Does Coursework Matter?
Uncovering the Role of Skills in the Returns to College

Yaling Xu*

April 24, 2026

Click here for the latest version.

Abstract

The continuing shift of the U.S. economy toward a high-skill base has increased the demand for college-educated workers. To understand how higher education prepares students for this evolving economy, a large body of literature in labor economics has focused on the causes and consequences of college enrollment, institutional selectivity, and major choice. Much less attention has been paid to a key dimension that shapes the skills students acquire in college—coursework. In this paper, I scrape and compile a new dataset of detailed course descriptions from Texas public universities. Using a large language model (GPT-4), I construct course-level measures of the cognitive skills, focusing on two widely taught and consistently identifiable domains: quantitative and writing skills. I then link these course-level skill measures to Texas administrative records that track students' educational histories and quarterly earnings. To estimate the returns to coursework-based skills, I implement an instrumental variables strategy that exploits variation in course offerings across cohorts within the same major. I find substantial early-career earnings returns to coursework-based quantitative skills, but no detectable returns to writing skills. These returns are especially large for underrepresented minority (URM) students and for students in less quantitatively intensive majors, suggesting that expanding access to quantitative coursework within majors may serve as a new lever for narrowing racial earnings gaps.

*Norwegian School of Economics; Department of Economics, Cornell University (email: yx558@cornell.edu).

I am very grateful for guidance and advice from Michael Lovenheim, Evan Riehl, and Nicholas Sanders. I also thank Francine Blau, Aline Bütikofer, Benjamin Goldman, Ian Lundberg, Hyewon Kim, Katrine Løken, Giulia Olivero, Zhuan Pei, Kjell Salvanes, Jason Sockin, Meredith Welch, Alexander Willén and participants at AEFP 2025, APPAM 2025 and seminars at Cornell and UCLA for helpful comments and suggestions. I thank the staff at the UT Dallas Research Center for assistance with the administrative data. The conclusions of this research do not necessarily reflect the opinions or official position of the Texas Education Agency, the Texas Higher Education Coordinating Board, Texas Workforce Commission, or the state of Texas. I acknowledge research support grant from the Cornell Population Center and the small labor grant from the Department of Economics at Cornell University. Any errors are my own.

1 Introduction

The continuing shift of the U.S. economy toward high-skill technologies has increased the relative demand for college-educated workers (Acemoglu and Autor, 2011; Autor, 2014). A large literature shows that pre-market cognitive skills, typically measured by math and verbal tests, predict both educational attainment and earnings and help explain wage inequality and racial gaps in the labor market (e.g. Aucejo and James, 2021; Murnane et al., 1995; Neal and Johnson, 1996). Most work on the role of higher education in skill formation focuses on coarse choices such as college enrollment, institutional selectivity, and major field (Lovenheim and Smith, 2023). These metrics treat what happens inside degree programs as a "black box", and measures of college-level cognitive skills remain scarce. Much less attention has been paid to a key dimension that affects which skills students actually acquire in college: coursework. Coursework is the primary channel through which colleges deliver skills, and both access to and selection among courses vary widely across students. Selective universities tend to offer more advanced, theory-intensive classes, whereas less selective institutions emphasize applied or vocational content (Biasi and Ma, 2022). Even within a given program, students choose from a broad menu of electives each term, and curricula evolve over time. These course-level decisions—shaped jointly by institutional offerings and student choice—determine the skills students ultimately acquire in college. Yet there is limited research on what skills are actually delivered and the consequences of these decisions.

The lack of research reflects three fundamental challenges in estimating the returns to college coursework and the skills it provides. First, it is difficult to identify which skills are taught in which courses. Second, few datasets link individual-level course-taking histories to both pre-collegiate academic profiles and post-college labor market outcomes. Third, course selection could be endogenous: students may choose courses based on unobserved traits such as motivation or earning potential, making it difficult to interpret correlations between coursework and earnings as causal.

In this paper, I develop new methods to identify quantitative and writing skills taught in college coursework and credibly estimate their causal impact on labor market outcomes. I first collect and compile a new dataset of detailed course descriptions from the course catalogs of Texas public universities. Using a large language model (GPT-4), I identify the skills students are likely to acquire in each course. I link these course-level skill measures to Texas administrative records that track students' educational histories and quarterly earnings. To estimate returns to coursework-based skills, I implement an instrumental variables (IV) strategy that exploits variation in course offerings across cohorts within the same program. I find large early-career earnings returns to coursework-based quantitative skills but no detectable returns to writing skills. These returns are particularly large for underrepresented minority (URM) students (Black and Hispanic) and students in less quantitatively intensive majors, suggesting that expanding access to quantitative coursework within majors may provide a promising policy lever for narrowing earnings gaps.

Specifically, I compile course descriptions from the catalogs of 27 Texas public universities from 2012 to 2020, covering more than 96 percent of all upper-division courses offered at these institutions.[1] I then identify whether each course teaches quantitative or writing skills. I focus on these two domains because they represent the core cognitive foundations of numeracy and literacy that decades of research link to educational attainment and earnings (e.g., Murnane et al., 1995, 2000; Bowles et al., 2001; Hanushek et al., 2015). They are also widely taught across majors and consistently described in catalogs. To identify these skills at scale, I use GPT-4, which is pretrained on vast textual data and can interpret nuanced context even when descriptions are brief or implicit (Wei et al., 2023; Bubeck et al., 2023). This approach is particularly useful for detecting skill content that is atypical for a course's disciplinary context or not evident from its title. I conduct multiple validation exercises to ensure the robustness and reliability of the GPT-generated labels.

After identifying skills at the course level, I link them to the Texas K–20 administrative data, creating the first panel that connects pre-collegiate information, college course-level skills, and labor market outcomes for a large statewide population. The Texas administrative data provide a uniquely rich setting to study the returns to coursework-based skills, offering detailed, large-scale information unavailable in other U.S.-based data sources (Andrews et al., 2024). First, it covers the universe of Texas public high school graduates and includes rich pre-collegiate measures—not only standard demographics but also statewide standardized test scores and high school course-taking histories that proxy for students' earning potential. Second, beyond major enrollment, it provides complete records for every course taken by every student in Texas public colleges since 2012, including course selections, credits, and grades. Third, it records quarterly earnings for all individuals covered by the state's Unemployment Insurance (UI) system. My analytic sample includes seven cohorts of students who completed college between 2014 and 2020, for whom I observe their labor market outcomes in their mid- to late twenties.

Next, I examine how students' observed characteristics correlate with course selection. Even within the same program, students who accumulate more quantitative skills from college coursework tend to have stronger math backgrounds (as measured by high school standardized math test scores), are more likely to be male, and are less likely to be Black or to have high standardized reading test scores in high school.[2] By contrast, students who accumulate more writing skills tend to have higher high school reading and math test scores, are more likely to have participated in gifted

[1] A course catalog is a list of courses offered by a postsecondary institution. Generally circulated online, it contains detailed descriptions of classes by subject, including credits awarded, rationale, key content, intended skills, and learning strategies.

[2] Throughout the text, I use the terms "school–major" and "program" interchangeably to refer to a specific four-digit Classification of Instructional Programs (CIP) major within a given institution—for example, the Biology major at UT Austin or the Economics major at Texas State University. This designation determines the set of courses available to students and the structure of the curriculum.

programs, and are less likely to be male or flagged as at risk of dropout. Although I observe extensive pre-collegiate achievement measures, it remains unclear whether OLS models yield causal estimates of the returns to coursework-based skills, as course selection may also be correlated with unobserved non-cognitive traits, preferences, and expectations.

I develop an instrumental variable that exploits term-by-term variation in course availability within the same program to address concerns about selection. Specifically, I use the share of upper-division, major-relevant courses offered within a given program that teach quantitative or writing skills as an instrument for the actual quantitative or writing skills students acquire.[3] The identifying variation arises from term-by-term fluctuations in departmental course offerings—for example, courses being added or removed, or temporary disruptions caused by faculty retirements and sabbaticals. The IV estimates should be interpreted as the return to quantitative or writing skills acquired through the major-relevant courses available to students, relative to peers in the same program, rather than the return to a common type or level of skill content across majors. The key identification assumption is that variation in offered skills affects post-graduation earnings only through the skills students acquire.

I first show that offered skills strongly predict actual skill exposure: a one standard deviation (SD) increase in the share of courses offering quantitative or writing skills increases the share of quantitative or writing courses completed by about 0.1 SD. To provide context, when a program adds about ten quantitative courses (or seven writing courses) — equivalent to a one SD increase in the share of offered courses — students take about 0.5 more quantitative or 0.4 more writing courses on average. Beyond establishing relevance, the first stage results highlight the role of course supply in shaping students' skill portfolios during college. I also conduct balance checks and find that variation in offered skills is not correlated with students' baseline characteristics.

I find economically meaningful returns to quantitative skills acquired through coursework but no detectable returns to coursework-based writing skills. A one SD increase in the share of upper-division courses a student completes that teach quantitative skills—roughly a 28 percentage-point (pp) rise—is associated with a 5.5 percent increase in quarterly earnings. This implies that, for the average student, taking one additional quantitative course raises quarterly earnings by about 1.1 percent, or roughly $150 per quarter, without changing their program. In contrast, a one SD increase in the share of courses teaching writing skills is associated with a 0.6 percent decrease in quarterly earnings, but the estimate is imprecise and statistically indistinguishable from zero.

Quantitative coursework yields larger economic returns for students who are often viewed as

[3] I focus on upper-division courses because these are where students diverge in their actual skill-building within the major, and lower-division courses are often general education or major requirements with little variation in supply. Major-relevant courses are those offered by the student's own program and related departments from which students in that program historically draw their coursework; these courses constitute the feasible course menu that students realistically face.

less quantitatively inclined. I find the strongest effects on earnings for URM students: a one SD increase in the share of quantitative courses completed raises quarterly earnings by 9.8 percent for URM students, compared to 4 percent for White students. The return for White students is not statistically significant at conventional levels, and the difference between the two groups is statistically significant. The URM–White gap in returns to quantitative skills is partly explained by the fact that URM students enter college with lower high-school math and reading achievement than peers in the same program and are more likely to be enrolled in less quantitatively intensive majors, while the remaining gap is consistent with heterogeneity in the returns to quantitative skills and may also reflect unobserved differences between URM and White students. Estimated returns to quantitative skills also are larger for women (6.5 percent), students attending less selective institutions (6.9 percent), and those whose high school math test scores are below the median of peers from the same college (7.8 percent), although the differences relative to their respective counterparts are not statistically significant.

Returns to quantitative skills are larger in majors that are less quantitatively focused. A one SD increase in quantitative skills raises quarterly earnings by 10.6 percent for students in less quantitative majors, compared to just 1.9 percent for those in more quantitatively intensive majors.[4] For example, the marginal benefit of taking one additional quantitative course is greater for a student majoring in communication than for a student majoring in engineering. This suggests that when majors already emphasize quantitative content, the marginal value of an additional quantitative course is limited. While estimated returns to coursework-based writing skills are generally negative or close to zero across subgroups, the point estimate is slightly larger for less writing-focused majors than for writing-intensive majors. Together, these patterns suggest potential skill complementarity within majors: students benefit more from acquiring skills that are underemphasized in their field of study, thereby supplementing their primary training with complementary competencies.

I show that the earnings effects primarily reflect returns to quantitative skill accumulation rather than other channels. To address concerns that changes in course offerings might reflect broader labor-market trends, I conduct an event-study analysis focusing on programs that experienced sharp and persistent increases in offered quantitative skills, often due to the introduction of new courses. I find no evidence of differential pre-trends in earnings between programs with large increases and those without. Earnings rise only after the increase in offered skills, with an estimated effect of about 6 percent—close in magnitude to the two-stage least squares (2SLS) estimate. Next, I construct an alternative instrument using only variation from temporarily closed and re-offered courses — which is less likely to reflect long-run departmental investments — and it yields qualitatively

[4] Here, "major" refers to CIP-4 fields of study (e.g., economics, biology). I classify a major as quantitatively intensive if its average share of courses teaching quantitative skills—pooled across all institutions within each CIP-4 field—is above the overall median across all fields.

similar results. I also discuss other potential channels, including peer effects, instructor effects, and signaling, and find that these mechanisms have limited ability to explain the main estimates.

The earnings effects are partly explained by increased productivity within industries and partly by improved access to higher-paying, quantitatively intensive industries, especially for URM students. When course offerings shift, students do not take more courses overall; instead, they reallocate their course mix toward more quantitative classes. GPA also does not change significantly, suggesting that the observed earnings effects are unlikely to be driven by increased academic intensity or GPA-based signaling. Students with more quantitative coursework are more likely to enter higher-paying, quantitatively intensive industries, particularly in finance and real estate, professional services, and technology sectors. These placement effects emerge earlier and are stronger for URM students, who have been historically underrepresented in these industries. Within industries, students who accumulate more quantitative skills exhibit a measurable advantage relative to their peers, and this advantage persists for at least five years after graduation, with no meaningful racial differences. Therefore, stronger quantitative preparation might enable students to take on more analytically demanding tasks, outperform peers, or move into higher-value roles within the same industry, allowing their earnings gains to persist beyond initial job placement. Taken together, the evidence suggests that expanding access to quantitative coursework, coupled with targeted academic advising, could serve as a practical and cost-effective strategy to improve early-career outcomes for URM students and reduce racial disparities, even without requiring a change of major.

This paper provides the first causal evidence on the labor market returns to college coursework and specific skills acquired from the coursework. Although coursework is a core component of higher education, its role in shaping labor market outcomes remains understudied. Previous research on the returns to college coursework often relies on metrics such as the number of credits earned or the distribution of credits across broad fields (Artz et al., 2014; Silos and Smith, 2015; Hamermesh and Donald, 2008; Light and Schreiner, 2019; Light and Wertz, 2022; Kane and Rouse, 1995). These credit-based approaches show that the quantity and type of courses students complete explain much of the variation in earnings across majors, but they cannot identify which aspects of those courses actually drive returns. In addition, these studies primarily rely on survey data or single-institution samples and use selection-on-observables designs, which do not fully address the endogeneity of course selection.[5] I address these gaps by exploiting variation in the skill content of courses offered across cohorts within the same program, showing that coursework itself—not just major or student background—plays a central role in shaping labor market outcomes. My paper also contributes to a growing literature that applies natural-language-processing tools to analyze curricular content. However, most of these studies are largely descriptive, focusing on outcomes

[5] An exception is Light and Wertz (2022), who use statewide administrative records from Ohio public universities.

such as graduation rates, aggregate earnings, or skill alignment at the institutional or major level. No prior study has linked college course content to individual earnings. For example, Biasi and Ma (2022) construct a measure of the "education–innovation gap" based on the similarity between course syllabi and frontier research and show that it predicts higher aggregate earnings using College Scorecard data. Chau et al. (2023) and Sabet et al. (2024) use word-embedding techniques to infer skill content and link it to average earnings by major, showing that similarities between syllabi and workplace tasks explain a substantial amount of early-career earnings variation, but they do not estimate returns to those skills. Moreover, existing syllabus-based samples may not be representative, as many institutions do not publish syllabi in accessible formats and not all courses are documented.[6] My study benefits from using official course-catalog descriptions—covering over 96% of courses across 27 Texas public universities—linked to population-level transcript and earnings data to examine individual labor market outcomes.

This paper contributes to research on the returns to skill by introducing a new approach to measuring skill acquisition during college. Most existing work in this area links adolescent cognitive skills—typically measured by standardized tests such as the AFQT or high school exams—to later earnings (Bowles et al., 2001; Murnane et al., 1995, 2000; Aucejo and James, 2021). Although college represents the single largest human-capital investment most individuals make after high school, no standardized assessments of general learning exist at the college level, and we lack comparable measures of skill acquisition. Prior research therefore typically relies on major choice as the proxy for college-acquired skills. I move beyond major choice and construct granular measures of quantitative and writing skills that mirror constructs in the cognitive-skills literature but reflect actual college learning. These measures also allow me to characterize majors by the skills students actually acquire through coursework.[7]

Relatedly, I show that differences in earnings—both across and within majors—are partly driven by the skills students acquire through coursework. While numerous studies document substantial variation in returns across college majors (Hamermesh and Donald, 2008; Altonji et al., 2012; Webber, 2014; Kirkeboen et al., 2016; Andrews et al., 2024), it remains unclear why some majors are more lucrative than others, and whether these differences reflect human capital accumulation or signaling. I find that these return differences are at least partly attributable to the skills embedded in coursework, not merely the credential itself. Recent work also documents considerable variation in returns within fields (Andrews et al., 2024), but offers limited explanation. I show

[6] For instance, Biasi and Ma (2022) show that the Open Syllabus Project slightly overrepresents Ivy-Plus schools.

