Title
How Instructors Regulate AI in College: Evidence from 31,000 Course Syllabi

Journal
Higher Education Working Paper Series, 26(1)

Author
Chirikov, Igor

Publication Date
2026-02-02

UC Berkeley Center for Studies in Higher Education

Higher Education Working Paper

Volume No. 26-1

How Instructors Regulate AI in College: Evidence from 31,000 Course Syllabi

Igor Chirikov
Center for Studies in Higher Education
University of California, Berkeley

Suggested Citation: Chirikov, I (2026). How Instructors Regulate AI in College: Evidence from 31,000 Course Syllabi. *CSHE Higher Education Working Paper Series*, Vol. 26-1. Retrieved from: escholarship.org/uc/item/9c51s3gs

How Instructors Regulate AI in College: Evidence from 31,000 Course Syllabi

Igor Chirikov*

Center for Studies in Higher Education,
Goldman School of Public Policy, UC Berkeley
chirikov@berkeley.edu

February 2026

Abstract

The way instructors support or restrict student use of generative artificial intelligence (AI) tools shapes which skills students develop in college and their preparation for AI-augmented jobs. This paper introduces a task-based approach to AI and skill formation to study how instructors regulate student use of AI, predicting that regulation depends on course task composition and the strength of AI capabilities in those tasks. I test these predictions using 31,000 course syllabi representing the full universe of courses at a large public research university from 2021 to 2025, providing novel large-scale longitudinal evidence on instructor responses to AI. I use computational methods to extract required course tasks and AI policies from each syllabus, tracking the same courses and instructors over time. I document three main results. First, instructors are warming toward AI over time, shifting from initial restrictive policies toward greater permissiveness, differentiating policies by task type, and increasingly introducing new AI-based tasks, with substantial disciplinary variation. Second, course task composition in 2021-22 (pre-ChatGPT) predicts AI regulation in 2025: courses requiring more practice in tasks with strong AI capabilities are more likely to regulate students' use of AI. Third, courses requiring more tasks in 2021-22 are more likely to differentiate permissions by task type in 2025, consistent with diverse task requirements creating both risks and opportunities for AI affecting task practice. Taken together, these findings indicate a move toward nuanced regulation aligned with task-level concerns about skill formation.

* I am grateful to Lynn Chien, Christopher Mach, Sophie McKenna, and Smrithi Senthilnathan for their exceptional research assistance throughout this project, supported by UC Berkeley's Data Discovery Program. I also thank Jesse Rothstein, Rene Kizilcec and Yan Tao for their thoughtful comments that improved this paper considerably. All errors are my own.

1 Introduction

As generative artificial intelligence (AI) reshapes professional work, instructors face decisions about whether and how students should use AI in their courses. These decisions are difficult because AI raises not only immediate academic integrity concerns about students misrepresenting their work, but also more fundamental questions about skill formation. AI can substitute for the complex cognitive tasks students must practice to develop skills such as writing, analysis, coding, and problem-solving. If AI substitutes for this practice, students may fail to develop skills; if AI is restricted entirely, students may miss opportunities to enhance learning or lack experience with tools expected in the labor market. How instructors navigate this tension shapes which tasks students perform, the skills they develop, and their preparation for AI-augmented workplaces.

To understand how instructors navigate this tension, I adapt task-based models of technological change (Autor et al. 2003; Acemoglu and Restrepo 2019) from production to education. These models conceptualize production as bundles of tasks allocated between workers and technology, and show that technology affects employment by displacing workers from some tasks while creating new tasks where workers retain comparative advantage. I apply this logic upstream to education. Students develop skills by practicing cognitively demanding tasks (Bjork and Bjork 2011; Freeman et al. 2014), with different courses requiring different task bundles. AI affects skill formation by changing which tasks students practice and how. Instructors thus face different pressures from AI depending on which tasks their courses require and AI capabilities in those tasks.

I develop a task-based approach to AI and skill formation that distinguishes three mechanisms through which AI affects task practice in education. *Task displacement* occurs when AI performs tasks instead of students, eliminating practice opportunities and risking skill erosion. *Task augmentation* occurs when AI supports task practice without displacing essential cognitive effort, potentially enhancing learning. *Task reinstatement* occurs when AI enables new tasks not previously feasible. Instructors respond to these mechanisms by restricting AI for tasks facing displacement risks, permitting AI for tasks with augmentation potential, or introducing new AI-based tasks. This framework predicts that instructors will increasingly regulate AI and differentiate permissions by task type, with courses requiring more practice in tasks where AI capabilities are strongest facing greater pressure to regulate.

Testing these predictions requires systematic data on instructor responses to AI, but existing evidence is limited and largely indirect. It comes primarily from research on institutional policies (McDonald et al. 2025), faculty surveys (Love and Blankstein 2024; Watson and Rainie 2026), and small-scale syllabi studies (Ali et al. 2025; Moore and Lookadoo 2024). Institutional guidance documents policy at the university level but may not reflect what instructors communicate to students. Surveys capture stated beliefs and attitudes but not actual course policies. Small-scale syllabi studies document AI policies within specific disciplines but cannot establish temporal trends or systematic patterns across fields.

To test these predictions systematically, I analyze over 31,000 course syllabi from 2021 to 2025 at a large research university. The longitudinal structure captures the pre-AI baseline (2021-2022), initial disruption (Spring 2023), and subsequent adaptation (2023-2025), and allows me to track the same courses taught by the same instructors across multiple years. I use computational methods to extract and classify AI policies and learning tasks from syllabi at scale, with LLM-assisted classification validated at 96% agreement with human coding.

By analyzing the full universe of courses, this study provides the first large-scale longitudinal evidence on how instructors are navigating students' use of AI in their courses. Course syllabi offer a direct window into instructors' decisions about students' use of AI. While syllabi may not fully reflect classroom practice or students' actual use of AI, they reveal which tasks instructors try to protect from displacement, which they view as candidates for augmentation, and which new AI-based tasks they reinstate.

The findings support the framework's predictions and advance understanding of how instructors navigate AI regulation by documenting substantial variation and evolution in their responses. First, while explicit AI regulation grew rapidly to 55% of courses by Fall 2025, the direction of that regulation is shifting. Initial policies were predominantly restrictive, but instructors are warming toward AI over time, introducing more permissive policies and incorporating new AI-based tasks into their courses. The framing of those policies has evolved as well, with declining emphasis on academic integrity and growing attention to AI's impact on learning. Second, courses requiring more practice in tasks where AI capabilities are strongest (writing and coding) in 2021-22 were more likely to adopt AI policies by Fall 2025, consistent with the prediction that courses face greater pressure when AI capabilities overlap with required task practice. Third, instructors increasingly differentiate AI permissions by task type, restricting AI for some tasks while permitting it for others within the same course. Fourth, courses with larger task bundles in 2021-22 were more likely to adopt differentiated policies by 2025, consistent with the prediction that diverse task requirements create more opportunities for instructors to perceive both displacement risks and augmentation/reinstatement potential.

This study makes three contributions. Theoretically, it demonstrates the potential of task-based models for understanding how AI affects skill formation, showing that the same displacement and reinstatement mechanisms that operate in production can be applied upstream to education. The framework may help reconcile divergent findings in the literature, where some studies document skill gains from AI use (Pardos and Bhandari 2024; Kumar et al. 2025) while others find skill erosion (Kosmyna et al. 2025; Budzyń et al. 2025). These findings are not contradictory but reflect different mechanisms operating on different tasks. Empirically, the study provides large-scale longitudinal evidence on instructors' responses to generative AI using actual course artifacts rather than surveys or institutional policies. The findings show that instructors are not uniformly considering AI as a threat to learning but actively navigating its implications, with systematic patterns in how they differentiate across tasks and disciplines. For policy, the substantial task and disciplinary variation documented here suggests that effective institutional AI policies require flexibility to accommodate disciplinary and task-level differences. The findings also highlight a potential feedback loop between AI in education and AI in production: if AI displaces skill-building tasks, students may graduate with weaker skills in precisely the areas where AI is strongest, further shifting comparative advantage toward AI and accelerating automation.

The paper proceeds as follows. Section 2 develops a conceptual framework extending task-based models to skill formation. Section 3 reviews existing evidence and presents research questions. Section 4 describes data and institutional context. Section 5 presents measurement and methods. Section 6 presents results. Section 7 discusses implications and concludes.

2 Conceptual Framework: Task-Based Approach to AI and Skill Formation

This section develops a conceptual framework to study instructors' responses to AI by adapting the task-based approach from labor economics. Task-based models have a rich empirical history of documenting the impact of technology, including AI, on work (Autor et al. 2003; Acemoglu and Restrepo 2019; Brynjolfsson et al. 2025), making it intriguing to explore whether they apply upstream in educational settings where skills are formed.

2.1. Task-Based Approach in Production

Task-based frameworks conceptualize production not as a direct function of labor and capital, but as the completion of a range of tasks that are allocated between these factors based on comparative advantage (Autor et al 2003; Acemoglu and Autor 2011; Autor 2013). In this framework, a task is a unit of work activity that produces output, while a skill is a worker's capability to perform various tasks. Capital takes the form of technology that can perform certain tasks. Technological change affects production by altering which tasks are performed by labor versus capital, not merely by making either factor more productive uniformly. Different industries and occupations require different bundles of tasks, creating variation in exposure to technological change based on task composition.

Acemoglu and Restrepo (2019) distinguish between two fundamental mechanisms through which technology reshapes the task content of production. First, automation enables capital to substitute for labor in tasks that can be accomplished by following explicit programmed rules. This substitution creates a *displacement effect* that reduces the share of tasks performed by labor. Second, technological advances can create entirely new tasks in which labor has comparative advantage. This generates a *reinstatement effect* that increases labor's share of tasks. The relative strength of these displacement and reinstatement effects determines the net impact of technological change on labor demand.

2.2 Task-Based Approach in Education

How does the task-based framework apply to education? The core insight is that the relationship between skills and tasks operates in reverse. In production, workers apply existing skills to perform tasks that produce output. In education, students perform tasks to develop skills. While learning involves multiple factors, task-based framework focuses on task practice as the primary mechanism of skill formation.

In this framework, a task is a unit of learning activity that requires cognitive effort. Tasks include, for example, reading papers, writing essays, solving problem sets, debugging code, or conducting research. A skill is a student's capability to perform such tasks independently. Skills develop primarily through task practice, with learning shaped by multiple factors including instruction, feedback, and the level of challenge. Research on learning demonstrates that skill acquisition requires learners to actively engage in task practice, not passively consume outputs (Bjork and Bjork 2011; Freeman et al. 2014; Soderstrom and Bjork 2015). The concept of "desirable difficulties" demonstrates that challenges which temporarily impede performance enhance long-term learning. Students develop skills through the cognitive effort of struggling with tasks, making errors, and refining approaches. Instructors develop courses by designing a task bundle (the specific mix of tasks) for students to practice. Different courses require different task bundles.

AI enters this framework as technology that can reshape the tasks students practice through the same two mechanisms identified in production: displacement and reinstatement. AI can displace tasks previously performed by students and reinstate new tasks not previously feasible.

However, task displacement in education has heterogeneous effects on skill formation that require distinguishing between two types. In production, when technology performs part of a task previously done by labor, this represents displacement that reduces labor's share of work. The effect on employment is similar regardless of which task components are affected. In education, displacement can have opposing effects on skill formation. In some cases, AI performing parts of tasks undermines learning by eliminating the cognitive practice through which skills develop. In other cases, AI performing parts of tasks enhances learning by enabling students to focus on core skill formation. I label this latter case of displacement as *task augmentation*, while recognizing that both represent AI performing work previously done by students. Together, task displacement, task augmentation, and task reinstatement determine which tasks students practice and consequently which skills they develop.

Task Displacement: Lost Practice Opportunities

Task displacement occurs when AI performs tasks instead of students. When AI completes a task instead of a student, the student loses the practice opportunity. This holds even when the AI-generated output is superior to what the student would have produced. Task displacement may be most consequential precisely for tasks where AI performance substantially exceeds student performance, since these are often the tasks where students most need practice.

Evidence from recent studies demonstrates that removing task practice through AI substitution leads to measurable skill erosion. Kosmyna et al. (2025) found that students who used LLMs for essay writing over four months consistently exhibited weaker cognitive engagement and reduced brain connectivity patterns compared to those writing without assistance. Effects persisted even when AI tools were subsequently removed. In healthcare, Budzyń et al. (2025) documented that endoscopists' adenoma detection rates in non-AI-assisted colonoscopies declined by 20% after several months of routine AI assistance, suggesting that delegating diagnostic tasks to AI degraded practitioners' independent diagnostic skills.

