

DISCUSSION PAPER SERIES

IZA DP No. 18055

**Generative AI in Higher Education:
Evidence from an Elite College**

Zara Contractor
Germán Reyes

AUGUST 2025

DISCUSSION PAPER SERIES

IZA DP No. 18055

Generative AI in Higher Education: Evidence from an Elite College

Zara Contractor

Middlebury College

Germán Reyes

Middlebury College and IZA

AUGUST 2025

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Generative AI in Higher Education: Evidence from an Elite College*

Generative AI is transforming higher education, yet systematic evidence on student adoption remains limited. Using novel survey data from a selective U.S. college, we document over 80 percent of students using AI academically within two years of ChatGPT's release. Adoption varies across disciplines, demographics, and achievement levels, highlighting AI's potential to reshape educational inequalities. Students predominantly use AI for augmenting learning (e.g., explanations, feedback), but also to automate tasks (e.g., essay generation). Positive perceptions of AI's educational benefits strongly predict adoption. Institutional policies can influence usage patterns but risk creating unintended disparate impacts across student groups due to uneven compliance.

JEL Classification: I23, O33, I21, J24, D83

Keywords: Generative AI, higher education, technology adoption, ChatGPT, student learning

Corresponding author:

Germán Reyes
Department of Economics
Middlebury College
014 Warner Hall
Middlebury, VT 05753
USA
E-mail: greyes@middlebury.edu

* We thank Jeff Carpenter, Amy Collier, Christa Deneault, Peter Matthews, Caitlin Myers, Andrea Robbett, and participants at the Cornell Behavioral Economics Group and LACBEE for helpful discussions and comments. Coco Kitai and Gavin Randolph provided excellent research assistance. Financial support from Middlebury College's Office of the Provost is gratefully acknowledged. This study was approved by the Middlebury College Institutional Review Board.

1 Introduction

The rapid advancement of generative artificial intelligence (AI) is transforming higher education at an unprecedented pace. The launch of ChatGPT and similar tools has introduced technologies capable of performing tasks central to academic assessment and learning—writing essays, solving complex problems, and explaining intricate concepts—instantly and at near-zero marginal cost. Yet systematic evidence on the nature and implications of this new technology remains scarce. How widespread is generative AI adoption among students, and what factors drive it? Do students primarily use AI to augment their learning or to automate coursework, potentially harming human capital development? Could disparities in access to premium AI resources amplify existing educational inequalities?

These questions are at the heart of ongoing debates about the role of AI in education, which have relied heavily on anecdotal evidence and speculation rather than rigorous evidence. As a result, universities have implemented policies that vary dramatically, from outright bans to permissive adoption, often without clear evidence on their effectiveness or unintended consequences (Nolan, 2023; Xiao et al., 2023; McDonald et al., 2025).

This paper addresses this gap by systematically examining generative AI adoption at Middlebury College, a highly selective liberal arts college in Vermont. With approximately 2,800 undergraduate students, Middlebury offers 49 majors across the arts, humanities, languages, social sciences, and natural sciences. Over the past decade, Middlebury has consistently ranked among the top 10 liberal arts colleges in U.S. News & World Report, with an average ranking of 8.5 between 2015 and 2025. The college’s selectivity is comparable to that of many elite research universities—Middlebury’s acceptance rate of approximately 10 percent is similar to that of Boston University (11 percent), Georgetown University (13 percent), and the University of Virginia (16 percent).

Our analysis draws on survey data collected from the student population between December 2024 and February 2025. The survey collected detailed information about AI usage, including the frequency and purpose of use, perceived impacts, and responses to institutional policies. To minimize selection bias, we framed the survey broadly as examining technology use and provided incentives for participation, achieving a 22.9 percent response rate with 634 responses representing a broad cross-section of the student population.

We document five main findings. First, generative AI is approaching near-universal adoption at an unprecedented speed. Over 80 percent of students use AI for academic purposes, up from less than 10 percent before Spring 2023. This represents one of the fastest

technology adoption episodes ever documented, dramatically exceeding the 40 percent adoption rate among U.S. workers (Bick et al., 2025) and the 23 percent among all U.S. adults (McClain, 2024). These levels are consistent with international findings of 50–70 percent adoption in university contexts (Nam, 2023; Stöhr et al., 2024; Ravšelj et al., 2025).

Second, AI adoption is markedly unequal across academic disciplines and demographic groups. Field of study is the strongest predictor of adoption, likely reflecting how well AI capabilities align with the academic tasks required across fields. Adoption ranges from 91.1 percent in Natural Sciences majors (including mathematics and computer science) to significantly lower levels in Literature (48.6 percent) and Languages (57.4 percent). Additionally, demographic disparities exist. For example, males adopt AI at higher rates than females (88.7 versus 78.4 percent), a pattern that is consistent with documented AI gender gaps in other contexts (Otis et al., 2024). Most notably, lower-achieving students have a higher adoption rate than their higher-achieving peers (87.1 versus 80.3 percent). As a result, AI could serve as an equalizing force if it enhances student learning—helping struggling students catch up to their peers—but could also widen achievement gaps if AI undermines skill development.

Third, generative AI transforms the students’ learning production function by *both* augmenting student capabilities and automating academic tasks. We classify academic tasks as augmentation when AI enhances human capabilities while maintaining student engagement (e.g., explaining concepts, proofreading) versus automation when AI directly produces outputs with minimal cognitive involvement (e.g., writing essays, creating images), and average usage rates across tasks within each category. We find that 61.2 percent of AI users employ these tool for augmentation purposes, while 41.9 percent use it for automation. Qualitative evidence reinforces these patterns: students describe AI as an “on-demand tutor” for augmentation purposes, particularly valuable when traditional resources like office hours are unavailable, while automation use centers on time savings during periods of overwhelming workload. These self-reported patterns align closely with actual usage data from Claude conversation logs (Handa et al., 2025a).

Fourth, institutional policies can significantly influence AI adoption, though their effectiveness can be undermined by informational frictions. We find that explicit prohibitions dramatically reduce self-reported intended use 39 percentage points (pp). Importantly, the effects of this policy vary substantially across student groups, with females showing larger reductions in usage under prohibition than males, suggesting that one-size-fits-all policies may inadvertently create disparate impacts across groups. Moreover, we also document

substantial informational gaps about institutional AI policies. Many students (19.2 percent) do not understand AI policy rules, few know about institution-provided premium AI resources (10.1 percent) or proper AI citation practices (32.6 percent)—a skill that is necessary for academic integrity.

Fifth, students believe that AI has a positive learning impact and these beliefs strongly predict their own AI usage. Most students believe that AI improves their understanding of course materials (70.2 percent) and learning ability (60.1 percent), though fewer believe it improves grades (41.1 percent). Perhaps unsurprisingly, students who believe that AI improves their academic outcomes are significantly more likely to use these tools. For example, a ten-percentage-point increase in the belief that AI improves learning corresponds to a 4.9 percentage-point increase in adoption, highlighting how perceptions can affect technology diffusion in educational settings.

Taken together, our findings reveal that generative AI has already reshaped students' college experience. The substantial information gaps that we document—regarding both permitted uses and available resources—suggest straightforward opportunities for policy intervention. Yet the observed heterogeneity across disciplines and student groups indicates that uniform policies risk unintended consequences: blanket prohibitions may disproportionately disadvantage students who benefit most from AI augmentation, while unrestricted use may encourage automation practices detrimental to skill development. Thus, effective policy requires distinguishing clearly between AI uses that enhance learning and those that undermine it.

While our findings are specific to the population we study—students at an elite liberal arts college—this setting offers valuable into AI's broader societal impact. The rapid uptake of generative AI in higher education may reflect a generalizable principle regarding the diffusion of generative AI: adoption may be fastest in contexts where AI consolidates fragmented tools into a unified platform. Students have long had access to various tools for the academic tasks for which they now use generative AI.¹ What distinguishes AI is its ability to provide all these services through a single, instantly accessible interface at virtually zero marginal cost. This consolidation likely drives students' perception of enhanced learning efficiency and contributes to the fast adoption. This principle may extend well beyond education: industries with fragmented, specialized tools could experience similar

¹Examples include Chegg and Course Hero for homework help and essay writing; Grammarly for proofreading and grammar checking; Khan Academy, Coursera, and YouTube tutorials for concept explanations; faculty office hours for personalized instruction; SparkNotes and CliffsNotes for text summaries; and Stack Overflow for coding assistance.

trajectories of rapid AI adoption, even if AI itself does not introduce fundamentally new capabilities.

Our findings contribute to a rapidly growing literature examining the adoption and impacts of generative AI. Recent work has focused on AI’s effects on worker productivity (Noy and Zhang, 2023; Dell’Acqua et al., 2023; Peng et al., 2023; Cui et al., 2024; Brynjolfsson et al., 2025) and its potential to transform occupations (Felten et al., 2021, 2023; Eloundou et al., 2024). While several papers document AI adoption in workplace settings (Bick et al., 2025; Hartley et al., 2025; Humlum and Vestergaard, 2025) and at the firm level (McElheran et al., 2023; Bonney et al., 2024; Kharazian, 2025), we examine adoption in higher education—a critical setting where future high-skilled workers develop human capital.

We add to an emerging literature examining generative AI in education. Some studies focus on the learning impacts of generative AI, finding mixed impacts depending on the setting (e.g., Lehmann et al., 2024; Bastani et al., 2025). A complementary strand of the literature documents adoption patterns across diverse educational contexts, including universities in Australia (Kelly et al., 2023), Ghana (Bonsu and Baffour-Koduah, 2023), Norway (Carvajal et al., 2024), Sweden (Stöhr et al., 2024), and multi-country studies (Ravšelj et al., 2025). To the best of our knowledge, this paper provides the first systematic evidence on generative AI adoption in U.S. higher education. Importantly, we focus specifically on how students use AI for academic purposes—unlike studies that measure more general usage including personal lives (e.g., Ravšelj et al., 2025)—and go beyond documenting adoption rates to examine how students integrate AI into their academic workflows, distinguishing between augmentation and automation, the role of student beliefs about AI’s educational impacts, and how institutional policies shape usage patterns.

Our findings also contribute to the literature on technology diffusion. Unlike the S-shaped adoption curves documented for most technologies (Griliches, 1957; Rogers, 1962), we observe extremely rapid adoption of generative AI, with 80 percent uptake within two years of ChatGPT’s release. This pattern contrasts sharply with the adoption trajectories of historical General Purpose Technologies (GPTs) like electricity, which required decades to achieve widespread use (David, 1990; Bresnahan and Trajtenberg, 1995). The rapid diffusion we document aligns with evidence that newer technologies diffuse faster than older ones (Comin and Hobijn, 2010a,b). Two factors likely explain AI’s exceptional adoption speed compared to other GPTs. First, unlike electricity or steam power, generative AI requires minimal physical infrastructure—students access it through existing devices—

and is available at no cost to consumers, eliminating the financial barriers that typically slow technology adoption. Second, as noted above, AI consolidates multiple specialized tools into a single platform, making its benefits immediately apparent without requiring specialized training or organizational restructuring, unlike previous GPTs (Brynjolfsson and Hitt, 2000).

2 Data: Novel Student Survey

2.1 Recruitment and Structure

We conducted the survey from December 2024 to February 2025. All Middlebury College students were contacted via email and received a reminder a few weeks and two months after the initial invitation. To minimize selection bias, the recruitment materials described the survey broadly as a study on students’ use of technology in their academic and personal lives. To incentivize participation, students who completed the survey qualified for entry into a lottery for Amazon gift cards ranging in value from \$50 to \$500. The full survey instrument is provided in Appendix E.

The survey contains three main sections (see Appendix Figure A1 for the survey flow). First, we gather demographic and academic information, including gender, race/ethnicity, type of high school attended (private or public), current academic year, and declared or intended major.² We also collect data on academic inputs and performance through students’ self-reported typical weekly hours spent studying and their first-year GPA.³

Second, we measure students’ experience with generative AI tools. We begin by asking whether students have ever used generative AI tools like ChatGPT or Claude. For those who have, we collect information about their usage patterns, including frequency of use during the academic semester, which specific AI models they use, and whether they pay for AI tools. We also gather data on how students use AI for different academic tasks, including writing assistance, learning support, and coding.

The final part of the survey elicits students’ beliefs about generative AI’s adoption and impacts. We ask students about their perceptions of AI’s effects on their academic

²We asked students to report their primary major, yet some reported two majors in the open-text box. In these cases, we keep the first major listed. Results are similar if we include each major-student pair in the dataset.

³We specifically asked about first-year GPA to obtain a measure of academic ability less influenced by AI tool usage, thus minimizing potential reverse causality or endogeneity between AI adoption and academic performance endogeneity.

experience across multiple dimensions, including learning, grades, time management, and understanding of course material. We then collect information on the role of institutional policies, asking students how different policy environments influence their likelihood of using AI. We also elicit students’ beliefs about AI usage among their peers, including their estimates of the fraction of Middlebury students who use AI for schoolwork and leisure. The survey concludes with two open-ended questions that allow students to share their experiences with AI in academic settings and provide feedback on the college’s AI policies and support services.

2.2 Sample and Summary Statistics

Out of Middlebury’s 2,760 enrolled students, 739 began the survey. We exclude 105 respondents who left the survey before reaching the generative AI usage module, leaving us with an analysis sample of 634 students (22.9 percent of the student body). This response rate is comparable to that of similar surveys (Wu et al., 2022). To make our sample more representative of Middlebury’s student body, we construct poststratification weights based on the distribution of declared majors from administrative records.⁴ Specifically, we weight observations by the ratio of each major’s share in the student population to its share in our survey responses.⁵

Table 1 presents summary statistics for our sample. Column 1 reports unweighted survey averages, column 2 shows averages after applying poststratification weights, and column 3 provides population benchmarks from administrative records where available. Our unweighted sample consists of 44.6 percent male and 50.8 percent females. The racial/ethnic composition includes 61.8 percent white, 15.5 percent Asian, 9.9 percent Hispanic, and 3.6 percent Black students. The majority of students (54.3 percent) attended a public school, while 42.0 percent attended a private high school. Our sample represents 43 different majors across seven fields of study, with 31.1 percent of respondents not yet having declared a major. In our analysis, we group these undeclared students by their intended field of study as reported in the survey.

Comparing our unweighted sample to administrative records reveals notable differences.

⁴We focus on achieving representativeness at the field of study level given substantial evidence that AI adoption varies systematically across academic disciplines and occupations, with usage patterns strongly tied to field-specific tasks (Stöhr et al., 2024; Bick et al., 2025; Humlum and Vestergaard, 2025; Ravšelj et al., 2025).

⁵We normalize the weights to sum to the total number of students (2,760) rather than the total number of declared majors, which differs due to some students having multiple majors.

Our sample overrepresents white students (61.8 versus 53.8 percent) and Asian students (15.5 versus 7.3 percent), while underrepresenting Black students (3.6 versus 5.2 percent) and Hispanic students (9.9 versus 12.4 percent). First-year students are overrepresented (31.1 versus 25.5 percent), while senior students are underrepresented (21.3 versus 28.7 percent). Our weighting procedure partially addresses these discrepancies—particularly in the distribution of academic fields, where the weighted figures closely approximate administrative records—but some demographic differences persist. Despite these differences, our weighted sample provides reasonably close approximations to the college population on most dimensions, particularly for academic characteristics and field distributions.