[7] The closest related work that characterizes majors using skills is Hemelt et al. (2023), who conceptualize majors as "bundles of skills." They assess the specificity of college majors by analyzing job postings and constructing a skill vector for each major based on the share of ads listing particular skills. This approach includes 11 skills—including cognitive and writing—but captures the skills employers associate with different fields of study. However, it reflects expected or perceived skill demand rather than the actual skills embedded in the coursework students complete.

that, even within a single program, cohort-to-cohort changes in course offerings generate meaningful differences in students' skill accumulation and earnings. Moreover, prior research shows that certain majors raise earnings by helping students sort into high-paying industries and occupations (Bleemer and Mehta, 2022). My findings suggest that similar sorting mechanisms also operate within majors, depending on the specific skills students acquire.

The remainder of the paper is organized as follows. Section 2 describes the data, the construction of course-level skill measures, and presents descriptive evidence on the distribution of skills. Section 3 discusses the empirical strategy. Section 4 presents the main results, including analyses of heterogeneity and mechanisms. Section 5 concludes.

2 Data and Measures

2.1 Course Catalog Data and Skill Measurement

To obtain the course content, I scrape and compile a new dataset from course catalogs published by 27 public universities in Texas between 2012 and 2020.[8] I collect course catalogs from university websites, retrieving structured HTML or XML when available and extracting text from PDFs when necessary. For each course, I extract the course prefix, number, title, credit hours, and a one- to two-paragraph description summarizing the course rationale, topics covered, and intended learning outcomes or skills. Figure 1 and Appendix Figure A.1 show several typical catalog entries.

I restrict the sample to fall and spring courses and focus on upper-division, lecture-based offerings. Upper-division courses offer students substantial discretion in course selection and exhibit meaningful within-major variation in skill content, whereas lower-division coursework is largely determined by general-education and major requirements. Because major declarations and transfers are typically settled by the junior year, I could observe complete upper-division course histories for most students. I process the raw scraped catalogs to construct a structured dataset suitable for analysis, as described in Appendix B.1. The final dataset includes approximately 240,000 courses and covers about 96% of all upper-division courses offered at the sampled institutions between 2012 and 2020. Since most excluded courses have small enrollments and are rarely taken, the dataset captures the vast majority of courses in which students enroll. Table 1, Panel B reports institutional coverage and course counts.

The course catalog dataset offers two key advantages for studying skill acquisition through college coursework. First, it provides detailed course descriptions that are not available in stan-

[8] I begin in 2012 because Texas administrative data first record student transcripts in that year. I exclude small or newly established public institutions with minimal enrollment, as well as any college for which I could not obtain full catalog coverage during the sample period. These 27 universities account for over 91% of graduates from four-year public institutions in Texas.

dard administrative data, which typically include only subject codes and, at most, course titles.[9] For example, two courses with the same title, such as *Research Methods*, may teach very different skills. One may emphasize statistical analysis, while the other focuses on academic writing. Second, course catalogs offer relatively comprehensive and consistent coverage of undergraduate coursework across Texas public institutions. Because they are maintained by universities and updated annually, they provide course content for nearly all classes, making large-scale analysis of curricular content feasible across institutions and fields.

To classify whether a course teaches a given skill, I use GPT-4 to analyze the text of each course description. For each course, I provide GPT with the prefix, number, title, and description, along with keywords and definitions for each skill category. Figure 1 presents the GPT prompt, the skill definitions used, and the expected output for a sample course. To mitigate the inherent variability in GPT responses, I adopt a multi-query prompting strategy that reduces the influence of randomness in any single output (Wang et al., 2023; Bsharat et al., 2024). For each course, I query GPT-4 with the same prompt ten times. Each query returns a binary response, with 1 indicating that the skill is taught and 0 otherwise. A course is classified as teaching a given skill if at least 6 out of 10 responses are affirmative.[10]

Using large language models (LLMs) to analyze course content is especially useful because they can interpret context even when information is limited. Most course descriptions consist of one or two paragraphs of free text and vary widely in style across institutions and departments. Traditional computational approaches, such as bag-of-words methods or word embeddings, rely on surface-level word frequencies and may fail to capture the contextual nuances through which skills are conveyed in course descriptions. In contrast, GPT-4 is trained on large textual corpora and can interpret meaning in context, even when descriptions are indirect or discipline-specific (Wei et al., 2023, 2022; Bubeck et al., 2023). This capability allows it to identify skills that are not clearly reflected in course titles or commonly associated with a given discipline. For example, as illustrated in Appendix Figure A.1, a course titled *Political Analysis* at Texas Tech emphasizes descriptive and inferential statistics, whereas a similarly titled course at UT Austin focuses on formal political theory. Appendix Figure A.1 also shows GPT-4 identifying skills atypical of a course's field, such as quantitative skills in a social science course and writing skills in a computer science

[9]Several efforts have been made to classify postsecondary courses. The National Center for Education Statistics developed the College Course Map (CCM), a taxonomy similar to the Classification of Instructional Programs (CIP) but with finer course-level detail. Its use has been limited because NCES's Postsecondary Education Transcript Studies (PETS) assign CCM codes manually, a process that is infrequent and applied to small samples. Paulson et al. (2024) extend this framework by training a natural language processing model to automate the mapping of course metadata to CCM codes, allowing for broader coverage. However, because this approach relies solely on metadata like titles and numbers, it may miss meaningful within-field variation described in full course texts.

[10]In robustness checks, I construct a continuous measure based on the share of responses labeled as 1, which represents the estimated probability that a course teaches the skill.

course.

I implement several strategies to assess the reliability of the GPT-based skill labeling. First, I cross-validate GPT's skill labeling with human labeling using a sample of 800 course descriptions.[11] Four undergraduate students from different academic backgrounds independently label the same set of courses. I compare their labels to the GPT-generated labels and define the agreement rate as the proportion of courses where both human and model labels match. The average agreement rate is 89% for quantitative skills and 82% for writing skills, with no evidence of systematic bias across fields. Second, to address concerns that course descriptions may be too brief to fully capture course content, I collect 50 course syllabi from UT Austin and compare GPT-generated labels based on the catalog descriptions versus the full syllabi. The agreement rate—the share of courses for which the description-based and syllabus-based labels coincide—is 94% for quantitative skills and 86% for writing skills. When catalog-based labeling does not flag a skill, especially for writing, it is typically because the relevant content only appears in assignments or grading rubrics of the course syllabi. Since my goal is to identify core curricular skills rather than instructor-specific practices, course descriptions provide a more consistent and stable source, particularly given that they remain largely unchanged over time (Light, 2024).

2.2 Texas Administrative Data and Key Variables

To estimate the returns to coursework-based skills, I link course-level skill measures to individual student records in Texas administrative data. These data combine pre-collegiate characteristics, detailed college educational histories, and post-graduation labor market outcomes. They come from three sources: the Texas Education Agency (TEA), the Texas Higher Education Coordinating Board (THECB), and the Texas Workforce Commission (TWC).

The TEA data cover all students who attended Texas public high schools and provide rich pre-collegiate covariates. These include demographic characteristics (gender and race/ethnicity), eligibility for free or reduced-price lunch as a proxy for economic disadvantage, English proficiency status, and indicators for whether a student was ever flagged as at risk of dropping out or ever participated in a gifted and talented program. I measure academic preparation using scores from state-mandated mathematics and reading exams administered in ninth grade and normalize these scores within each test year. To proxy for pre-collegiate preparation and subject orientation, I control for the number of regular and Advanced Placement (AP) courses completed in middle school

[11] Based on a power analysis targeting a 95% confidence level and a ±5% margin of error—and assuming an expected agreement rate of 85%—the minimum required sample size is approximately 200 course descriptions. To ensure broad representation across academic disciplines, I expand the sample to 800 course descriptions and apply a stratified sampling strategy that guarantees at least three courses from each discipline.

and high school by subject area.[12]

The THECB data track all students enrolled in Texas public postsecondary institutions since 1992, including semester-level records on college enrollment, declared major, academic standing, and degree completion. Detailed course-taking histories are available from 2012 onward. Each course record includes the course prefix, number, credits, GPA, instruction mode (e.g., lecture, practicum, seminar), and the offering institution. Although I do not observe coursework from private or out-of-state institutions, Texas public universities enroll the vast majority of the state's college students.

From the THECB data, I construct a sample of students who graduated with a bachelor's degree from one of 27 Texas public universities between 2014 and 2020. These universities account for roughly 91% of four-year public college graduates in the state. I define a student's major as the four-digit CIP code of their degree and assign cohorts based on the calendar year in which the degree is completed.[13] I link these students to their K–12 and workforce records using unique individual identifiers. Following Andrews et al. (2024), I aggregate majors into ten broader fields for heterogeneity analysis. Appendix Table A.1 lists the detailed majors included in each field.

My main outcome variable is the log of average quarterly earnings from one to five years after college graduation. Earnings data come from the TWC, which records quarterly wages for all Texas workers covered by the state's unemployment insurance (UI) system.[14] I deflate all earnings to 2019 dollars. To account for sporadic unemployment episodes and to focus on returns for full-time work, I follow Altonji and Zhu (2025) and only keep earnings observations that (1) are part of a sequence of at least three consecutive quarters, (2) occur after graduation and outside any term of undergraduate or graduate enrollment at a Texas public institution, (3) fall between \$2,000 and \$250,000 per quarter, and (4) occur at least three quarters after degree attainment. I use North American Industry Classification System (NAICS) industry codes from TWC records to study industry sorting as a potential mediator of skill returns.

To quantify the skills each student acquires through coursework, I construct a measure of coursework-based skill intensity. For each student i and skill domain s, I first calculate the share of courses that teach that skill. Let $\mathscr{C}_i$ denote the set of upper-division, lecture-based courses taken by student i during fall and spring terms. The raw skill intensity is defined as $\text{RawSkill}_i^{(s)} =$

[12] Subject categories include English language arts, mathematics, science, social studies, foreign languages, fine arts, business education, career and technical education (CTE) and technology applications, and an "other" category. AP counts are available for the first six categories; for AP mathematics, technology applications are included.

[13] In robustness checks, I define a student's major as the major they declare before their junior year, and cohort as the year they first appear as juniors or seniors at that institution. Results are similar under this definition.

[14] TWC data do not capture self-employed workers, federal employees, or those working outside of Texas. As a result, I cannot distinguish between individuals who are unemployed, out of the labor force, or employed in settings not covered by the UI system. However, Texas has a relatively low out-migration rate, and prior research suggests that out-of-state attrition does not meaningfully bias earnings comparisons across fields (Andrews et al., 2016, 2014; Mountjoy and Hickman, 2021; Foote and Stange, 2022).

$\frac{\sum_{k \in \mathscr{C}_i} \mathbf{1}\{\text{Course } k \text{ teaches skill } s\}}{|\mathscr{C}_i|}$, where the indicator $\mathbf{1}\{\cdot\}$ equals one if course k teaches skill s according to the GPT-based classification. Because course-taking patterns differ systematically across institutions and majors, I standardize this measure within each school–major cell. This prevents high-variance programs from disproportionately influencing pooled estimates and allows the measure to be interpreted as a student's relative skill accumulation within their own program.

2.3 Descriptive Overview

2.3.1 Summary Statistics

Table 1 summarizes my sample. The sample includes seven graduating cohorts from Texas public four-year universities, comprising 239,577 students who earned a degree from one of 27 institutions. Of these students, 59 percent are female and 24 percent are economically disadvantaged. There is sizable representation among Hispanic, African American, and Asian students: 52 percent are White, 9 percent are Black, 29 percent are Hispanic, and 8 percent are Asian. As expected, these four-year students are positively selected in terms of academic preparation; 22 percent were enrolled in some form of gifted or talented program before college. The average early-career quarterly earnings for graduates in this sample is \$13,068.

I classify institutions into three tiers. Tier 1 includes the two flagship universities—UT Austin and Texas A&M—both of which are among the largest and most selective public institutions in the country. Tier 2 includes other large, moderately selective public universities: Texas Tech, University of Houston, University of North Texas, UT Dallas, UT Arlington, UT San Antonio, and UT El Paso. Tier 3 includes the remaining, less selective four-year institutions. Roughly 35 percent of students in the sample graduated from Tier 3 schools. Additionally, 26 percent of students had previously enrolled in either a two-year college or another four-year institution other than their graduation college before junior year.

On average, students complete 18 courses during their junior and senior years. Of these, 37 percent are identified as teaching quantitative skills and 24 percent as teaching writing skills. The remaining courses either focus on other skill areas not confidently classified using course descriptions or lack sufficient information to make a reliable classification.

2.3.2 Course-Taking and Skill Patterns

Quantitative and writing skills are taught across a wide range of fields, but the courses delivering these skills vary systematically by field. Table 2 provides illustrative examples of high-enrollment courses that teach quantitative or writing skills. Quantitative skills are not limited to traditional mathematics, engineering, or computer science courses. In quantitatively intensive fields, courses classified as teaching quantitative skills often involve advanced or field-specific technical material. By contrast, in less quantitatively intensive fields, quantitative skills are often taught through applied, empirical, or methods-oriented courses. Examples include *Statistics for the Behavioral*

Sciences, *Advertising Media Planning Foundations*, and *Conduct of Communication Inquiry* in the communications field, as well as *Experimental Psychology*, *Game Theoretic Methods in Political Science*, and *Quantitative Data Analysis* in the social sciences. A similar pattern appears for writing skills. In communications, social sciences, and humanities-related fields, courses classified as teaching writing skills often focus on discipline-specific analysis, theory, campaigns, literacy, or topic-based writing. By contrast, in more technical or professional fields, writing skills are often embedded in courses emphasizing professional communication, field-specific reporting, design, ethics, or law. Examples include *Professional and Technical Communication* in computer science, *Engineering Communication* and *Introduction to Engineering Design* in engineering, *Biomedical Writing* in biology and health, and *Introduction to Construction Law* in vocational fields.[15]

Figure 2, Panel A, shows the share of courses that teach quantitative and writing skills across fields. Quantitative skills are most frequently taught in mathematics, engineering, and the physical sciences, while writing skills are most prevalent in the humanities and communication fields. I then cluster selected majors based on the share of courses that teach quantitative versus writing skills. Panel B shows a negative correlation between the two: majors with a high concentration of quantitative skills, such as mathematics and engineering, tend to offer relatively few writing-intensive courses. Conversely, majors such as English, journalism, and public relations emphasize writing while offering limited quantitative training. Importantly, substantial heterogeneity exists even within broad fields. As shown in Figure 2, Panel B, and Appendix Figure A.2, for example, within the business field, finance, economics, and accounting are considerably more quantitatively intensive than marketing or general business administration. Similarly, within health-related majors, public health includes a greater share of writing-intensive coursework, whereas biology places stronger emphasis on quantitative skills. These patterns suggest that students' skill accumulation varies systematically across broad fields of study and depends critically on the specific major within each field.

3 Empirical Framework

3.1 Research Design

I estimate how the skills students acquire through their college coursework affect early-career earnings. The primary outcome of interest, Y_i, is the log of average quarterly earnings from one to five years after college graduation. The key explanatory variable is the share of upper-division courses that teach a given skill s taken by student i, standardized within school–major cells, denoted $\text{Skill}_i^{(s)}$.

[15] In my sample, within the Liberal Arts + Humanities field, most students are enrolled in Multi-/Interdisciplinary Studies majors (CIP code 30). In Texas, students preparing for K–12 teaching careers often select these majors (Deneault, 2024). Consequently, many courses in this field emphasize curriculum design and classroom practice.