Task Augmentation: Enhanced Practice Opportunities

Task augmentation occurs when AI enhances student performance of existing tasks without displacing the essential practice. Rather than replacing student work or creating entirely new tasks, AI enables students to practice existing tasks more effectively or at higher levels of complexity. Multiple mechanisms enable this augmentation, with two prominent examples being cognitive load reduction and scaffolding, though others are possible.

Cognitive load theory demonstrates that learners have limited working memory capacity (Sweller 1994). AI can handle peripheral aspects of tasks that impose extraneous cognitive load, freeing mental resources for germane load focused on core learning objectives. Technologies have long served this function in education: calculators freed students from arithmetic computation to focus on mathematical reasoning and problem-solving. Similarly, AI tools that format citations or organize research sources allow students to devote cognitive effort to the core task of synthesis and argumentation, augmenting rather than displacing the practice of analytical writing.

AI can also augment task practice through scaffolding by providing support at the edge of student capabilities without eliminating cognitive challenge (Wood, Bruner, and Ross 1976; Pea 2018). AI tutoring systems can offer just-in-time feedback and hints that guide students through tasks while preserving the problem-solving process. AI coding assistants can explain error messages without rewriting code, augmenting debugging practice without displacing it.

Early experimental evidence suggests LLM-based explanations can enhance learning (Pardos and Bhandari 2024; Kumar et al. 2025). However, because general-purpose AI tools can equally provide scaffolding or generate complete solutions, realizing task augmentation rather than task displacement requires intentional course design to guide students toward complementary rather than substitutive use.

Task Reinstatement: New Practice Opportunities

Task reinstatement occurs when AI enables entirely new task types that create practice opportunities not previously available. Students can now practice formulating effective prompts, critically evaluating AI-generated outputs for accuracy and appropriateness, and iteratively refining AI assistance to achieve learning goals. When students must compare multiple AI-generated solutions, verify AI reasoning, or identify errors in AI outputs, they engage in novel tasks that develop analytical and metacognitive skills distinct from traditional skills. These new tasks represent genuine additions to the course task space rather than substitutions for existing practice.

2.3 Instructors' Responses to the Impact of AI on Tasks and Skills

Task displacement, task augmentation, and task reinstatement determine which tasks students practice and consequently which skills they develop. Importantly, displaced, augmented, and reinstated tasks do not lead to equivalent skills. Task displacement reduces practice in traditional tasks and the skills they build. Task augmentation preserves practice in these same tasks while potentially improving learning. Task reinstatement creates practice in new tasks that develop distinct AI-based skills.

Balancing the three mechanisms is not simply a matter of offsetting losses with gains. Even if task augmentation improves practice and task reinstatement creates as many learning opportunities as task displacement eliminates, the net effect on skill portfolios could be negative if displaced skills are more foundational or more valued in labor markets than augmented or reinstated skills.

Instructors navigate these trade-offs by making decisions about which tasks students will practice in their courses. AI's capability to affect student task practice creates displacement pressures as well as augmentation and reinstatement opportunities. Instructors respond to these pressures and opportunities by making decisions about how to structure AI use in their courses. They can restrict AI for tasks facing displacement, permit or encourage AI use for tasks where augmentation can enhance skills, or introduce new AI-related tasks. Figure 1 visualizes these three mechanisms and instructors' responses to AI capabilities.

This framework assumes that instructors understand that students develop skills by practicing tasks, recognize that AI can affect task practice, and respond to these pressures when making course design decisions. These assumptions are plausible. Instructors structure courses around assignments and activities intended to develop skills, observe student work directly, and have professional incentives to support student learning. While instructors may not think explicitly in terms of task displacement or augmentation, they can recognize when AI threatens essential practice or offers opportunities to enhance learning.

Figure 1. Framework: Effects of AI on Tasks, Skills and Instructor Responses

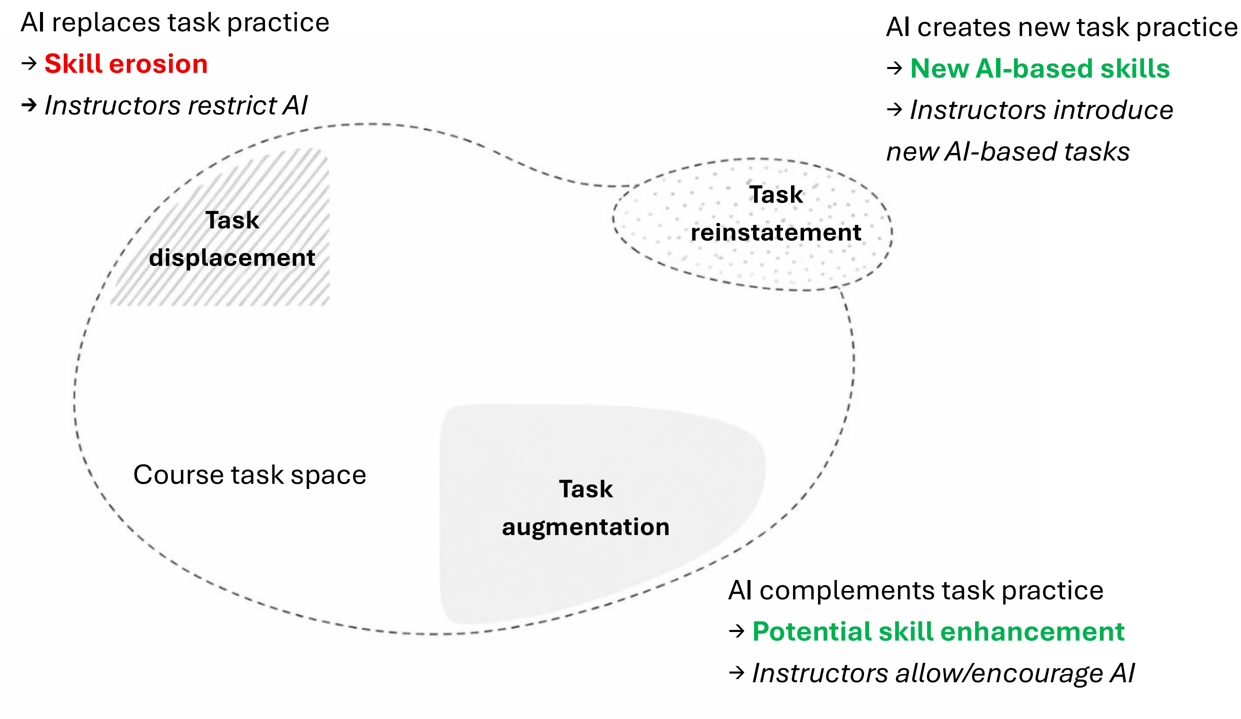

If this framework accurately describes how AI affects learning tasks and how instructors respond, several patterns can be observed:

Prediction 1a: Instructors increasingly regulate AI in response to its capability to affect student task practice. The direction of regulation (whether increasingly restrictive or permissive) reveals whether instructors perceive displacement risks or augmentation/reinstatement opportunities as dominant.

Prediction 1b: Courses requiring more practice in tasks where AI capabilities are strongest (e.g., writing, coding) face greater pressures and opportunities from AI's capability to displace, augment, and reinstate tasks. Consequently, these courses are more likely to regulate AI.

Prediction 2a: Instructors increasingly differentiate AI permissions by task type, restricting tasks where displacement threatens essential practice while permitting tasks where augmentation/reinstatement potential exists.

Prediction 2b: Courses with larger task bundles create more opportunities for displacement, augmentation, and reinstatement to occur within the same course. Consequently, these courses are more likely to differentiate AI regulation by task type (i.e., to specify permissions and restrictions for particular tasks rather than applying a single uniform rule).

Several clarifications about these predictions are important. The predictions concern instructor responses to AI capabilities, not direct effects on student learning. Instructor responses may not determine actual student AI use, as students can access AI tools outside the classroom regardless of instructor decisions. Similarly, instructor beliefs about displacement and augmentation may not accurately reflect AI's actual effects on learning: instructors might restrict tasks that don't actually face displacement risks or permit tasks where augmentation doesn't materialize. Finally, observing patterns consistent with these predictions cannot establish whether instructor responses actually improve student skill formation. The framework predicts how instructors respond to perceived pressures and opportunities, not whether their responses achieve intended learning outcomes.

Nevertheless, understanding instructor responses is valuable for several reasons. Instructors shape the learning environment through course design decisions that directly influence which practice opportunities students encounter. They possess domain expertise about disciplinary skill requirements and observe student work in real time. Patterns in instructor responses therefore provide systematic evidence of professional judgments forming during the critical early period of widespread AI adoption, informing both pedagogical practice and institutional policy development.

3 Existing Evidence and Research Questions

3.1 Existing Evidence and Knowledge Gaps

The four predictions about instructor responses find partial support in existing research, but systematic evidence examining these patterns at scale is limited.

Research relevant to AI regulation (Prediction 1a) exists primarily at the institutional level or faculty surveys. McDonald et al. (2025) found 63% of R1 universities developed AI guidance within one year of ChatGPT's release. However, the American Association of University Professors reports that despite 90% of institutions introducing AI initiatives, only 20% published formal policies (Paris et al. 2025). The translation from institutional guidance to course-level regulation remains uncertain. An American Association of Colleges and Universities' survey of 1,057 U.S. faculty members in November 2025 found that 87% of them reported they created policies regulating the use of AI by their students (Watson and Rainie 2026), while some other institutional surveys found lower levels of AI policy adoption (e.g., Szelenyi 2024). Beyond aggregate adoption rates from surveys, systematic evidence on course-level regulatory prevalence and temporal evolution is limited.

The direction of regulation comes primarily from small-scale studies with limited samples. Ali et al. (2025) found computer science syllabi focused predominantly on prohibition, while other studies documented approaches ranging from punitive to enthusiastic across disciplines (Moore and Lookadoo 2024; Hooper and Lunn 2024; Tong et al. 2025). These studies suggest variation exists but cannot establish temporal trends. The evolution toward greater restrictiveness or permissiveness over time remains unknown. Observed differences may reflect systematic patterns or idiosyncratic choices.

Whether courses requiring more practice in tasks with strong AI capabilities are more likely to regulate AI (Prediction 1b) is particularly understudied. While existing studies document AI policies in specific disciplines or course types, they do not systematically examine whether regulation patterns vary with the intensity of required task practice in domains where AI capabilities are strongest. To the best of my knowledge, research has not tested whether courses requiring extensive writing or coding practice regulate AI at higher rates than courses with different task compositions.

The evidence on task differentiation patterns in AI regulation (Predictions 2a and 2b) is emerging but still limited. Some qualitative case studies document instructors permitting AI for specific activities while restricting it for others (Tsao 2025). Faculty surveys show substantial disagreement about what constitutes legitimate AI use for common tasks like following AI-generated outlines or using AI to edit drafts (Watson and Rainie 2026). This lack of consensus may drive instructors to specify task-level permissions rather than adopt uniform policies. Key gaps remain: whether instructors systematically distinguish tasks where AI might support versus eliminate essential practice, whether differentiation includes new AI-engagement activities, and whether such patterns increased over time.

Across these areas, existing evidence is suggestive but inconclusive. Institutional policies may not reflect course-level practice and faculty surveys capture attitudes or aggregate experimentation rather than specific pedagogical decisions. Small-scale syllabi studies provide more direct evidence but cannot establish whether observed patterns are systematic or idiosyncratic. Temporal evidence is largely cross-sectional, documenting snapshots rather than tracking evolution. Addressing these gaps requires large-scale, longitudinal analysis of course-level policies with explicit attention to temporal evolution of AI regulation and task composition.

3.2 Research Questions

I address these gaps by analyzing the full universe of course syllabi from a large public research university for courses taught between 2021 and 2025. The sample includes over 31,000 unique syllabi and captures courses before widespread AI availability (2021-2022), the first semester after ChatGPT's release (Spring 2023), and subsequent adaptation (2023-2025). For each syllabus, I code whether it includes language around AI use regulation, the restrictiveness and reasoning features of that regulation, and, if available, permitted and restricted tasks, and new AI-based tasks introduced. I also measure overall course task composition by identifying the number of tasks where AI has strong capabilities (writing and coding) among all required course tasks.

I examine four research questions directly corresponding to predictions:

Research Question 1: How do instructors respond to students' use of AI in their courses?