3 Generative AI Usage Patterns Among Students

3.1 Adoption of Generative AI

Generative AI has achieved a remarkably high adoption rate among Middlebury College students. Figure 1 presents the distribution of AI usage frequency during the academic semester, categorized into four levels: “Rarely” (a few times per semester), “Occasionally” (a few times per month), “Frequently” (a few times per week), and “Very Frequently” (daily or almost daily). Overall, 82.5 percent of students report using generative AI for academic purposes, with substantial variation in usage intensity: 23.5 percent use it rarely, 22.2 percent occasionally, 26.2 percent frequently, and 10.6 percent very frequently.

The adoption rate at Middlebury aligns closely with patterns observed across other higher education institutions. A *BestColleges* survey of U.S. undergraduate and graduate students in fall 2023 found that 56 percent had used AI on assignments or exams (Nam, 2023). Ravšelj et al. (2025) surveyed higher education students globally between late 2023 and early 2024, finding that 71.4 percent had ever used ChatGPT—though this figure encompasses all usage rather than academic use specifically. Similarly, Stöhr et al. (2024) document that 63 percent of Swedish university students had used ChatGPT by spring 2023, with 35.4 percent reporting regular use and 27.6 percent rare use. Among Norwegian university students, Carvajal et al. (2024) find that 68.9 percent use AI tools occasionally or more frequently.

These adoption rates in higher education far exceed those in the general population and workforce. Pew Research finds that only 23 percent of U.S. adults have ever used ChatGPT (McClain, 2024), while Gallup reports that just one-third of U.S. workers have used AI for work (Gallup, 2024). Bick et al. (2025) and Hartley et al. (2025) estimate

that about 40 percent of the U.S. working-age population used generative AI for work as of late 2024 and early 2025. Both studies document substantial heterogeneity across industries, with information services showing the highest adoption rates at 56–62 percent—still considerably below the 80 percent adoption rate at Middlebury. Even among workers in AI-exposed occupations, [Humlum and Vestergaard \(2025\)](#) find adoption rates of only 41 percent.⁶

3.2 Adoption by Student Characteristics and Field of Study

AI adoption varies considerably across demographic groups and academic disciplines (Figure 1 and Appendix Table A1). Males report higher usage rates than females (88.7 versus 78.4 percent). Usage patterns differ markedly by race/ethnicity: Black students (92.3 percent) and Asian students (91.3 percent) exhibit the highest adoption rates, while white students (80.2 percent) and Hispanic students (77.9 percent) report lower usage. Students from private high schools use AI more frequently than those from public schools (84.1 versus 80.4 percent). Notably, students with below-median GPAs report higher usage rates than their higher-achieving peers (87.1 versus 80.3 percent). Adoption varies widely by field of study: Natural Sciences leads with 91.1 percent usage, followed by Social Sciences at 84.6 percent, while Languages (57.4 percent) and Literature (48.6 percent) show substantially lower adoption rates.

To examine how student characteristics jointly relate to AI adoption, we estimate multivariate OLS regressions that include all observed characteristics simultaneously.⁷ Table 2 presents regression estimates using four usage thresholds as outcomes. Each column represents the probability of meeting progressively higher frequency thresholds: any AI use (column 1), at least monthly use (column 2), at least weekly use (column 3), and daily use (column 4).

The regression results confirm the patterns observed in the descriptive statistics. Holding other characteristics constant, males are 10.3 pp more likely than females to use AI (column 1, $p < 0.05$), with this gender gap typically widening at higher usage frequencies

⁶The higher adoption rates in higher education may partly reflect demographic composition. Younger and more educated individuals consistently show greater AI adoption: [McClain \(2024\)](#) find that 43 percent of adults under 30 have used ChatGPT compared to 23 percent overall. [Bick et al. \(2025\)](#) document that workers aged 18–29 are twice as likely to use AI at work as those aged 50–64, and college-educated workers are twice as likely to use AI as those without degrees. Similar age and education gradients appear in [Humlum and Vestergaard \(2025\)](#) and [Liu and Wang \(2024\)](#).

⁷We exclude first-year GPA from these regressions because this variable is unavailable for current first-year students.

(columns 2–4, all $p < 0.01$). Black and Asian students show substantially higher adoption rates than white students, at 11.8 and 10.7 pp respectively (column 1, both $p < 0.01$). Students from public high schools are 3.0 pp less likely to use AI than those from private schools, but this difference is not statistically significant. Field of study emerges as the strongest predictor of adoption. Compared to Natural Sciences majors, students in Literature, Languages, Arts, and Humanities all show lower usage rates, with the differences being statistically significant at high usage frequencies (columns 3 and 4). Social Sciences majors exhibit adoption rates similar in magnitude to Natural Sciences majors across all usage frequency thresholds.

Our findings on heterogeneity in AI adoption align with patterns documented in other settings. The gender gap in AI adoption at Middlebury—10.3 pp higher for males—is consistent with evidence across multiple studies. A meta-analysis of 18 studies by [Otis et al. \(2024\)](#) finds that males are 10–20 percentage points more likely to use generative AI than females. This gender gap in AI adoption appears in all educational studies ([Nam, 2023](#); [Carvajal et al., 2024](#); [Stöhr et al., 2024](#); [Ravšelj et al., 2025](#)). Our finding that students with below-median GPAs are more likely to use AI aligns with [Carvajal et al. \(2024\)](#), who document higher adoption rates among students with lower admission grades. This pattern of greater adoption among lower-achieving students suggests that AI could narrow achievement gaps if it enhances learning and skill development, but could widen these gaps if it undermines the acquisition of fundamental skills.

The differences in adoption rates across academic fields at Middlebury mirror patterns documented in other educational settings. [Stöhr et al. \(2024\)](#) find that technology and engineering students exhibit significantly higher ChatGPT usage compared to students in humanities. Similarly, [Nam \(2023\)](#) report that 62 percent of business majors and 59 percent of STEM majors have used AI tools for coursework, compared to 52 percent of humanities majors. [Ravšelj et al. \(2025\)](#) document comparable disciplinary differences, with applied sciences students showing substantially higher usage rates than arts and humanities students.

These academic differences persist into the workforce. [Bick et al. \(2025\)](#) find stark variation by college major: STEM graduates have the highest AI adoption rates (46.0 percent), followed by Business/Economics graduates (40.0 percent), while Liberal Arts graduates show substantially lower rates (22.4 percent). [Humlum and Vestergaard \(2025\)](#) document similar patterns by occupation, finding that roles requiring strong writing and technical skills—such as marketing specialists and journalists—have the highest adoption

rates. These consistent patterns across educational and professional contexts suggest that field-specific factors, particularly the applicability of AI tools to different types of tasks, shape adoption in systematic ways.

3.3 Timing of Generative AI Adoption

The speed of technology diffusion is a critical determinant of its economic and social impact (David, 1990; Hall and Khan, 2003; Stokey, 2021). In educational contexts, rapid technology adoption can create or exacerbate inequalities between early and late adopters, particularly if the technology confers significant learning advantages (World Bank, 2016). To track the timing of AI adoption among Middlebury students, we asked them when they first began using generative AI for academic purposes, with options ranging from “This semester (Fall 2024)” to “Before Spring 2023” (as a reference, ChatGPT’s public launch was in November 2022).

Students adopted generative AI at an extraordinary pace. Figure 2 shows that the cumulative adoption rate grew dramatically from less than 10 percent before Spring 2023 to slightly above 80 percent by Fall 2024. The pace of adoption has accelerated over time, likely reflecting improvements in AI capabilities. Among current users, 25.7 percent adopted AI in Fall 2024 alone, compared to 19.9 percent in Spring 2024 and 16.3 percent in Fall 2023.

The adoption rate among Middlebury students far exceeds that observed in other populations and represents one of the fastest technology adoption episodes ever documented. For comparison, Bick et al. (2025) show that it took over 20 years for computers to reach an 80 percent adoption rate in the U.S. working-age population, and about 15 years for internet adoption to reach similar levels. Even generative AI adoption in the broader population has been markedly slower: Pew Research found that just 23 percent of U.S. adults had ever used ChatGPT as of February 2024, up from 18 percent in July 2023 (McClain, 2024). The dramatically faster adoption rate we document—reaching over 80 percent in less than two years—suggests academic settings may uniquely accelerate AI diffusion.

To identify early versus late adopters of generative AI, in Table 3 and Appendix Figure A2, we analyze how adoption timing varies across student characteristics. The rate of adoption varies substantially across student characteristics. Male students led adoption, with a 8.9 pp higher probability of using AI before Spring 2023 compared to females (column 1, $p < 0.01$)—a gender gap that persists across all periods (columns 2–5). Black and Asian students also adopted earlier than white students, though these differences reach

statistical significance only in later periods (columns 4–5). Field of study is a strong predictor of adoption timing. For example, students in Languages consistently lagged behind Natural and Social Sciences majors, with significantly lower adoption rates across nearly all time periods.

3.4 Choice of Generative AI Models

Major AI companies operate on a freemium model, where free versions coexist with premium subscriptions that offer higher usage limits and access to more advanced models. This tiered structure creates potential for a new form of educational disparities: if paid versions confer substantial academic advantages, students who cannot afford subscriptions may be systematically disadvantaged. To investigate these potential disparities, we presented respondents with a list of options including both free and paid versions of popular models. We also collected information on monthly subscription expenditures, with response options ranging from no active subscription to spending more than \$40 monthly (standard subscriptions to major models like ChatGPT cost \$20 per month at the time).

OpenAI’s ChatGPT dominates AI usage among Middlebury AI users, with the free version capturing the largest market share. Figure 3 shows that 89.3 percent of AI users rely on the free version of ChatGPT, making it dramatically more popular than any alternative. Google Gemini (13.5 percent) and Microsoft Copilot (7.7 percent) are distant competitors, while other platforms each capture less than 5 percent of users. This dominance mirrors patterns in other academic settings, where Stöhr et al. (2024) find substantially higher familiarity with ChatGPT compared to alternative platforms, and aligns with broader workforce trends documented by Bick et al. (2025).⁸

Despite near-universal AI adoption, only 11.3 percent of AI users pay for any AI service.⁹ This figure is remarkably similar to the 8 percent found in Ravšelj et al. (2025)’s multi-country survey. This low payment rate suggests that for most students, the premium features of paid versions—primarily higher usage limits and access to more advanced models—do not justify the subscription cost. However, payment patterns reveal significant disparities: males and Asian students are substantially more likely to purchase AI subscrip-

⁸ChatGPT’s dominance is less pronounced in the general population. Bick et al. (2025) find that ChatGPT leads with 28.5 percent adoption among U.S. adults, followed by Google Gemini at 16.3 percent—a much smaller gap than we observe at Middlebury.

⁹The number of students who report using paid ChatGPT (12.8 percent) is slightly larger than the number who report paying for any AI service (11.3 percent). This discrepancy may arise because some students access paid versions through shared accounts and therefore do not personally pay for the subscription.

tions (Appendix Table A2 and Appendix Figure A3), potentially reflecting differences in usage intensity across demographic groups.

4 Generative AI and the Production of Learning

4.1 The Use of Generative AI across Academic Tasks

How is generative AI transforming the traditional inputs to student learning? The educational production function includes inputs like time spent studying (e.g., [Stinebrickner and Stinebrickner, 2008](#)), faculty instruction (e.g., [Fairlie et al., 2014](#)), peer interactions (e.g., [Sacerdote, 2001](#)), and academic support services (e.g., [Angrist et al., 2009](#)). AI tools have the potential to complement or substitute for these traditional inputs. For example, using AI to explain concepts might substitute for faculty office hours, while using it for proofreading might reduce time needed for academic support services.

To understand the role of AI in students' learning production function, we collected information about students' AI usage across ten common academic tasks: proofreading, generating ideas, writing essays, editing essay drafts, coding assistance, creating images, explaining concepts, composing emails, summarizing materials, and finding information. For each task, students indicated their frequency of AI use on a five-point scale ranging from never to daily. We supplemented this quantitative data with open-ended responses about how AI influences their academic work process.

Students use generative AI for a wide range of academic tasks, with the highest adoption rates for learning support and text-processing activities. Figure 4, Panel A shows that explaining concepts is the most common use case, with 80.3 percent of AI users using it for this task. Summarizing texts follows as the second most common task (74.0 percent), followed by finding information and generating ideas (63.1 and 61.9 percent). Writing assistance tasks like proofreading and editing essays are also common, used by 54.1 and 47.3 percent of AI users, respectively. Technical applications like coding help are significant (34.4 percent), considering that many academic degrees involve no programming. Notably, while 23.5 percent of AI users report using it for writing essays, this represents a relatively low adoption rate compared to other academic uses, suggesting students may be more hesitant to use AI for primary content creation. The lowest adoption rate is for creating images at 20.4 percent, likely reflecting fewer academic use cases for this capability.

An important limitation of our survey is that it relies on self-reported usage, which may introduce measurement error. For example, students might underreport uses they

perceive as academically inappropriate—such as essay writing—while overreporting those viewed as legitimate learning tools (Ling and Imas, 2025). To assess the validity of our self-reported measures, we compare our findings with Handa et al. (2025a), who analyze actual Claude usage patterns among users with university email addresses. Several caveats apply: their data captures conversation-level interactions rather than student-level usage, and most students in our sample use ChatGPT rather than Claude (Section 3.4). Despite these limitations, the comparison provides a useful benchmark for evaluating our survey responses.

Reassuringly, our results are consistent with Anthropic’s data. Both studies identify explaining concepts and technical problem-solving as primary use cases. Handa et al. (2025a) report that the second largest use case (33.5 percent of conversations) involves “technical explanations or solutions for academic assignments,” while we find that 80.3 percent of AI users use it for explaining concepts—a difference likely attributable to our student-level versus their conversation-level measurement. Similarly, the most common usage category in Claude involves “designing practice questions, editing essays, or summarizing academic material” (39.3 percent of conversations), aligning with our findings that 74.0 percent of AI users use it for summarizing texts and 47.3 percent for editing essays. The disciplinary patterns also converge: Handa et al. (2025a) finds that computer science, natural sciences, and mathematics conversations are overrepresented, which mirrors our finding that Natural Science majors show significantly higher AI adoption rates (Figure 1). Overall, these convergent findings from self-reported survey data and actual usage logs suggest that our results capture genuine patterns of student AI engagement rather than merely reflecting social desirability in responses.

4.2 Automation versus Augmentation

Are students using generative AI primarily to *augment* their learning or to *automate* their coursework? This distinction is crucial for understanding AI’s impacts (Autor and Thompson, 2025)—augmentation may enhance students’ learning processes while maintaining their active engagement and critical thinking, whereas automation produces fully-formed outputs that could be submitted with minimal student input, potentially harming learning.

To examine this empirically, Figure 4, Panel B categorizes the ten measured tasks based on whether they augment or automate academic tasks. We classify tasks as augmentation when they enhance human capabilities (explaining concepts, finding information, proof-reading, and editing drafts) and as automation when they directly produce outputs (writ-

ing essays, creating images, composing emails, summarizing texts, generating ideas, and coding assistance). We then calculate the percentage of AI users who employ each category at various frequencies.¹⁰

Students use generative AI for both augmentation and automation, but with markedly different frequency. While 61.2 percent of AI users report using AI for augmentation tasks, 41.9 percent report using it for automation—thus, there is a substantial 19.3 pp difference. This gap is mainly driven by occasional use (19.7 percent for augmentation versus 12.7 percent for automation) and frequent use (16.3 versus 9.0 percent). Most strikingly, students are more than twice as likely to use AI daily for augmentation (5.0 percent) compared to automation (2.5 percent). The higher frequency of augmentation use also suggests that students find augmentation uses more valuable for their day-to-day academic activities. Our findings align closely with patterns observed in actual AI usage data: [Handa et al. \(2025b\)](#) analyze real conversation logs from Claude and find that 57 percent of workplace AI interactions involve augmentation while 43 percent involve automation.