The basic strategy relates student i's earnings Y_i to their realized skill exposure:

$$Y_i = \theta \, \text{Skill}_i^{(s)} + \alpha_{jm} + \mu_c + \psi^\top \mathbf{X}_i + \xi_i. \tag{1}$$

I include school–major fixed effects (α_{jm}) to absorb time-invariant differences across programs, and cohort fixed effects (μ_c) to control for graduation-year shocks. I also condition on a rich set of pre-collegiate covariates $\mathbf{X}_i$, described in the data section, to account for observable differences in academic preparation and socioeconomic status. Standard errors are clustered at the school–major level throughout the analysis. The school–major fixed effects are crucial. In Appendix Table A.3, I show that adding controls and other fixed effects beyond the school–major fixed effects does not change the estimated coefficient on $\text{Skill}_i^{(s)}$ very much. However, the R^2 of the fully controlled specification remains modest, so there is still substantial unexplained variation in earnings, and it is possible that unobserved characteristics, preferences, or constraints also affect course selection. Therefore, I use an IV approach to address remaining concerns about endogenous course selection within programs.

To address endogeneity in course selection, I construct an IV that exploits term-by-term variation in the availability of major-relevant courses that teach a given skill. The key idea is that students' exposure to specific skills depends partly on the set of courses available to them, which changes over time as new courses are added, old courses are dropped permanently, or courses are suspended temporarily, for example when faculty go on sabbatical or leave. As a result, two otherwise similar students in the same program but in different cohorts (and thus different upper-division term windows) may face different opportunities to take skill-teaching courses.

To construct the instrument, I first identify the set of courses available to students in each term, focusing on major-relevant courses. I define major-relevant courses as those offered by departments from which students in a given program historically draw a substantial share of their coursework. Appendix Figure A.3 shows the departments I classify as major-relevant for selected majors at UT Austin. For each major, the size (and label) of each bubble represents the share of courses taken by students in that major from the corresponding department. For example, finance majors at UT Austin typically take courses in the statistics and business-related departments. I focus on major-relevant courses because students complete most of their degree requirements within a consistent subset of related departments, and access to courses in other departments is often limited by prerequisites or major restrictions. This set also represents the part of the curriculum that academic programs can directly influence through course planning.

For every program in each academic term, I compute the proportion of major-relevant courses that teach a given skill. For example, I calculate what fraction of courses open to junior economics majors at UT Austin in Fall 2015 teach quantitative skills. I then average this proportion over the upper-division terms during which each student is enrolled, which serves as my instrumental variable.[16] The instrument is standardized within school–major cells, similar to the treatment variable. The endogenous treatment—the share of courses a student actually takes that teach a given skill—is therefore instrumented with the average availability of such courses in their major-relevant departments during college. Because the instrument is based on course offerings, it provides plausibly exogenous variation in exposure to each skill. For instance, if a faculty member who normally teaches a quantitative course goes on sabbatical in 2015, juniors and seniors in that year face lower exposure to quantitative skills. The detailed construction steps can be found in Appendix B.2.

Figure 3 visualizes the source of IV variation. For each major, I decompose the set of major-relevant, upper-division courses that teach the focal skill into four types: always offered, permanently closed, temporarily closed, and newly introduced. The IV exploits variation from courses that are not always offered. Across majors, many programs have non-trivial shares of such courses, indicating meaningful changes in the menu of skill-bearing offerings over time. Quantitative fields (e.g., Engineering and Physical Sciences) combine a large, stable core of always-offered quantitative courses with a sizable non-always-offered margin. In quantitatively light fields, the overall quantitative intensity is much smaller, and the quantitative content is relatively more concentrated in non-always-offered offerings. For writing, field differences are driven mainly by differences in overall writing intensity, while the composition across availability types is relatively similar across fields.

I begin with a series of reduced-form regressions to estimate the intent-to-treat effect of skill-teaching course offerings:

$$Y_i = \rho \, \text{OfferedSkill}_i^{(s)} + \alpha_{jm} + \mu_c + \delta^\top \mathbf{X}_i + \nu_i, \tag{2}$$

where $\text{OfferedSkill}_i^{(s)}$ is the share of upper-division courses available in student i's program during their junior and senior years that teach skill s, standardized within school–major cells.

My identification strategy exploits within-program, across-cohort variation in the availability of skill-teaching courses. This setup is conceptually similar to a difference-in-differences design: I compare students in the same major and institution (i.e., the same "unit") whose exposure to course supply varies across cohorts (i.e., over "time"). Some cohorts happen to be enrolled in years when more quantitative courses are offered, while others face a more limited set of options. The coefficient ρ measures how earnings change when students are exposed to more skill-intensive courses in their program. This is a policy-relevant parameter for schools or departments deciding

[16] Courses are weighted by their average class size, because students are more likely to enroll in larger classes. I use average class size to remove endogeneity arising from short-run swings in course popularity; variation in the IV therefore only reflects whether a course is offered. Within each program, courses are weighted by the historical share of that program's enrollments taken in each department—for example, economics department courses receive more weight for economics majors than for students in other majors.

whether to expand course offerings in areas emphasizing writing or quantitative skills. Specifically, a one SD increase in the share of courses teaching skill s leads to a $100 \times (e^{\rho} - 1)$ percent change in average earnings for graduates of that program.

Next, I estimate a series of 2SLS regressions. The first- and second-stage equations are:

$$\text{Skill}_i^{(s)} = \pi \, \text{OfferedSkill}_i^{(s)} + \alpha_{jm} + \mu_c + \kappa^{\top}\mathbf{X}_i + \varepsilon_i. \tag{3}$$

$$Y_i = \beta \, \widehat{\text{Skill}}_i^{(s)} + \alpha_{jm} + \mu_c + \lambda^{\top}\mathbf{X}_i + \eta_i. \tag{4}$$

The first-stage coefficient, π, captures how students adjust their actual skill accumulation through course choices when course offerings in their program emphasize skill s change. The second-stage coefficient, β, identifies the causal effect of coursework-based skill on earnings for students whose skill accumulation responds to supply-driven changes in course availability. Specifically, a one SD increase in the share of upper-division courses a student takes that teach skill s—relative to peers in the same program—increases earnings by $100 \times \left(e^{\beta} - 1\right)$ percent, on average.

It is worth noting that a course classified as teaching quantitative or writing skills does not correspond to a common course type, difficulty level, or curriculum across majors, as illustrated in Table 2. However, students do not choose from the full university course catalog. For example, when a sociology student takes quantitative courses, it is rare for them to choose *Introduction to Number Theory* from the mathematics department; instead, they are more likely to take courses such as *Quantitative Data Analysis*. The variation used in my IV therefore comes from changes in the availability of skill-bearing courses within realistic, program-specific choice sets. Thus, the estimates should be interpreted as the return to quantitative or writing skills acquired through program-related coursework available to students, relative to peers in the same program, rather than the return to a single homogeneous skill bundle across majors.

3.2 Validity of the Instrument

For identification, the instrument must satisfy the assumptions of relevance, independence, monotonicity, and the exclusion restriction (Angrist et al., 1996). The relevance assumption requires that the instrument be strongly correlated with the endogenous regressor. Column (2) of Table 3 reports results from the first-stage regression in equation 3. I find that the skills offered through available coursework are strongly predictive of actual skill accumulation: F-statistics for both writing and quantitative skills exceed 80, well above conventional thresholds (Stock and Yogo, 2005). Specifically, a one SD increase in the share of courses offered that teach quantitative or writing skills leads to approximately a 0.1 SD increase in the share of corresponding courses completed by students. To provide context: a one SD increase in share of quantitative courses roughly corresponds to adding about 10 quantitative courses to a typical program's curriculum. This shift leads students to take roughly 0.5 additional quantitative courses, on average. For writing skills, a one SD increase—equivalent to about 7 more writing-intensive courses—results in an average increase of 0.4 writing courses taken per student. In practice, the within-program, across-cohort variation in skill-relevant offerings is more modest: the average change is just 4 pp, or about 1.5 additional courses. This implies that roughly 7–8% of students respond by taking the marginal skill-teaching course.[17]

The monotonicity assumption requires that the increase in the availability of courses teaching skill s cannot cause any student to take fewer of those courses. Although this assumption cannot be tested directly, it is plausible in my context. An increase in the instrument expands (or tilts) the menu of available options: students who value skill s now have more opportunities to enroll, while students who wish to avoid it are at worst unaffected. There may be rare instances of scheduling conflicts that prevent a student from enrolling in a skill-teaching course, but such cases are likely idiosyncratic and not systematically related to the instrument. Therefore, the individual probability of taking a course teaching skill s should weakly increase. Moreover, I show in Appendix Table A.4 and Table A.5 that the first-stage is positive and statistically significant across all subsamples, further supporting the plausibility of monotonicity.

The independence assumption requires that, after controlling for fixed effects and observed covariates, the share of courses teaching skill s to which a student is exposed is unrelated to potential earnings outcomes. I assess its plausibility using the balance tests proposed by Pei et al. (2019). Table 4 reports regressions of key individual characteristics on the treatment (actual skills) in columns (1) and (3), and on the instrument (offered skills) in columns (2) and (4). Appendix Table A.6 presents results for remaining covariates, specifically the number of high school courses completed in each field. Columns (1) and (3) show that, even within the same program, students who accumulate more quantitative skills tend to have stronger math backgrounds, are more likely to be male, and are less likely to be Black or to have high reading scores in high school. By contrast, students who accumulate more writing skills tend to have higher high school reading and math test scores, are more likely to have participated in gifted programs, and are less likely to be male or to

[17] Programs offer an average of 37 upper-division courses. For quantitative skills, the SD of the offering share is 0.28, so a one SD increase corresponds to $0.28 \times 37 = 10.36$ additional quantitative courses. With a first-stage coefficient of 0.10 and an SD of the quantitative-taken share of 0.28, this implies a $0.10 \times 0.28 = 0.028$ (2.8 pp) increase in the share of quantitative courses taken, or $0.028 \times 18 = 0.50$ additional quantitative courses per student (18 courses taken on average). For writing skills, the SD of the offering share is 0.20, so a one SD increase corresponds to $0.20 \times 37 = 7.4$ additional writing-intensive courses. With a first-stage coefficient of 0.11, this implies a $0.11 \times 0.20 = 0.022$ (2.2 pp) increase in the share of writing courses taken, or $0.022 \times 18 = 0.40$ additional writing-intensive courses per student. In practice, the within-program change is about 4 pp: $0.04 \times 37 = 1.48$ courses. In SD units, this is $0.04/0.28 = 0.143$, implying $0.10 \times 0.143 = 0.0143$ SD in quantitative course-taking, i.e., $0.0143 \times 0.28 = 0.004$ (0.4 pp), or $0.004 \times 18 = 0.072$ courses per student, consistent with about 7–8% of students taking one additional course at the margin.

be flagged as at risk of dropout. However, the instruments are not correlated with observed student characteristics: almost all coefficients in columns (2) and (4) are small, and only two characteristics are significant at the 10% level. The joint F-statistics for the full set of covariates are less than 1.3 for both skills, indicating that variation in offered skills—driven by course availability—is unrelated to student characteristics. To assess overall balance, I use all covariates to predict earnings and construct a "predicted earnings" variable. This predicted value, shown at the bottom of Table 4, is uncorrelated with the instrument, further suggesting that the instrument is balanced with respect to observed student characteristics.

The exclusion restriction requires that changes in the skills available through coursework affect post-college earnings only through the skills students actually accumulate. This is a strong assumption and cannot be tested directly. In Section 4.5, I examine several potential violations. First, I test whether within-program variation in offered skills proxies for time-varying departmental inputs or broader program changes. Second, I assess whether course supply responds to broader shifts, such as rising student demand or anticipated labor market trends, that could independently affect student outcomes. Third, I evaluate whether peer and network channels, instructor quality, or signaling could affect the interpretation. Fourth, I assess whether other skill content could drive the results. Taken together, the evidence suggests that course availability primarily operates by shifting students' quantitative skill accumulation, and that the associated earnings effects largely reflect these shifts.

4 Results

4.1 Main Results

The OLS estimates in Table 3, Column (1), show a positive association between coursework-based quantitative skills and early-career log quarterly earnings, and a negative association between coursework-based writing skills and earnings. Specifically, a one SD increase in the share of courses teaching quantitative skills is associated with approximately a 1.5% increase in quarterly earnings, while a one SD increase in coursework-based writing skills is associated with a 0.6% decrease.

The reduced-form results in Column (4) indicate that expanding quantitative course offerings within a given program leads to average earnings gains, while expanding access to writing courses has no meaningful impact on average earnings. For a given program, a one SD increase in the share of courses offered that teach quantitative skills is associated with a 0.5% increase in average quarterly earnings for students in that program. A one SD increase in quantitative skills corresponds to a 28 pp rise in the availability of quantitative courses—about 10 additional courses out of an average of 37 upper-division offerings. In practical terms, this implies that adding a single quantitative course within a program (approximately a 2.7 pp increase in course share) is associated with a 0.05% increase in average earnings.

The 2SLS estimates in Column (5) provide a causal link between coursework-based skills and earnings. For the average student, a one SD increase in the share of upper-division courses that teach quantitative skills—roughly a 28 pp increase—is associated with a 5.5% increase in quarterly earnings. Since students take an average of 18 upper-division courses, a one SD increase corresponds to approximately $0.28 \times 18 = 5$ additional quantitative courses. This implies that taking one additional quantitative course raises quarterly earnings by about 1.1%, or roughly $150 per quarter. Conversely, a one SD increase in writing-intensive coursework is associated with a 0.6% decrease in earnings, though the estimate is small and statistically indistinguishable from zero.

I benchmark the magnitude of these returns against prior work on earnings differences across majors and show that coursework-based skill accumulation explains a non-trivial and economically plausible share of those gaps. In my sample, Business and Economics majors complete about 50% of their coursework in quantitative subjects, compared to 15% for Social Science majors—a 35 percentage-point difference, or 1.25 standard deviations (the SD of quantitative course share is 0.28). If a Social Science major adopted the quantitative course mix typical of Business or Economics without changing majors, my estimates imply that their earnings would increase by $0.055 \times \left(\frac{0.35}{0.28}\right) \approx 0.07$ log points, or about 7%. This gain explains roughly one-quarter of the observed Business–Social Science earnings gap in my sample (0.27 log points). This magnitude is comparable to findings in Andrews et al. (2024), who use similar field categories and estimate that majoring in Economics or Finance rather than a Social Science raises quarterly earnings by $3,334 (about 0.399 log points) six to ten years after high school. Likewise, Bleemer and Mehta (2021) show that students who major in Economics earn about 46% higher annual earnings than those who pursue second-choice majors (typically in the social sciences). My estimated 7% return from adjusting skill accumulation within the major accounts for roughly 15% of that economics major premium. Benchmarks using other major comparisons produce similar results, with coursework-based skill differences explaining on the order of 10–30% of observed earnings gaps across fields.

4.2 Heterogeneity

Coursework-based quantitative skills yield larger economic returns for students traditionally considered less quantitatively inclined. Figure 4, Panel A and Appendix Table A.4 report the effects of coursework-based quantitative skills on earnings by race, gender, prior academic achievement, institutional selectivity, family income, and major type. I find the strongest effects for URM students: a one SD increase in the share of courses with quantitative skills increases quarterly earnings by 9.8 percent for Black and Hispanic students. The estimated effect for White students is 4 percent, but it is not statistically significant at conventional levels. For Asian students, the coefficient is small and not statistically distinguishable from zero.[18]

Women benefit more than men: a one SD increase in coursework-based quantitative skills is associated with a statistically significant 6.5 percent increase in quarterly earnings for women, compared to a 3.5 percent increase for men, which is not statistically significant at conventional levels.

Returns also vary by prior academic achievement and institution selectivity. For students not identified as gifted during K–12, a one SD increase in coursework-based quantitative skills leads to a 7 percent increase in quarterly earnings. In contrast, the estimated effect for gifted students is close to zero. Additionally, students with lower prior math preparation—defined as having high school math scores below the median among peers who attended the same college—experience substantial and statistically significant earnings returns (7.8 percent), whereas the effect for students with higher math scores is much smaller (3.5 percent) and not statistically significant. Students at less-selective Tier 3 institutions gain 6.9 percent per one SD increase in quantitative skills, while students at more selective Tier 1/2 institutions experience smaller and statistically insignificant effects (3.6 percent). One exception arises when examining family income, proxied by eligibility for free or reduced-price lunch. Here, I find earnings returns for non-economically disadvantaged students are 5.3 percent; although the returns for economically disadvantaged students are 6.7 percent, they are not statistically significant, and the difference between these two groups is not statistically significant.