I focus on two dimensions of instructors' response:

1.1. Regulation: Do courses increasingly include AI policies over time and across disciplines?

1.2 Restrictiveness: Does regulation become increasingly restrictive or permissive, and how does this vary across fields? I also examine policy characteristics including the reasoning instructors provide to understand what drives directional changes.

Research Question 2: Does pre-generative AI course task composition predict AI regulation in 2025? Focusing on tasks where current AI capabilities are strongest (writing and coding), I test whether courses that required more practice in these tasks before ChatGPT's release (November 2022) are more likely to include AI policies by Fall 2025.

Research Question 3: Do instructors differentiate AI permissions by task type (restricting certain tasks while permitting others), and what new AI-based tasks emerge? These patterns show whether instructor responses are consistent with task-based logic about displacement, augmentation and reinstatement.

Research Question 4: Does pre-generative AI course task bundle size (overall number of tasks) predict differentiated approach to AI regulation? I test whether courses that required more task practice before ChatGPT's release (November 2022) are more likely to include a differentiated approach to regulating AI use by students in their syllabi (explicitly restricting some tasks while allowing others) in 2025.

This study addresses several limitations in existing research. Analyzing syllabi at the course level provides systematic evidence on actual pedagogical decisions rather than institutional policies or self-reported survey responses. The scale enables examination of temporal evolution, disciplinary variation, and task-level patterns that other studies cannot establish.

Recent advances in computational text analysis have made large-scale study of teaching practices possible. Researchers have increasingly treated course syllabi as rich data sources that document pedagogical decisions and curricular priorities. Studies using natural language processing on hundreds of thousands of syllabi have examined curriculum trends (Hong et al. 2025), alignment between academic content and workforce skills (Javadian Sabet et al. 2024), and incorporation of digital competencies (Yang 2023). This work demonstrates the viability of syllabi as artifacts of pedagogical thinking that can be analyzed computationally at institutional scale.

However, syllabi data have important limitations. Syllabi capture stated policies, not classroom enforcement or student compliance. Some instructors regulate AI through verbal instructions, assignment prompts, or learning management systems rather than syllabi, meaning this analysis likely underestimates total regulation. Syllabi reveal instructor intentions but cannot measure actual student AI use or learning outcomes. Despite these limitations, course syllabi provide the most direct, systematic evidence available on how instructors structure AI use in educational settings during the critical early period of widespread AI adoption.

4 Data

Syllabi Collection

My undergraduate research assistants and I collected course syllabi from a large, selective public research university in Texas enrolling over 50,000 students across all major academic fields. Texas law requires all public institutions to publish undergraduate course syllabi online. The university did not mandate AI policies or require specific language on syllabi, allowing me to observe organic instructor responses to AI. We web-scraped all available syllabi from Spring 2017 through Fall 2025 from the university's public course catalog website.

For each course listing, we extracted metadata including year, semester, department, course number, course title, instructor names, and syllabus PDF link. This process yielded 77,507 initial syllabus records. We used optical character recognition (OCR) to extract text from scanned content when needed. We excluded 63 password-protected or corrupted files (less than 0.1% of the sample).

Deduplication

We observed three types of duplicate syllabi. First, instructors often teach multiple sections of the same course using identical syllabi. We removed these duplicates by keeping only the first occurrence within each combination of year, semester, department, course number, title, and instructor. Second, multiple instructors sometimes teach different sections using shared course templates. We identified these by computing pairwise cosine similarity between TF-IDF representations of syllabus text within each year-semester combination, flagging syllabi pairs exceeding 0.98 similarity as duplicates. Third, some courses appear under multiple departments (cross-listings). We identified these using the same text similarity approach. After deduplication, the dataset contained 65,326 unique syllabi.

Disciplinary Classification

The university organizes courses under 107 departments. We manually classified these departments into nine broad academic fields: Arts, Business, Engineering, Humanities, Life Sciences, Math/CS, Physical Sciences, Social Sciences and Other. This classification enables examining disciplinary variation in AI policy adoption.

Analytical Sample

I included only syllabi for courses offered in Spring or Fall semesters. I excluded Summer sessions because they typically enroll fewer students and include a higher proportion of abbreviated courses with atypical syllabi structures. While we collected syllabi from 2017 through 2025, I limited the analytical sample to 2021-2025. I selected 2021-2022 as the pre-ChatGPT baseline period because AI policies were essentially nonexistent in these years.

My final cross-sectional sample includes 31,692 unique syllabi from Spring 2021 through Fall 2025 spanning all nine academic fields. The temporal range captures four semesters establishing the pre-ChatGPT baseline (Spring 2021-Fall 2022) and five semesters documenting adaptation following ChatGPT's November 2022 launch (Spring 2023-Fall 2025). Tables A1-A3 in the Appendix provide descriptive statistics for syllabi counts by year, semester, discipline, and course level.

To examine within-instructor evolution, I constructed a longitudinal subsample of courses taught by the same instructor or instructor team from 2022-2025. For each course-term, I identified the instructor roster and retained panels observed in at least two years. The resulting subsample contains 4,325 unique course-instructor panels (see Table A4 for longitudinal sample characteristics).

5 Measures and Methods

5.1 AI Regulations and Task Measurement

Computational Text Analysis with LLM Validation

I use large language models (LLMs) to extract and classify AI policies and learning tasks from course syllabi at scale. I began by developing classification schemes through manual coding. Four undergraduate research assistants and I independently coded 500 randomly sampled syllabi snippets mentioning AI, achieving 93% inter-rater reliability, and then reconciling the differences through discussion. I then tested whether OpenAI's GPT-4.1-mini could replicate human classifications. The LLM-based approach achieved 96% accuracy compared to human coding, substantially outperforming traditional machine learning models.[2] Based on this validation, I adopted the LLM-based approach for full-scale analysis.

I implemented an adaptive consensus protocol to improve reliability. For each classification task, I ran the model three times independently with temperature set to 0. If all runs produced identical results, I accepted that classification. If results differed, I ran two additional classifications and selected the majority vote across all five attempts. This approach is similar to having five independent reviewers code each item and taking the majority decision.

AI Policy Extraction and Classification

Course syllabi average 3,300 words and cover many topics beyond AI policies. I used a two-stage process to identify and classify AI-related content. First, I extracted AI-related text from each syllabus, focusing the model on portions discussing AI tools, generative AI, or related

[2] Models based on random forest or support vector machines achieved 75-80% accuracy when predicting human classifications of the permissiveness of AI policy on held-out test data.

policies. Second, I classified the extracted content along multiple dimensions: policy type (fully prohibited, prohibited with exceptions, conditionally permitted, fully permitted, or no policy), whether AI is referenced as a learning tool, whether the policy emphasizes academic integrity concerns, whether attribution of AI use is required, and whether it discusses AI's impact on learning.

Task Extraction and Classification

To examine task differentiation in AI policies (RQ3 and RQ4), I extracted learning tasks mentioned within AI policy text and classified each task's relationship to AI. The classification prompt instructed the model to identify whether the policy explicitly permits AI use (task augmentation), explicitly prohibits AI use (task displacement), or does not clearly specify (unspecified). Tasks were then grouped into categories based on initial hand coding of 100 extracted tasks: ideation/planning, drafting/revising, editing/proofreading, reasoning/problem-solving, study support/synthesis, coding/technical work, and other.

To test whether course task composition predicts AI regulation (RQ2) and to test whether task bundle size predicts task differentiation in AI regulation (RQ4), I extracted learning tasks that students are required to perform from course syllabi. The extraction prompt instructed the model to identify concrete academic activities using verb-object structure (e.g., "write essay," "solve problems," or "read papers"). The prompt explicitly excluded administrative actions, abstract goals, and institutional warnings.

I extracted tasks from two sources with different purposes. For courses taught in 2021-2022 (before ChatGPT's release), I extracted tasks from the same courses to establish pre-AI baseline task requirements, enabling tests of whether baseline task composition predicts subsequent AI regulation. For Fall 2025 syllabi, I extracted tasks from full course syllabi to measure most recent task composition for robustness checks.

Tables A6 and A7 in the Appendix summarize task extraction results. Nearly all syllabi (98%) yielded at least one task. Writing tasks appear in approximately 80% of courses, while coding tasks appear in about 8%. Task composition varies across disciplines in expected ways: Humanities and Social Sciences courses average over four writing tasks, Engineering and Math/CS courses have higher concentrations of coding tasks, and Arts courses show the highest concentration of design tasks. This variation matches expected disciplinary differences, lending credibility to the extraction method.

Advantages and Limitations

This computational approach enabled analyzing over 31,000 syllabi, a scale requiring approximately 3,000 person-hours with manual coding. The adaptive consensus protocol and two-stage extraction process improved reliability and accuracy while ensuring consistency across all syllabi. The approach is replicable using documented prompts and model versions.

The approach has important limitations. While validation showed 96% accuracy and spot checks indicated high accuracy on expanded schemes, some classifications may contain errors that cannot be identified without manual review. Classification schemes require judgment, and some policies fall ambiguously between categories. Classification accuracy depends on prompt design, and model updates may classify identical content differently, complicating exact replication. Finally, I did not conduct full human validation of all expanded classification dimensions, though spot checks showed high accuracy.

5.2 Analytical Approach

RQ1: Adoption and Direction of AI Regulation

To examine how instructors respond to students' use of AI in their courses over time, I estimate two complementary specifications:

$$Y_{ct} = \beta_0 + \beta_1 T_t + \alpha_c + \varepsilon_{ct} \quad \text{(1) Within-course-instructor}$$

$$Y_{it} = \beta_0 + \beta_1 T_t + \gamma X_{it} + \delta_d + \varepsilon_{it} \quad \text{(2) Cross-sectional}$$

where T_t measures time (years from baseline in equation (1); semester count in equation (2)). Controls X include course type (lower division, upper division, graduate), credit hours, and log syllabus word count. Course-instructor fixed effects α_c absorb all time-invariant course characteristics in equation (1), while discipline fixed effects δ_d control for field-specific differences in equation (2). Standard errors are clustered at the department level.

Equation (1) follows the same courses taught by the same instructor or instructor teams over multiple years and represents my main specification. Equation (2) uses unique syllabi from Spring/Fall semesters and serves as a robustness check. To examine disciplinary differences, I also estimate equation (2) on Fall 2025 data only, testing for joint significance of discipline coefficients using a Wald test.

Outcome Variables:

- **Adoption**: Binary indicator for any AI policy.
- **Restrictiveness**: Binary indicator for fully restrictive policies (primary); permissiveness index (for robustness/visualization): -2 (fully prohibited) to +2 (fully permitted), with intermediate values for prohibited with exceptions (-1) and conditionally permitted (+1).
- **Policy characteristics**: binary indicators for references to academic dishonesty, learning outcomes, AI as a learning tool, and attribution requirements in AI policies.

Sample: 2021-2025 for adoption; 2023-2025 for restrictiveness analysis.

RQ2: Task Composition and Regulation

To test whether courses requiring more practice in high AI capability tasks before ChatGPT are more likely to regulate AI, I estimate:

$$Y_i = \beta_0 + \beta_1 S_{i\_2021\text{-}22} + \beta_2 N_{i\_2021\text{-}22} + \gamma X_i + \delta_d + \varepsilon_i \quad (3)$$

where $S_{i\_2021\text{-}22}$ measures the number of required high AI capability tasks (writing and coding) for the course in 2021-2022 (before ChatGPT's November 2022 release) and $N_{i\_2021\text{-}22}$ is the total task count. Controls X include course type, credit hours, and log word count. Discipline fixed effects δ_d are included, and standard errors are clustered at the department level.

Task Variables: I extract learning tasks from full course syllabi using verb-object structure. The extraction identifies concrete academic activities with verbs including write, code, read, present, design, solve, discuss, etc. I extract tasks across all verb categories to measure total task counts ($N_{i\_2021\text{-}22}$), but for analyzing task composition, I construct counts for three specific verb categories: write, code, and present (for a placebo test).

For each verb category, I construct two versions of counts. The main specification uses all task statements, where each mention is counted separately. As a robustness check, I use unique task counts, where the LLM groups semantically similar or identical task statements together and counts each group once. For example, if a syllabus lists "write a 2-page essay," "write a 5-page essay," and "write report," the all-mention write count equals three. However, the unique write count equals two, as the two essay-related statements are grouped into one unique task, while the report remains a separate unique task. I control for log syllabus word count to account for differences in verbosity across syllabi.

I define high AI capability tasks as the sum of write and code task counts, representing domains where current AI capabilities are strongest. I measure task composition from both 2021-2022 syllabi (establishing temporal precedence) and Fall 2025 syllabi (providing a larger sample).