Qualitative evidence from open-ended responses provides additional insights into students’ motivations for using AI (see Appendix C for additional results and validation of the open-ended measure). Students’ descriptions align closely with the augmentation-automation framework. For augmentation, many characterize AI as an “on-demand tutor,” particularly valuable when traditional resources like office hours are unavailable. Non-native English speakers frequently mention using AI for proofreading to overcome language barriers, while students from technical majors describe using it to debug code and understand error messages. For automation, time savings emerged as the dominant motivation, with 21.7 percent of open-ended responses explicitly mentioning efficiency benefits (Appendix Figure C3). Students describe turning to AI during periods of overwhelming workload or looming deadlines, using it to generate initial drafts or complete routine assignments. That students automate tasks mostly when time-pressed helps explain why automation is less common than augmentation.¹¹

¹⁰We acknowledge that this boundary is not always clear-cut, as usage patterns matter. For instance, coding assistance augments learning when students use it to understand concepts and debug their work, but becomes automation if they merely copy solutions without comprehension. Similarly, summarizing texts could augment learning by helping students identify key points, or automate by replacing their own reading. Our categorization reflects the most likely use case for each task based on how the questions were framed in our survey (Appendix E).

¹¹This is consistent with anecdotal evidence suggesting that students tend to rely on automation when under time pressure, particularly in courses outside their major. See, for example: “[What Happens After A.I. Destroys College Writing?](#)” Hua Hsu, *The New Yorker*, June 30, 2025.

Our finding that students favor augmentation over automation extends beyond Middlebury. In Appendix Figure A4, we re-analyze data from Ravšelj et al. (2025)—a multi-country survey of higher education students—to examine how this balance varies across institutional quality.¹² We classify universities into quintiles based on their Times Higher Education World University Rankings. Overall, students who use AI worldwide show similar adoption rates for augmentation (64.6 percent) and automation (63.1 percent) tasks—a much smaller gap than at Middlebury (Panel A). However, this aggregate pattern masks heterogeneity by institutional quality. Students at top-quintile universities show a modest preference for augmentation over automation, with this gap narrowing monotonically down the institution quality distribution. By the bottom quintile, augmentation and automation usage are virtually identical (Panel B).

4.3 Heterogeneity in Augmentation versus Automation Usage

Understanding whether augmentation and automation patterns vary across student populations may help to design targeted support policies. To assess this, we construct four measures of augmentation and automation usage. First, we create binary indicators for whether students use AI for any augmentation or automation task. Second, we calculate the proportion of tasks in each category for which students employ AI. Third, using the Likert-scale responses, we create intensity measures that capture how frequently students use AI for augmentation and automation purposes. Finally, we compute the difference between augmentation and automation variables, as a measure that directly compares students’ relative preference for augmentation versus automation. Table 4 presents regression estimates using these measures as outcomes.

The balance between augmentation and automation varies substantially by student characteristics. Males show higher adoption of both augmentation and automation tasks, with a similar magnitude in each case—they are 6.0 pp more likely to use augmentation prompts and 5.5 pp more likely to use automation prompts compared to females (columns 1 and 4, $p < 0.05$). Black students have a substantially higher augmentation intensity than white students (0.871 points higher, equivalent to 38.1 percent of the outcome mean, $p < 0.01$), but show no statistically significant differences in automation usage. This results in Black students having the strongest preference for augmentation over automation, with

¹²While Ravšelj et al. (2025) elicit a different set of academic tasks than our survey, we categorize them similarly. The tasks we classify as augmentation include proofreading, translating, study assistance, and research assistance. Automation tasks include academic writing, professional writing, creative writing, brainstorming, summarizing, calculating, coding assistance, and personal assistance.

a 19.5 pp higher share of augmentation tasks relative to automation (column 7, $p < 0.01$). Asian students show high usage of both kinds relative to white students. Differences across field of study effects are sizable: humanities, languages, and literature majors tend to show lower usage of both augmentation and automation tools compared to natural science majors.

5 Institutional Policies and AI Adoption

5.1 The Role of Institutional Policies in Shaping Student Behavior

Institutional policies are central to shaping the adoption and diffusion of new technologies (Acemoglu, 2025). In the context of generative AI in higher education, understanding how policies affect student behavior is essential for guiding evidence-based decision-making. To examine this, we asked students to report their likelihood of using AI under various policy scenarios, ranging from complete prohibition to unrestricted use. This analysis is particularly relevant in light of the ongoing debate about how universities should regulate generative AI use in academic settings (Nolan, 2023), and the widespread variation in institutional policies across colleges (Xiao et al., 2023).

Institutional policies substantially influence students' reported likelihood of using generative AI. Figure 5 shows that when generative AI use is unrestricted, 52.4 percent of students report being likely or extremely likely to use it. This likelihood decreases modestly when policies require citation (40.9 percent) or when no explicit policy exists (42.3 percent). However, explicit prohibition creates a dramatic shift: only 13.4 percent of students report they would be likely or extremely likely to use AI when it is banned, while 72.9 percent say they would be unlikely or extremely unlikely to do so (note that these figures include both current AI users and students who do not use generative AI for academic purposes). These results suggest that institutional policies can significantly influence generative AI usage patterns, though a small fraction of students report they would likely use AI even when explicitly prohibited.

The magnitude of policy effects in our study aligns closely with findings from other contexts. Carvajal et al. (2024) estimate that banning AI reduces usage by 37.2 pp among females and 20.6 pp among males, for an overall drop of about 28.9 pp. In our survey, a ban leads to a 37.8 pp decline in usage, with a larger decrease among females (49.6 pp) than males (40.1 pp). These parallel results underscore how institutional policies can unintentionally produce disparate effects across gender. Notably, these differential

policy effects extend beyond gender. Other demographic and academic characteristics—such as race and field of study—also moderate how students respond to policy restrictions (Appendix Table A3), underscoring that institutional policies can produce non-neutral impacts along multiple dimensions.

5.2 Understanding of Institutional Policies and Resources

Information gaps and inattention can significantly affect technology adoption decisions (Duflo et al., 2011; Hanna et al., 2014). For example, imperfect information about rules, available resources, or proper usage guidelines could lead to underadoption of beneficial technologies or inadvertent policy violations. To test for the existence of information gaps, we examine three dimensions of policy understanding. First, we asked students whether they find AI policies in their current classes clear. Second, we measured awareness of free access to premium AI tools through the college—a resource that could reduce inequality but only if students know it exists. Third, we assessed whether students know how to properly cite AI when required, a mechanical skill necessary for academic integrity. We supplemented these measures with open-ended feedback about the college’s AI policies and support services.

Student understanding of institutional policies is high, though significant gaps remain. Figure 6, Panel A shows that most students (79.1 percent) understand when and where they are allowed to use AI in their classes, but a nontrivial minority (19.2 percent) find AI policies unclear. Moreover, critical knowledge gaps persist elsewhere. Only 10.1 percent know they have free access to Microsoft Copilot through the college (Panel B), and just 32.6 percent understand how to properly cite AI use (Panel C). These gaps vary systematically: females show better policy understanding than males (81.5 versus 75.8 percent), and non-white students demonstrate higher awareness across all three dimensions compared to white students.¹³

The qualitative evidence from open-ended responses reinforces these information frictions and reveals implementation challenges (Appendix D). Students express frustration with vague guidelines, requesting specific examples of acceptable versus unacceptable use cases. Many advocate for formal training, noting that simply knowing policies exist dif-

¹³Similar patterns of limited awareness and inconsistent enforcement of AI policies have been documented in other educational contexts. Stöhr et al. (2024) find that only 19.1 percent of Swedish students report that their teachers or universities have rules or guidelines on responsible AI use, suggesting widespread policy ambiguity. Similarly, Nam (2023) reports that while 58 percent of U.S. students say their school has an AI policy, 28 percent indicate that policies vary by course or professor, potentially creating confusion.

fers from understanding how to effectively integrate AI into their workflow. A particularly striking theme is students’ perception of blanket prohibitions as both ineffective and unfair, as they create a prisoner’s dilemma situation where compliant students are disadvantaged relative to those who secretly violate restrictions. Many responses call for a balanced approach—permitting AI use that supports learning while restricting uses that replace it—though the boundary between the two remains contested.

6 Beliefs About AI’s Educational Impact and Peer Usage

6.1 Student Beliefs of AI’s Impact on Educational Outcomes

Technology adoption decisions are shaped by beliefs about potential returns (Foster and Rosenzweig, 2010). Students’ perceptions of how AI affects their learning may influence whether they adopt these tools and how they integrate them into their academic workflows. To elicit these beliefs, we asked students to evaluate AI’s impact across four dimensions of their academic performance: understanding of course materials, overall learning ability, time management, and course grades. For each dimension, students rated AI’s effect on a five-point scale ranging from “significantly reduces” to “significantly improves.” These subjective assessments provide insight into the perceived value of AI tools from the student perspective, which may differ from their actual effects on learning outcomes.

Students tend to believe that generative AI is beneficial for their academic performance, though the perceived benefits vary across dimensions. Figure 7 shows that the majority of students (70.2 percent) believe that generative AI improves their understanding of course materials, and 60.1 percent report improvement in their ability to grasp concepts, retain information, or learn new skills. Similarly, 59.4 percent report that AI improves their ability to complete assignments on time. Notably, while students believe AI helps their learning and assignment completion, they are less confident about its impact on course grades—41.1 percent believe it improves their grades, while 55.4 percent report no effect and 3.5 percent report negative effects. This pattern suggests that while students perceive generative AI as improving their learning process and workflow—through better understanding, skill development, and timely completion of work—these benefits do not necessarily translate into better course grades.¹⁴

¹⁴Perceived benefits vary across student groups (Appendix Table A4). Black students report the most positive perceptions, being 36.2 pp more likely than white students to believe AI improves learning ability. Male students consistently perceive greater benefits than females across most dimensions (10-15 pp higher).

Beliefs about AI’s benefits strongly predict adoption. In Figure 8, we plot the relationship between the percent of students who use AI (x -axis) against the percentage who believe AI improves a specific outcome (y -axis) for different subgroups of students (e.g., males, white students, public-school students, etc.). Across all four measured academic dimensions, there is a strong and statistically significant relationship between perceiving positive AI effects and AI adoption. The relationship is strongest for beliefs about the ability to improve course grades (Panel D): a 10-pp increase in the belief that AI improves grades is associated with a 5.6 pp increase in AI adoption ($p < 0.01$). Similar positive relationships exist for beliefs about learning ability (4.3 pp, $p < 0.05$), understanding of course materials (4.5 pp, $p < 0.01$), and timely assignment completion (5.3 pp, $p < 0.05$). These findings suggest that student beliefs about AI’s academic benefits—irrespective of the actual benefits—may play a crucial role in shaping adoption decisions.

These positive beliefs align with findings from other contexts. Ravšelj et al. (2025) report that the majority of students in their sample believe that ChatGPT improves their general knowledge (68.8 percent) and specific knowledge (62.7 percent)—remarkably similar to our finding that most students believe AI improves understanding of course materials (68.5 percent) and learning ability (57.6 percent). Similarly, 57.4 percent of students in Ravšelj et al. (2025)’s sample believe that ChatGPT helps meet assignment deadlines, while 59.4 percent in our data report AI improves timely assignment completion. Stöhr et al. (2024) provide complementary evidence from Swedish universities, where 47.7 percent of students believe that AI makes them more effective learners, yet only 17.3 percent believe these tools improve their grades—mirroring our finding that perceived learning benefits from AI use exceed perceived effects on course performance.

6.2 Student Beliefs About Peer Use of Generative AI

Students’ beliefs about their peers’ AI usage may influence their own adoption through multiple channels, including social norms (e.g., Giaccherini et al., 2019), social learning (e.g., Foster and Rosenzweig, 1995; Beaman et al., 2021), peer effects (e.g., Bailey et al., 2022), and competitive pressure to avoid falling behind (e.g., Goehring et al., 2024). To measure these beliefs, we asked students to estimate what fraction of their peers use generative AI for different purposes and would use under different policy environments. Figure 9

Students in humanities, languages, and literature report significantly less optimistic views about AI’s academic benefits compared to natural sciences majors, with literature majors being 26.2 pp less likely to believe AI improves grades.

presents the distribution of these beliefs. Panels A-C show students’ estimates of peer AI usage for schoolwork, leisure, and any purpose respectively. Panels D-F display students’ beliefs about AI usage under three policy environments: when classes have no explicit AI policy, when classes allow AI use, and when classes prohibit AI use.

Students systematically underestimate their peers’ AI usage. On average, they believe 65.2 percent of peers regularly use generative AI for schoolwork (Panel A), while our survey reveals an actual usage rate of 82.5 percent—a 17.3 pp gap. This underestimation appears consistent across educational contexts: [Stöhr et al. \(2024\)](#) find that only 38.7 percent of Swedish students believe AI chatbot use is common among peers, despite 63 percent actual usage. When we examine beliefs about policy-contingent behavior, students estimate that 62.7 percent of peers use AI when no explicit policy exists (Panel D), rising to 72.2 percent under permissive policies (Panel E) and falling to 43.5 percent under prohibition (Panel F). These beliefs align directionally with self-reported intentions, though with notable magnitude differences (Appendix Figure [A5](#)). For instance, while students believe 43.5 percent of their peers use AI in classes that prohibit it, only 28.4 percent report they themselves would be likely to do so.¹⁵

Notably, students’ perceptions of peer AI adoption are closely linked to their own adoption behavior. Appendix Figure [A6](#) shows a strong positive association between beliefs about peer usage and actual usage rates across student groups. This behavior is consistent with several psychological mechanisms, including the “false consensus effect” ([Ross et al., 1977](#)), selection neglect ([Jehiel, 2018](#)), or interpersonal projection bias ([Bushong and Gagnon-Bartsch, 2024](#)). Still, virtually all groups systematically underestimate peer usage—actual usage rates exceed believed usage for every demographic and academic group we examine. This figure also shows that underestimation varies substantially across student characteristics: males use AI at 88.7 percent but believe only 66.6 percent of peers do so (22.1 pp gap), while females show a smaller gap (77.6 percent actual versus 65.1 percent believed, 13.3 pp gap).

Taken together, students’ systematic underestimation of peer AI usage coupled with the strong relationship between beliefs and adoption suggests that misperceptions about social norms may shape technology diffusion. If AI enhances learning, then correcting these misperceptions through information provision could accelerate beneficial adoption;

¹⁵Interestingly, students believe that disallowing generative AI in classes has a lower deterrent impact than suggested by self-reported behavior. Based on students’ beliefs, disallowing AI reduces usage by 19.2 pp (from 62.7 percent to 43.5 percent), whereas based on self-reported behavior, it would reduce usage by 39.5 pp (from 68.4 percent to 28.9 percent).

conversely, if AI undermines skill development, then students’ underestimation of peer usage may serve as an unintentional safeguard against harmful overadoption.

7 Discussion

This paper presents the first systematic evidence on generative AI adoption at a highly selective U.S. college. Using novel survey data, we document exceptionally rapid and widespread adoption, substantial shifts in the educational production function through augmentation and automation, and the significant roles of students’ beliefs and institutional policies in shaping AI use.