Returns to coursework-based quantitative skills are larger in majors that are not traditionally quantitative. Within less quantitatively intensive majors, when students increase their quantitative skills by one SD, their quarterly earnings rise by 10.6 percent.[19] In contrast, in more quantitatively intensive majors, a one SD increase in quantitative skills is associated with a 1.9 percent increase in earnings, which is not statistically distinguishable from zero.[20] I further illustrate this pattern in Figure 5, which plots the average share of courses with quantitative skills in each field against the corresponding 2SLS earnings estimates. While not all field-level estimates are statistically significant due to limited power, a consistent pattern emerges: within majors where students generally take fewer quantitative courses (e.g., social sciences), taking more quantitative coursework yields larger earnings gains, whereas in majors with already high baseline quantitative exposure (e.g.,

[18] The first-stage relationship is weaker for Asians (coefficient = 0.065, F-statistic = 11.2), and their sample size is smaller (N ≈ 19,000). Therefore, I include Asian students in the main sample but compare White versus Black/Hispanic students, where the effects are more precisely identified.

[19] I define quant-light majors as those whose average student quantitative skill accumulation falls below the median across all majors in the sample, using CIP-4 classifications pooled across institutions. Similarly, writing-light majors are defined as those below the median in average writing skill accumulation. Appendix Table A.2 provides detailed information on major classifications.

[20] To formally test subgroup differences, I estimate a fully saturated version of Equations (3) and (4), interacting all regressors with subgroup indicators. Results presented in Appendix Table A.7 confirm that returns to coursework-based quantitative skills are statistically significantly larger for URM students relative to White students, for non-gifted students relative to gifted peers, and for students in quant-light majors relative to those in quant-heavy majors.

engineering), the incremental returns to additional quantitative coursework are more limited.[21]

Returns to coursework-based writing skills are generally negative or close to zero across subgroups and not statistically significant at conventional levels, as reported in Figure 4, Panel B and Appendix Table A.5. However, the point estimates are generally larger for groups traditionally viewed as less writing-inclined. Specifically, returns per SD increase in writing skills are higher for men (0.1 percent) than for women (–1.3 percent), for economically disadvantaged students (3 percent) than for non-disadvantaged students (–1.9 percent), and for students with lower reading test scores (2.4 percent) than for students with higher reading scores (–3.7 percent). While field-level estimates are not statistically different from zero, returns for students in writing-light majors (0.2 percent) are larger than for those in writing-heavy majors (–0.6 percent).

The results have two key implications. First, accumulating more quantitative skills through coursework even within a student's existing program translates into substantial earnings gains, especially for URM students. This suggests that expanding access to quantitative courses within majors may offer a practical and potentially low-cost strategy for state higher-education systems and university leaders aiming to reduce earnings disparities. Second, the findings provide suggestive evidence of skill complementarity within majors: students appear to benefit more from acquiring skills that are underemphasized in their primary field of study. This is particularly important given the paper's focus on upper-division courses, as the core curriculum might not be sufficient to provide the full breadth of skills that benefit students in the labor market. Thus, broadening skill sets through complementary coursework can improve labor-market outcomes.

4.3 Complier Characteristics and Interpretation of the IV Estimates

The IV estimates capture a group-specific local average treatment effect (LATE) for compliers—students whose acquired skill changes when offered skill changes. To interpret the findings, I examine how the identifying variation is distributed across predetermined characteristics. Since both my treatment D_i (acquired skill) and instrument Z_i (offered skill) are continuous, I adapt the complier characterization in Frandsen et al. (2023) to this setting. Specifically, for a predetermined characteristic X_i, I estimate a 2SLS regression in which the dependent variable is the interaction $X_i D_i$, instrumenting D_i with Z_i and controlling for school–major fixed effects and cohort fixed effects. The resulting estimand can be written as $\hat{\delta} = \frac{\mathrm{Cov}(Z_i, X_i D_i)}{\mathrm{Cov}(Z_i, D_i)}$, where the covariances are computed after partialling out the same fixed effects used in the main specification. This estimand is a responsiveness-weighted average of the pre-treatment characteristic X_i, placing greater weight on

[21] A potential concern is that one SD change in quantitative skill may correspond to different numbers of quantitative courses across fields. Thus, field heterogeneity in SD units could partly reflect differences in scaling rather than differences in marginal returns. I rescale the field-specific estimates into approximate per-course units using each field's average upper-division course load and raw quantitative-share dispersion. The qualitative pattern is unchanged, with larger implied effects in quantitatively light fields.

students whose D_i is more sensitive to changes in Z_i.[22] In Appendix Table A.12, I report both the raw population share of each characteristic X in group g (i.e., $\Pr(X = 1 \mid g)$ for indicator X, where g is the full sample, White, or URM) and the responsiveness-weighted share $\Pr_\omega(X = 1 \mid g)$. If the weighted share exceeds the raw share, that characteristic is over-represented among compliers relative to the baseline population in group g; if it is below the raw share, it is under-represented.

I examine the complier characteristics for quantitative skill in the pooled sample. First, the identifying variation is disproportionately generated by White students: 62.2% of the identifying variation comes from White students, relative to 52.0% in the full population, while 33.6% comes from URM students, relative to 38.5% in the full population. Second, compliers are slightly less likely to be economically disadvantaged and slightly more likely to be in Tier-3 schools, although the magnitudes are small, and gender differences are negligible. Third, less identifying variation comes from quant-light majors. Fourth, within programs, the distribution of within-program high-school math and reading ranks among compliers is close to the baseline distribution, implying that, within a program, students with different pre-collegiate academic preparation respond similarly to offered quantitative coursework. Together, these patterns suggest that the pooled IV estimate is not primarily driven by some of the highest-return margins, particularly URM students and students in quant-light majors. The pooled estimate may therefore understate the returns to policies that expand quantitative coursework specifically in those higher-return contexts.

Appendix Table A.12, Panel B, reports the complier characterization for writing skill. Similar to quantitative skill, the identifying variation is disproportionately generated by White students, is more concentrated in Tier-3 schools, and is more likely to come from female students. Within programs, students with different math and reading preparation respond similarly to changes in offered writing coursework.

Implications for the URM–White earnings-return gap. I find that returns to quantitative skills are higher for URM students than for White students. Differences in estimated effects across race groups could reflect either (1) heterogeneity in treatment effects or (2) differences in the composition of compliers across groups. I therefore examine compliers' characteristics across race groups. I first compare the responsiveness-weighted share to the baseline population share within each race group g. The within-group complier–baseline gaps are similar for URM and White samples, suggesting that the URM estimate is not larger simply because the instrument moves a different type of student within each race group (e.g., if compliers within the URM group were concentrated among low-math-preparation students while compliers within the White group were concentrated among high-

[22]With continuous D and Z, $\hat{\delta}$ can be interpreted as a responsiveness-weighted average of X_i, $\hat{\delta} = \mathbb{E}[\omega_i X_i]/\mathbb{E}[\omega_i]$, where ω_i measures each student's first-stage responsiveness (e.g., $\omega_i = \partial D_i/\partial Z_i$ in a local linear approximation). For indicator X, this is the corresponding responsiveness-weighted share of $X = 1$ in group g, which I denote by $\Pr_\omega(X = 1 \mid g) \equiv \mathbb{E}[\omega_i X_i \mid g]/\mathbb{E}[\omega_i \mid g]$. In the binary LATE setting, this weighted share coincides with $\Pr(X = 1 \mid \text{complier}, g)$.

math-preparation students).

However, URM compliers still differ from White compliers because the baseline composition differs across races. URM students are more heavily concentrated in the lowest within-program quartile of high-school math scores (30.5%) than White students (22.4%), and correspondingly have a higher complier share (32.3% vs. 24.1%).[23] Likewise, URM students are more likely to be in quant-light majors in the baseline population (60.1% vs. 53.3% for Whites), and this difference persists among compliers (51.1% vs. 40.4%). URM students are also more likely to be female in the baseline population (62.8% vs. 56.7% for Whites), with a similar pattern among compliers (59.8% vs. 55.3%). These baseline differences in composition place URM compliers disproportionately in contexts where returns to quantitative skills are higher and can therefore account for part of the URM–White return gap.

To assess whether these composition differences can explain the full racial gap, I reweight URM compliers to match the pre-collegiate characteristics of White compliers, following the propensity-score approach in Dodini et al. (2025). The results are reported in Appendix Table A.13. The reweighted estimates remain statistically significant and are modestly smaller than the baseline estimates. The attenuation is primarily driven by reweighting on within-program high-school math and reading preparation, which reduces the URM–White gap but does not eliminate it. Therefore, the URM–White gap in returns to quantitative skills is partly explained by the fact that URM students enter college with lower math and reading preparation relative to peers in the same program and are more likely to be enrolled in quant-light majors. The remaining difference is consistent with heterogeneity in the returns to quantitative coursework or with unobserved differences between URM and White students.

4.4 Mechanisms

In this section, I show that the earnings effects of quantitative skills acquired through coursework are partly driven by enhanced productivity within industries, and partly by access to higher-paying, quantitatively intensive industries, especially for URM students.

Before analyzing the mechanisms behind the earnings effects, it is important to understand what the estimates identify. Is the increase in quantitative skills driven by a heavier overall course load, or by a reallocation of existing courses toward more quantitative ones? I examine how students' course-taking behaviors change during college. Table 5, Panel A, shows that a one SD increase in the share of courses teaching quantitative skills translates to approximately 2.8 additional upper-division quantitative courses—a 40% increase over the baseline mean of 7 courses. In contrast, the total number of upper-division courses taken remains unchanged across racial groups, except for

[23] Consistent with this, Appendix Table A.11 shows that, within a school–major–cohort, URM students have lower high-school math and reading test scores than White students.

Asian students.[24] Similarly, I find no significant changes in GPA. These results indicate that the estimated returns are not driven by increased overall workload, academic performance, or GPA-based signaling, but rather by a reallocation of coursework. For URM students, increased quantitative coursework is associated with a reduction in in-state public graduate school enrollment. One possible explanation is that higher early-career earnings among URM students increase the opportunity cost of immediate graduate study. Alternatively, the decline may reflect out-of-state, private, or delayed graduate program enrollment; thus, we should be cautious in interpreting this pattern as a reduction in educational investment per se.

The earnings effects are partly driven by early-career placement into higher-paying industries where quantitative skills are particularly rewarded, especially for URM students. While the Texas administrative data do not include detailed occupational classifications, they report industries of employment at the quarterly level. Using these data, I assess whether quantitative skills increase students' likelihood of working in specific industries. Results reported in Table 5, Panel C, indicate that quantitative skills increase the probability of employment in Finance, Insurance, Real Estate (FIRE) and Professional, Scientific, and Technical (PST) industries.[25] A one SD increase in quantitative skills raises the likelihood of working in FIRE/PST by 4.1 pp in Years 1–2 (a 14.9% increase over the 0.275 baseline) and 5.3 pp in Years 3–5 (a 17.2% increase over the 0.308 baseline). I find no meaningful shifts in most other industries (Appendix Table A.14). For URM graduates, this placement effect appears earlier: a one SD increase in quantitative skills raises FIRE/PST employment by 5.6 pp in Years 1–2 and is statistically significant at the 10% level, which represents a 22.6% increase over the 0.248 baseline. In Years 3–5, effects are statistically significant at the 10% level for both URM and White graduates (6.6 pp and 3.8 pp, respectively). The point estimate is larger for URM students, but differences are not statistically significant.[26]

To further characterize destination industries, I construct an industry-level quantitative intensity index, capturing how quantitatively skilled the typical worker is in each industry.[27] A one SD

[24] The total number of upper-division courses significantly increases for Asian students, as they have a higher baseline share of quantitative courses (46 percent versus 34 to 38 percent for other groups), making them more likely to respond by adding courses.

[25] Defined using NAICS industry codes: (i) Finance and Insurance (NAICS 52), (ii) Real Estate and Rental and Leasing (NAICS 53), and (iii) Professional, Scientific, and Technical Services (NAICS 54). An indicator equals one if any UI-covered quarter in the window is spent in these industries. I refer to this grouping as "FIRE/PST"

[26] The later FIRE/PST placement effect for White graduates likely reflects a mediated response to quantitative coursework—via (i) delayed labor-market entry (e.g., attending non-TX public graduate schools, self-employment, or out-of-state jobs), consistent with the small employment dip in Years 1–2, and (ii) post-entry industry switching. When restricting the sample to White graduates with positive Texas UI earnings in Years 1–2 (i.e., early entrants), the Years 3–5 FIRE/PST effect remains positive and statistically significant.

[27] For each industry g and time window t, I compute the average quantitative coursework share among all other graduates working in that industry: $\text{QuantIndex}_{gt} = \frac{1}{N_{gt}-1}\sum_{j\in g,t,j\neq i}\text{QuantSkill}_j$. I assign each individual to the NAICS 4-digit industry in which they have at least three quarters of earnings and the highest average earnings within the window.

increase in coursework-based quantitative skills increases the industry quant index by 1.5 pp for URM students in Years 1–2, about 4.0% of their 0.378 baseline, but this estimate is only statistically significant at the 10% level. I find no comparable effect for White students. The means of the dependent variable rows show that URM students often have lower average earnings and are underrepresented in higher-paying industries and in industries with more quantitatively skilled workforces.

The earnings effects are also partly driven by increased productivity within the industry. I construct a "within-industry quantitative-skill advantage " index, defined as the difference between a student's quantitative-skill level and the average among employees in the same industry.[28] A one SD increase in quantitative skills raises this relative advantage by 13.3 pp in Years 1–2 and 12.4 pp in Years 3–5. I find no meaningful racial heterogeneity in this within-industry advantage. This indicates that additional quantitative coursework may allow college graduates to undertake more analytically demanding tasks or outperform peers even within the same industry.

Quantitative skills acquired through college coursework help close early-career racial gaps in labor market outcomes. Within two years of college graduation, at baseline, URM graduates earn about 13% less than their White peers ($11,020 vs. $12,641), work slightly fewer quarters (6.38 vs. 6.48 out of 8), are less likely to enter FIRE/PST industries (24.8% vs. 28.3%), and are employed in industries with lower average quantitative intensity (37.8% vs. 40.3%). Quantitative coursework helps mitigate these disparities: for URM students, a one SD increase in quantitative skills is associated with a 9.7% increase in quarterly earnings, a 5.6 pp increase in FIRE/PST entry (approximately 22.6% relative to the 24.8% baseline), and a 1.5 pp increase in average industry quantitative-skill level (about 4.0% relative). It also raises the number of working quarters by 0.36 (roughly 6% increase relative to baseline). For White students, the corresponding earnings and employment coefficients are small and statistically indistinguishable from zero, though FIRE/PST entry increases in Years 3–5 by 3.8 pp. In the third to fifth years following college graduation, these gains largely persist: the URM earnings effect remains at 7.5%, and the FIRE/PST entry effect is 6.6 pp. Therefore, expanding access to quantitative coursework within the major could be a low-cost strategy to reduce baseline URM-White disparities in both earnings and representation, particularly in the first few years after graduation.

How does quantitative coursework enhance initial job placement for URM students? One possibility is that quantitative coursework may mitigate initial employer skepticism by demonstrating relevant technical skills. Employers often face greater uncertainty about productivity for minority candidates and, as a result, tend to rely more heavily on observable credentials such as education

[28] I implement a leave-one-out strategy. The within-industry quantitative advantage is calculated as the difference between an individual's quantitative skill and the average quantitative skill of graduates in the same industry from all except themselves. $\text{QuantAdv}_{it} = \text{QuantSkill}_{it} - \text{QuantIndex}_{g(i)t}$.

when evaluating Black applicants (e.g., Lang and Manove (2011)). Additionally, coursework on transcripts might help URM applicants pass initial employer screening processes that traditionally rely on informal signals—such as professional networks, internships, or recommendations—which minority students typically have less access to. Finally, improved quantitative skills could also alter URM students' preferences, increasing their likelihood of applying to quantitative roles or industries.