Outcome Variable: Binary indicator for any AI policy in Fall 2025. The key parameter β1 tests whether courses with more high AI capability tasks are more likely to adopt regulation.

Robustness Checks: I conduct five robustness checks. First, I use unique task counts instead of all task mentions. Second, I use binary indicators for whether any high AI capability tasks appear, rather than continuous counts. Third, I estimate separate coefficients for writing and coding tasks rather than combining them. Fourth, I measure task composition from Fall 2025 syllabi instead of 2021-2022 syllabi, providing a larger sample though missing temporal precedence.[3] Fifth, I conduct a placebo test using oral presentation tasks, where AI capabilities are weaker and should not predict regulation if the mechanism operates through AI capability strength.

Sample: Fall 2025 syllabi that have an identical course match in 2021-2022 for temporal precedence tests; all Fall 2025 syllabi for robustness checks using Fall 2025 task composition.

[3] The 2021-2022 measurement may reflect pandemic-era course structures that changed by 2025, which the Fall 2025 specification addresses. The correlation between 2021-2022 and 2025 tasks with high AI capability counts is 0.7, indicating substantial but not perfect stability in course task structures over time.

RQ3: Task Differentiation

To examine whether instructors differentiate AI permissions by task type and what new AI-based tasks emerge (RQ3), I estimate the same specifications as RQ1 (equations (1) and (2)).

Outcome Variables: Binary indicators: Any tasks explicitly permitted; any tasks explicitly prohibited; any new AI-based tasks; any task differentiation (any permitted or prohibited or new AI-based tasks in same policy).

Sample: 2023-2025 (courses with AI policies).

RQ4: Task Bundle Size and Task Differentiation

To test whether courses with larger task bundles before ChatGPT are more likely to differentiate AI regulation by task type (RQ4), I estimate:

$$Y_i = \beta_0 + \beta_1 N_{i\_2021\text{-}22} + \gamma X_i + \delta_d + \varepsilon_i \quad (4)$$

where $N_{i\_2021\text{-}22}$ measures the total number of required tasks from 2021-2022 syllabi (before ChatGPT's November 2022 release). Controls X include course type, credit hours, and log syllabus word count. Discipline fixed effects δ_d are included, and standard errors are clustered at the department level.

Task Bundle Size Variable: I measure task bundle size using the same verb-object extraction approach described in RQ2, counting all required learning tasks across all verb categories. The main specification uses all task statements, where each mention is counted separately. As a robustness check, I use unique task counts, where semantically similar statements are grouped together. I control for log syllabus word count to account for differences in verbosity across syllabi.

Outcome Variable: Binary indicator for whether the AI policy in Fall 2025 differentiates permissions by task type (whether it explicitly restricts some tasks while permitting others, or introduces new AI-based tasks alongside restrictions or permissions). The key parameter β_1 tests whether courses with larger pre-AI task bundles are more likely to adopt differentiated AI regulation. The prediction is that courses with more diverse task requirements create more opportunities for instructors to perceive both displacement risks and augmentation/reinstatement potential within the same course, making task-level differentiation more likely.

Robustness Checks: I conduct two robustness checks. First, I use unique task counts from 2021-2022 syllabi instead of all task mentions. Second, I measure task bundle size from Fall 2025 syllabi instead of 2021-2022 syllabi, providing a larger sample though missing temporal precedence.

Sample: Fall 2025 syllabi with AI policies that have an identical course match in 2021-2022 for temporal precedence tests; all Fall 2025 syllabi with AI policies for robustness checks using Fall 2025 task composition.

6 Results

6.1. Adoption and Direction of AI Regulation (RQ1)

6.1.1 AI Regulation

Figure 2 presents the share of syllabi with explicit policies regulating students' use of AI from 2021 through Fall 2025. Adoption of AI policies increased rapidly since Spring 2023 (the first semester after ChatGPT's release), reaching 55% by Fall 2025.

Disciplinary variation is substantial. Business, Humanities, Social Sciences, and Life Sciences show the highest adoption rates between 58% and 63% by Fall 2025. Engineering shows the lowest adoption at 37%. The Wald test in Appendix Table B2 confirms that these disciplinary differences are statistically significant.

Figure 2. AI Policy Adoption Trends by Discipline

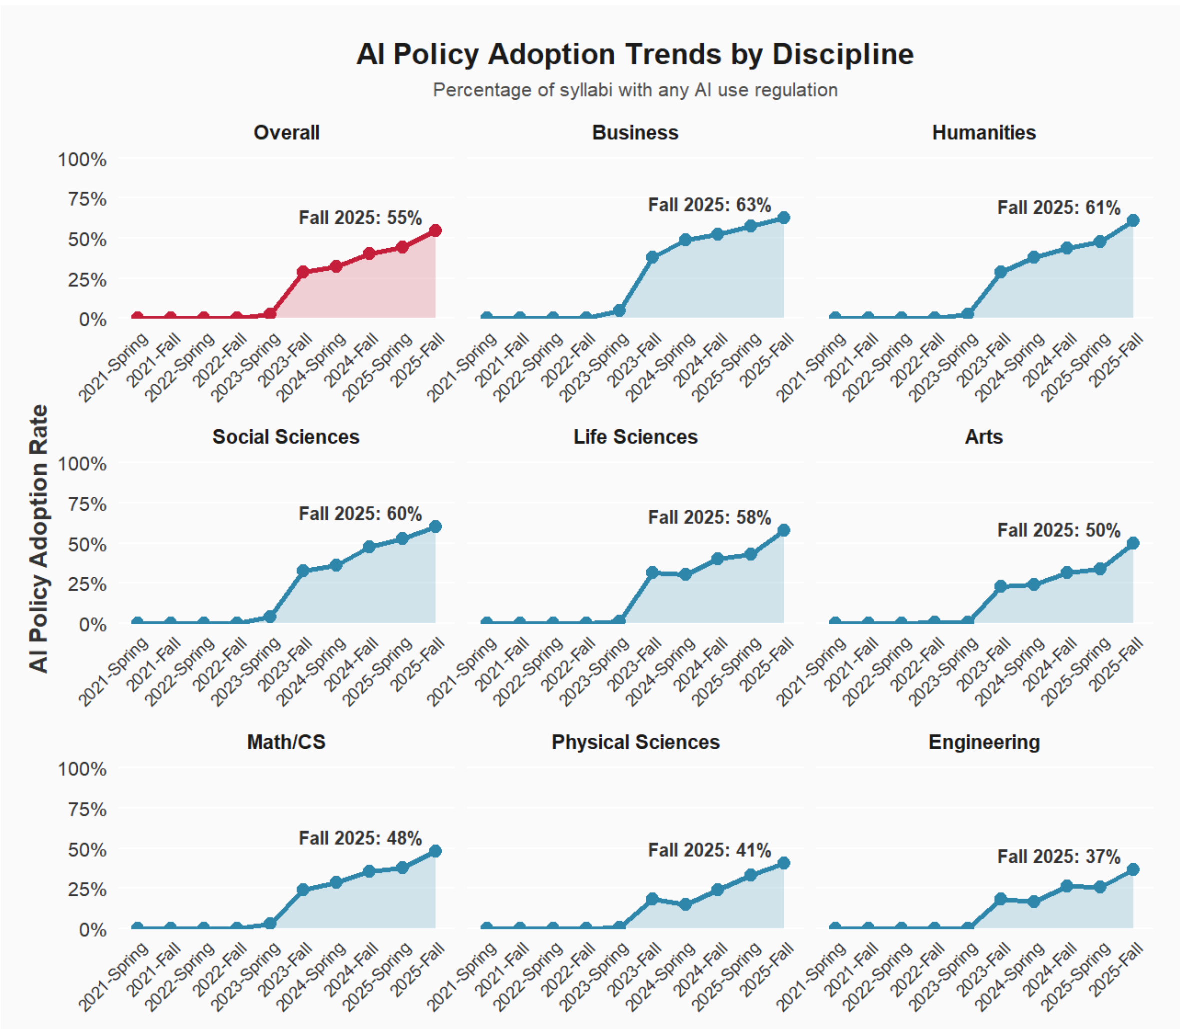

Table 1 presents within-course-instructor estimates tracking the same courses taught by the same instructors over time. AI regulation is increasing by 15.3 percentage points per year. Because this estimate holds course and instructor fixed, it reflects instructors actively updating their syllabi rather than compositional shifts in which courses are offered. The repeated cross-section estimates in Appendix Table B1 show a consistent magnitude of 6.9 percentage points per semester.

Table 1. AI Regulation and Restrictiveness Over Time (Within Courses)

	(1) AI Regulation (all courses)	(2) AI Fully Restricted (courses with policies)	(3) AI Permissiveness Index (courses with policies)
Time (years)	0.1534***	-0.0496***	0.1215***
	(0.0074)	(0.0074)	(0.0196)
Log Word Count	0.0840***	-0.0453	-0.0393
	(0.0204)	(0.0427)	(0.0904)
Observations	12,254	2,948	2,948
Panels (course-instructors)	4,325	1,326	1,326
R-squared	0.265	0.036	0.030

Notes: SE clustered at department level. ***p<0.001, **p<0.01, *p<0.05.
AI Regulation and AI Fully Restricted are binary. AI Permissiveness Index ranges from -2 (fully restricted) to +2 (fully permitted).

6.1.2 Direction of AI Regulation

Figure 3 shows the evolution of policy permissiveness from 2021 to Fall 2025, with the index ranging from -2 (fully prohibited) to +2 (fully permitted). Initial policies in 2023 were restrictive on average. Over subsequent semesters, the overall trend moved toward greater permissiveness, reaching approximately -0.4 points by Fall 2025.

Disciplinary trajectories diverge. Business moved most rapidly toward permissive policies and by Fall 2024 was the only discipline with average permissiveness above zero. Humanities followed the opposite path, adopting the most restrictive policies in 2023 and remaining the most restrictive through Fall 2025. Other disciplines cluster between these extremes, all showing gradual movement toward greater permissiveness over time.

Table 1 quantifies these trends within courses. The proportion of fully restrictive policies decreased by almost 5 percentage points per year among courses with AI policies (Column 2). The continuous permissiveness index increased by 0.12 points per year (Column 3). Both estimates come from within-course-instructor specifications, indicating that individual instructors shifted toward more permissive AI regulation over time.

Appendix Table B2 documents disciplinary differences in Fall 2025. Among courses with AI policies, Business courses are 43 percentage points less likely than Humanities courses to be fully restrictive. Math/CS, Engineering, Arts, and Physical Sciences are 29 to 36 percentage points less likely to be fully restrictive than Humanities. These disciplinary differences are statistically significant.