Our results offer three implications for institutional policy and ongoing debates about AI in education. First, we identify low-cost opportunities to improve institutional policy effectiveness through targeted information provision. The significant gaps we document in students’ understanding of institutional AI policies, citation practices, and available AI resources suggest that simple interventions—such as clear guidelines, illustrative examples of acceptable uses, and AI literacy programs—can reduce unintentional academic integrity violations and support beneficial AI integration. Qualitative feedback strongly indicates student demand for more explicit guidance on responsible AI use.

Second, our evidence challenges alarmist narratives that conflate widespread AI adoption with universal academic dishonesty based on anecdotal accounts.¹⁶ Although AI use is indeed near-universal, we find clear evidence that students primarily employ AI as a tool for strategic task management: to enhance learning (augmentation) and selectively automate tasks when facing high time opportunity cost—not solely to circumvent academic effort. This distinction matters: by normalizing academic dishonesty as inevitable and universal, these narratives may shift social norms and encourage students who would otherwise use AI responsibly to engage in prohibited behaviors, believing “everyone else is doing it.”


Third, our findings caution against policy extremes of either blanket prohibition or unrestricted AI use (Merchant, 2024; McDonald et al., 2025). Blanket prohibitions disproportionately harm students who benefit most from AI’s augmentation functions—particularly lower-achieving students—while also creating uneven compliance, placing conscientious students at a disadvantage relative to rule-breakers. Conversely, unrestricted AI use based solely on revealed preference arguments ignores important market failures in educational

¹⁶See, for example: “Everyone Is Cheating Their Way Through College. ChatGPT has unraveled the entire academic project.” John Herrman, *New York Magazine*, May 7, 2025.

settings. Students often hold overly optimistic beliefs about AI’s learning benefits despite mixed empirical evidence, potentially leading to unintended negative learning outcomes. Most concerningly, permissive policies risk creating competitive dynamics where students feel compelled to adopt AI not for its learning benefits but simply to avoid falling behind in an educational “arms race” ([Goehring et al., 2024](#)).

Figures and Tables


Figure 1: The Adoption of Generative AI among Middlebury College Students



Notes: This figure shows the fraction of students who report using AI during the academic semester, categorized by demographic characteristics, high school type, academic cohort, GPA, and field of study. Usage frequency is divided into four levels: “Rarely” (a few times a semester), “Occasionally” (a few times a month), “Frequently” (a few times a week), and “Very Frequently” (daily or almost daily).


The category “All students” provides the baseline usage rate for the full sample. Gender categories are based on self-identification, with non-binary responses excluded due to a small sample size. “Private HS” refers to students who attended private high schools, while “Public HS” includes public institutions. “Cohort” denotes the student’s academic year, ranging from first-year (“Freshman”) to fourth-year and beyond (“Senior”). GPA categories (“GPA > p50” and “GPA < p50”) split students into groups above or below the median first-year GPA, as self-reported on a 4.0 scale. See Appendix B.1 for the classification of majors into fields of study.

Figure 2: The Evolution of Generative AI Adoption among Middlebury College Students



Notes: This figure shows the cumulative percent of students who reported using generative AI tools for academic purposes over time. The data is based on retrospective self-reports collected in our December 2024 survey, where students were asked “When did you first start using any form of Generative AI for academic purposes?” Response options ranged from “Before Spring 2023” to “This semester (Fall 2024).” The x -axis represents academic semesters, while the y -axis represents the cumulative adoption rate. Vertical lines represent 95 percent confidence intervals calculated with heteroskedasticity-robust standard errors clustered at the student level.


Figure 3: Adoption of Generative AI Models Among College Students




Notes: This figure shows the adoption rates of various AI models as of Fall 2024. The horizontal axis shows the percent of students who reported using each tool, and the vertical axis lists the tools in descending order of adoption rates.

Figure 4: Academic Uses of Generative AI

Panel A. Across Common Academic Tasks




Panel B. Across Tasks that Augment versus Automate Student Effort




Notes: This figure shows the percent of AI students who use generative AI for different academic tasks. For each task, usage frequency is divided into four levels: “Rarely” (a few times a semester), “Occasionally” (a few times a month), “Frequently” (a few times a week), and “Very Frequently” (daily or almost daily). The number at the end of each bar represents the total percent of students who use AI for that purpose at any frequency. Tasks are ordered by total usage, from highest to lowest. Results are based on responses to the question: “For academic purposes, which of the following tasks do you typically use generative AI for?” Sample includes all students who reported using AI during the academic semester.

Figure 5: Student Reported Likelihood of Using AI under Different Policies




Notes: This figure shows the percent of students who report different likelihoods of using AI under various policy scenarios. For each policy, responses are categorized on a five-point scale from “Extremely unlikely to use AI” to “Extremely likely to use AI.” The sample includes all survey respondents. The question asked was: “How likely are you to use generative AI in a class with each of the following AI policies?”


Figure 6: Understanding of Generative AI Policies and Resources



Panel B. Awareness of Copilot Access




Panel C. Knowledge of Citation Requirements



Notes: This figure shows students' understanding of institutional AI policies and resources. Panel A displays the percent of students who report understanding AI policies in their classes, those who report having no explicit policy, and those who find policies unclear, broken down by demographic characteristics. Panel B shows the percent of students who are aware of their free access to Microsoft Copilot through Middlebury College. Panel C presents the percent of students who report knowing how to properly cite AI use in their academic work when required. For Panels B and C, horizontal lines represent 95 percent confidence intervals. Sample includes all survey respondents.


Figure 7: Student Beliefs about the Impact of AI on their Academic Performance




Notes: This figure shows the percent of Middlebury students who believe that AI improves, reduces, or has no effect on different aspects of their academic experience. For each outcome, responses are categorized into three groups: “Improves” combines “significantly improves” and “somewhat improves” responses, “Reduces” combines “significantly reduces” and “somewhat reduces” responses, and “No effect” represents neutral responses. Sample includes all students who report using AI during the academic semester. “Don’t know” responses are excluded.

Figure 8: Relationship Between AI Adoption and Beliefs About AI's Academic Benefits


Panel A. Learning ability (e.g., ability to grasp concepts, learn new skills, etc.)




Panel B. Understanding of course materials



Panel C. Ability to complete assignments on time




Panel D. Course grades



Notes: This figure shows the relationship between AI adoption rates and beliefs that AI improves various academic outcomes across different student groups. Each panel plots the percent of students who use AI (x -axis) against the percentage who believe AI improves a specific outcome (y -axis). Points represent different student groups categorized by demographics (circles), academic characteristics (triangles), and field of study (squares). The dashed line shows the linear fit across all groups. Groups with fewer than 10 students are excluded.

Figure 9: Student Beliefs about Generative AI Usage at Middlebury College



Notes: This figure shows the distribution of students' beliefs about generative AI usage among their peers at Middlebury College. Panels A-C display students' estimates of the percent of their peers who regularly use AI for schoolwork, leisure activities, and any purpose, respectively. Panels D-F show students' beliefs about AI usage in classes with different AI policies: those without an explicit policy (Panel D), those that allow AI use (Panel E), and those that prohibit AI use (Panel F). Each panel shows a histogram with bins of width ten percentage points (e.g., responses between 1-10 fall in the 10 bin, 11-20 in the 20 bin, etc.). The red dashed line indicates the mean response. Sample excludes respondents with missing values or who selected the default response for all six categories (which equals zero).

Table 1: Summary Statistics of Survey Participants

	Survey Sample		Admin records
	Unweighted (1)	Weighted (2)	(3)
Panel A. Demographics			
Male	0.446	0.433	0.463
Female	0.508	0.516	0.533
White	0.618	0.603	0.538
Black	0.036	0.033	0.052
Hispanic	0.099	0.106	0.124
Asian	0.155	0.162	0.073
Private high school	0.420	0.399	–
Public high school	0.543	0.556	–
Panel B. Academic Characteristics			
GPA	3.740	3.736	3.670
Hours spent on academics per week	17.899	17.889	–
Freshman	0.311	0.355	0.255
Sophomore	0.273	0.272	0.257
Junior	0.202	0.179	0.201
Senior	0.213	0.194	0.287
Panel C. Field of Study			
Arts	0.011	0.021	0.023
Humanities	0.052	0.068	0.073
Languages	0.021	0.020	0.025
Literature	0.035	0.025	0.025
Natural Sciences	0.218	0.249	0.244
Social Sciences	0.353	0.247	0.243
Has not declared major	0.311	0.371	0.364
N (# degrees)	43	43	49
N (# students)	634	2,760	2,760

Notes: This table presents summary statistics from our survey of college students. Panel A reports demographic characteristics, including the proportion of participants identifying as male, female, white, Black, Hispanic, Asian, or who attended a private or public high school. Panel B provides academic characteristics, such as GPA (only available for non-freshmen), average weekly hours spent on academics, and academic year distribution (Freshman, Sophomore, Junior, and Senior). Note that in column 1–2, GPA refers to self-reported first-year GPA while in column 3 it is the overall GPA during Spring 2024. Panel C summarizes the distribution of participants across different fields of study. Major groups are mutually exclusive.

Table 2: Student Characteristics Associated with Frequency of Generative AI Use

	Outcome: Uses AI during the semester with frequency of at least...			
	A few times a semester	A few times a month	A few times a week	Daily or almost daily
	(1)	(2)	(3)	(4)
Male	0.103*** (0.033)	0.168*** (0.042)	0.208*** (0.042)	0.096*** (0.029)
Black	0.118*** (0.042)	0.214** (0.099)	0.195* (0.111)	0.022 (0.075)
Latino	-0.012 (0.062)	0.054 (0.073)	-0.019 (0.066)	-0.020 (0.037)
Asian	0.107*** (0.037)	0.164*** (0.055)	0.084 (0.057)	0.054 (0.041)
Public HS	-0.030 (0.033)	-0.050 (0.043)	-0.033 (0.042)	-0.011 (0.026)
Sophomores	0.076* (0.042)	0.065 (0.057)	0.081 (0.055)	0.053 (0.038)
Juniors	0.107** (0.047)	0.144** (0.065)	0.154** (0.068)	0.058 (0.040)
Seniors	0.052 (0.052)	0.081 (0.068)	0.136** (0.066)	0.040 (0.044)
Arts	-0.205 (0.148)	-0.214 (0.187)	-0.009 (0.200)	-0.167*** (0.054)
Humanities	-0.137 (0.083)	-0.291*** (0.097)	-0.303*** (0.080)	-0.106*** (0.032)
Languages	-0.245 (0.156)	-0.238 (0.155)	-0.337*** (0.050)	-0.079** (0.032)
Literature	-0.368*** (0.116)	-0.218* (0.124)	-0.339*** (0.078)	-0.093*** (0.028)
Social Sci.	-0.043 (0.041)	0.014 (0.053)	0.019 (0.055)	0.044 (0.038)
Mean Dep. Var.	0.825	0.589	0.368	0.106
R-squared	0.077	0.089	0.116	0.058
N (Students)	616	616	616	616

Notes: This table assesses the relationship between AI adoption and student characteristics. We estimate:

$$Y_i = \alpha + \beta X_i + \varepsilon_i,$$

where Y_i is a binary indicator of AI usage frequency threshold and X_i is a vector of student characteristics including gender, race/ethnicity, high school type, cohort indicators, and academic division. Students who have not declared their major are classified into fields of study based on their intended major.

Each column uses a different threshold for AI usage frequency, categorized as: “Rarely” (a few times a semester), “Occasionally” (a few times a month), “Frequently” (a few times a week), and “Very Frequently” (daily or almost daily). Column 1 defines usage as any nonzero frequency; column 2 includes at least occasional use; column 3 includes frequent or higher use; and column 4 captures only very frequent use.

The omitted categories are: Natural Sciences for academic division, white students for race/ethnicity, freshmen for cohort, female for gender, and private high school for school type. Heteroskedasticity-robust standard errors clustered at the student level in parentheses. Observations are weighted to adjust for sampling. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Student Characteristics Associated with Timing of Generative AI Adoption

	Outcome: Started using generative AI...				
	Before Spring 2023 (1)	Spring 2023 or before (2)	Fall 2023 or before (3)	Spring 2024 or before (4)	Fall 2024 or before (5)
Male	0.089*** (0.024)	0.156*** (0.035)	0.222*** (0.042)	0.134*** (0.043)	0.085** (0.035)
Black	-0.016 (0.059)	0.031 (0.097)	0.064 (0.110)	0.184* (0.103)	0.142*** (0.042)
Latino	-0.010 (0.036)	0.036 (0.055)	0.036 (0.063)	-0.043 (0.069)	-0.026 (0.063)
Asian	0.041 (0.039)	0.029 (0.049)	0.127** (0.059)	0.144** (0.057)	0.132*** (0.038)
Public HS	-0.014 (0.024)	-0.057 (0.035)	-0.076* (0.041)	-0.010 (0.043)	-0.058* (0.034)
Sophomores	-0.070** (0.028)	-0.083* (0.043)	0.031 (0.053)	0.178*** (0.058)	0.054 (0.046)
Juniors	-0.067* (0.035)	-0.082 (0.055)	0.080 (0.067)	0.255*** (0.063)	0.087* (0.050)
Seniors	-0.077** (0.033)	-0.020 (0.060)	0.076 (0.066)	0.145** (0.068)	0.045 (0.054)
Arts	-0.013 (0.105)	-0.181 (0.131)	0.156 (0.146)	0.054 (0.146)	-0.182 (0.143)
Humanities	-0.011 (0.036)	0.019 (0.077)	-0.039 (0.096)	-0.085 (0.097)	-0.096 (0.085)
Languages	-0.015 (0.019)	-0.116*** (0.042)	-0.286*** (0.052)	-0.312** (0.151)	-0.235 (0.155)
Literature	0.124 (0.082)	0.117 (0.106)	-0.023 (0.115)	-0.098 (0.125)	-0.333*** (0.117)
Social Sci.	0.035 (0.026)	0.032 (0.042)	-0.029 (0.050)	-0.001 (0.053)	-0.018 (0.044)
Mean Dep. Var.	0.089	0.199	0.357	0.551	0.801
R-squared	0.045	0.061	0.088	0.080	0.069
N (Students)	633	633	633	633	633

Notes: This table assesses the relationship between AI adoption and student characteristics. We estimate:

$$Y_i = \alpha + \beta X_i + \varepsilon_i,$$

where Y_i is a binary indicator of AI adoption date and X_i is a vector of student characteristics including gender, race/ethnicity, high school type, cohort indicators, and academic division. Students who have not declared their major are classified into fields of study based on their intended major.

Each column presents results for a different threshold of AI adoption. Column 1 shows the probability of adopting AI before Spring 2023; column 2 by Spring 2023; column 3 by Fall 2023; column 4 by Spring 2024; and column 5 by Fall 2024. The dependent variable in each regression is a binary indicator equal to one if the student had adopted AI by the specified time period.

