Gruver, “Teacher-student interaction: A game theoretic extension of the economic theory of education,” *Mathematical Social Sciences*, vol. 13, no. 1, pp. 19–47, Sep 1987. DOI:10.1016/0165-4896(87)90057-6. [Online]. Available: https://www.sciencedirect.com/science/article/pii/0165489687900576

[8] H. Bonesrønning, “Do the teachers’ grading practices affect student achievement?” *Education Economics*, vol. 12, no. 2, pp. 151–167, Aug. 2004. DOI:10.1080/0964529042000239168. [Online]. Available: https://www.tandfonline.com/doi/full/10.1080/0964529042000239168

[9] L. Wittie, “Incorporating a Teach for Mastery System in a Computer Science Course,” in *2023 ASEE Annual Conference & Exposition*, no. 10.18260/1-2–43624. Baltimore , Maryland: ASEE Conferences, Jun 2023. DOI:10.18260/1-2–43624. [Online]. Available: https://peer.asee.org/43624

[10] C. W. Cobb and P. H. Douglas, “A theory of production,” *The American Economic Review*, vol. 18, no. 1, pp. 139–165, Mar 1928. [Online]. Available: http://www.jstor.org/stable/1811556

[11] K. J. Arrow, H. B. Chenery, B. S. Minhas, and R. M. Solow, “Capital-Labor Substitution and Economic Efficiency,” *The Review of Economics and Statistics*, vol. 43, no. 3, pp. 225–250, Aug 1961. DOI:10.2307/1927286. [Online]. Available: http://www.jstor.org/stable/1927286

[12] T. J. Holmes and H. Sieg, “Chapter 2 - Structural Estimation in Urban Economics,” in *Handbook of Regional and Urban Economics*, ser. Handbook of Regional and Urban Economics, G. Duranton, J. V. Henderson, and W. C. Strange, Eds. Elsevier, May 2015, vol. 5, pp. 69–114. DOI:10.1016/B978-0-444-59517-1.00002-7. [Online]. Available: https://www.sciencedirect.com/science/article/pii/B9780444595171000027

[13] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, A. Vijaykumar, A. P. Bardelli, A. Rothberg, A. Hilboll, A. Kloeckner, A. Scopatz, A. Lee, A. Rokem, C. N. Woods, C. Fulton, C. Masson, C. Häggström, C. Fitzgerald, D. A. Nicholson, D. R. Hagen, D. V. Pasechnik, E. Olivetti, E. Martin, E. Wieser, F. Silva, F. Lenders, F. Wilhelm, G. Young, G. A. Price, G.-L. Ingold, G. E. Allen, G. R. Lee, H. Audren, I. Probst, J. P. Dietrich, J. Silterra, J. T. Webber, J. Slavič, J. Nothman, J. Buchner, J. Kulick, J. L. Schönberger, J. V. de Miranda Cardoso, J. Reimer, J. Harrington, J. L. C. Rodríguez, J. Nunez-Iglesias, J. Kuczynski, K. Tritz, M. Thoma, M. Newville, M. Kümmerer, M. Bolingbroke, M. Tartre, M. Pak, N. J. Smith, N. Nowaczyk, N. Shebanov, O. Pavlyk, P. A. Brodtkorb, P. Lee, R. T. McGibbon, R. Feldbauer, S. Lewis, S. Tygier, S. Sievert, S. Vigna, S. Peterson, S. More, T. Pudlik, T. Oshima, T. J. Pingel, T. P. Robitaille, T. Spura, T. R. Jones, T. Cera, T. Leslie, T. Zito, T. Krauss, U. Upadhyay, Y. O. Halchenko, Y. Vázquez-Baeza, and S. 1.0 Contributors, “SciPy 1.0: fundamental algorithms for scientific computing in

Python," *Nature Methods*, vol. 17, no. 3, pp. 261–272, Mar 2020. DOI:10.1038/s41592-019-0686-2. [Online]. Available: https://doi.org/10.1038/s41592-019-0686-2

[14] A. Foster and E. Gehrke, "Start What You Finish! Ex Ante Risk and Schooling Investments in the Presence of Dynamic Complementarities," National Bureau of Economic Research, Cambridge, MA, Tech. Rep. w24041, Nov. 2017. DOI:10.3386/w24041. [Online]. Available: http://www.nber.org/papers/w24041.pdf

[15] F. E. Kydland and E. C. Prescott, "Time to Build and Aggregate Fluctuations," *Econometrica*, vol. 50, no. 6, p. 1345, Nov. 1982. DOI:10.2307/1913386. [Online]. Available: https://www.jstor.org/stable/1913386?origin=crossref