Figure 3. Evolution of AI Policy Permissiveness by Discipline

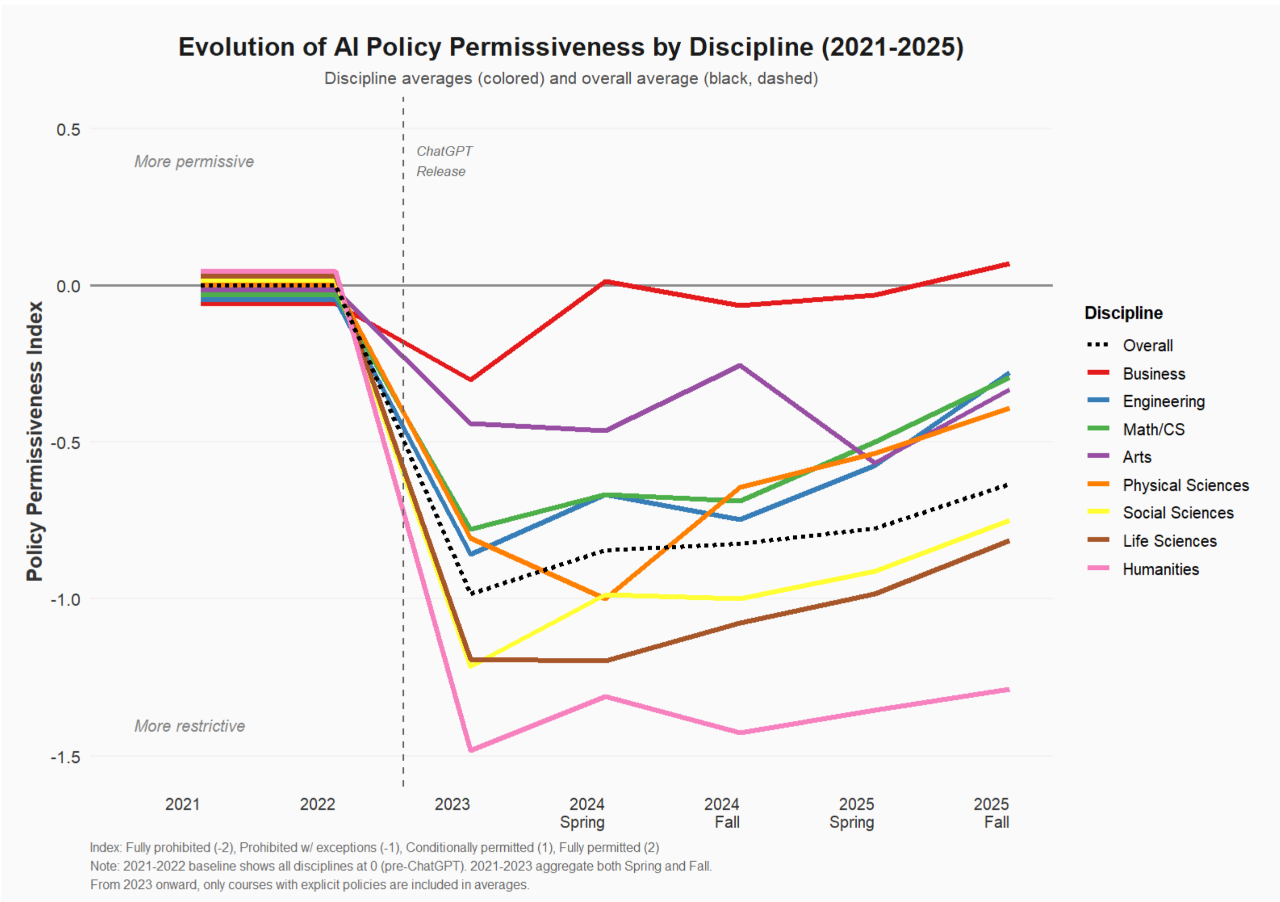

6.1.3 Policy Characteristics

Beyond the presence and direction of AI policies, I examine how instructors frame AI regulation. Figure 4 tracks four features of AI policies from Spring 2023 through Fall 2025: references to academic integrity, references to AI's impact on learning, framing AI as a learning tool, and requirements for attribution of AI use.

Academic integrity concerns dominated early policies but have declined in prevalence, from 63% in Spring 2023 to 49% by Fall 2025. In contrast, references to AI's impact on learning increased from 1% to 29% over the same period. Attribution requirements showed the strongest growth, rising from 16% to 43%. References to AI as a learning tool remain relatively rare at 11% by Fall 2025, though this represents growth from near zero.

The within-course-instructor estimates in Appendix Table B4 confirm that individual instructors are shifting how they frame AI regulation over time. References to learning impact increased by 5.1 percentage points per year, learning tool framing by 2.4 percentage points per year, and attribution requirements by 2.6 percentage points per year. The decline in academic integrity framing is statistically significant in the repeated cross-section (Appendix Table B3) but not within courses, suggesting the aggregate decline reflects newer policies being less likely to emphasize academic integrity rather than instructors removing such language from existing policies.

Figure 4. AI Policy Characteristics

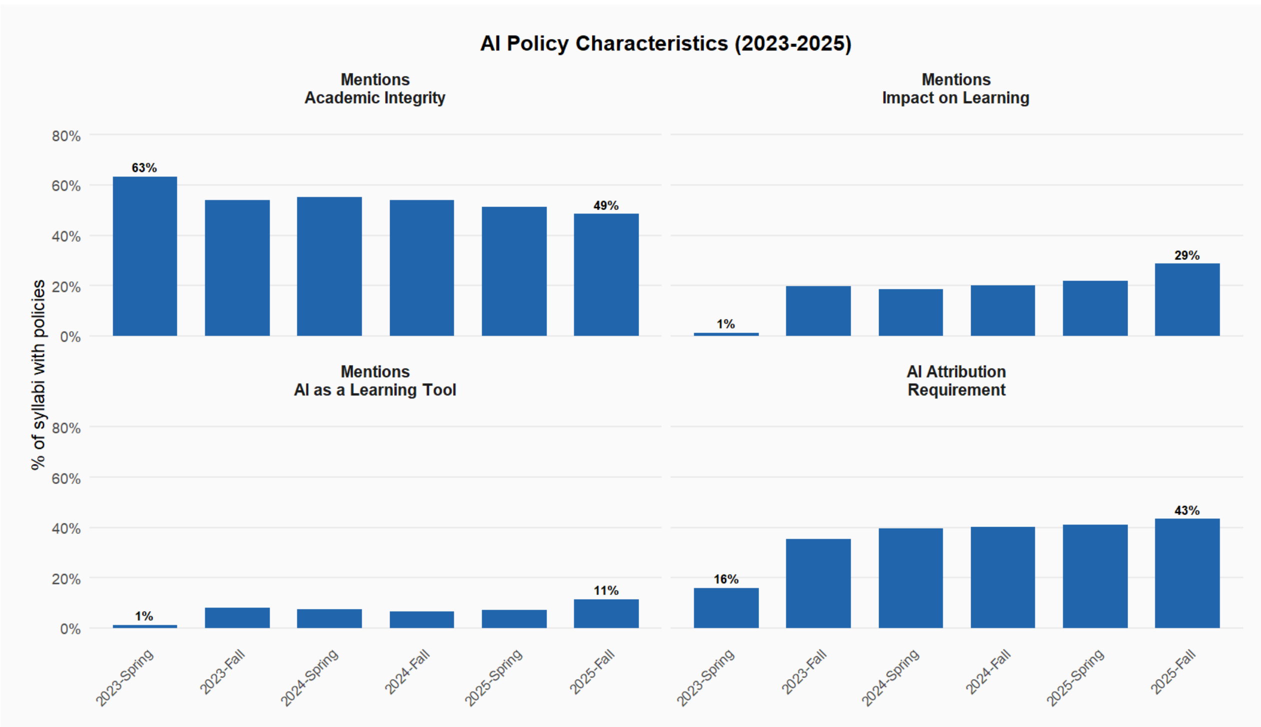

6.2. Task Composition and AI Regulation (RQ2)

The framework predicts that courses requiring more practice in tasks where AI capabilities are strongest face greater pressures from AI and are therefore more likely to regulate. I test this by examining whether the number of writing and coding tasks required before ChatGPT's release predicts AI policy adoption by Fall 2025.

Table 2 presents the main results. Each additional high AI capability task (writing or coding) in 2021-22 is associated with about a 3 percentage point increase in the likelihood of having an AI policy by Fall 2025. This relationship holds after controlling for total task count, course characteristics, and discipline fixed effects (Column 3). The estimate implies that a course requiring five writing or coding tasks is approximately 15 percentage points more likely to adopt AI regulation than a course requiring none, holding other factors constant.

Table 2. Number of High AI Capability Tasks in 2021-22 and AI Regulation Adoption in 2025

	Outcome: AI Regulation Adoption by Fall 2025		
	(1)	**(2)**	**(3)**
Number of high AI capability tasks (writing, coding) in 2021-22	0.0322*** (0.0069)	0.0317*** (0.0067)	0.0264*** (0.0073)
Total tasks	Yes	Yes	Yes
Controls	-	Yes	Yes
Discipline FE	-	-	Yes
Observations (courses)	907	907	907
R-squared	0.039	0.144	0.158

Notes: SE clustered at department level. ***p<0.001, **p<0.01, *p<0.05.
Controls include credit hours, log word count, and course level (lower/upper/graduate)

Appendix Tables C1 through C5 present robustness checks. The relationship holds when using unique task counts rather than all mentions (Table C1) and when using a binary indicator for any high AI capability tasks (Table C2). Estimating separate coefficients for writing and coding tasks shows that both independently predict regulation, with coding tasks showing a somewhat larger coefficient (Table C3). Using task composition measured in 2025 rather than 2021-22 yields similar estimates with a larger sample (Table C4), though this specification cannot establish temporal precedence.

Table C5 presents a placebo test using oral presentation tasks, where AI capabilities are weaker. Oral presentation tasks do not significantly predict AI regulation adoption. This pattern is consistent with the prediction that task composition matters specifically for tasks where AI capabilities create displacement pressures or augmentation/reinstatement opportunities.

6.3. Task Differentiation in AI Regulation (RQ3)

6.3.1 Trends in Task Differentiation

The framework predicts that instructors increasingly differentiate AI permissions by task type, restricting tasks where displacement threatens essential practice while permitting tasks where augmentation/reinstatement potential exists. Table 3 presents within-course-instructor estimates for courses with AI policies. The share of policies that explicitly permit AI for some tasks increased by 4.1 percentage points per year. The share that explicitly restrict AI for some tasks increased by 3.1 percentage points per year. Task differentiation overall, defined as specifying permissions or restrictions for particular tasks rather than applying a uniform rule, increased by 4.5 percentage points per year. All estimates are statistically significant and robust to the repeated cross-section specification in Appendix Table D1.

Table 3. Task Differentiation in AI Regulation Over Time (Within Courses with AI Policy)

	(1) Any Permitted Tasks	(2) Any Restricted Tasks	(3) New AI-based Tasks	(4) Task Differentiation
Time (years)	0.0409***	0.0310**	0.0178**	0.0453***
	(0.0079)	(0.0094)	(0.0059)	(0.0079)
Log Word Count	0.0492	0.0571	0.0515	0.0769
	(0.0408)	(0.0511)	(0.0347)	(0.0476)
Observations	2,948	2,948	2,948	2,948
Panels (course-instructors)	1,326	1,326	1,326	1,326
R-squared	0.023	0.010	0.007	0.025

Notes: SE clustered at department level. ***p<0.001, **p<0.01, *p<0.05.

6.3.2 Which Tasks Do Instructors Permit and Restrict?

Figure 5 shows task-level AI regulation in Fall 2025 among policies that mention specific tasks. The left panel displays the share of policies that restrict versus permit AI use for each task type. The right panel shows how frequently each task is mentioned.

Instructors most commonly restrict AI for drafting and revising (79% restricted) and reasoning and problem-solving (65%). They most commonly permit AI for editing and proofreading (83% permitted), study support and synthesis (80%), and coding and technical work (75%). Ideation and planning is the most contested, with 46% of policies permitting and 54% restricting AI use for this task.

Figure 5. Task-level AI Regulation in Fall 2025

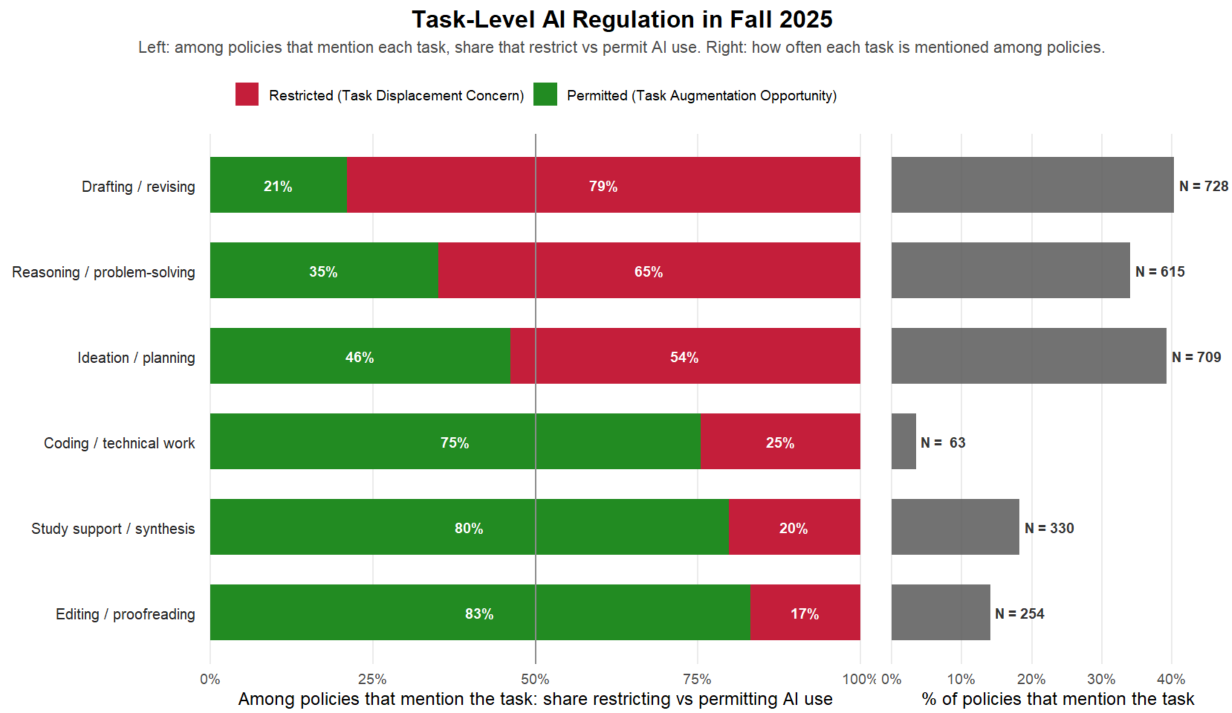

6.3.3 New AI-Based Tasks

Figure 6 presents the share of syllabi mentioning at least one new AI-based task, such as AI-integrated assignments, prompting exercises, or output verification activities. Overall adoption reached 11% by Fall 2025, growing from near zero before ChatGPT's release.