The omitted categories are: Natural Sciences for academic division, white students for race/ethnicity, freshmen for cohort, female for gender, and private high school for school type. Heteroskedasticity-robust standard errors clustered at the student level in parentheses. Observations are weighted to adjust for sampling. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Student Characteristics Associated with Task Augmentation and Automation

	Augmentation Tasks			Automation Tasks			Difference: Augm. - Autom.	
	Any > 0 (1)	Share > 0 (2)	Intensity (3)	Any > 0 (4)	Share > 0 (5)	Intensity (6)	Share (7)	Intensity (8)
Male	0.060** (0.027)	0.098*** (0.030)	0.337*** (0.083)	0.055** (0.028)	0.091*** (0.024)	0.227*** (0.054)	0.007 (0.026)	0.110* (0.064)
Black	0.041 (0.072)	0.196** (0.082)	0.871*** (0.285)	-0.015 (0.080)	0.001 (0.066)	0.183 (0.157)	0.195*** (0.073)	0.687*** (0.199)
Latino	0.041 (0.040)	0.106** (0.047)	0.378*** (0.130)	0.081** (0.035)	0.082** (0.040)	0.238** (0.097)	0.025 (0.046)	0.140 (0.104)
Asian	0.076** (0.033)	0.068* (0.038)	0.226** (0.106)	0.046 (0.033)	0.051* (0.028)	0.120* (0.063)	0.016 (0.035)	0.107 (0.083)
Public HS	0.005 (0.030)	0.010 (0.031)	0.037 (0.083)	-0.038 (0.028)	-0.019 (0.023)	-0.012 (0.054)	0.028 (0.026)	0.050 (0.061)
Sophomores	-0.009 (0.034)	-0.013 (0.039)	-0.042 (0.107)	0.013 (0.038)	-0.002 (0.032)	-0.012 (0.068)	-0.011 (0.035)	-0.030 (0.088)
Juniors	-0.012 (0.041)	0.020 (0.049)	0.098 (0.129)	0.058 (0.039)	0.059 (0.037)	0.116 (0.084)	-0.039 (0.040)	-0.018 (0.096)
Seniors	-0.050 (0.046)	-0.021 (0.048)	0.090 (0.133)	0.066* (0.037)	0.067* (0.035)	0.227** (0.090)	-0.088** (0.042)	-0.136 (0.098)
Arts	-0.129 (0.157)	-0.188** (0.085)	-0.510*** (0.187)	-0.056 (0.090)	-0.221*** (0.045)	-0.330** (0.141)	0.033 (0.069)	-0.180* (0.108)
Humanities	-0.000 (0.065)	-0.050 (0.059)	-0.386*** (0.147)	-0.141 (0.092)	-0.150*** (0.052)	-0.300** (0.119)	0.100 (0.062)	-0.086 (0.118)
Languages	0.143*** (0.034)	0.012 (0.096)	-0.180 (0.296)	0.112*** (0.033)	-0.140* (0.084)	-0.394** (0.154)	0.152*** (0.054)	0.214 (0.207)
Literature	-0.063 (0.124)	-0.109 (0.111)	-0.355 (0.256)	0.102*** (0.027)	-0.027 (0.068)	-0.187 (0.137)	-0.082 (0.080)	-0.168 (0.184)
Social Sci.	0.046 (0.033)	0.080** (0.039)	0.293*** (0.105)	0.035 (0.028)	0.062** (0.029)	0.184*** (0.070)	0.019 (0.032)	0.110 (0.076)
Mean Dep. Var.	0.912	0.612	2.285	0.913	0.419	1.801	0.193	0.484
R-squared	0.036	0.073	0.126	0.053	0.117	0.129	0.042	0.063
N (Students)	515	515	515	515	515	515	515	515

Notes: This table reports estimated associations between student characteristics and their use of generative AI for academic tasks. In columns 1 and 4, the outcome is a dummy that equals one if a student reports using AI with any frequency for at least one augmentation or automation task, respectively. In columns 2 and 5, the outcome is the share of tasks within each category for which the student reports any use. In columns 3 and 6, the outcome is a continuous measure capturing average usage frequency for each task category, based on raw Likert-style responses. In columns 7 and 8, the outcome is the difference in average task share and usage intensity between augmentation and automation, respectively. Regressions are weighted and report robust standard errors clustered at the student level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

References

- Acemoglu, D. (2025). Nobel lecture: Institutions, technology, and prosperity. *American Economic Review*, 115(6):1709–1748.
- Angrist, J., Lang, D., and Oreopoulos, P. (2009). Incentives and services for college achievement: Evidence from a randomized trial. *American Economic Journal: Applied Economics*, 1(1):136–163.
- Autor, D. and Thompson, N. (2025). Expertise.
- Bailey, M., Johnston, D., Kuchler, T., Stroebel, J., and Wong, A. (2022). Peer effects in product adoption. *American Economic Journal: Applied Economics*, 14(3):488–526.
- Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., and Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26):e2422633122.
- Beaman, L., BenYishay, A., Magruder, J., and Mobarak, A. M. (2021). Can network theory-based targeting increase technology adoption? *American Economic Review*, 111(6):1918–1943.
- Bick, A., Blandin, A., and Deming, D. J. (2025). The rapid adoption of generative AI. Technical report, National Bureau of Economic Research.
- Bonney, K., Breaux, C., Buffington, C., Dinlersoz, E., Foster, L. S., Goldschlag, N., Haltiwanger, J. C., Kroff, Z., and Savage, K. (2024). Tracking Firm Use of AI in Real Time: A Snapshot from the Business Trends and Outlook Survey.
- Bonsu, E. M. and Baffour-Koduah, D. (2023). From the consumers’ side: Determining students’ perception and intention to use chatgpt in ghanaian higher education. *Journal of Education, Society & Multiculturalism*, 4(1):1–29.
- Bresnahan, T. F. and Trajtenberg, M. (1995). General purpose technologies ‘engines of growth’? *Journal of econometrics*, 65(1):83–108.
- Brynjolfsson, E. and Hitt, L. M. (2000). Beyond computation: Information technology, organizational transformation and business performance. *Journal of Economic Perspectives*, 14(4):23–48.
- Brynjolfsson, E., Li, D., and Raymond, L. (2025). Generative AI at Work. *The Quarterly Journal of Economics*, 140(2):889–942.
- Bushong, B. and Gagnon-Bartsch, T. (2024). Failures in forecasting: An experiment on interpersonal projection bias. *Management Science*, 70(12):8735–8752.

- Carvajal, D., Franco, C., and Isaksson, S. (2024). Will artificial intelligence get in the way of achieving gender equality? Discussion Paper 03, NHH Norwegian School of Economics.
- Comin, D. and Hobijn, B. (2010a). An exploration of technology diffusion. *American Economic Review*, 100(5):2031–2059.
- Comin, D. and Hobijn, B. (2010b). *Technology Diffusion and Postwar Growth*, volume 25, pages 209–246. University of Chicago Press.
- Cui, Z. K., Demirer, M., Jaffe, S., Musolff, L., Peng, S., and Salz, T. (2024). The Effects of Generative AI on High Skilled Work: Evidence from Three Field Experiments with Software Developers.
- David, P. A. (1990). The dynamo and the computer: an historical perspective on the modern productivity paradox. *The American economic review*, 80(2):355–361.
- Dell’Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraymer, L., Candelon, F., and Lakhani, K. R. (2023). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality.
- Duflo, E., Kremer, M., and Robinson, J. (2011). Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya. *American Economic Review*, 101(6):2350–2390.
- Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702):1306–1308.
- Fairlie, R. W., Hoffmann, F., and Oreopoulos, P. (2014). A community college instructor like me: Race and ethnicity interactions in the classroom. *American Economic Review*, 104(8):2567–2591.
- Felten, E. W., Raj, M., and Seamans, R. (2021). Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses.
- Felten, E. W., Raj, M., and Seamans, R. (2023). How will Language Modelers like Chat-GPT Affect Occupations and Industries?
- Foster, A. D. and Rosenzweig, M. R. (1995). Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of political Economy*, 103(6):1176–1209.
- Foster, A. D. and Rosenzweig, M. R. (2010). Microeconomics of technology adoption. *Annu. Rev. Econ.*, 2(1):395–424.
- Gallup (2024). AI in the Workplace: Answering 3 Big Questions.

- Giaccherini, M., Herberich, D. H., Jimenez-Gomez, D., List, J. A., Ponti, G., and Price, M. K. (2019). The behavioralist goes door-to-door: Understanding household technological diffusion using a theory-driven natural field experiment. Technical report, National Bureau of Economic Research.
- Goehring, G., Mezzanotti, F., and Ravid, S. A. A. (2024). Technology adoption and career concerns: Evidence from the adoption of digital technology in motion pictures. Technical report, National Bureau of Economic Research.
- Griliches, Z. (1957). Hybrid corn: An exploration in the economics of technological change. *Econometrica*, 25(4):501–522.
- Hall, B. H. and Khan, B. (2003). Adoption of new technology.
- Handa, K., Bent, D., Tamkin, A., McCain, M., Durmus, E., Stern, M., Schiraldi, M., Huang, S., Ritchie, S., Syverud, S., Jagadish, K., Vo, M., Bell, M., and Ganguli, D. (2025a). Anthropic education report: How university students use claude.
- Handa, K., Tamkin, A., McCain, M., Huang, S., Durmus, E., Heck, S., Mueller, J., Hong, J., Ritchie, S., Belonax, T., et al. (2025b). Which economic tasks are performed with ai? evidence from millions of claude conversations. *arXiv preprint arXiv:2503.04761*.
- Hanna, R., Mullainathan, S., and Schwartzstein, J. (2014). Learning Through Noticing: Theory and Evidence from a Field Experiment. *The Quarterly Journal of Economics*, 129(3):1311–1353.
- Hartley, J., Jolevski, F., Melo, V., and Moore, B. (2025). The Labor Market Effects of Generative Artificial Intelligence.
- Humlum, A. and Vestergaard, E. (2025). Large Language Models, Small Labor Market Effects.
- Hutto, C. and Gilbert, E. (2014). Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the international AAAI conference on web and social media*, volume 8, pages 216–225.
- Jehiel, P. (2018). Investment strategy and selection bias: An equilibrium perspective on overoptimism. *American Economic Review*, 108(6):1582–1597.
- Kelly, A., Sullivan, M., and Strampel, K. (2023). Generative artificial intelligence: University student awareness, experience, and confidence in use across disciplines. *Journal of University Teaching and Learning Practice*, 20(6):1–16.
- Kharazian, A. (2025). Ramp ai index: Monthly measurement of ai adoption by american businesses. Technical report, Ramp Economics Lab.


- Lehmann, M., Cornelius, P. B., and Sting, F. J. (2024). AI Meets the Classroom: When Does ChatGPT Harm Learning?
- Ling, Y. and Imas, A. (2025). Underreporting of AI use: The role of social desirability bias. *Available at SSRN*.
- Liu, Y. and Wang, H. (2024). Who on earth is using generative ai? Policy Research Working Paper WPS10870, World Bank Group, Washington, D.C.
- McClain, C. (2024). Americans’ use of ChatGPT is ticking up, but few trust its election information.
- McDonald, N., Johri, A., Ali, A., and 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.
- McElheran, K., Li, J. F., Brynjolfsson, E., Kroff, Z., Dinlersoz, E., Foster, L. S., and Zolas, N. (2023). AI Adoption in America: Who, What, and Where.
- Merchant, B. (2024). The new luddites aren’t backing down. *The Atlantic*. Accessed: 2025-06-16.
- Nam, J. (2023). 56% of college students have used ai on assignments or exams. Technical report. Edited by Lyss Welding. Accessed June 11, 2025.
- Nolan, B. (2023). Here are the schools and colleges that have banned chatgpt. *Business Insider*. Accessed: 2025-01-27.
- Noy, S. and Zhang, W. (2023). Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence.
- Otis, N. G., Delecourt, S., Cranney, K., and Koning, R. (2024). *Global Evidence on Gender Gaps and Generative AI*. Harvard Business School.
- Peng, S., Kalliamvakou, E., Cihon, P., and Demirer, M. (2023). The Impact of AI on Developer Productivity: Evidence from GitHub Copilot.
- Ravšelj, D., Keržič, D., Tomažević, N., Umek, L., Brezovar, N., A. Iahad, N., Abdulla, A. A., Akopyan, A., Aldana Segura, M. W., AlHumaid, J., Allam, M. F., Alló, M., Andoh, R. P. K., Andronic, O., Arthur, Y. D., Aydın, F., Badran, A., Balbontín-Alvarado, R., Ben Saad, H., Bencsik, A., Benning, I., Besimi, A., Bezerra, D. D. S., Buizza, C., Burro, R., Bwalya, A., Cachero, C., Castillo-Briceno, P., Castro, H., Chai, C. S., Charalambous, C., Chiu, T. K. F., Clipa, O., Colombari, R., Corral Escobedo, L. J. H., Costa, E., Crețulescu, R. G., Crispino, M., Cucari, N., Dalton, F., Demir Kaya, M., Dumić-Čule, I., Dwidienawati, D., Ebardo, R., Egbenya, D. L., Faris, M. E., Fečko, M., Ferrinho, P., Florea, A., Fong, C. Y., Francis, Z., Ghilardi, A., González-Fernández, B., Hau, D., Hossain, M. S., Hug, T., Inasius, F., Ismail, M. J., Jahić, H., Jessa, M. O., Kapanadze,

- M., Kar, S. K., Kateeb, E. T., Kaya, F., Khadri, H. O., Kikuchi, M., Kobets, V. M., Kostova, K. M., Krasmane, E., Lau, J., Law, W. H. C., Lazăr, F., Lazović-Pita, L., Lee, V. W. Y., Li, J., López-Aguilar, D. V., Luca, A., Luciano, R. G., Machin-Mastromatteo, J. D., Madi, M., Manguela, A. L., Manrique, R. F., Mapulanga, T., Marimon, F., Marinova, G. I., Mas-Machuca, M., Mejía-Rodríguez, O., Meletiou-Mavrotheris, M., Méndez-Prado, S. M., Meza-Cano, J. M., Mirke, E., Mishra, A., Mital, O., Mollica, C., Morariu, D. I., Mospan, N., Mukuka, A., Navarro Jiménez, S. G., Nikaj, I., Nisheva, M. M., Nisiforou, E., Njiku, J., Nomnian, S., Nuredini-Mehmedi, L., Nyamekye, E., Obadić, A., Okela, A. H., Olenik-Shemesh, D., Ostoj, I., Peralta-Rizzo, K. J., Peštek, A., Pilav-Velić, A., Pires, D. R. M., Rabin, E., Raccanello, D., Ramie, A., Rashid, M. M. U., Reuter, R. A. P., Reyes, V., Rodrigues, A. S., Rodway, P., Ručinská, S., Sadzaglishvili, S., Salem, A. A. M. S., Savić, G., Schepman, A., Shahpo, S. M., Snouber, A., Soler, E., Sonyel, B., Stefanova, E., Stone, A., Strzelecki, A., Tanaka, T., Tapia Cortes, C., Teira-Fachado, A., Tilga, H., Titko, J., Tolmach, M., Turmudi, D., Varela-Candamio, L., Vekiri, I., Vicentini, G., Woyo, E., Yorulmaz, Ö., Yunus, S. A. S., Zamfir, A.-M., Zhou, M., and Aristovnik, A. (2025). Higher education students’ perceptions of ChatGPT: A global study of early reactions. *PLOS ONE*, 20(2):e0315011.
- Rogers, E. M. (1962). *Diffusion of Innovations*. Free Press, New York.
- Ross, L., Greene, D., and House, P. (1977). The “false consensus effect”: An egocentric bias in social perception and attribution processes. *Journal of experimental social psychology*, 13(3):279–301.
- Sacerdote, B. (2001). Peer Effects with Random Assignment: Results for Dartmouth Roommates. *The Quarterly Journal of Economics*, 116(2):681–704.
- Stinebrickner, R. and Stinebrickner, T. R. (2008). The causal effect of studying on academic performance. *The BE Journal of Economic Analysis & Policy*, 8(1).
- Stöhr, C., Ou, A. W., and Malmström, H. (2024). Perceptions and usage of AI chatbots among students in higher education across genders, academic levels and fields of study. *Computers and Education: Artificial Intelligence*, 7:100259.
- Stokey, N. L. (2021). Technology diffusion. *Review of Economic Dynamics*, 42:15–36.
- World Bank (2016). *World development report 2016: Digital dividends*. World Bank Publications.
- Wu, M.-J., Zhao, K., and Fils-Aime, F. (2022). Response rates of online surveys in published research: A meta-analysis. *Computers in Human Behavior Reports*, 7.
- Xiao, P., Chen, Y., and Bao, W. (2023). Waiting, banning, and embracing: An empirical analysis of adapting policies for generative ai in higher education. *arXiv preprint arXiv:2305.18617*.