4.5 Alternative Channels and Robustness

4.5.1 Addressing Alternative Channels

In this section, I show that a range of alternative channels have limited power to explain the estimated earnings effects. Overall, the evidence is most consistent with an interpretation in which course availability primarily operates by shifting students' quantitative skill accumulation, and the associated earnings effects largely reflect these shifts.

Labor market demand. To address concerns that course offerings may respond endogenously to labor market demand, I conduct an event-study analysis focusing on programs that experienced sharp and persistent increases (at least 50 pp) in the share of courses teaching quantitative skills, typically due to the introduction of new quantitative courses. I define the event year as the first academic year in which these new courses become available to students. Each student is assigned an event time based on their graduation year relative to the event: students who graduate in the event year itself (i.e., seniors with limited exposure to the new courses) are considered partially treated and coded as event time 0. Students who graduate the year after the event are the fully treated cohort (event time ≥ 1). Students in earlier cohorts (event time ≤ -1), who completed all coursework before the curricular change, are considered untreated. I estimate a two-way fixed effects model for post-graduation earnings, controlling for program and cohort fixed effects. I find no evidence of differential pre-trends in earnings between programs with large increases in quantitative course offerings and those without, but I do find rising post-treatment earnings gains for cohorts graduating after the year of the curricular change, with an estimated effect of around 6 percent—close in magnitude to the 2SLS estimates. Figure 6 plots the estimated event-time coefficients, and Appendix Table A.10 reports the full regression results. To assess robustness to treatment-effect heterogeneity, I also implement alternative estimators proposed by Callaway and Sant'Anna (2021), Sun and Abraham (2021) and Borusyak et al. (2024). Appendix Figure A.4 overlays these three estimators alongside the TWFE estimator; all show a clear and consistent increase in earnings after the event, reinforcing the robustness of the results.

Departmental inputs and curriculum strategy. Another concern is that variation in the instrument might be driven by broader changes in departmental resources or strategy, such as expansions in faculty, classroom space, or advising that directly affect earnings independent of skill content. I first aggregate the data to the program-by-term level and regress average class size on the instrument. Table 6, Panel A shows that an increase in the share of courses teaching quantitative skills does not significantly change average class size within the program, suggesting that the instructional capacity of the program does not change much.

Second, I exploit the fact that a substantial share of within-program variation arises from courses that were temporarily closed and later re-offered. These changes account for about one-third of all within-program changes in offered skills and are less likely to be systematically correlated with long-run departmental investments or labor-market expectations. I construct an alternative instrument based only on this "temporary closure" variation. Table 6, Panel A, Column (2) presents the reduced-form results for quantitative skills, and Column (3) presents the 2SLS estimates. As expected, the first-stage coefficient is smaller, with an F-statistic of 18, but still above the conventional weak-instrument threshold (Stock and Yogo, 2005). Notably, the reduced-form estimate remains around 0.005, similar to the main specification, indicating that a one SD increase in the share of offered courses teaching quantitative skills raises cohort-level earnings by approximately 0.5 percent, regardless of whether the course-supply change is temporary or permanent. The 2SLS estimate rises to 0.12 and is statistically significant at the 10-percent level. This likely reflects a higher marginal return among students whose course-taking responds to short-run availability (compliers). Overall, these results show that even temporary fluctuations in course offerings can meaningfully affect student outcomes and help alleviate concerns that the main findings are driven fully by long-run shifts in departmental inputs or labor market anticipation.

Peer effects and networks. The instrument might change not only which courses students take, but also with whom they take them. If cohorts exposed to more quantitative courses systematically end up in classes with "stronger" peers, earnings could rise through peer effects rather than through skill accumulation. To examine this, I proxy peer quality using the mean high school math score and lower-division GPA of classmates.[29] Table 6, Panel B shows no systematic relationship between offered skills and peer composition in terms of observable preparation. I then add the peer controls and re-estimate the 2SLS model in Table 6, Panel C, Columns (1) and (2); the estimated return to quantitative coursework-based skills changes very little. Another concern is that curriculum changes could alter student composition—for example, if stronger students sort into programs when quantitative offerings expand. To assess this, I conduct an event-study analysis of peer quality for programs experiencing sharp and persistent increases in offered quantitative skills. Appendix Figure A.5 shows that peer quality is flat in the years leading up to the curriculum change and

[29]For each upper-division course section, I compute the mean high school math score and lower-division GPA of all other students in the section. I then average these section-level peer measures across each student's upper-division schedule to obtain that student's average exposure to peer quality.

remains stable afterward, with no evidence of differential pre-trends or post-event shifts. Moreover, to assess whether networking operates across cohorts—for example, if employers update beliefs after interviewing a "treated" cohort, or if seniors pass job opportunities to juniors—I construct a lagged version of the instrument based on offered skills for prior cohorts in the same program. When I instrument actual quantitative skills with lagged offered skills, the returns to quantitative skills are statistically indistinguishable from zero (Table 6, Panel C, Column (3)).

Other skills and bundled content. Courses labeled as "quantitative" or "writing" may also teach other valuable skills—such as computer, financial, or other skills that are not captured by my measures. If the instrument shifted these other dimensions rather than quantitative or writing skills per se, the IV estimates could be picking up a broader bundle of content. To explore this possibility, I estimate 2SLS models with multiple endogenous skill measures and multiple instruments, including quantitative, writing, financial, and computer skills. Although the financial and computer skill measures are not perfectly measured by GPT, the results in Table 6, Panel D show that the estimated return to quantitative skills, conditional on writing, financial, and computer exposure, is close to the baseline estimate. The coefficients on the financial and computer skills are imprecisely estimated once quantitative skills are included, suggesting that the earnings effects are mainly driven by students accumulating more quantitative skills.

Instructor effects. Another concern is that quantitative courses are systematically taught by instructors with higher value added. Unfortunately, I do not yet have instructor identifiers, so I cannot directly estimate this channel. However, to overturn a skill-based interpretation, one would have to argue that instructors teaching quantitative courses systematically raise earnings through channels other than improving students' skills, for example, by providing networking or mentoring that are independent of skills. Several patterns make this interpretation less compelling. First, the estimated returns to quantitative coursework-based skills are largest in less selective institutions, where systematic assignment of "star" instructors to particular classes and dense faculty–employer networks are less common. Second, existing work on college instructor value added typically finds modest effects on student outcomes, even in settings with quasi-experimental variation in instructor assignment (e.g., Carrell and West (2010)). These findings suggest that a purely instructor-driven story would have to be unusually strong to match the magnitudes observed here. If anything, to the extent that stronger instructors are more likely to be assigned to quantitative courses, this would reinforce the interpretation that higher returns arise because students acquire more of the skills.

Signaling. One might also worry that the estimated effects reflect signaling rather than human capital. For example, additional quantitative courses might help students by changing what appears on the transcript and improving interview chances, even if students do not actually learn more. While it is not possible to fully rule out this channel, it is difficult to reconcile with a purely signaling-based story. First, the earnings gains persist for several years after graduation and do not collapse after the initial job match; if the effects were driven only by signaling, they would be expected to fade as employers learn about workers' true productivity over time (Altonji and Pierret, 2001). Second, the returns are largest at less selective institutions, where employers are less likely to have targeted recruiting pipelines tied to specific course sequences. Third, even if additional courses help students secure interviews, job retention and career progression typically depend on actual competence.

4.5.2 Other Robustness Checks

In this section, I implement a series of robustness checks to address remaining concerns related to identification and to assess the robustness of the estimated effects.

Redefining cohort and major. In the main specification, I define major fixed effects using the degree major at completion and construct the instrument based on the courses offered in that final major. Although most students declare a major before taking upper-division coursework, some switch majors during junior or senior years. As a result, the instrument may reflect student decisions that are correlated with unobserved earnings potential, potentially threatening the independence assumption. To address this concern, I (i) redefine the major fixed effects using the major declared before the junior year, (ii) redefine the cohort fixed effect as the calendar year in which the student first attains junior standing, and (iii) reconstruct the course-availability instrument using the set of courses linked to the pre-junior major, ensuring that each course is mapped to the major the student intended at the time of enrollment. Column (1) of Table 7 shows that the estimated return to quantitative skills falls only slightly—from 0.055 to 0.048—and remains statistically significant at the 10-percent level. This modest change suggests that any bias arising from post-treatment major switching or graduation-timing decisions is limited.

Controlling for school-specific linear time trends. The baseline specification includes school fixed effects to account for time-invariant characteristics of institutions. However, unobserved factors such as gradual improvements in institutional quality may cause earnings trajectories to evolve differently across schools. To address this concern, I include school-specific linear time trends in the main specification. Column (2) of Table 7 shows that the estimates remain stable.

Alternative skill measures. LLMs such as GPT may produce variable responses to the same prompt due to randomness in their output generation. Prior research shows that using a multi-query prompting strategy, followed by averaging the responses or selecting the most consistent ones, can reduce this variability and improve labeling accuracy (Wang et al., 2023; Bsharat et al., 2024). In the main specification, I classify a course as teaching a given skill if at least 6 out of 10 GPT prompts return a positive label. As a robustness check, I use an alternative approach that treats the proportion of positive responses as a continuous measure of the probability that a course teaches the skill. I also construct an alternative version of the skill exposure measure that uses credit-hour weights rather than simple course counts, so that a four-credit course contributes more to a student's skill profile than a one-credit course. Columns (3) and (4) of Table 7 show that the results are robust using both alternative skill measures.

Excluding graduate-school attendees. Because the baseline outcome averages quarterly earnings from one to five years after graduation, some students in the sample may enroll in graduate school and then return to the labor market. To ensure that the results are not driven by this group, I exclude all individuals who are ever recorded in a Texas public graduate program. Column (5) of Table 7 shows that the 2SLS coefficient on coursework-based quantitative skills falls only modestly—from 0.055 to 0.040—yet remains statistically significant at the 10-percent level.

5 Conclusion

Rapid technological change and shifting economic landscapes have raised the demand for high-skilled college workers (Acemoglu and Autor, 2011; Autor et al., 2003). Yet we still know little about how such skills are taught and accumulated in higher education. Coursework is the primary channel through which colleges deliver knowledge and students acquire skills, but measuring skills at scale and credibly identifying their labor-market returns is difficult, as course content is rarely observed and course choices are endogenous. In this paper, I address these challenges by extracting skill content from course descriptions using LLMs and linking those skill measures to individual-level transcripts and post-college earnings. I then use an instrumental variables strategy that exploits within-program, across-cohort variation in the availability of skill-bearing courses to examine whether and how coursework-based skills explain variation in earnings within and across majors.

I find that coursework-based quantitative skills lead to economically meaningful early-career earnings gains, even conditional on major. In contrast, I do not detect average returns to writing skills acquired through coursework. Quantitative skills particularly benefit students who are traditionally viewed as less quantitatively inclined: marginal returns are statistically higher for students in less quantitative majors and larger for URM students relative to white peers.

The earnings effects operate through two channels. Students with more quantitative coursework are more likely to enter higher-paying, quantitatively intensive industries, particularly in finance and real estate, professional services, and technology sectors. These placement effects appear immediately after college and are stronger for URM students, suggesting that quantitative training may help reduce disparities in economic outcomes. I also find a persistent within-industry quantitative-skill advantage, implying that additional quantitative coursework enables students to take on more analytically demanding tasks or outperform peers in similar roles.

These findings have several policy implications. First, curriculum design may benefit from paying attention not only to majors but also to the distribution and accessibility of high-return quantitative content within majors. Changes in the availability of courses teaching certain skills significantly affect the skills students actually accumulate and have measurable impacts on labor-market outcomes. Importantly, "quantitative" skills extend beyond traditional math courses like calculus. Therefore, embedding applied quantitative modules in coursework, especially in quantitatively light fields or at less-selective institutions, can improve returns without requiring students to switch majors. Moreover, the returns to quantitative coursework are particularly large for URM students, women, and those with lower pre-collegiate math preparation. This suggests that, in addition to expanding course supply, universities should adopt advising strategies and reduce enrollment frictions to nudge these students toward quantitative coursework and help close gaps in economic outcomes. Lastly, this paper introduces a practical catalog-skill framework for identifying course-level skill exposure. Such an approach can serve as an audit tool for institutions to track skill development and align curricula with evolving labor-market demand.

This paper suggests several directions for future research. The data currently cover only early-career (mid-20s) outcomes. Future work using longer-run earnings trajectories, more granular occupational detail, or qualitative surveys could help reveal how quantitative skills shape productivity, task allocation, and career progression. Additionally, while this paper focuses on two fundamental domains—quantitative and writing—the framework can be extended to a broader set of skill categories as richer course content and better classification tools become available. Understanding and shaping the granular skill content of college education is critical for informing education policy, guiding institutional design, and improving equity and efficiency in postsecondary education.

References

Acemoglu, D. and Autor, D. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. In Card, D. and Ashenfelter, O., editors, *Handbook of Labor Economics*, volume 4, pages 1043–1171. Elsevier.

Altonji, J. G., Blom, E., and Meghir, C. (2012). Heterogeneity in Human Capital Investments: High School Curriculum, College Major, and Careers. *Annual Review of Economics*, 4(Volume 4, 2012):185–223.

Altonji, J. G. and Pierret, C. R. (2001). Employer Learning and Statistical Discrimination. *The Quarterly Journal of Economics*, 116(1):313–350.

Altonji, J. G. and Zhu, Z. (2025). Returns to Specific Graduate Degrees: Estimates Using Texas Administrative Records.

Andrews, R., Li, J., and Lovenheim, M. F. (2014). Heterogeneous paths through college: Detailed patterns and relationships with graduation and earnings. *Economics of Education Review*, 42:93–108.

Andrews, R. J., Imberman, S. A., Lovenheim, M. F., and Stange, K. (2024). The Returns to College Major Choice: Average and Distributional Effects, Career Trajectories, and Earnings Variability. *The Review of Economics and Statistics*, pages 1–45.

Andrews, R. J., Li, J., and Lovenheim, M. F. (2016). Quantile Treatment Effects of College Quality on Earnings. *Journal of Human Resources*, 51(1):200–238.

Angrist, J. D., Imbens, G. W., and Rubin, D. B. (1996). Identification of Causal Effects Using Instrumental Variables: Rejoinder. *Journal of the American Statistical Association*, 91(434):468–472.

Artz, G. M., Kimle, K. L., and Orazem, P. F. (2014). Does the Jack of All Trades Hold the Winning Hand? Comparing the Role of Specialized versus General Skills in the Returns to an Agricultural Degree. *American Journal of Agricultural Economics*, 96(1):193–212. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1093/ajae/aat063.

Aucejo, E. and James, J. (2021). The Path to College Education: The Role of Math and Verbal Skills. *Journal of Political Economy*, 129(10):2905–2946.

Autor, D. H. (2014). Skills, education, and the rise of earnings inequality among the “other 99 percent”. *Science*, 344(6186):843–851.

Autor, D. H., Levy, F., and Murnane, R. J. (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, 118(4):1279–1333.

Biasi, B. and Ma, S. (2022). The Education-Innovation Gap.

Bleemer, Z. and Mehta, A. (2022). Will Studying Economics Make You Rich? A Regression Discontinuity Analysis of the Returns to College Major. *American Economic Journal: Applied*

Economics, 14(2):1–22.

Bleemer, Z. and Mehta, A. S. (2021). College Major Restrictions and Student Stratification. *SSRN Electronic Journal*.

Borusyak, K., Jaravel, X., and Spiess, J. (2024). Revisiting event-study designs: Robust and efficient estimation. *The Review of Economic Studies*, 91(6):3253–3285.

Bowles, S., Gintis, H., and Osborne, M. (2001). The Determinants of Earnings: A Behavioral Approach. *Journal of Economic Literature*, 39(4):1137–1176.

Bsharat, S. M., Myrzakhan, A., and Shen, Z. (2024). Principled Instructions Are All You Need for Questioning LLaMA-1/2, GPT-3.5/4. arXiv:2312.16171 [cs].

Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y. T., Li, Y., Lundberg, S., Nori, H., Palangi, H., Ribeiro, M. T., and Zhang, Y. (2023). Sparks of Artificial General Intelligence: Early experiments with GPT-4. arXiv:2303.12712 [cs].

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-Differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Carrell, S. E. and West, J. E. (2010). Does professor quality matter? evidence from random assignment of students to professors. *Journal of Political Economy*, 118(3):409–432.