Disciplinary variation is pronounced. Business leads at 27% by Fall 2025, followed by Social Sciences and Arts at 11%. Math/CS, Engineering, and Life Sciences cluster around 9%. Physical Sciences and Humanities show the lowest adoption at 6% and 5%. The within-course-instructor estimates in Table 3 confirm that new AI-based tasks are increasing over time, at 1.8 percentage points per year.

Figure 6. New AI-Based Tasks

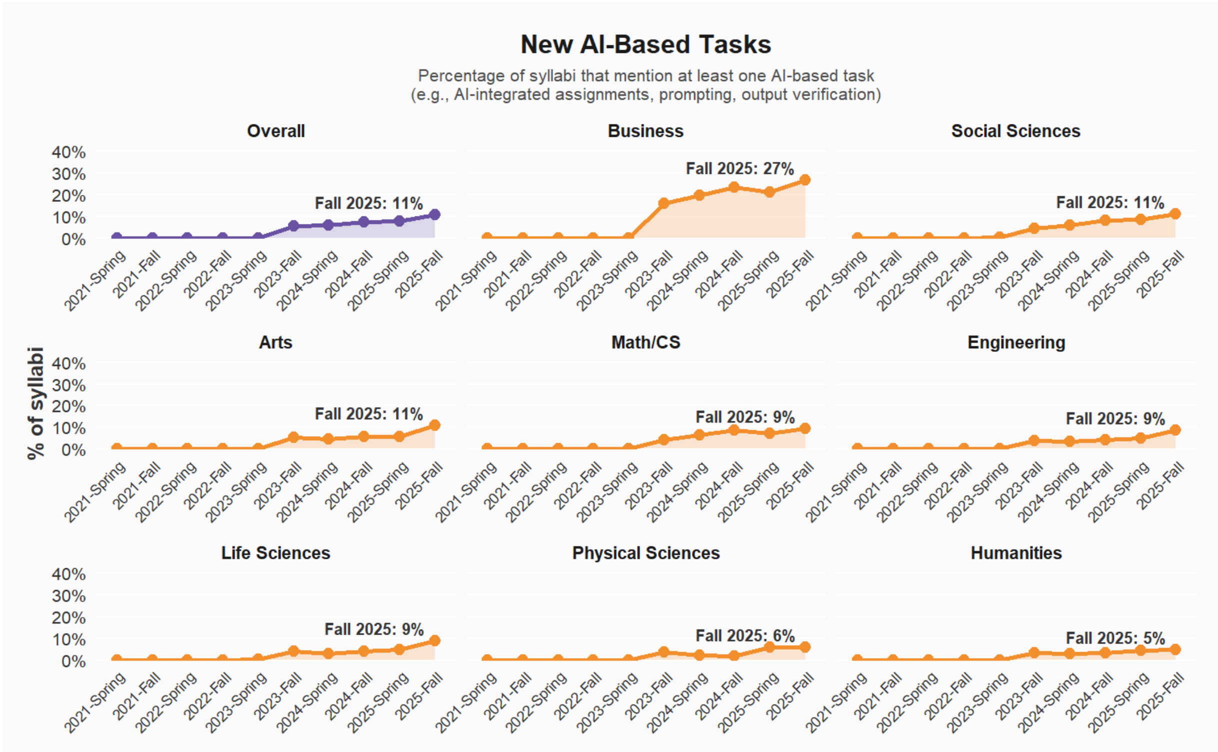

6.4 Task Bundle Size and Task Differentiation in AI Regulation (RQ4)

The framework predicts that courses with larger task bundles create more opportunities for both displacement and augmentation to occur within the same course, making task-level differentiation more likely. I test this by examining whether the total number of tasks required before ChatGPT's release predicts differentiated AI regulation by Fall 2025.

Table 4 presents the main results. Each additional task in 2021-22 is associated with a 1.3 percentage point increase in the likelihood of differentiated AI regulation by Fall 2025. This relationship holds after controlling for course characteristics and discipline fixed effects (Column 3). A course requiring ten tasks is approximately 11 percentage points more likely to differentiate AI permissions than a course requiring two.

Appendix Tables E1 and E2 present robustness checks. The relationship holds when using unique task counts rather than all mentions (Table E1). However, when measuring task bundle size in 2025 rather than 2021-22, the relationship is not statistically significant with full controls (Table E2). This difference suggests that pre-existing course structure drives differentiation rather than contemporaneous task composition, though the 2025 measure may also reflect instructors adjusting task bundles in response to AI.

Table 4. Overall Number of Tasks in 2021-22 and Task Differentiation in AI Regulation in 2025

	Outcome: Task Differentiation in AI Regulation in 2025 (Courses with AI Policies)		
	(1)	**(2)**	**(3)**
Overall number of tasks (task bundle size) in 2021-22	0.0148*** (0.0030)	0.0123*** (0.0030)	0.0134*** (0.0030)
Controls	-	Yes	Yes
Discipline FE	-	-	Yes
Observations (courses)	486	486	486
R-squared	0.037	0.118	0.148

Notes: SE clustered at department level. ***p<0.001, **p<0.01, *p<0.05.
Controls include credit hours, log word count, and course level (lower/upper/graduate)

7 Discussion

This study examined how instructors respond to AI in their courses using over 31,000 syllabi from a large public research university. I document four main findings.

First, AI use regulation grew rapidly from near zero before ChatGPT's release to 55% of courses by Fall 2025. Initial policies were predominantly restrictive, but have shifted toward greater permissiveness over time. The framing of policies has also evolved, with declining emphasis on academic integrity and growing attention to AI's impact on learning.

Second, courses requiring more writing and coding practice before ChatGPT were more likely to adopt AI policies by Fall 2025, consistent with the prediction that courses face greater pressure to respond when AI capabilities overlap with required task practice.

Third, instructors increasingly differentiate AI permissions by task type. They tend to permit AI for tasks like editing and proofreading while restricting it for tasks like drafting and revision. New AI-based tasks such as prompt formulation and output verification are emerging, though adoption remains modest at 11% of courses.

Fourth, courses with larger task bundles before ChatGPT were more likely to adopt differentiated AI policies, suggesting that diverse task requirements create more opportunities for instructors to perceive both displacement risks and augmentation potential within the same course.

Taken together, these findings provide novel empirical evidence on how instructors are responding to AI in educational settings, and the patterns observed are consistent with the task-based theoretical framework developed in this paper.

7.1 Empirical Contributions

A central finding of this study is that instructors' responses to AI are warming over time. Initial policies in 2023 were predominantly restrictive, but the direction has shifted toward greater permissiveness over subsequent semesters, both among instructors updating existing policies and in newly adopted policies. Two pieces of evidence support this pattern: the within-course-instructor increase in permissiveness and the growing incorporation of AI into task practice through new AI-based assignments.

This trajectory may not be apparent in prior research for several reasons. Cross-sectional designs cannot establish temporal trends, and studies conducted in 2023 or early 2024 captured the initial restrictive phase before accommodation emerged (e.g., Ali et al. 2025). Survey-based studies may also reflect sample composition. If respondents disproportionately represent disciplines like Humanities, where restrictiveness remains high, aggregate estimates will overstate overall restrictiveness and understate the shift toward permissiveness occurring in fields like Business (e.g., Watson and Rainie 2026).

These findings suggest a reframing of the narrative around instructors' responses to AI. Rather than viewing AI primarily as a threat to learning, instructors appear to be actively adapting and experimenting with how it can be incorporated into their courses. Instructors appear to be learning about AI capabilities and adjusting their approaches accordingly, moving from blanket restrictions toward more nuanced policies that permit some uses while restricting others.

A second empirical contribution is documenting systematic differentiation in how instructors regulate AI. Instructors do not treat AI uniformly. They distinguish among tasks, restricting AI for some activities while permitting or even encouraging it for others. This task-level variation is obscured in studies that measure AI policy as a single binary or categorical variable.

Disciplinary variation interacts with task composition. Fields with intensive writing requirements, like Humanities, remain more restrictive. Fields oriented toward professional preparation, like Business, have moved more rapidly toward permissive policies and new AI-based tasks. These patterns suggest that aggregate measures of instructors' AI policy mask meaningful heterogeneity that reflects differences in what students are expected to practice and learn.

7.2 Theoretical Contributions

The empirical patterns observed in this study align with predictions derived from the task-based framework. Instructors increasingly regulate AI, with the shift toward permissiveness suggesting that initial displacement concerns are giving way to recognition of augmentation potential (Prediction 1a). Courses requiring more writing and coding practice before ChatGPT were more likely to adopt AI policies, consistent with the prediction that courses face greater pressure when AI capabilities overlap with required tasks (Prediction 1b). Instructors differentiate AI permissions by task type (Prediction 2a), and courses with larger task bundles are more likely to adopt differentiated policies (Prediction 2b).

The task-based framework may help reconcile seemingly contradictory findings in the emerging literature on AI and learning. Some studies document learning gains from AI use (Pardos and Bhandari 2024; Kumar et al. 2025), while others find skill erosion (Kosmyna et al. 2025; Budzyń et al. 2025). The framework suggests these findings are not contradictory but reflect different mechanisms operating on different tasks.

When AI performs tasks instead of students, displacement occurs and skill formation suffers. When AI supports task performance without eliminating essential practice, augmentation occurs and learning may improve. The same AI tool used on the same assignment can produce either effect depending on how students engage with it.

This perspective implies that aggregate measures of AI use are unlikely to yield consistent effects on learning. The effect depends on the task, the nature of AI engagement, and whether that engagement constitutes displacement or augmentation. Instructor efforts to differentiate AI permissions by task type represent an attempt to navigate these distinctions, though whether their judgments accurately reflect underlying learning effects remains an open question.

7.3. Limitations

It's worth restating and addressing the limitations of this study. First, syllabi capture stated policies, not classroom enforcement or student compliance. Instructors may regulate AI through verbal instructions or assignment prompts rather than syllabi, meaning this analysis likely underestimates total regulation. Students may also use AI regardless of stated policies. However, syllabi represent deliberate, documented decisions that instructors make when designing their courses. As such, they reveal instructor preferences about how students should engage with AI, even if actual behavior diverges.

Second, I cannot observe actual student AI use or learning outcomes. The framework predicts how instructors respond to perceived displacement and augmentation pressures, but whether their responses effectively protect skill formation remains untested. Instructor judgments about which tasks face displacement versus augmentation may not align with actual effects on learning. That said, instructors possess domain expertise about disciplinary skill requirements, observe student work directly, and have professional incentives to support learning. Their revealed preferences about AI use provide systematic evidence of professional judgments forming during a critical period of AI adoption.

Third, the data come from a single large public research university in Texas. Instructors' responses may differ at other institution types, in other regions, or in other countries. However, the sample includes the full universe of courses at a large, research-intensive institution spanning all major academic fields, providing substantial variation in disciplines, course levels, and task compositions. Moreover, the findings align with predictions derived from a general theoretical framework, which increases confidence that similar patterns would emerge in other contexts where the same mechanisms operate.

Finally, LLM-based classification may introduce measurement error. While validation showed high accuracy, noise in classification would likely attenuate estimates toward zero, making the reported relationships conservative.

7.4 Implications

7.4.1 For Instructors

The findings suggest that task-level thinking may be a useful framework for AI policy design. Rather than adopting blanket restrictions or permissions, instructors might consider which specific tasks face displacement risks and which offer augmentation potential. The patterns observed here, where instructors restrict AI for core compositional tasks while permitting it for mechanical editing, represent one approach to this differentiation. Whether this particular pattern optimizes learning is an empirical question, but the logic of distinguishing among tasks appears to be gaining traction among instructors.