Appendix


A Appendix Figures and Tables

Figure A1: Generative AI Usage Survey Design Overview




Note: This figure illustrates the structure of the AI usage survey conducted at Middlebury College in December 2024. The survey collected information across three main sections: (1) demographic and academic background, (2) patterns of generative AI usage including adoption timing, frequency, and specific applications, and (3) perceptions of AI's impact on learning and responses to institutional policies.

Figure A2: Cumulative Generative AI Use by Student Characteristic



Notes: This figure presents cumulative AI use based on different student characteristics. Each panel displays the cumulative distribution of AI use based on a specific characteristic: gender, race, school type, first-year GPA, cohort, or field of study. The cumulative percent of students is plotted against usage categories. The legends and colors correspond to subgroups within each demographic variable.


Figure A3: Percent of Students Who Pay for Generative AI Tools




Notes: This figure shows the percent of AI users who pay for AI tools (through any platform) across different demographic groups. Horizontal lines represent 95 percent confidence intervals calculated with heteroskedasticity-robust standard errors clustered at the student level.

Figure A4: Academic Uses of ChatGPT: Evidence from Global Survey

Panel A. Across Tasks that Augment versus Automate Student Effort




Panel B. Across the College Quality Distribution




Notes: This figure shows the percent of students who use ChatGPT for different academic tasks based on data from [Ravšelj et al. \(2025\)](#). Panel A displays usage patterns across tasks categorized as augmenting (proofreading, translating, study assistance, research assistance) versus automating (academic writing, professional writing, creative writing, brainstorming, summarizing, calculating, coding assistance, personal assistance) student effort. Panel B shows usage patterns across university quality quintiles based on World University Rankings, with universities ranked in the top 20% (top quintile) showing slightly higher rates of augmentation relative to automation compared to bottom quintile institutions. The analysis includes universities with at least 30 student responses and excludes observations with missing usage data.

Figure A5: Student Beliefs about AI Usage versus Actual Usage under Different Policies



Notes: This figure compares students' beliefs about AI usage with actual usage rates under different AI policies. For each policy type, blue bars show the percent of students who report being "Neutral," "Likely," or "Very Likely" to use AI, while red bars show students' mean beliefs about what percent of their peers use AI in classes with that policy. The sample includes all survey respondents for both actual usage and beliefs measures. "AI Use is Unrestricted" refers to classes with no restrictions on AI use, "No Explicit AI Policy" refers to classes without a stated policy, and "AI Use is Prohibited" refers to classes that ban AI use entirely.

Figure A6: Relationship Between Beliefs About AI Usage and Actual AI Usage



Notes: This figure shows the relationship between students' beliefs about AI usage among their peers and actual AI usage rates across different demographic groups. Each point represents a different group of students (by demographics, academic characteristics, or field of study). The *y*-axis shows the percent of students in each group who report using AI for academic purposes. The *x*-axis shows the mean belief within each group about what percent of Middlebury students use AI. The dashed line shows the linear fit. Sample includes all survey respondents with at least ten observations per group. Students' beliefs about AI usage are positively correlated with actual usage patterns, suggesting that students have relatively accurate perceptions of AI adoption among their peers.

Table A1: Generative AI Usage Frequency by Student Characteristics

	By Usage Frequency				
	Any use (1)	Rarely (2)	Occasionally (3)	Frequently (4)	Very Frequently (5)
Panel A. Demographics					
Male	88.7	20.2	19.1	33.4	16.1
Female	78.4	25.7	25.1	21.0	6.6
White	80.2	25.6	19.8	24.0	10.8
Black	92.3	15.1	23.5	41.4	12.3
Hispanic	77.9	19.9	28.7	22.8	6.6
Asian	91.3	19.5	27.0	30.1	14.8
Private HS	84.1	23.7	22.2	27.2	10.9
Public HS	80.4	24.7	21.4	24.4	9.9
Panel B. Academic Characteristics					
GPA > p50	80.3	26.3	17.4	23.2	13.4
GPA < p50	87.1	18.7	26.3	32.2	9.9
Freshman	81.3	24.3	24.1	24.7	8.1
Sophomore	84.4	25.6	22.6	23.7	12.4
Junior	86.2	21.6	23.1	28.6	12.8
Senior	78.7	21.0	17.4	29.9	10.4
Panel C. Field of Study					
Arts	73.3	21.0	0.0	52.3	0.0
Humanities	74.8	39.3	21.8	12.1	1.7
Languages	57.4	27.0	30.3	0.0	0.0
Literature	48.6	10.0	33.7	5.0	0.0
Natural Sciences	91.1	22.0	21.9	33.8	13.3
Social Sciences	84.6	17.9	20.4	29.1	17.2

Notes: This table presents the percent of students in each demographic group who report using AI at different frequencies during the academic semester. Each cell shows the percent of students within that group. Column 1 reports the total percent who use AI at any frequency. Columns 2 to 5 represent increasing usage frequencies: rarely (1–2 times per semester), occasionally (monthly), frequently (weekly), and very frequently (multiple times per week). Panel A reports percentages by demographic characteristics. Panel B shows percentages by academic characteristics. Panel C presents percentages by field of study.

Table A2: Student Characteristics Associated with Choice of Generative AI Models

	Outcome: =1 if student uses				
	OpenAI's ChatGPT (1)	Google's Gemini (2)	Microsoft Copilot (3)	Other Model (4)	Pays for GenAI (5)
Male	0.007 (0.017)	0.086** (0.035)	0.010 (0.026)	-0.003 (0.032)	0.107*** (0.030)
Black	-0.032 (0.062)	-0.010 (0.097)	-0.013 (0.061)	-0.014 (0.077)	-0.023 (0.044)
Latino	0.010 (0.025)	-0.033 (0.051)	0.033 (0.051)	-0.004 (0.051)	0.023 (0.043)
Asian	0.009 (0.017)	-0.077** (0.035)	-0.014 (0.034)	-0.067** (0.032)	0.116*** (0.043)
Public HS	-0.010 (0.014)	0.005 (0.033)	0.021 (0.026)	-0.012 (0.030)	-0.001 (0.028)
Sophomores	-0.000 (0.028)	0.047 (0.044)	-0.064* (0.034)	-0.003 (0.039)	-0.008 (0.037)
Juniors	0.025 (0.018)	0.119* (0.063)	-0.061 (0.042)	0.043 (0.048)	-0.001 (0.040)
Seniors	0.023 (0.018)	0.108* (0.058)	-0.066* (0.036)	0.045 (0.048)	0.066 (0.047)
Arts	0.011 (0.019)	-0.200*** (0.050)	0.060 (0.126)	-0.084** (0.041)	0.489*** (0.183)
Humanities	-0.009 (0.032)	-0.186*** (0.054)	-0.067*** (0.024)	0.022 (0.074)	-0.041 (0.049)
Languages	0.033* (0.020)	-0.078 (0.096)	-0.064** (0.027)	-0.108*** (0.035)	-0.049 (0.034)
Literature	0.029* (0.015)	-0.162*** (0.039)	-0.065*** (0.023)	0.100 (0.130)	-0.059* (0.030)
Social Sci.	0.008 (0.017)	-0.074 (0.048)	-0.004 (0.029)	0.039 (0.038)	0.064* (0.036)
Mean Dep. Var.	0.973	0.135	0.077	0.110	0.114
R-squared	0.012	0.051	0.027	0.024	0.138
N (Students)	516	516	516	516	516

Notes: This table assesses the relationship between AI model adoption and student characteristics. We estimate:

$$Y_i = \alpha + \beta X_i + \varepsilon_i,$$

where Y_i is a binary indicator of AI model usage (columns 1-4) or payment for AI services (column 5), and X_i is a vector of student characteristics including gender, race/ethnicity, high school type, cohort indicators, and academic division.

Each column presents results for a different model or payment outcome. Column 1 shows usage of OpenAI's ChatGPT, column 2 Google Gemini, column 3 Microsoft Copilot, column 4 any other AI model, and column 5 whether the student pays for any generative AI service.

The omitted categories are Natural Sciences for academic division, white students for race/ethnicity, and freshmen for cohort. Heteroskedasticity-robust standard errors clustered at the student level in parentheses. Observations are weighted to adjust for sampling. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A3: Student Characteristics Associated with AI Usage Likelihood Under Different Institutional Policies

	Outcome: Would use generative AI in a given policy scenario...				
	AI Use is Unrestricted (1)	No Explicit AI Policy (2)	AI Allowed if Cited (3)	AI Use is Prohibited (4)	Prohibition Impact (5)
Male	0.105*** (0.038)	0.158*** (0.040)	0.101** (0.042)	0.166*** (0.041)	-0.061 (0.049)
Black	0.132** (0.057)	0.082 (0.095)	0.034 (0.105)	-0.061 (0.097)	0.193* (0.110)
Latino	0.098* (0.057)	0.019 (0.069)	0.114* (0.062)	0.037 (0.066)	0.060 (0.079)
Asian	0.105** (0.048)	0.131*** (0.049)	0.042 (0.057)	0.003 (0.052)	0.102 (0.067)
Public HS	-0.064* (0.038)	-0.041 (0.041)	-0.046 (0.042)	0.023 (0.039)	-0.087* (0.047)
Sophomores	0.102** (0.048)	0.120** (0.053)	0.066 (0.056)	0.033 (0.055)	0.069 (0.064)
Juniors	0.115** (0.055)	0.103* (0.062)	0.052 (0.064)	0.036 (0.066)	0.079 (0.070)
Seniors	0.085 (0.061)	0.043 (0.066)	0.025 (0.067)	-0.017 (0.063)	0.102 (0.074)
Arts	-0.248 (0.182)	-0.161 (0.179)	-0.100 (0.206)	-0.213* (0.116)	-0.035 (0.236)
Humanities	-0.182* (0.094)	-0.253** (0.100)	-0.279*** (0.100)	-0.166** (0.077)	-0.016 (0.105)
Languages	-0.442*** (0.156)	-0.557*** (0.095)	-0.540*** (0.120)	-0.211*** (0.048)	-0.231 (0.154)
Literature	-0.389*** (0.117)	-0.288** (0.117)	-0.271** (0.116)	-0.021 (0.101)	-0.368*** (0.100)
Social Sci.	-0.047 (0.046)	0.030 (0.049)	-0.003 (0.052)	0.086 (0.054)	-0.133** (0.056)
Mean Dep. Var.	0.751	0.676	0.661	0.271	0.479
R-squared	0.087	0.113	0.076	0.068	0.041
N (Students)	599	599	599	599	599

Notes: This table examines how student characteristics relate to self-reported likelihood of using generative AI under different policy scenarios. Each column presents results for different policy scenarios. In columns 1-4, the dependent variable equals one if the student reports being “neutral,” “likely” or “extremely likely” to use AI under the specified policy, and zero if they report being “unlikely,” or “extremely unlikely.” Column 5 represents the impact of moving from unrestricted use to complete prohibition.

The omitted categories are: Natural Sciences for academic division, white students for race/ethnicity, freshmen for cohort, female for gender, and private high school for school type. Heteroskedasticity-robust standard errors clustered at the student level in parentheses. Observations are weighted to adjust for sampling. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A4: Student Characteristics Associated with Perceived Learning Benefits

	Outcome: Believes that generative AI improves...				
	Learning Ability (1)	Understand Materials (2)	Course Grades (3)	Assignment Completion (4)	Time on Academics (5)
Male	0.109** (0.045)	0.089** (0.041)	0.143*** (0.046)	0.019 (0.047)	0.176*** (0.045)
Black	0.362*** (0.060)	0.234*** (0.051)	0.257* (0.133)	-0.053 (0.123)	0.108 (0.104)
Latino	0.176** (0.069)	0.066 (0.065)	0.198*** (0.076)	0.019 (0.075)	0.077 (0.072)
Asian	0.106* (0.060)	0.122** (0.053)	0.063 (0.061)	0.013 (0.064)	0.128** (0.061)
Public HS	-0.056 (0.045)	-0.043 (0.041)	0.016 (0.045)	-0.045 (0.047)	-0.042 (0.045)
Sophomores	0.056 (0.057)	0.040 (0.053)	0.015 (0.061)	-0.142** (0.062)	0.025 (0.061)
Juniors	-0.026 (0.073)	0.008 (0.064)	0.048 (0.072)	-0.153** (0.074)	0.029 (0.071)
Seniors	-0.122* (0.069)	-0.069 (0.068)	-0.116* (0.068)	-0.157** (0.073)	-0.060 (0.072)
Arts	0.316*** (0.119)	0.098 (0.130)	-0.364*** (0.065)	0.250 (0.238)	0.131 (0.165)
Humanities	-0.104 (0.101)	-0.387*** (0.097)	-0.059 (0.094)	-0.042 (0.103)	-0.245** (0.095)
Languages	-0.229* (0.128)	-0.441*** (0.131)	-0.105 (0.120)	-0.235* (0.133)	-0.246** (0.122)
Literature	-0.277*** (0.107)	-0.295** (0.124)	-0.262*** (0.069)	-0.101 (0.120)	-0.130 (0.126)
Social Sci.	0.107* (0.055)	-0.022 (0.053)	0.075 (0.057)	0.022 (0.059)	0.070 (0.057)
Mean Dep. Var.	0.576	0.685	0.376	0.489	0.564
R-squared	0.102	0.119	0.084	0.037	0.087
N (Students)	561	565	543	560	560

Notes: This table assesses the relationship between AI adoption and student characteristics. Each column presents results for beliefs about different academic outcomes: learning ability (e.g., ability to grasp concepts, retain information, or learn new skills) in column 1, understanding of course materials in column 2, course grades in column 3, ability to complete assignments on time in column 4, and time spent on academics in column 5. The dependent variable in each regression equals one if the student believes AI “somewhat improves” or “significantly improves” the outcome, and zero if they believe it has no effect, reduces, or significantly reduces the outcome. “Don’t know” responses are excluded.

The omitted categories are: Natural Sciences for academic division, white students for race/ethnicity, freshmen for cohort, female for gender, and private high school for school type. Heteroskedasticity-robust standard errors clustered at the student level in parentheses. Observations are weighted to adjust for sampling. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

B Empirical Appendix

B.1 Field of Study Classifications

This appendix details the classification of majors into broad fields of study used in Figure 1:

- **Natural Sciences:** Includes Biology, Biochemistry, Chemistry, Computer Science, Earth and Climate Sciences/Geology, Environmental Studies, Mathematics, Molecular Biology & Biochemistry, Neuroscience, Physics, and Statistics.
- **Social Sciences:** Includes Anthropology, Economics, Education, Geography, International & Global Studies, International Politics & Economics, Political Science, Psychology, and Sociology.
- **Humanities:** Includes American Studies, Black Studies, Classics, History, Architectural Studies, Art History & Museum Studies, History of Art & Architecture, Philosophy, Religion, and Classical Studies.
- **Literature:** Includes Comparative Literature, English/English & American Literatures, and Literary Studies.
- **Languages:** Includes Arabic, Chinese, French & Francophone Studies, German, Japanese Studies, Russian, and Spanish.
- **Arts:** Includes Film & Media Culture, Music, Studio Art, and Theatre.