Chau, H., Bana, S. H., Bouvier, B., and Frank, M. R. (2023). Connecting higher education to workplace activities and earnings. *PLOS ONE*, 18(3):e0282323.

Deneault, C. (2024). Local Labor Markets and Selection into the Teaching Profession.

Dodini, S., Willén, A., and Zhu, J. L. (2025). Union Membership and Native-Immigrant Labor Market Gaps.

Foote, A. and Stange, K. M. (2022). Attrition from Administrative Data: Problems and Solutions with an Application to Postsecondary Education. Technical Report w30232, National Bureau of Economic Research.

Frandsen, B., Lefgren, L., and Leslie, E. (2023). Judging Judge Fixed Effects. *American Economic Review*, 113(1):253–277.

Hamermesh, D. S. and Donald, S. G. (2008). The effect of college curriculum on earnings: An affinity identifier for non-ignorable non-response bias. *Journal of Econometrics*, 144(2):479–491.

Hanushek, E. A., Schwerdt, G., Wiederhold, S., and Woessmann, L. (2015). Returns to skills around the world: Evidence from PIAAC. *European Economic Review*, 73:103–130.

Hemelt, S. W., Hershbein, B., Martin, S., and Stange, K. M. (2023). College majors and skills: Evidence from the universe of online job ads. *Labour Economics*, 85:102429.

Kane, T. J. and Rouse, C. E. (1995). Labor-Market Returns to Two- and Four-Year College. *The American Economic Review*, 85(3):600–614.

Kirkeboen, L. J., Leuven, E., and Mogstad, M. (2016). Field of Study, Earnings, and Self-Selection. *The Quarterly Journal of Economics*, 131(3):1057–1111.

Lang, K. and Manove, M. (2011). Education and Labor Market Discrimination. *American Economic Review*, 101(4):1467–1496.

Light, A. and Schreiner, S. (2019). College major, college coursework, and post-college wages. *Economics of Education Review*, 73:101935.

Light, A. and Wertz, S. S. (2022). Should English majors take computer science courses? Labor market benefits of the occupational specificity of major and nonmajor college credits. *Economics of Education Review*, 88:102263.

Light, J. (2024). Student Demand and the Supply of College Courses.

Lovenheim, M. and Smith, J. (2023). Returns to different postsecondary investments: Institution type, academic programs, and credentials. In *Handbook of the Economics of Education*, volume 6, pages 187–318. Elsevier.

Mountjoy, J. and Hickman, B. R. (2021). The Returns to College(s): Relative Value-Added and Match Effects in Higher Education. Technical Report w29276, National Bureau of Economic Research.

Murnane, R. J., Willett, J. B., Duhaldeborde, Y., and Tyler, J. H. (2000). How important are the cognitive skills of teenagers in predicting subsequent earnings? *Journal of Policy Analysis and Management*, 19(4):547–568. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/1520-6688%28200023%2919%3A4%3C547%3A%3AAID-PAM2%3E3.0.CO%3B2-%23.

Murnane, R. J., Willett, J. B., and Levy, F. (1995). The Growing Importance of Cognitive Skills in Wage Determination. *The Review of Economics and Statistics*, 77(2):251–266.

Neal, D. A. and Johnson, W. R. (1996). The Role of Premarket Factors in Black-White Wage Differences. *Journal of Political Economy*, 104(5):869–895.

Paulson, A., Stange, K., and Flaster, A. (2024). Classifying Courses at Scale: a Text as Data Approach to Characterizing Student Course-Taking Trends with Administrative Transcripts.

Pei, Z., Pischke, J.-S., and Schwandt, H. (2019). Poorly Measured Confounders are More Useful on the Left than on the Right. *Journal of Business & Economic Statistics*, 37(2):205–216.

Sabet, A. J., Bana, S. H., Yu, R., and Frank, M. R. (2024). Course-Skill Atlas: A national longitudinal dataset of skills taught in U.S. higher education curricula. arXiv:2404.13163 [cs, econ, q-fin].

Silos, P. and Smith, E. (2015). Human capital portfolios. *Review of Economic Dynamics*, 18(3):635–652.

Stock, J. H. and Yogo, M. (2005). Testing for Weak Instruments in Linear IV Regression. Technical Working Paper 0284, National Bureau of Economic Research.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.

Wang, X., Wei, J., Schuurmans, D., Le, Q., Chi, E., Narang, S., Chowdhery, A., and Zhou, D. (2023). Self-Consistency Improves Chain of Thought Reasoning in Language Models. arXiv:2203.11171 [cs].

Webber, D. A. (2014). The lifetime earnings premia of different majors: Correcting for selection based on cognitive, noncognitive, and unobserved factors. *Labour Economics*, 28:14–23.

Wei, J., Tay, Y., Bommasani, R., Raffel, C., Zoph, B., Borgeaud, S., Yogatama, D., Bosma, M., Zhou, D., Metzler, D., Chi, E. H., Hashimoto, T., Vinyals, O., Liang, P., Dean, J., and Fedus, W. (2022). Emergent Abilities of Large Language Models. arXiv:2206.07682 [cs].

Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q., and Zhou, D. (2023). Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. arXiv:2201.11903 [cs].

Tables

Table 1: Summary Statistics

Panel A: Student-level Characteristics			**Panel B: Institutional Coverage of Course Catalog**			
Variable	Mean	SD	Institution	Courses Covered in Catalog Data	Total Courses Offered	Coverage Rate (%)
Covariates:			*Tier 1:*			
Male	0.413	0.492	UT-Austin	18195	19309	94.23%
White	0.520	0.500	TAMU	22513	23087	97.51%
Black	0.094	0.291				
Hispanic	0.291	0.454	*Tier 2:*			
Asian	0.080	0.271	TX Tech	13519	13935	97.01%
Econ. Disadvantaged	0.240	0.427	Houston	16649	17291	96.29%
At Risk for Dropout	0.136	0.343	North TX	17024	17588	96.79%
Gifted Program	0.221	0.415	UT-Dallas	7302	7732	94.44%
Limited English Proficiency	0.012	0.172	UT-Arlington	12230	12699	96.31%
Math Test Score	0.838	0.667	UT-San Antonio	11673	12016	97.15%
Reading Test Score	0.661	0.506	UT-El Paso	10252	10585	96.85%
Skill Measures:			*Tier 3:*			
Share of Courses w/ Quant Skills	0.369	0.279	UT-Tyler	5267	5519	95.43%
Share of Courses w/ Writing Skills	0.243	0.198	Midwestern State	5343	5519	96.81%
			Angelo State	4966	5210	95.32%
Earnings Outcomes:			TAMU-International	4188	4540	92.25%
Log Quarterly Earnings	9.367	0.482	TAMU-Commerce	6103	6899	88.46%
Quarterly Earnings	13068.25	6417.19	Prairie View A&M	4527	4753	95.25%
			UT-Permian Basin	3079	3159	97.47%
Other Student Characteristics:			Houston-Downtown	5855	6111	95.81%
# Upper-division Courses	17.698	5.374	TAMU-Central TX	2545	2628	96.84%
Tier 1 School	0.265	0.442	Houston-Victoria	2643	2783	94.97%
Tier 2 School	0.385	0.487	TAMU-Corpus Christi	6151	6456	95.28%
Tier 3 School	0.350	0.477	Sam Houston State	9903	10180	97.28%
			Stephen F. Austin	7640	7905	96.65%
			TX State-San Marcos	15289	15756	97.04%
			TAMU-Kingsville	5623	5888	95.50%
			Lamar	5365	5677	94.50%
			Tarleton State	7225	7428	97.27%
			West TX A&M	6685	6935	96.40%
Number of Students	239577		Total	237754	247588	96.03%

Notes: Panel B includes all upper-division courses offered by the listed institutions from 2012 to 2020. The sample in Panel A consists of students who graduated from a Texas public four-year university listed in Panel B between 2014 and 2020. "Economically Disadvantaged" is defined as eligibility for free or reduced-price lunch in high school. Math and reading test scores are based on 9th-grade state assessments and are standardized to have a mean of zero and a standard deviation of one within the test year. Skill measures report the raw mean and SD of the share of courses an individual student took that taught quantitative or writing skills. Earnings are measured as average quarterly earnings 1–5 years after graduation, in 2019 dollars.

Table 2: Top-Enrolled Courses Teaching Quantitative and Writing Skills by Field

Field	**Courses with Quantitative Skills**	**Courses with Writing Skills**
Agriculture + Natural Resources	Economic Analysis for Agribusiness Management	Ethics in Agribusiness and Agricultural Economics
Agriculture + Natural Resources	Marketing Agricultural and Food Products	Food and Agricultural Sales
Agriculture + Natural Resources	Comprehensive Genetics	U.S. Environmental Regulations
Communications	Advertising Media Planning Foundations	Integrated Communications Management
Communications	Statistics for the Behavioral Sciences	Advanced Studies in Advertising
Communications	Conduct of Communication Inquiry	Integrated Communications Campaigns
CS and Information Sciences	Algorithms and Complexity	Professional and Technical Communication
CS and Information Sciences	Principles of Computer Systems	Contemporary Issues in Computer Science
CS and Information Sciences	Design and Analysis of Algorithms	Data Base Management Systems
Vocational	Exercise Physiology	Intro to Technical Writing
Vocational	Biomechanics	Contemporary Legal Issues in Law Enforcement
Vocational	Statistical Methods	Introduction to Construction Law
Engineering + Architecture	Dynamics and Vibrations	Engineering Communication
Engineering + Architecture	Engineering Analysis for Mechanical Engineers	Intermediate Design
Engineering + Architecture	Materials and Manufacturing Selection in Design	Introduction to Engineering Design
Liberal Arts + Humanities	General Science	Assessing Literacy: Early Childhood Through Grade Six
Liberal Arts + Humanities	Mathematics in the Elementary School	Literacy Instruction for Early Childhood Through Grade Six
Liberal Arts + Humanities	Problem Solving in Mathematics	The Teaching of Reading
Biology + Health	Comprehensive Biochemistry II	Biomedical Writing
Biology + Health	Evolution	Biomedical Explorations through Narrative
Biology + Health	Biostatistics	Pharmacological Basis for Nursing
Physical Sciences + Math	Introduction to Number Theory	Principles of Geological Writing
Physical Sciences + Math	Probability I	Introduction to Mathematical Reasoning and Proof
Physical Sciences + Math	Applied Statistics	History of Mathematics
Social Sciences	Experimental Psychology	History and Theory
Social Sciences	Game Theoretic Methods in Political Science	Topics in American Government and Politics
Social Sciences	Quantitative Data Analysis	Experimental Psychology
Business + Economics	Business Finance	Strategic Management and Business Policy
Business + Economics	Marketing	Business Law I
Business + Economics	Accounting and Financial Information Systems	Business Law and Ethics

Notes: The table lists the top enrolled courses within each field that teach either writing or quantitative skills. Courses are offered by departments within the respective field and are selected based on enrollment by students majoring in that field.

Table 3: Effects of Coursework-Based Skills on Log Quarterly Earnings: OLS, First Stage, Reduced Form, and 2SLS Estimates

	(1) OLS	(2) First Stage	(3) First Stage F-stats	(4) Reduced Form	(5) 2SLS	(6) Share of Courses w/ Skill (mean, [SD])
Quant Skill	0.015***	0.099***	82.281	0.005**	0.055**	36.9%
	(0.002)	(0.011)		(0.002)	(0.022)	[27.9%]
Writing Skill	-0.006***	0.111***	99.321	-0.001	-0.006	24.3%
	(0.002)	(0.011)		(0.002)	(0.019)	[19.8%]
Individual Controls	Yes	Yes		Yes	Yes	
Cohort FE	Yes	Yes		Yes	Yes	
School × Major FE	Yes	Yes		Yes	Yes	
N	239577	239577	239577	239577	239577	239577

Notes: Column (1) reports OLS estimates of the association between acquired skills and log quarterly earnings from estimating Equation 1. Column (2) reports first-stage coefficients of the *offered* skill on the *acquired* skill from estimating Equation 3; column (3) reports the Kleibergen-Paap Wald F statistic. Column (4) reports reduced-form coefficients of offered skills on earnings from estimating Equation 2. Column (5) reports 2SLS estimates from estimating Equation 4. Column (6) reports raw means (and standard deviations in brackets) of the share of courses taken by students with the relevant skill. Treatments and instruments are standardized (mean 0, SD 1) within school–major (program), so coefficients in columns (1)–(5) represent per–standard-deviation effects. All regressions include individual controls (as described in the data section), school-by-major fixed effects, and cohort fixed effects. Parentheses contain standard errors clustered at the school-by-major level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Balance Test

Dep. var.	(1) Acquired Quant Skill	(2) Offered Quant Skill	(3) Acquired Writing Skill	(4) Offered Writing Skill
White	0.002	-0.000	0.002	0.003
	(0.002)	(0.002)	(0.002)	(0.002)
Black	-0.002*	-0.002	-0.002	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)
Hispanic	0.000	0.000	-0.001	0.001
	(0.001)	(0.002)	(0.002)	(0.002)
Asian	0.001	0.000	0.001	-0.002*
	(0.001)	(0.001)	(0.001)	(0.001)
Male	0.031***	0.002	-0.014**	0.002
	(0.002)	(0.002)	(0.005)	(0.002)
High School Math Test	0.011**	0.000	0.012***	-0.001
	(0.004)	(0.003)	(0.002)	(0.003)
High School Reading Test	-0.006***	-0.010*	0.012***	-0.007
	(0.002)	(0.005)	(0.002)	(0.004)
Gifted	-0.000	-0.001	0.006***	0.001
	(0.001)	(0.002)	(0.001)	(0.002)
At risk	-0.000	-0.001	-0.003***	-0.000
	(0.001)	(0.002)	(0.001)	(0.001)
Economically Disadvantaged	-0.002*	-0.002	-0.001	0.001
	(0.001)	(0.002)	(0.001)	(0.002)
English Proficiency	-0.000	-0.001	-0.000	-0.001
	(0.000)	(0.001)	(0.000)	(0.001)
Predicted Earnings	0.008***	0.001	-0.002*	0.000
	(0.001)	(0.001)	(0.001)	(0.001)
F statistic: All coefficients zero	32.789	1.100	38.167	1.248
Cohort FE	Yes	Yes	Yes	Yes
School × Major FE	Yes	Yes	Yes	Yes
Observations	239577	239577	239577	239577

Notes: This table examines how individual characteristics relate to the treatments (acquired skills) in columns (1) and (3), and to the instruments (offered skills) in columns (2) and (4). For each row characteristic X_i, I estimate the regression $X_i = \beta \cdot T_i$ or $Z_i + \alpha_{j,m} + \mu_c + \varepsilon_i$, where X_i is a single covariate, T_i denotes acquired (actual) skills, Z_i denotes offered skills, $\alpha_{j,m}$ denotes school-by-major fixed effects, and μ_c denotes cohort fixed effects. The covariate set includes all variables shown in the table, as well as high school course-taking measures by field; coefficients on the latter are omitted here for brevity and reported in Appendix Table A.6. I also report the F-statistic for the joint significance of all coefficients on the covariates. "Predicted earnings" are predicted using the full set of covariates in the main specification. Parentheses contain standard errors clustered at the school-by-major level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Effects of Coursework-based Quantitative Skills on Academic and Labor Market Outcomes