7.4.2 For Institutions

The substantial disciplinary variation documented here suggests that effective institutional AI policies require flexibility to accommodate disciplinary and task-level differences. Humanities courses and Business courses face different task compositions and different professional preparation goals. A policy well-suited to one context may be poorly suited to another. Institutions might consider providing frameworks and principles while allowing instructors autonomy to navigate task-specific tradeoffs within their domains of expertise.

7.4.3 For Labor Markets

The task-based framework highlights a potential feedback loop between AI in education and AI in production. If AI displaces skill-building tasks in educational settings, students may graduate with weaker capabilities in precisely the areas where AI is most capable. This erosion of human skill would further shift comparative advantage toward AI, potentially accelerating automation in those tasks.

Consider writing as an example. If students practice writing less because AI can draft text, they may enter the labor market with weaker writing skills. Weaker human writing skills make AI-generated writing relatively more attractive to employers, reinforcing the displacement of human labor in writing tasks. The same logic applies to coding and other domains where AI capabilities are strong.

This feedback loop suggests that decisions made in classrooms today may shape the trajectory of automation tomorrow. Instructor responses to AI are not merely pedagogical choices but upstream determinants of skill supply that feed into labor market dynamics. The task-based framework developed by Acemoglu and Restrepo emphasizes reinstatement of new tasks as a counterbalance to displacement in production. A parallel point applies in education: reinstatement of new learning tasks may be essential to maintaining human comparative advantage in an AI-augmented economy.

7.4.4 For Future Research

Several directions for future research emerge from this study. First, studies should measure task-level AI use rather than overall AI use. Aggregate measures obscure the distinction between displacement and augmentation that the framework identifies as central to understanding effects on learning.

Second, research is needed to test whether instructor judgments about displacement and augmentation align with actual learning effects. Instructors may restrict tasks that do not actually face displacement risks or permit tasks where augmentation does not materialize. Experimental designs that vary AI access by task type could address this question.

Third, longitudinal studies tracking skill formation under different AI policies would provide direct evidence on learning effects. The current study documents instructor responses but cannot assess whether those responses achieve their intended goals.

Fourth, replication at other institutions and in other educational contexts would establish generalizability. Community colleges, elite private universities, and international settings may show different patterns of instructors' responses. In addition, future work should examine whether instructors' responses vary by their characteristics and employment status, which may shape both the incentives and constraints instructors face when regulating student AI use.

7.5 Conclusion

This study provides novel, large-scale longitudinal evidence on how instructors are responding to AI in their courses. Instructors are not simply resisting or fearing AI but actively navigating its implications for student learning. The patterns observed, including increasing regulation, a shift toward permissiveness, and differentiation by task type, are consistent with instructors weighing displacement risks against augmentation/reinstatement opportunities. Whether these responses effectively protect skill formation remains to be seen. But understanding how educators are adapting to AI is a necessary first step, both for improving pedagogical practice and for anticipating how changes in skill formation may shape labor markets in an AI-augmented economy.

References

1. Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In Handbook of labor economics (Vol. 4, pp. 1043-1171). Elsevier.
2. Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. Journal of economic perspectives, 33(2), 3-30.
3. Ali, A., Collier, A. H., Dewan, U., McDonald, N., & Johri, A. (2025, February). Analysis of generative AI policies in computing course syllabi. In Proceedings of the 56th ACM Technical Symposium on Computer Science Education V. 1 (pp. 18-24).
4. Autor, D. H. (2013). The "task approach" to labor markets: an overview. Journal for Labour Market Research, 46(3), 185-199.
5. Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. The Quarterly journal of economics, 118(4), 1279-1333.
6. Bjork, E. L., & Bjork, R. A. (2011). Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning. Psychology and the real world: Essays illustrating fundamental contributions to society, 2(59-68), 56-64.
7. Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. The Quarterly Journal of Economics, 140(2), 889-942.
8. Budzyń, K., Romańczyk, M., Kitala, D., Kołodziej, P., Bugajski, M., Adami, H. O., ... & Mori, Y. (2025). Endoscopist deskilling risk after exposure to artificial intelligence in colonoscopy: a multicentre, observational study. The Lancet Gastroenterology & Hepatology, 10(10), 896-903.
9. Freeman, S., Eddy, S. L., McDonough, M., Smith, M. K., Okoroafor, N., Jordt, H., & Wenderoth, M. P. (2014). Active learning increases student performance in science, engineering, and mathematics. Proceedings of the national academy of sciences, 111(23), 8410-8415.
10. Hong, Y., Kim, B., Jeon, J., & Kim, L. (2025). Has higher education become more interdisciplinary? a longitudinal analysis of syllabi using natural language processing. Humanities and Social Sciences Communications, 12(1), 1841.
11. Hooper, K., & Lunn, S. J. (2024, October). Values in Education: Exploration of Artificial Intelligence Ethics Syllabi Using Natural Language Processing Analyses. In 2024 IEEE Frontiers in Education Conference (FIE) (pp. 1-8). IEEE.
12. Javadian Sabet, A., Bana, S. H., Yu, R., & Frank, M. R. (2024). Course-Skill Atlas: A national longitudinal dataset of skills taught in US higher education curricula. Scientific Data, 11(1), 1086.
13. Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X. H., Beresnitzky, A. V., ... & Maes, P. (2025). Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task. arXiv preprint arXiv:2506.08872, 4.

14. Kumar, H., Rothschild, D. M., Goldstein, D. G., & Hofman, J. M. (2025, July). Math Education With Large Language Models: Peril or Promise?. In International Conference on Artificial Intelligence in Education (pp. 60-75). Cham: Springer Nature Switzerland.
15. Love, S. J., & Blankstein, M. (2024, August 22). US Instructor Survey 2024: Findings from a National Survey.
16. McDonald, N., Johri, A., Ali, A., & Collier, A. H. (2025). Generative artificial intelligence in higher education: Evidence from an analysis of institutional policies and guidelines. Computers in Human Behavior: Artificial Humans, 3, 100121.
17. Moore, S., & Lookadoo, K. (2024). Communicating clear guidance: Advice for generative AI policy development in higher education. Business and Professional Communication Quarterly, 87(4), 610-629.
18. Pardos, Z. A., & Bhandari, S. (2024). ChatGPT-generated help produces learning gains equivalent to human tutor-authored help on mathematics skills. Plos one, 19(5), e0304013.
19. Paris, B., Conti-Cook, C., Greene, D., Jones, K. M., Justie, B., Kirschenbaum, M., ... & Widder, D. G. (2025). Artificial Intelligence and Academic Professions. Available at SSRN 5376147.
20. Pea, R. D. (2018). The social and technological dimensions of scaffolding and related theoretical concepts for learning, education, and human activity. In Scaffolding (pp. 423-451). Psychology Press.
21. Soderstrom, N. C., & Bjork, R. A. (2015). Learning versus performance: An integrative review. Perspectives on Psychological Science, 10(2), 176-199.
22. Sweller, J. (1994). Cognitive load theory, learning difficulty, and instructional design. Learning and instruction, 4(4), 295-312.
23. Szelenyi, B. (2024). Adopt or Avoid: Faculty Dilemmas and Decisions on Generative AI in Teaching and Research. Unpublished Manuscript. Senate Information Technology Policy Committee, Northeastern University.
24. Tong, S. T., DeTone, A., Frederick, A., & Odebiyi, S. (2025). What are we telling our students about AI? An exploratory analysis of university instructors' generative AI syllabi policies. Communication Education, 1-22.
25. Tsao, J. (2025). Trajectories of AI policy in higher education: Interpretations, discourses, and enactments of students and teachers. Computers and Education: Artificial Intelligence, 100496.
26. Watson, C.E. and Rainie, L. (2026). The AI Challenge: How College Faculty Assess the Present and Future of Higher Education in the Age of AI. AAC&U Survey Report.
27. Wood, D., Bruner, J. S., & Ross, G. (1976). The role of tutoring in problem solving. Journal of child psychology and psychiatry, 17(2), 89-100.
28. Yang, T. C. (2023). Application of artificial intelligence techniques in analysis and assessment of digital competence in university courses. Educational Technology & Society, 26(1), 232-243.

Appendix

Table A1. Syllabi Count by Year and Semester (2021-2025)

Semester	N Unique Syllabi
2021 Spring	3,020
2021 Fall	3,107
2022 Spring	3,110
2022 Fall	3,096
2023 Spring	3,178
2023 Fall	3,114
2024 Spring	3,156
2024 Fall	3,220
2025 Spring	3,393
2025 Fall	3,298
Total (2021-2025)	31,692

Table A2. Syllabi Count by Discipline

Discipline	N (Overall)	% (Overall)	N (Fall 2025)	% (Fall 2025)
Social Sciences	7,561	23.9	775	23.5
Humanities	5,702	18.0	548	16.6
Arts	4,074	12.9	455	13.8
Business	3,078	9.7	318	9.6
Engineering	2,953	9.3	295	8.9
Life Sciences	2,592	8.2	251	7.6
Math/CS	2,378	7.5	286	8.7
Physical Sciences	1,868	5.9	202	6.1
Other	1,486	4.7	168	5.1

Table A3. Syllabi Count by Course Type

Course Type	N (Overall)	% (Overall)	N (Fall 2025)	% (Fall 2025)
Lower division	9,895	31.2	1,033	31.3
Upper division	17,487	55.2	1,717	52.1
Graduate	4,288	13.5	547	16.6

Table A4. Longitudinal Sample Characteristics

Metric	Value
Total course-instructor pairs	4,325
Total observations	12,254
Average observations per panel	2.83
Year range	2022-2025
Distribution by years observed:	
2 years	1781 (41.2%)
3 years	1484 (34.3%)
4 years	1060 (24.5%)

Table A5. Task Composition Summary

Variable	2022 All	2022 Unique	2025 All	2025 Unique
Total tasks	9.67 (6.44)	5.04 (2.65)	10.38 (7.23)	5.29 (2.61)
	[N = 907]		[N = 3,296]	
Write tasks	3.08 (3.35)	1.57 (1.36)	3.35 (3.53)	1.69 (1.34)
Code tasks	0.14 (0.56)	0.09 (0.3)	0.13 (0.56)	0.08 (0.28)
Tasks where AI capability is strong (Write + Code)	3.22 (3.35)	1.66 (1.36)	3.49 (3.52)	1.77 (1.33)
Present tasks	0.85 (1.29)	0.51 (0.58)	1.05 (1.42)	0.59 (0.6)

Notes: Table shows mean and SD (in parentheses) for task counts. “All” columns count all task mentions; “Unique” columns count deduplicated tasks.

Table A6. Percentage of Courses with Any Tasks

Variable	2022 Baseline (%)	Fall 2025 (%)
Any tasks (non-zero extraction)	98.3	98.6
Any write tasks	77.9	81.2
Any code tasks	8.3	8.0
Any high AI capability tasks	81.4	83.8
Any present tasks	46.7	54.0

Table A7. Average Tasks by Discipline (Fall 2025)

Discipline	Write	Code	Present	Design
Arts	3.06	0.09	1.44	2.65
Business	2.96	0.12	1.16	0.75
Engineering	1.54	0.46	0.54	0.56
Humanities	4.30	0.01	1.14	0.35
Life Sciences	3.49	0.07	1.08	0.42
Math/CS	1.86	0.44	0.67	0.53
Physical Sciences	1.95	0.16	0.60	0.20
Social Sciences	4.21	0.04	1.03	0.50

Table B1. AI Regulation and Restrictiveness Over Time (Repeated Cross-Section)

	(1) AI Regulation	(2) AI Fully Restricted	(3) AI Permissiveness Index
Time (semester count)	0.0687***	-0.0339***	0.0821***
	(0.0031)	(0.0049)	(0.0126)
Credit Hours	-0.0064	0.0459**	-0.0949*
	(0.0066)	(0.0162)	(0.0479)
Log Word Count	0.0993***	-0.0903***	0.1591**
	(0.0075)	(0.0211)	(0.0541)
Upper Division Course	-0.0001	-0.1151***	0.2962***
(reference: Lower Division)	(0.0110)	(0.0238)	(0.0602)
Graduate Course	-0.0442	-0.1577***	0.4332***
(reference: Lower Division)	(0.0237)	(0.0322)	(0.1032)
Discipline FE	Yes	Yes	Yes
Observations	31,670	6,600	6,600
R-squared	0.277	0.128	0.114

Notes: SE clustered at department level. ***p<0.001, **p<0.01, *p<0.05.
Regulation and Fully Restricted are binary. Permissiveness ranges from -2 (fully restricted) to +2 (fully permitted).