B.2 Task-Specific Use of Generative AI

The aggregate patterns documented in Section 4 may mask important heterogeneity across student groups. Different students may adopt AI tools at varying rates depending on their academic needs and field-specific norms. To investigate this heterogeneity systematically, Appendix Table B1 presents regression estimates examining how student characteristics predict AI usage for each of our ten measured academic tasks.

Usage patterns differ markedly by student characteristics. Male students consistently exhibit higher adoption rates across most applications, with particularly pronounced differences for finding information (16.4 pp; $p < 0.01$), summarizing texts (12.5 pp; $p < 0.01$), and creating images (16.9 pp; $p < 0.01$). Black students show substantially higher usage for information-gathering tasks, being 23.8 pp more likely to use AI for finding information compared to white students ($p < 0.05$). They also exhibit dramatically higher adoption of writing assistance tools, with 27.6 pp higher usage for editing text ($p < 0.05$) and 23.6 pp for writing emails ($p < 0.10$). Latino students demonstrate significantly higher usage for generating ideas (24.0 pp; $p < 0.01$) and writing emails (20.7 pp; $p < 0.01$), while Asian students show markedly higher adoption for writing emails (24.5 pp; $p < 0.01$) and explaining concepts (10.9 pp, $p < 0.05$). Students from public high schools report lower usage rates for several applications, particularly for concept explanation (7.3 pp lower; $p < 0.10$) and writing emails (7.4 pp lower; $p < 0.10$), but higher usage for proofreading (8.7 pp higher; $p < 0.10$).

Field of study emerges as a particularly strong predictor of usage patterns. Arts majors show significantly lower adoption across multiple tasks compared to Natural Science students, with gaps of 49.7 pp for finding information ($p < 0.01$), 44.0 pp for generating ideas ($p < 0.05$), and 62.1 pp for coding assistance ($p < 0.01$). Similarly, humanities majors are 36.8 pp less likely to use AI for generating ideas ($p < 0.01$) and 30.8 pp less likely for coding assistance ($p < 0.01$). Languages majors exhibit the most pronounced differences, being 52.4 pp less likely to use AI for summarizing texts ($p < 0.01$) and 46.9 pp less likely for generating ideas ($p < 0.01$). By contrast, social science students show higher usage for explaining concepts (8.0 pp; $p < 0.10$) and summarizing texts (14.7 pp; $p < 0.01$).

Table B1: Student Characteristics Associated with Task-Specific Use of Generative AI

	=1 if student uses generative AI with any frequency during the academic semester to...									
	Explain Cnpts. (1)	Summ. Texts (2)	Find Info. (3)	Gen. Ideas (4)	Proof- read (5)	Edit Text (6)	Write Emails (7)	Code Assist. (8)	Write Essays (9)	Create Images (10)
Male	0.070*	0.125***	0.164***	0.053	0.052	0.105**	0.057	0.089**	0.051	0.169***
	(0.037)	(0.042)	(0.045)	(0.046)	(0.049)	(0.048)	(0.046)	(0.044)	(0.043)	(0.039)
Black	0.112	0.065	0.238**	-0.119	0.159	0.276**	0.236*	-0.129	0.090	-0.138*
	(0.074)	(0.108)	(0.095)	(0.123)	(0.126)	(0.123)	(0.127)	(0.101)	(0.121)	(0.077)
Latino	0.168***	0.085	0.093	0.240***	0.078	0.086	0.207**	-0.155**	-0.000	0.113
	(0.050)	(0.066)	(0.076)	(0.063)	(0.082)	(0.081)	(0.083)	(0.061)	(0.070)	(0.075)
Asian	0.109**	0.033	0.018	0.037	0.076	0.067	0.245***	0.055	0.009	-0.072*
	(0.044)	(0.056)	(0.060)	(0.060)	(0.064)	(0.064)	(0.062)	(0.060)	(0.054)	(0.041)
Public HS	-0.073*	-0.044	-0.001	0.018	0.086*	0.025	-0.074*	-0.023	0.002	0.009
	(0.038)	(0.042)	(0.046)	(0.046)	(0.049)	(0.048)	(0.045)	(0.043)	(0.041)	(0.037)
Sophomores	0.021	0.019	0.039	-0.004	-0.010	-0.103	-0.065	0.148**	-0.035	-0.074
	(0.048)	(0.056)	(0.062)	(0.060)	(0.064)	(0.064)	(0.060)	(0.058)	(0.057)	(0.049)
Juniors	0.028	-0.051	0.059	0.064	0.036	-0.042	0.001	0.344***	0.017	-0.021
	(0.055)	(0.067)	(0.071)	(0.067)	(0.076)	(0.076)	(0.074)	(0.069)	(0.072)	(0.058)
Seniors	0.001	0.021	-0.017	0.041	0.004	-0.072	-0.039	0.359***	-0.054	0.073
	(0.060)	(0.064)	(0.073)	(0.072)	(0.075)	(0.076)	(0.069)	(0.072)	(0.060)	(0.062)
Arts	-0.066	-0.005	-0.497***	-0.440**	0.197	-0.385***	0.214	-0.621***	-0.202***	-0.270***
	(0.144)	(0.155)	(0.160)	(0.198)	(0.201)	(0.128)	(0.163)	(0.089)	(0.054)	(0.061)
Humanities	-0.051	-0.151	-0.146	-0.368***	-0.061	0.056	-0.075	-0.308***	-0.048	0.049
	(0.090)	(0.108)	(0.102)	(0.096)	(0.114)	(0.109)	(0.091)	(0.081)	(0.090)	(0.092)
Languages	-0.444**	-0.524***	0.007	-0.469***	0.176	0.307	0.167	-0.279	-0.167***	0.431**
	(0.181)	(0.154)	(0.221)	(0.162)	(0.216)	(0.213)	(0.219)	(0.173)	(0.038)	(0.215)
Literature	-0.232	-0.035	-0.139	0.197	0.102	-0.167	0.134	-0.363***	-0.207***	0.114
	(0.157)	(0.158)	(0.153)	(0.120)	(0.161)	(0.149)	(0.167)	(0.118)	(0.040)	(0.144)
Social Sci.	0.080*	0.147***	0.061	0.134**	0.096	0.085	0.112*	-0.094	0.114**	-0.044
	(0.043)	(0.048)	(0.057)	(0.054)	(0.060)	(0.061)	(0.058)	(0.058)	(0.058)	(0.045)
Mean Dep. Var.	0.803	0.740	0.631	0.619	0.541	0.473	0.371	0.344	0.235	0.204
R-squared	0.078	0.087	0.077	0.116	0.028	0.054	0.083	0.135	0.040	0.094
N (Students)	515	515	515	515	515	515	515	515	515	515

Notes: This table reports estimated associations between student characteristics and use of AI for specific academic tasks. Each column shows the result for a different academic task. The omitted categories are: Natural Sciences for academic division, white students for race/ethnicity, freshmen for cohort, female for gender, and private high school for school type. Regressions are weighted and use heteroskedasticity-robust standard errors clustered at the student level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

C Qualitative Evidence on Student Perspectives on AI Use

In this section, we analyze student responses to an open-ended question about their use of generative AI. The survey question asked: “Please describe the factors that have personally influenced your use of generative AI in your academic work. What initially led you to try it, what has motivated you to use it or caused you to hesitate?” This was an optional question that 48.3 percent of survey respondents answered. Appendix Figure C1 presents a word cloud of the most frequent words in student responses.

C.1 Validating the Open-Ended Response Measure

We begin by validating our open-ended response measure to show that it contains meaningful signal. To validate this measure, we employ sentiment analysis using VADER (Valence Aware Dictionary and sEntiment Reasoner), a widely-used lexicon and rule-based sentiment analysis tool (Hutto and Gilbert, 2014). We calculate sentiment scores for each response and examine whether these scores correlate with students’ actual AI adoption behaviors. Intuitively, students who express more positive sentiment toward AI in their open-ended responses should be more likely to have adopted AI tools in practice. We test this hypothesis by examining two measures of AI adoption: whether students have ever used generative AI and whether they currently use AI for academic purposes. Appendix Figure C2 presents binned scatterplots showing the relationship between AI adoption rates and standardized sentiment scores.

Sentiment towards generative AI strongly predicts AI adoption. For both outcomes, students with negative sentiment scores (below zero) show substantially lower adoption rates—around 70-85 percent for ever using AI and 70-75 percent for academic use. In contrast, students with positive sentiment scores display markedly higher adoption rates, reaching nearly 100 percent for general use and 95-100 percent for academic purposes among those with the most positive sentiment. The relationship is particularly pronounced for academic AI use (Panel B), where the coefficient of $\hat{\beta} = 0.168$ is twice as large as for general adoption.¹⁷ These systematic patterns confirm that our open-ended responses capture meaningful variation in student attitudes that corresponds to real behavioral differences.

¹⁷This stronger association for academic use makes intuitive sense, as our open-ended question specifically asked about academic AI use, making the sentiment scores particularly relevant for this domain.

C.2 How Students use Generative AI

To analyze these responses systematically, we classified each response using keywords based on their content. For example, if a student mentioned using AI to save time, we tagged the response with the keyword “time-saver.” If a student expressed concerns about AI’s impact on learning, we tagged it with “negative learning.” Responses could receive multiple keywords if they discussed several themes. Appendix Figure C3 shows the frequency of keywords in our classification. The responses reveal how students use AI tools, what motivates this use, and what causes some to avoid or limit their use.

The most common use of AI is as an explanatory tool. Nearly 30 percent of responses mentioned using AI to understand course material. Students frequently ask AI to break down complex concepts from readings and lectures, particularly when they find the material difficult to understand. For example, one student reported: “I can ask AI to explain concepts to me that I have a hard time grasping. [...] I can keep asking ‘simplify’ or ‘break down even more.’” Students also use AI to summarize dense academic readings, which they argue helps them manage heavy reading loads.

Students employ AI throughout different stages of the writing process. Some use AI to generate initial drafts that serve as starting points. One student explained: “Helps me get started with a base for most of my essays. It feels easier to edit something already written and make it my own than to write from scratch.” Others use AI more narrowly for brainstorming when stuck on specific problems. As one student noted: “I use it if I am feeling stuck to push me to the right direction (whether a mathematical problem or an essay idea).” Many also report using AI as an editing tool to improve grammar, sentence structure, and overall writing flow. This is especially the case for non-native English speakers. As one student explained: “English is not my first language and it frustrates me sometimes that I cannot find the best way to phrase a certain idea and AI is a useful tool to have to find alternate expressions.”

Students frequently mentioned using AI for specific academic tasks. In courses that require coding, students often use AI for debugging code and understanding programming concepts. Students also employ AI for administrative tasks like formatting citations and drafting routine emails. Finally, many students use AI as an enhanced search engine. One student reported: “It has significantly reduced the time it takes to conduct research on new topics and ideas, and helps me by giving me a thorough selection of sources to use for projects of any kind.”

C.3 Why Students Adopt AI

Time savings was the most commonly cited reason for using AI. Nearly 30 percent of responses mentioned using AI to complete work more efficiently. Students often viewed AI as a way to manage demanding course loads. Many students particularly embrace AI assistance for tasks viewed as mechanical or administrative. A student noted they use AI for “Writing emails quick and creating resume/ cover letter templates.” But AI assistance goes beyond grunt work. Some students use it to “spend less time doing assignments and homework.” This is particularly true if students don’t view the work as central to their academic experience. As one student explained: “when I come across work I deem as ineffective, I want to spend as little time as possible doing it.”

Having an on-demand tutor for academic support was another key motivation. One-quarter of responses described using AI as an “explainer” when other resources were unavailable or inconvenient. As one student noted, “I use it as a last resort (if there are no office hours, after looking up videos, etc.) if I need extra help. I’d like to think that the way I use it is similar to going to office hours or TA hours.”

Peer influence also drove AI adoption. Some students reported feeling pressure to use AI to remain competitive with their classmates. One student explained: “I noticed others use it, are getting better grades than me, and they say they learn better with the help of AI, so I gave it a try.” Others worried about being at a competitive disadvantage: “Other people were using it and told me about it. I felt like I would be at a disadvantage if I wasn’t also using it.”

C.4 Concerns and Limitations

Students expressed several concerns about AI use in academic work. The most frequent worry was about negative impacts on learning. One student who initially used AI extensively reported: “In the past, I have simply plugged and chugged homework assignments into ChatGPT and submitted it. Those assignments feedback from teachers was positive and I was getting good grades, but I definitely felt that my own learning outcomes to be significantly worse.” Other students viewed AI use as fundamentally incompatible with their educational goals. As one humanities student explained: “my task is as a humanities student is to think, not calculate; why should I let AI do the thinking for me? It would defeat the purpose of pursuing my education.”

Many students described ethical concerns about AI use. Responses suggested uncer-

tainty about appropriate boundaries. One student noted: “I never use it to explicitly write something because that feels like overt cheating, but sometimes I hesitate when it completely solves Econ problems. I understand how it does it, and it helps me to learn, but it still sometimes feels a little morally gray.” Another expressed similar ambivalence: “I tend to only use it when [it] will save me time in a moral way.”

Students also emphasized the importance of maintaining ownership of their work. Many expressed pride in producing original work and hesitation about diluting that ownership through AI use. As one student explained: “I don’t have interest in using generative AI for my academic work because I want my work to reflect my own ideas.” Another noted: “It usually would not even occur to me to turn to AI to substitute writing because I want to take credit for my work, and using AI seems to diminish that.”

Technical limitations deterred some students from using AI tools. Students reported concerns about inaccurate outputs (“hallucinations”) and poor output quality, particularly for creative writing or complex analytical tasks. For example, one student noted that “In my poetry class we were instructed to use it to come up with poems and they were awful, so that kinda turned me away from using it to do my work for me.”

C.5 Discussion


We conclude with two overarching themes that emerged from the responses.

First, students vary substantially in how they incorporate AI into their academic lives. Crucially, this heterogeneity largely depends on what students perceive as “appropriate” uses of AI. For activities that they perceive as core to their academic journey—like writing essays or solving problems—many students hesitate to use AI. A student articulated this clearly: “Most of my work is writing or reading. If I’m not doing the writing, what is the purpose of me taking the class?” Yet, students draw different boundaries between central tasks and grunt work. Some use AI extensively, viewing their role as akin to a manager that provides high-level direction while AI handles implementation. Others restrict AI use to specific tasks like brainstorming, editing, or drafting emails. Still others avoid AI entirely for academic work, often for ethical reasons. Even among AI users, adoption patterns reflect individual trade-offs between time savings, learning goals, and academic integrity concerns.

Second, there appears to be a fundamental tension between efficiency and learning. The time-saving benefits are easy to observe, quantifiable, and tangible. But these time savings are unlikely to be a free lunch. Some benefits may come at the cost of spending less time


with material that requires deeper engagement to digest. As one student noted: “There may be a negative effect in that it eliminates much of the ‘struggle’ in learning.” Yet, having an on-demand tutor that explains concepts in relatable ways can also improve learning. As one student explained: “It can explain concepts to me in a way that is tailored to my learning style.” This suggests that the impact of AI use on learning outcomes depends not on whether students use AI—almost all do, to some extent—but rather on how they employ these tools.

Figure C1: Word Cloud of Student Motivations for Generative AI Use




Notes: Word cloud displaying words that appear at least five times in 147 student responses after removing common English stop words and the word “AI”. Text size is proportional to word frequency. The visualization is based on responses to the question: “Please describe the factors that have personally influenced your use of generative AI in your academic work. What initially led you to try it, what has motivated you to use it or caused you to hesitate?”