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.	All	White	URM			
Panel A. Academic Outcomes						
# of Upper-division Courses	1.018	0.487	0.419			
	(0.682)	(0.514)	(0.790)			
Dependent mean	17.698	17.435	17.909			
# of Upper-division Courses w/ Quant Skills	2.849***	2.652***	2.343***			
	(0.406)	(0.299)	(0.385)			
Dependent mean	7.013	7.016	6.624			
GPA	0.007	-0.004	-0.015			
	(0.037)	(0.034)	(0.051)			
Dependent mean	3.080	3.134	2.985			
Attending TX Grad School	-0.010	0.023	-0.076**			
	(0.024)	(0.022)	(0.038)			
Dependent mean	0.209	0.197	0.227			
N	239577	124450	91989			
		Years 1–2			Years 3–5	
	All	White	URM	All	White	URM
Panel B. Earnings and Employment						
Ln(Avg. Qtr. Earnings)	0.039*	0.013	0.097***	0.037*	0.022	0.075**
	(0.020)	(0.024)	(0.035)	(0.021)	(0.025)	(0.035)
Dependent mean	9.387	9.441	9.304	9.460	9.510	9.376
N	224433	116899	86771	221963	115436	85573
# of Quarters Employed	0.027	-0.149	0.364*	0.213	0.013	0.616*
	(0.119)	(0.134)	(0.195)	(0.180)	(0.175)	(0.331)
Dependent mean	6.379	6.479	6.375	10.149	10.315	10.078
N	234922	121954	90427	158443	85000	58630
Panel C. Destination-Industry Characteristics						
Prob. of Ever Working in FIRE/PST Industry	0.041**	0.033	0.056*	0.053***	0.038*	0.066*
	(0.019)	(0.022)	(0.033)	(0.019)	(0.021)	(0.034)
Dependent mean	0.275	0.283	0.248	0.308	0.313	0.284
Industry Average Quant Skill	0.002	-0.005	0.015*	0.005	-0.000	0.011
	(0.006)	(0.007)	(0.009)	(0.006)	(0.006)	(0.009)
Dependent mean	0.395	0.403	0.378	0.387	0.396	0.369
Panel D. Within-Industry Quant Skill Advantage						
Quant Skill Relative to the Industry Average	0.133***	0.135***	0.131***	0.124***	0.129***	0.118***
	(0.007)	(0.008)	(0.010)	(0.007)	(0.008)	(0.010)
Dependent mean	0.011	0.009	0.004	0.016	0.014	0.007
N	224310	116841	86721	221815	115372	85496

Notes: Each cell reports 2SLS estimates of the effect of quantitative skills on the listed outcome. Rows labeled "Dependent mean" report the sample mean of each outcome. Panel A presents academic outcomes: the total number of upper-division courses completed, the number of upper-division quantitative courses, cumulative college GPA at graduation, and an indicator for enrollment in a graduate program at a Texas public university (equal to 1 if any such record is observed; 0 otherwise). Panels B–D present labor market outcomes measured in two post-BA windows: Years 1–2 and Years 3–5. Ln(Avg. Qtr. Earnings) refers to the log of average quarterly UI-covered earnings in the respective window, following the sample restrictions described in Section 2. Employment is measured by the number of quarters employed during each period. FIRE/PST industry equals 1 if the individual worked in any quarter within Finance & Insurance (NAICS 52), Real Estate and Rental and Leasing (NAICS 53), or Professional, Scientific & Technical Services (NAICS 54) during the window. For industry-level quantitative measures, each graduate is assigned to the NAICS 4-digit industry in which they had at least three quarters of employment and the highest average earnings within the window. Industry average quant skill is defined as the (leave-one-out) mean quantitative-coursework share among other graduates working in the same NAICS-4 industry during the window. Within-industry quant advantage is calculated as the individual's quantitative coursework share minus the corresponding industry average. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Evaluating Alternative Channels to the Quantitative Skill Channel

	(1)	(2)	(3)
Panel A: Departmental inputs and temporary-closure IV			
Dep. var.	Avg class size	Log earnings	Log earnings
Specification	*program-term regression*	*IV from temporary-closure courses*	*IV from temporary-closure courses*
Offered Quant Skill	0.017	0.005*	
	(0.064)	(0.003)	
Quant Skill			0.117*
			(0.068)
First Stage			0.046***
			(0.011)
First stage F-statistic			17.814
N (program-terms / students)	25596	239577	239577
Panel B: Effect of Offered Quant Skill on peer composition			
Dep. var.	Peer HS math	Peer lower-div GPA	
Offered Quant Skill	-0.002	-0.003	
	(0.002)	(0.004)	
N	239503	239295	
Panel C: IV with peer controls and lagged IV			
Dep. var.	Log earnings	Log earnings	Log earnings
Specification	*+ peer HS math control*	*+ peer lower-div GPA control*	*Lagged-IV spec*
Quant Skill	0.053***	0.057***	0.028
	(0.021)	(0.022)	(0.028)
N	239503	239295	239069
Panel D: Multiple-skill IV estimates			
Dep. var.	Log earnings		
Specification	*Multi-skill IV*		
Quant Skill	0.069**		
	(0.033)		
Writing Skill	-0.031		
	(0.030)		
Computer Skill	-0.035		
	(0.025)		
Financial Skill	0.062		
	(0.048)		
N	232943		

Notes: Each column reports coefficients from a separate regression. Panel A uses program–term data in column (1) to regress average class size on Offered Quant Skill, and student-level data in columns (2)–(3) to report reduced-form and 2SLS estimates of log quarterly earnings using Offered Quant Skill from temporary-closure courses as an instrument for Quant Skill. In Panel A, the temporary-closure instrument is constructed from courses that temporarily disappear and later reappear in the schedule. Panel B uses student-level data to regress peer composition measures—Peer HS math and Peer lower-division GPA, defined as leave-one-out averages of classmates' high school math scores and lower-division GPA across a student's upper-division courses—on Offered Quant Skill. Panel C reports 2SLS estimates of the effect of Quant Skill on log quarterly earnings: columns (1)–(2) use Offered Quant Skill as the instrument and add peer controls (Peer HS math or Peer lower-division GPA), while column (3) uses Lagged Offered Quant Skill, constructed from the previous cohort's offered-skill exposure in the same program, as an alternative instrument. Panel D reports 2SLS estimates from a multiple-skill specification in which quant skill, writing skill, computer skill, and financial skill are treated as endogenous and are instrumented by their corresponding offered-skill measures. All student-level specifications include school–major and cohort fixed effects and the baseline set of individual covariates described in data section. Program-term-level specification includes school–major and term fixed effects. Standard errors, in parentheses, are clustered at the school–major level (student-level regressions) and at the program level (program–term regressions). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Robustness Checks for Estimates of Quantitative Skills

	(1) Redefine Major and Cohort	(2) Add School-specific Linear Trends	(3) Probabilistic Skill Measure	(4) Credit-Weighted Skill Measure	(5) Excluding Grad-school Attendees
First Stage	0.089***	0.092***	0.092***	0.094***	0.097***
	(0.010)	(0.010)	(0.011)	(0.012)	(0.012)
First Stage F-stats	83.836	79.242	76.167	66.627	71.593
Reduced Form	0.004*	0.006***	0.005**	0.005**	0.004*
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
2SLS	0.048*	0.061**	0.053**	0.054**	0.040*
	(0.026)	(0.025)	(0.023)	(0.022)	(0.023)
Individual controls	Yes	Yes	Yes	Yes	Yes
Cohort FE		Yes	Yes	Yes	Yes
School × Major FE		Yes	Yes	Yes	Yes
School × Major (Jr. entry) FE	Yes				
Year of Jr. Entry FE	Yes				
School Lin. Trends		Yes			
N	224493	239577	239577	239577	189443

Notes: Each column reports a variant of the baseline specification with log quarterly earnings as the dependent variable. Column (1) redefines the fixed effects and instrument using information prior to junior year: school×major (pre-junior) fixed effects and year-of-junior-entry cohort fixed effects; the instrument is constructed using variation in courses offered within the student's pre-junior major. Column (2) augments the baseline with school-specific linear time trends. Column (3) replaces the binary course-skill indicator with a probabilistic measure (the share of GPT prompts indicating that a course teaches the skill). Column (4) weights course exposures by credit hours. Column (5) excludes students who ever attend a Texas public graduate program. Treatments and instruments are standardized (mean 0, SD 1) within program. Standard errors (in parentheses) are clustered at the corresponding school×major level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figures

Figure 1: GPT Prompt Example and Skill Definitions

Prompt Shot #1: Skill Definitions and Instructions

You are an educational expert. Below are the JSON definitions of two skills:

```
{
  "Mathematical or Quantitative": "Understand, manipulate, and apply numerical and
mathematical concepts. Relies on numerical data to make decisions, analyze situations, and
predict outcomes, utilizing structured methodologies from quantitative disciplines. Example
skills include: Problem Solving, Research, Analytical Thinking, Statistical Reasoning,
Decision-making, Math, and Statistics.",

  "Writing": "Composing, drafting, and editing books, papers, reports, releases, scripts, and
other text-based documents. Example skills include: Writing, Written Communication, Editing,
Synthesis, and Reporting."
}
```

Based on the provided list of skills and their definitions, determine whether each skill is likely to be taught in the course described below. For each skill, output 1 if it is likely taught, and 0 if it is not or if there is insufficient information. Use the following JSON format:

```
{"Mathematical or Quantitative": 0 or 1, "Writing": 0 or 1}
```

Prompt Shot #2: Course Inputs

Example Course:

```
{

  “ECON 3338 Economics of Crime (3 semester credit hours): This course empirically analyzes
the causes and consequences of crime and the criminal justice system using the tools of
economics. A main focus of the course is reading and discussing empirical research papers in
the economics of crime literature. Topics include the effects of incarceration on the
incarcerated, the impacts of policies intended to deter crime or reduce recidivism, racial
disparities in the criminal justice system, and the causes and consequences of domestic
violence. Major projects include creation of a data portfolio examining one of several
sources of national crime data using tables, graphs, and statistical relationships and a
group presentation on a major episode or issue in U.S. crime policy.”

}
```

GPT Responses

Example Course Output:

```
{"Mathematical or Quantitative": 1, "Writing": 1}
```

Notes: This figure shows the skill definitions and the JSON prompt used to label course descriptions with quantitative and writing skills. For each course, GPT-4 is queried 10 times; a course is coded as teaching a given skill if at least 6 of 10 responses are positive.

Figure 2: Skill Distribution Across Fields/Majors

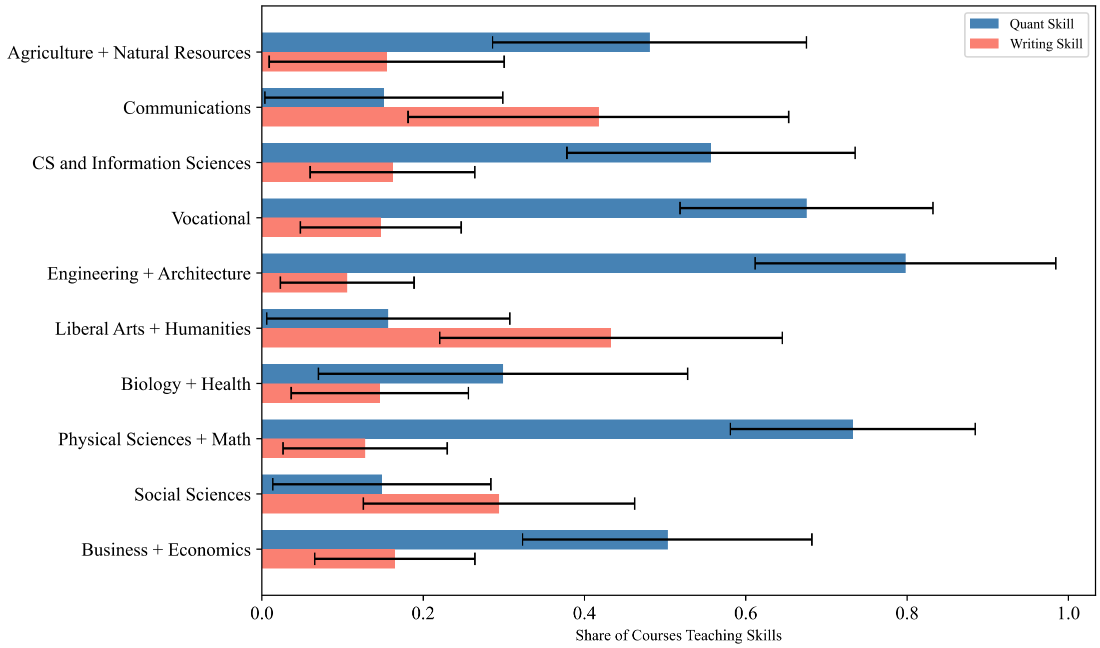

Panel A.

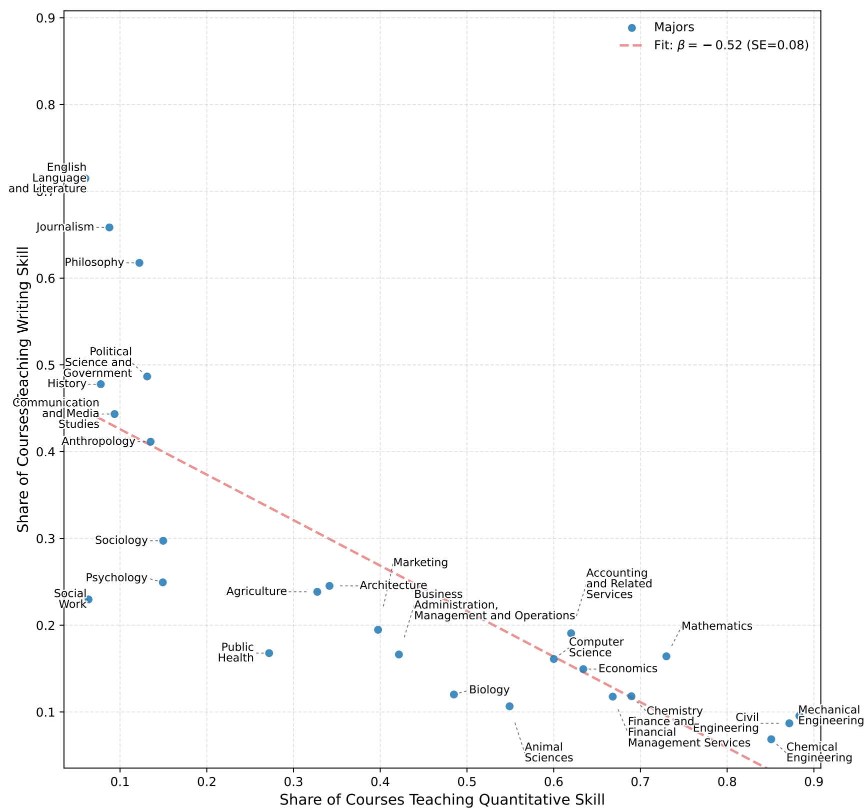

Panel B.

Notes: Panel A shows the average share of courses teaching quantitative and writing skills across fields. Error bars represent standard deviations. Panel B plots selected CIP-4 majors by their mean share of courses teaching quantitative (x-axis) and writing (y-axis) skills.

Figure 3: Breakdown of IV-Relevant Course Offerings by Field

Notes: This figure decomposes major-relevant upper-division courses within each selected major into four categories based on their availability over time: *always offered*, *permanently closed*, *temporarily closed* (courses removed and later reintroduced), and *newly introduced*. Instruments are constructed from variation in courses that are not always offered. Colored bars represent the share of skill-relevant courses in each category (e.g., # of quant courses that are newly introduced / total # of courses in the major). The grey bar represents the share of courses in the major that do not teach the given skill.