Table B2. Disciplinary Differences in AI Regulation and Restrictiveness (Fall 2025)

	(1) AI Regulation	(2) AI Fully Restricted	(3) AI Permissiveness Index
Arts	-0.1437*	-0.3133***	0.8336**
	(0.0588)	(0.0765)	(0.2584)
Business	-0.0665	-0.4252***	1.1319***
	(0.0603)	(0.0590)	(0.1496)
Engineering	-0.1902***	-0.3410***	0.8763***
	(0.0386)	(0.0792)	(0.1886)
Life Sciences	-0.1007	-0.1641*	0.3003
	(0.0565)	(0.0785)	(0.1956)
Math/CS	-0.0931	-0.3596***	0.8831***
	(0.0992)	(0.0703)	(0.1733)
Other	-0.0123	-0.2594***	0.6472***
	(0.0535)	(0.0520)	(0.1258)
Physical Sciences	-0.1877*	-0.2902***	0.8269***
	(0.0751)	(0.0580)	(0.1589)
Social Sciences	-0.0436	-0.1979**	0.4056**
	(0.0552)	(0.0608)	(0.1349)
Credit Hours	-0.0196	0.0150	-0.0464
	(0.0187)	(0.0206)	(0.0546)
Log Word Count	0.2252***	-0.0950***	0.1526*
	(0.0191)	(0.0241)	(0.0752)
Upper Division Course	0.0109	-0.1278***	0.3568***
(reference: Lower Division)	(0.0239)	(0.0327)	(0.0924)
Graduate Course	-0.0595	-0.1580***	0.4055**
(reference: Lower Division)	(0.0435)	(0.0420)	(0.1320)
Observations	3,297	1,807	1,807
R-squared	0.127	0.122	0.107
Wald test (disciplines)	p = 0.0002	p = 0.0000	p = 0.0000
Reference	Humanities	Humanities	Humanities

Notes: SE clustered at department level. ***p<0.001, **p<0.01, *p<0.05.
Regulation and Fully Restricted are binary. Permissiveness ranges from -2 (fully restricted) to +2 (fully permitted).

Table B3. AI Policy Characteristics (Repeated Cross-Section)

	(1) Academic Integrity	(2) Impact on Learning	(3) AI as a Learning Tool	(4) AI Attribution Requirement
Time (semester count)	-0.0142***	0.0308***	0.0101***	0.0237***
	(0.0040)	(0.0045)	(0.0028)	(0.0048)
Credit Hours	0.0061	-0.0080	0.0016	-0.0536***
	(0.0125)	(0.0117)	(0.0097)	(0.0137)
Log Word Count	0.0964***	0.1150***	0.0117	0.1100***
	(0.0230)	(0.0227)	(0.0131)	(0.0203)
Upper Division Course	-0.0625**	0.0253	0.0102	0.1102***
(reference: Lower Division)	(0.0196)	(0.0187)	(0.0130)	(0.0204)
Graduate Course	-0.0687	0.0288	0.0341	0.1287***
(reference: Lower Division)	(0.0380)	(0.0265)	(0.0212)	(0.0369)
Discipline FE	Yes	Yes	Yes	Yes
Observations	6,600	6,600	6,600	6,600
R-squared	0.070	0.046	0.032	0.085

Notes: SE clustered at department level. ***p<0.001, **p<0.01, *p<0.05.

Table B4. AI Policy Characteristics (Within Course-Instructor FE)

	(1) Academic Integrity	(2) Impact on Learning	(3) AI as a Learning Tool	(4) AI Attribution Requirement
Time (years)	-0.0023	0.0511***	0.0241***	0.0263***
	(0.0096)	(0.0079)	(0.0061)	(0.0062)
Log Word Count	0.1612*	0.0384	0.0713	0.0331
	(0.0676)	(0.0381)	(0.0342)	(0.0272)
Observations	2,948	2,948	2,948	2,948
Panels (course-instructors)	1,326	1,326	1,326	1,326
R-squared	0.008	0.032	0.021	0.013

Notes: SE clustered at department level. ***p<0.001, **p<0.01, *p<0.05.

Table C1. Number of <u>Unique</u> High AI Capability Tasks in 2021-22 and AI Regulation Adoption in 2025

	Outcome: AI Regulation Adoption by Fall 2025		
	(1)	**(2)**	**(3)**
Number of unique high AI capability tasks (writing, coding) in 2021-22	0.0848*** (0.0184)	0.0843*** (0.0174)	0.0721*** (0.0181)
Total unique task count	Yes	Yes	Yes
Controls	-	Yes	Yes
Discipline FE	-	-	Yes
Observations (courses)	907	907	907
R-squared	0.064	0.154	0.167

Notes: SE clustered at department level. ***p<0.001, **p<0.01, *p<0.05.
Controls include credit hours, log word count, and course level (lower/upper/graduate)

Table C2. Number of <u>Any</u> High AI Capability Tasks in 2021-22 and AI Regulation Adoption in 2025

	Outcome: AI Regulation Adoption by Fall 2025		
	(1)	**(2)**	**(3)**
Number of any high AI capability tasks (writing, coding) in 2021-22	0.0848*** (0.0184)	0.0843*** (0.0174)	0.0721*** (0.0181)
Total unique task count	Yes	Yes	Yes
Controls	-	Yes	Yes
Discipline FE	-	-	Yes
Observations (courses)	907	907	907
R-squared	0.064	0.154	0.167

Notes: SE clustered at department level. ***p<0.001, **p<0.01, *p<0.05.
Controls include credit hours, log word count, and course level (lower/upper/graduate)

Table C3. Number of Writing and Coding tasks in 2021-22 and AI Regulation Adoption in 2025

	Outcome: AI Regulation Adoption by Fall 2025		
	(1)	**(2)**	**(3)**
Number of writing tasks in 2021-22	0.0318***	0.0314***	0.0255***
	(0.0070)	(0.0067)	(0.0074)
Number of coding tasks in 2021-22	0.0570**	0.0536**	0.0564**
	(0.0190)	(0.0184)	(0.0209)
Total unique task count	Yes	Yes	Yes
Controls	-	Yes	Yes
Discipline FE	-	-	Yes
Observations (courses)	907	907	907
R-squared	0.040	0.144	0.159

Notes: SE clustered at department level. ***p<0.001, **p<0.01, *p<0.05.
Controls include credit hours, log word count, and course level (lower/upper/graduate)

Table C4. Number of High AI Capability Tasks in 2025 and AI Regulation Adoption in 2025

	Outcome: AI Regulation Adoption by Fall 2025		
	(1)	**(2)**	**(3)**
Number of high AI capability tasks (writing, coding) in 2025	0.0267***	0.0266***	0.0225***
	(0.0049)	(0.0044)	(0.0038)
Total tasks	Yes	Yes	Yes
Controls	-	Yes	Yes
Discipline FE	-	-	Yes
Observations (courses)	3,296	3,296	3,296
R-squared	0.042	0.133	0.142

Notes: SE clustered at department level. ***p<0.001, **p<0.01, *p<0.05.
Controls include credit hours, log word count, and course level (lower/upper/graduate)

Table C5. Placebo Test: Number of Oral Presentation Tasks in 2021-22 and AI Regulation Adoption in 2025

	Outcome: AI Regulation Adoption by Fall 2025		
	(1)	**(2)**	**(3)**
Number of oral presentation tasks in 2021-22	0.0182 (0.0172)	0.0221 (0.0163)	0.0229 (0.0167)
Total tasks	Yes	Yes	Yes
Controls	-	Yes	Yes
Discipline FE	-	-	Yes
Observations (courses)	907	907	907
R-squared	0.017	0.123	0.147

Notes: SE clustered at department level. ***p<0.001, **p<0.01, *p<0.05.
Controls include credit hours, log word count, and course level (lower/upper/graduate)

Table D1. Task Differentiation in AI Regulation Over Time (Repeated Cross-Section; Courses with AI Policy)

	(1) Any Permitted Tasks	(2) Any Restricted Tasks	(3) New AI-based Tasks	(4) Task Differentiation
Time (semester count)	0.0259*** (0.0044)	0.0121* (0.0058)	0.0087* (0.0039)	0.0261*** (0.0052)
Credit Hours	-0.0299* (0.0140)	0.0017 (0.0191)	-0.0395*** (0.0098)	-0.0453** (0.0143)
Log Word Count	0.1074*** (0.0190)	0.0612** (0.0187)	0.0994*** (0.0159)	0.1343*** (0.0204)
Upper Division Course (reference: Lower Division)	0.0573** (0.0214)	-0.0377 (0.0278)	0.0576** (0.0177)	0.0723** (0.0233)
Graduate Course (reference: Lower Division)	0.0718* (0.0320)	-0.0528 (0.0393)	0.0466 (0.0348)	0.0701* (0.0304)
Discipline FE	Yes	Yes	Yes	Yes
Observations	6,600	6,600	6,600	6,600
R-squared	0.092	0.017	0.078	0.110

Notes: SE clustered at department level. ***p<0.001, **p<0.01, *p<0.05.

Table D2. Disciplinary Differences in Task Differentiation in AI Regulation (Fall 2025)

	(1) Any Permitted Tasks	(2) Any Restricted Tasks	(3) New AI-based Tasks	(4) Task Differentiation
Arts	0.0167	-0.0475	0.0955	0.1343
	(0.0551)	(0.0702)	(0.0790)	(0.0844)
Business	0.2903***	-0.1038	0.2521***	0.3814***
	(0.0655)	(0.0729)	(0.0481)	(0.0718)
Engineering	0.1329*	-0.0742	0.1271**	0.2385**
	(0.0559)	(0.0644)	(0.0440)	(0.0844)
Life Sciences	-0.0215	-0.0490	-0.0113	0.0382
	(0.0629)	(0.0822)	(0.0384)	(0.0710)
Math/CS	0.1405*	-0.1158	0.0879	0.2346**
	(0.0674)	(0.0646)	(0.0536)	(0.0729)
Other	0.0949	-0.0198	0.0776*	0.1432**
	(0.0529)	(0.0824)	(0.0326)	(0.0517)
Physical Sciences	0.1391*	-0.0871	0.0369	0.1804**
	(0.0677)	(0.0811)	(0.0545)	(0.0629)
Social Sciences	0.0773	0.0690	0.0512	0.1329*
	(0.0538)	(0.0615)	(0.0412)	(0.0607)
Credit Hours	-0.0065	-0.0266	-0.0385**	-0.0209
	(0.0186)	(0.0206)	(0.0125)	(0.0205)
Log Word Count	0.0904***	0.0661*	0.1298***	0.1183***
	(0.0273)	(0.0270)	(0.0202)	(0.0308)
Upper Division Course	0.0799**	-0.0287	0.0894***	0.0844*
(reference: Lower Division)	(0.0289)	(0.0301)	(0.0241)	(0.0351)
Graduate Course	0.0752	-0.0709	0.0732	0.0758
(reference: Lower Division)	(0.0386)	(0.0429)	(0.0436)	(0.0429)
Discipline FE	Yes	Yes	Yes	Yes
Observations	1,807	1,807	1,807	1,807
R-squared	0.061	0.025	0.089	0.083

Notes: SE clustered at department level. ***p<0.001, **p<0.01, *p<0.05.

Table E1. Overall Number of Unique Tasks in 2021-22 and Task Differentiation in AI Regulation in 2025

	Outcome: Task Differentiation in AI Regulation in 2025 (Courses with AI Policies)		
	(1)	(2)	(3)
Overall number of unique tasks (task bundle size) in 2021-22	0.0360*** (0.0066)	0.0293*** (0.0065)	0.0296*** (0.0064)
Controls	-	Yes	Yes
Discipline FE	-	-	Yes
Observations (courses)	486	486	486
R-squared	0.036	0.118	0.144

Notes: SE clustered at department level. ***p<0.001, **p<0.01, *p<0.05.
Controls include credit hours, log word count, and course level (lower/upper/graduate)

Table E2. Overall Number of Tasks in 2025 and Task Differentiation in AI Regulation in 2025

	Outcome: Task Differentiation in AI Regulation in 2025 (Courses with AI Policies)		
	(1)	(2)	(3)
Overall number of tasks (task bundle size) in 2025	0.0045* (0.0020)	0.0012 (0.0019)	0.0024 (0.0018)
Controls	-	Yes	Yes
Discipline FE	-	-	Yes
Observations (courses)	486	486	486
R-squared	0.004	0.038	0.084

Notes: SE clustered at department level. ***p<0.001, **p<0.01, *p<0.05.
Controls include credit hours, log word count, and course level (lower/upper/graduate)