Figure C2: Relationship Between AI Sentiment and AI Adoption



Notes: This figure presents the relationship between AI sentiment and AI adoption. Panel A shows the proportion of respondents who have ever used generative AI, while Panel B shows the proportion who use AI for academic purposes. Each point represents the mean adoption rate for respondents within sentiment score bins of width 0.2. Sentiment scores are standardized compound scores computed using [Hutto and Gilbert \(2014\)](#)'s VADER algorithm applied to responses to an open-ended question about generative AI. Positive values indicate positive sentiment and negative values indicate negative sentiment. The dashed lines show OLS best-fit lines estimated on the microdata, with coefficients and standard errors (in parentheses) displayed in the top-left corner of each panel.

Figure C3: Frequency of Keywords in Student Motivations for AI Use



Notes: The figure shows the share of open-ended responses that mentioned different themes related to AI use. The responses come from the question “Please describe the factors that have personally influenced your use of generative AI in your academic work. What initially led you to try it, what has motivated you to use it or caused you to hesitate?” Color coding indicates the category of each theme. Usage type refers to how students use AI tools. Motivation captures what drove students to try AI. Concerns include mentions of course policies and academic integrity, individual reservations about AI use, worries about AI’s impact on education, and AI’s technical limitations.

D Qualitative Evidence on Student Views of AI Policies

In this section, we analyze student responses to an open-ended question about Middlebury’s AI policies. The survey asked: “Do you have any specific feedback or suggestions about Middlebury’s generative AI policies, resources, or support services?” Appendix Figure [D1](#) presents a word cloud of the most frequent words in student responses. To analyze these responses systematically, we classified each response using keywords based on their content. Appendix Figure [D2](#) shows the frequency of keywords in our classification.

D.1 Polarized Views on Generative AI Policy Approaches

Students expressed markedly different views about appropriate AI policies, revealing fundamental disagreement about the path forward. Some strongly advocated for embracing AI technology. As one student argued, “The tool is there, there is supply and there is demand. Don’t fight another war on drugs. Don’t live in a fake reality”. Others called for significant restrictions, arguing that “the use of generative AI is dishonest and corrosive” and that it “prohibits these organic processes and divorces students from true learning”.

However, the most common position advocated for a balanced approach that would allow beneficial uses while restricting harmful ones. Students distinguished between uses that enhance learning (like concept explanation) and those that substitute for learning (like generating entire essays). One student articulated this nuanced view particularly well: “AI also can really be helpful at explaining a textbook problem that doesn’t make sense, or guiding slightly with homework, or creating study materials, or editing/tightening up your prose. All of those things are good, and universities should figure out how to maximize AI use for those reasons and to minimize students just feeding their problem sets into ChatGPT”

A recurring theme was the futility of blanket bans. Many students emphasized that prohibition would be ineffective, with one noting “I don’t think anyone really cares what the policy of any given class is. If professors want people to not use it, they need to structure assessments in a way that will discourage use.” Another compared AI bans to restricting internet use, arguing “AI policies seem to be totally irrelevant. It’s like telling people they can’t use the internet as a resource for the class.” This ineffectiveness of bans creates fairness concerns. As one student explained: “I think if it is banned in a class, that should be enforced (and right now it absolutely is not)... As with any form of cheating, those who don’t cheat are put at a disadvantage.” Another student expressed similar frustration: “I

find it discouraging when I hear classmates saying they use AI for things such as essays when they use it in dishonest ways.”

D.2 Need for Clear Guidelines

The most frequently expressed concern was the need for clear guidelines about generative AI use. Several students reported confusion about what constitutes acceptable use, with one noting “I think it should be more clear whether we can use it and how and how to cite it since most professors rarely mention it at all”. Other students emphasized the importance of professors explaining their policies upfront and their rationale, with one stating “I think that Professor’s should be very specific about what is allowed and their reasoning behind their policy”.

Many students advocated for standardization across classes, observing that “Sometimes its confusing when one class allows it and another doesn’t and the other encourages it and so on so if there was a school wide or department wide policy that could help”. Yet there is also disagreement regarding standardization. Some students preferred leaving the decision to individual professors, arguing that “GenAI is more effective in some classes/majors than others. Making sure professors understand how students use GenAI and how useful GenAI is in their class (given the course structure, nature of assignments/material, etc.) is very important for the class policy.”

D.3 Training and Support Services

Students strongly emphasized the need for training in appropriate generative AI use. Many suggested that the college should provide guidance on using AI tools effectively while maintaining academic integrity. One student proposed “a workshop that teaches you to effectively use GenAI without violating the honor code”. Students also expressed interest in learning how to leverage AI to enhance their learning experience rather than circumvent it. As one student explained: “I think it could be useful to develop some sort of training. How do we use AI in a way that actually benefits our learning? I tried out some things on my own but I feel that I need more guidance.”

This desire for training was often linked to workplace preparedness. Students recognized that AI proficiency would be valuable in their careers, with one noting “As the world uses more and more AI, I think it is an important tool that students should know how to leverage”. Another emphasized: “The moment us students leave campus, we will be using

it in the professional world, and when used in combination with one's own skills, it is merely a tool to maximize efficiency”.

D.4 Discussion

The responses reveal several key tensions in AI policy implementation. First, students express a desire to use AI in ways that enhance rather than substitute for learning, yet they recognize that blanket bans are ineffective and worry about being disadvantaged if they follow restrictions while others do not.

Second, while students desire clear guidelines, they also want flexibility to accommodate legitimate uses that vary by discipline and assignment type. Different courses and majors may find different AI uses appropriate based on their learning objectives and assessment types.

Third, there is tension between faculty autonomy in setting course policies and students' desire for consistent institutional standards. While some support letting professors determine appropriate AI use for their specific courses, others argue that varying policies across classes create confusion and enforcement challenges.


Finally, students report significant variation in faculty attitudes toward AI. Some students perceived faculty fear or misunderstanding of AI tools, noting that categorical bans often reflect a lack of understanding about AI's capabilities and limitations rather than pedagogical considerations.

Figure D1: Word Cloud of Student Feedback on Generative AI Policies



Notes: Word cloud displaying words that appear at least five times in 133 student responses after removing common English stop words and the word “AI”. Text size is proportional to word frequency. The visualization is based on responses to the question: “Do you have any specific feedback or suggestions about Middlebury’s generative AI policies, resources, or support services?”

Figure D2: Frequency of Keywords in Student Feedback on AI Policies



Notes: The figure shows the share of open-ended responses that mentioned different themes related to Middlebury’s AI policies. The responses come from the question “Do you have any specific feedback or suggestions about Middlebury’s generative AI policies, resources, or support services?” Color coding indicates the category of each theme. Policy Approach captures suggestions about how AI should be regulated at the college. Adoption Views reflect positions on whether and how AI should be integrated into academic work. Training Needs indicates requests for guidance and support. Other Concerns include issues of workplace relevance and fairness.

E Survey Instrument

Default Question Block

Click below to confirm
you are human.



This is a consent form. Please read it carefully, and click below to accept and continue.

You are invited to participate in a research study on on students' use of technology in their academic and personal lives. The study should take around 5 minutes on average to complete. Participation in this survey is voluntary, and you can end your participation at any time by exiting the browser window. If you agree to participate in this study, you will answer questions about your technology use, preferences, and attitudes.

Eligibility: You must be over 18 to participate.

Compensation: As compensation for your time and effort, you will have the option to enter a drawing for multiple gift

cards ranging in value from \$50 to \$500. To enter the drawing, you will be directed to a separate form where you can provide your email address. Your email address cannot be connected to your survey responses.

Risks and Benefits: Your participation in this survey presents no greater risk than everyday Internet use. We cannot and do not guarantee or promise that you will receive any benefits from this study. Your participation may benefit society by improving our understanding of technology usage and its impacts.

Confidentiality: We will make no attempt to identify participants and will keep the data private by storing it securely in a password-protected file on Middlebury's secure servers. The anonymized responses of all survey participants may be shared with other researchers for academic research purposes. We will never share your name or any identifying information with anyone. While we are not collecting any direct identifying information, please be aware that in a small community like Middlebury, there is a very small possibility that some individuals' identities could be ascertained based on their responses.

Contact Information: This survey is being conducted for academic research purposes. The principal investigators are Professor Zara Contractor (zcontractor@middlebury.edu) and Professor Germán Reyes (greyes@middlebury.edu), whom you may contact for specific questions about the research study. For questions about your rights as a research

participant, you may contact the Middlebury College IRB at irb@middlebury.edu.

Agreement to Participate: By clicking to continue, you indicate that you have read this consent form and voluntarily agree to participate in the study.

- ☐ I AGREE TO PARTICIPATE IN THIS STUDY
- ☐ I DO NOT AGREE TO PARTICIPATE IN THIS STUDY

Thank you for your interest in this study. Since you did not agree to participate, you are ineligible to proceed.

First, we would like to ask you some questions about yourself.

Please answer each of the following questions:

What is your gender?

- ☐ Male
- ☐ Female

- ☐ Non-binary / third gender
- ☐ Prefer not to say

Which of the following categories best describes your race/ethnicity? (Please select all that apply)

- ☐ White
- ☐ Black or African American
- ☐ Hispanic or Latino
- ☐ Asian
- ☐ Other (Please specify)
- ☐ Prefer not to say

For your final year of high school (or equivalent), what type of educational setting did you attend?

- ☐ Public high school
- ☐ Private high school
- ☐ Other (Please specify)
- ☐ Prefer not to say

What is your current academic year?

- ☐ First year (Freshman)
- ☐ Second year (Sophomore)

- ☐ Third year (Junior)
- ☐ Fourth year (Senior)
- ☐ Fifth year or beyond
- ☐ Prefer not to say

What was your cumulative GPA in your first year at Middlebury? Please round to one decimal place. If you do not remember your exact first-year GPA, provide your best estimate.

On a typical week during the academic year, how many hours do you spend studying or working on assignments outside of class?

What is your primary major?

- ☐ My primary major is:
- ☐ I have not declared a major
- ☐ Prefer not to say

What is your intended primary major?

Intended Primary Major

AI Use

Next, we will ask some questions about your experiences with Artificial Intelligence tools.

There are no right or wrong answers – we are interested in understanding your honest perspectives and experiences.

Please remember that your responses are **completely anonymous** and this study has been reviewed and approved by Middlebury College's Institutional Review Board. The research team cannot link your responses to your identity in any way.

Have you ever used any form of Generative AI such as ChatGPT, GitHub Copilot, Claude, etc.?

☐ Yes

☐ No

How often do you use Generative AI tools during the academic semester?

- ☐ Never
- ☐ Rarely (a few times a semester)
- ☐ Occasionally (a few times a month)
- ☐ Frequently (a few times a week)
- ☐ Very frequently (daily or almost daily)

When did you first start using any form of Generative AI for academic purposes?

- ☐ This semester (Fall 2024)
- ☐ Last semester (Spring 2024)
- ☐ Fall 2023
- ☐ Spring 2023
- ☐ Before Spring 2023

Which of the following AI tools do you currently use on a regular basis during the academic year? (Select all that apply)

- ☐ ChatGPT (Free version)
- ☐ ChatGPT (Paid version)

- ☐ Claude (Free version)
- ☐ Claude (Paid version)
- ☐ Meta LLaMA
- ☐ Google Gemini
- ☐ Midjourney
- ☐ Microsoft Copilot
- ☐ GitHub Copilot
- ☐ Other (please specify)

For academic purposes, which of the following tasks do you typically use Generative AI for? (Select all that apply)

	Never	Rarely (few times per semester)	Occasionally (few times per month)	Frequently (few times per week)	Very frequently (daily or almost daily)
Proofreading	☐	☐	☐	☐	☐
Creating images	☐	☐	☐	☐	☐
Editing essay drafts	☐	☐	☐	☐	☐
Finding information	☐	☐	☐	☐	☐
Explaining concepts	☐	☐	☐	☐	☐
Generating ideas	☐	☐	☐	☐	☐
Coding assistance	☐	☐	☐	☐	☐
Composing emails	☐	☐	☐	☐	☐
Writing essays	☐	☐	☐	☐	☐

	Never	Rarely (few times per semester)	Occasionally (few times per month)	Frequently (few times per week)	Very frequently (daily or almost daily)
Summarizing materials	☐	☐	☐	☐	☐

How much do you currently pay per month for Generative AI subscription services?

- ☐ \$0 (I don't have an active paid subscription)
- ☐ \$0 (but I used to pay)
- ☐ Between \$1 and \$20
- ☐ Between \$21 and \$40
- ☐ More than \$40

Effects on academics and learning

We'd like to understand how Generative AI affects different aspects of your academic experience.

For each item below, please indicate whether you think AI use has a positive, negative, or has no effect.

Remember, there are no right or wrong answers - we're interested in your personal experience.

	Significantly improves	Somewhat improves	No effect	Somewhat reduces	Significantly reduces	I don't know
Your learning (e.g., your ability to grasp concepts, retain information, or learn new skills)	☐	☐	☐	☐	☐	☐
The amount of time you spend on academics (e.g., assignments, studying)	☐	☐	☐	☐	☐	☐
Your understanding of course material	☐	☐	☐	☐	☐	☐
Your grades	☐	☐	☐	☐	☐	☐
Your ability to complete assignments on time	☐	☐	☐	☐	☐	☐

Institutional policies

Next, we would like to understand your experience with Middlebury College AI policies.

Do you find the Generative AI use policy in the majority of your current classes to be clear?

- ☐ Yes: I understand when and where I'm allowed to use AI.
- ☐ No: There is no AI policy in place.

☐ No: The AI policy is unclear or vague.

If a class's AI policy says you need to cite Generative AI use, would you know how to do so?

☐ Yes

☐ No

How likely are you to use Generative AI in a class with each of the following AI policies? (Please rate from extremely unlikely to extremely likely.)

	Extremely unlikely to use AI	Unlikely	Neutral	Likely	Extremely likely to use AI
AI use is prohibited entirely.	☐	☐	☐	☐	☐
AI use is allowed if cited.	☐	☐	☐	☐	☐
AI use is unrestricted.	☐	☐	☐	☐	☐
No explicit AI policy is provided.	☐	☐	☐	☐	☐

Are you aware that you have access to the premium version of Microsoft Copilot through Middlebury College?


☐ Yes

☐ No


Others use

Next, we would like to understand your perceptions about Generative AI usage among other Middlebury students.

What fraction of Middlebury students do you think regularly use Generative AI tools? (Please provide your best estimate for each)



What percent of Middlebury students do you think regularly use Generative AI tools for classes with the following AI policies:



Factors influence AI

Please describe the factors that have personally influenced your use of Generative AI in your academic work. What initially led you to try it, what has motivated you to use it or caused you to hesitate? (Optional)

Do you have any specific feedback or suggestions about Middlebury's Generative AI policies, resources, or support

services? (Optional)



Completion

Thank you for completing our survey! Your responses have been recorded and will help us better understand how students engage with AI technology.

To enter the lottery for gift cards ranging from \$50 to \$500:

1. Click the link below to submit your email address
2. Enter the following unique code **`${e://Field/Random%20ID}`**
3. Your email submission will be collected separately from this survey.

[Submission link](#)

Important Privacy Note:

- The Google Form collecting emails is completely separate from this survey
- Your survey responses remain anonymous

- The research team cannot link emails to survey responses
- Emails will only be used for the lottery and will be deleted after the drawing

You may now close this window.

Powered by Qualtrics