Figure 4: Effect of Coursework-based Skills on Log Quarterly Earnings by Subgroup

Panel A. Quantitative Skills

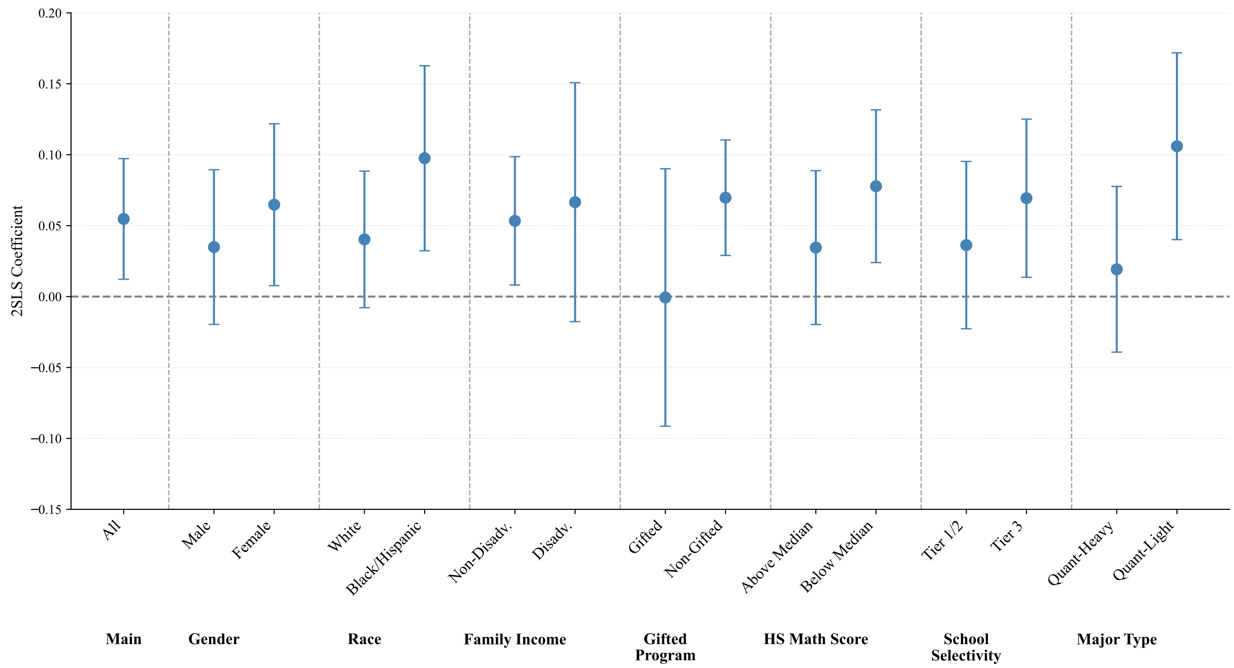

Panel B. Writing Skills

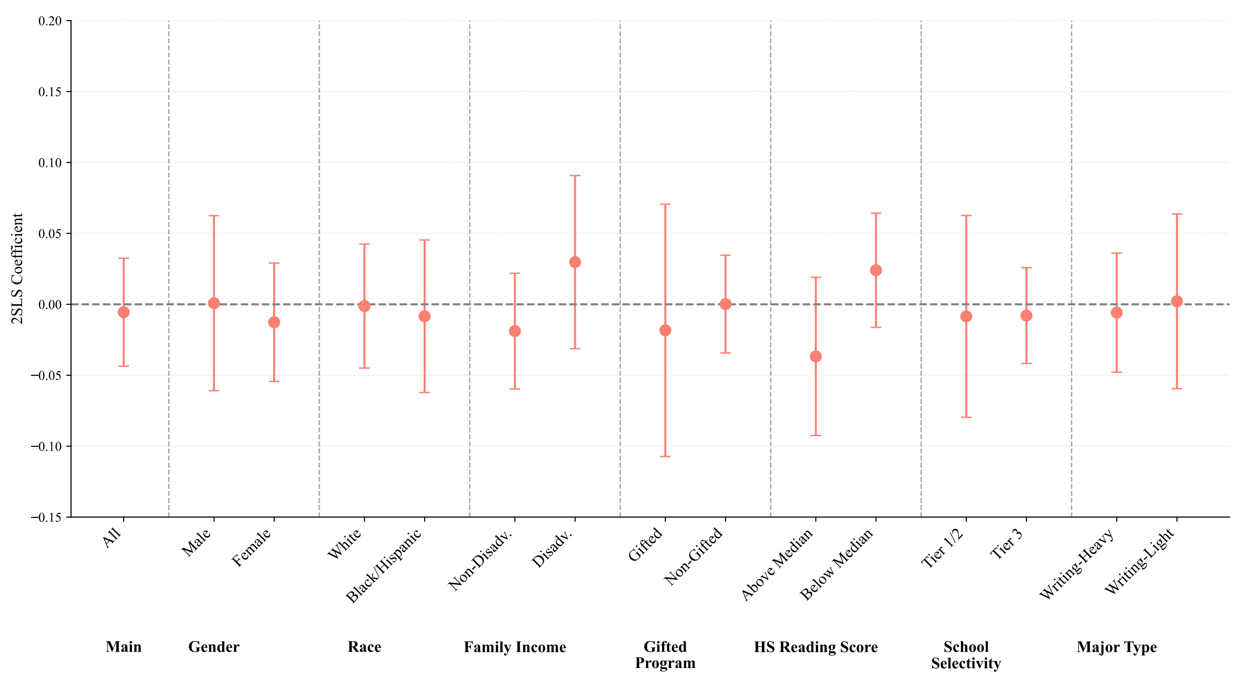

Notes: This figure displays heterogeneity in the effect of coursework-based skills on log quarterly earnings. Each dot reports a 2SLS point estimate of the effect of a one SD increase in the share of upper-division courses taken that teach quantitative/writing skills; vertical lines represent 95% confidence intervals. Estimates are based on Equation 4, run separately for each subgroup listed on the left. "HS math above median" refers to students whose standardized state math test scores in high school are above the median within their degree-granting college. "HS reading above median" refers to students whose standardized state reading test scores in high school are above the median within their degree-granting college. "Major type" is defined based on whether the average quantitative/writing skill share of a CIP-4 major is above or below the median across all CIP-4 majors pooled across schools (see Appendix Table A.2). School selectivity tiers follow the categorization in Table 1. All regressions include individual-level covariates (excluding the grouping variable for the subsample), school-by-major fixed effects, and cohort fixed effects. Standard errors are clustered at the school-by-major level.

Figure 5: Effect of Coursework-based Quantitative Skills on Log Quarterly Earnings across Fields

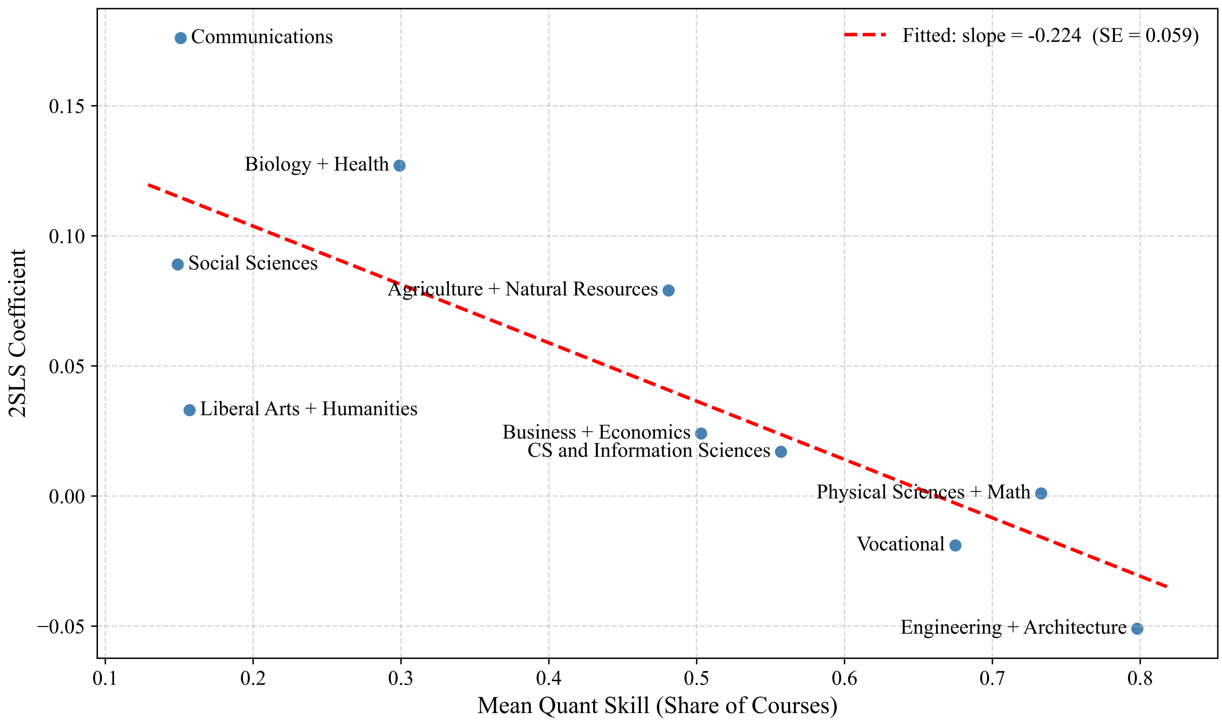

Notes: This figure displays heterogeneity in the effect of coursework-based quantitative skills on log quarterly earnings across fields. Each dot represents a broad academic field. The vertical axis plots the field-specific 2SLS estimate from Equation 4 of the effect of a one SD increase in a student's share of upper-division courses taken that teach quantitative skills on log quarterly earnings. The horizontal axis reports the mean share of courses in the field that teach quantitative skills, averaged across programs and years.

Figure 6: Event Study

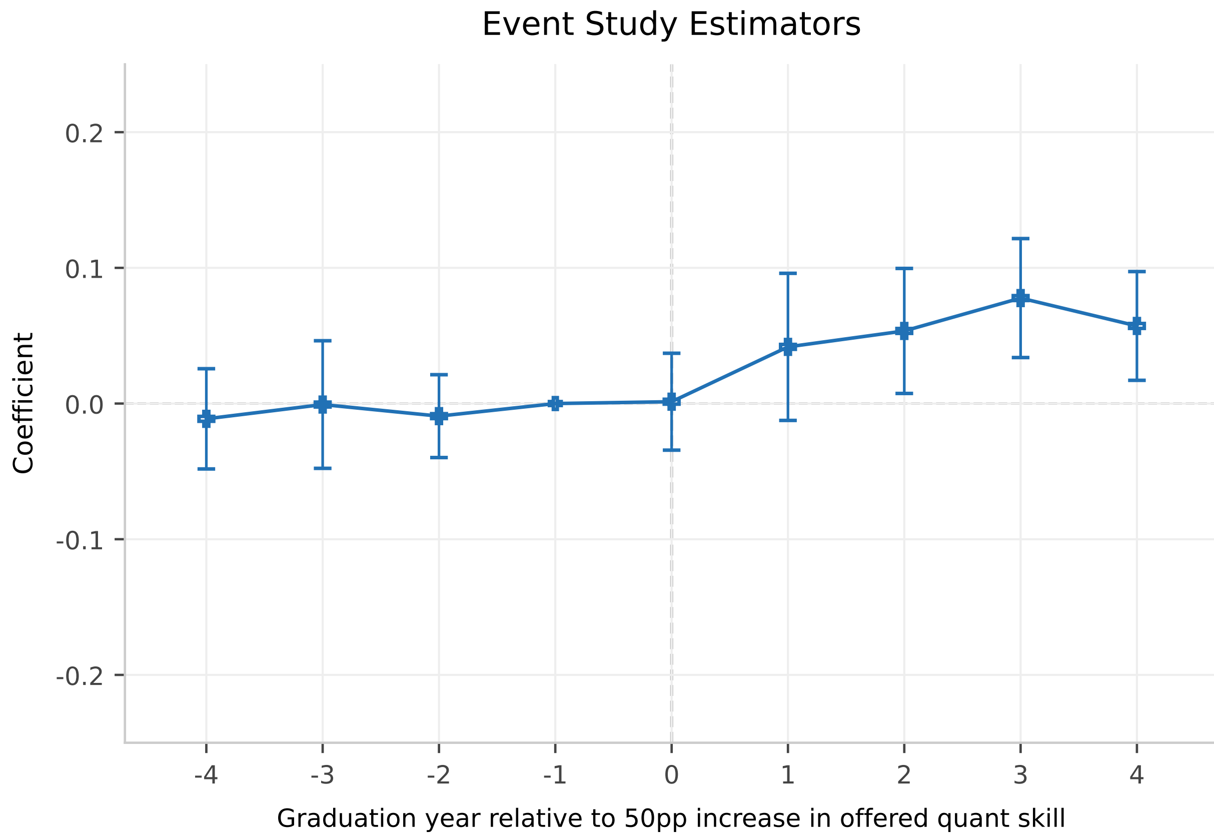

Notes: This figure plots event-time coefficients from a two-way fixed-effects regression of log quarterly earnings on indicators for time since a large curricular change within a program (school×major). Treated programs are those experiencing a sharp and persistent increase (≥50 percentage points, relative to the prior year) in the share of courses teaching quantitative skills, typically due to the introduction of new quantitative courses. Event time is defined by graduation year relative to the event: students graduating in the event year are partially treated ($t = 0$; seniors during the change), and those graduating the year after are fully treated ($t \geq 1$). The omitted category is $t = -1$. Regressions include school-by-major and cohort fixed effects; standard errors are clustered at the school-by-major level. Error bars show 95% confidence intervals.

Appendices

A Appendix Tables and Figures

Table A.1: Aggregate Major Groups

Aggregate Major Group	Specific Major	CIP Code
Agriculture + Natural Resources	Agriculture, Agriculture Operations, and Related Sciences	01, 02
	Natural Resources and Conservation	03
Communications	Communication, Journalism, and Related Programs	09
	Communications Technologies/Technicians and Support Services	10
	Family and Consumer Sciences/Human Sciences	19
CS and Information Sciences	Computer and Information Sciences and Support Services	11
Vocational	Personal and Culinary Services	12
	Engineering Technologies/Technicians	15
	Vocational Home Economics	20
	Parks, Recreation, Leisure, and Fitness Studies	31
	Basic Skills	32
	Leisure and Recreational Activities	36
	Science Technologies/Technicians	41
	Security and Protective Services	43
	Construction Trades	46
	Mechanic and Repair Technologies/Technicians	47
	Precision Production	48
	Transportation and Materials Moving	49
	Reserve Officer Training Corps	28
	Military Technologies	29
	Citizenship Activities	33
	Health-Related Knowledge and Skills	34
	Interpersonal and Social Skills	35
	Personal Awareness and Self-Improvement	37
Engineering + Architecture	Architecture and Related Services	04
	Engineering	14
Liberal Arts + Humanities	Area, Ethnic, Cultural, and Gender Studies	05
	Foreign Languages, Literatures, and Linguistics	16
	English Language and Literature/Letters	23
	Liberal Arts and Sciences, General Studies and Humanities	24
	Library Science	25
	Multi/Interdisciplinary Studies	30
	Philosophy and Religious Studies	38
	Theology and Religious Vocations	39
	Visual and Performing Arts	50
	History	4508, 54
Biology + Health	Biological and Biomedical Sciences	26
	Health Professions and Related Clinical Sciences	51
	Residency Programs	60
Physical Sciences + Math	Physical Sciences	40
	Mathematics and Statistics	27
Social Sciences	Legal Professions and Studies	22
	Psychology	42
	Public Administration and Social Service Professions	44
	Social Sciences, General	4501
	Anthropology	4502
	Archeology	4503
	Criminology	4504
	Demography and Population Studies	4505
	Geography and Cartography	4507
	International Relations and Affairs	4509
	Political Science and Government	4510
	Sociology	4511
	Urban Studies/Affairs	4512
	Sociology and Anthropology	4513
	Rural Sociology	4514
	Social Sciences, Other	4599
Business + Economics	Business, Management, Marketing, and Related Support Services	52, 08
	Economics	4506
Undeclared	Undeclared	99

Source: Texas Higher Education Coordinating Board data as described in the text.

Table A.2: Classification of Majors by Quantitative and Writing Skill Intensity

Quant-Heavy Majors	Avg. Share of Quant-Skill Courses Taken
Civil Engineering Technologies/Technicians	94.69%
Manufacturing Engineering	92.27%
Materials Engineering	91.69%
Engineering-Related Fields	90.80%
Statistics	89.38%
Aerospace, Aeronautical and Astronautical Engineering	88.62%
Ocean Engineering	88.50%
Systems Engineering	88.11%
Mechanical Engineering	88.01%
Nuclear Engineering	87.17%
Civil Engineering	86.85%
Agricultural Engineering	86.04%
Engineering Physics	85.13%
Chemical Engineering	84.64%
Petroleum Engineering	84.32%
Electromechanical Instrumentation and Maintenance Technologies/Technicians	83.74%
Physics	83.56%
Geological/Geophysical Engineering	83.13%
Environmental/Environmental Health Engineering	82.92%
Electrical, Electronics and Communications Engineering	82.82%
Applied Mathematics	82.24%
Atmospheric Sciences and Meteorology	82.16%
Metallurgical Engineering	81.67%
Mechanical Engineering Related Technologies/Technicians	81.54%
Computational Science	80.74%
Industrial Engineering	80.52%
Architectural Engineering	77.89%
Engineering Technology, General	77.80%
Biomedical/Medical Engineering	77.24%
Engineering-Related Technologies	76.94%
Computer Engineering Technologies/Technicians	75.29%
Construction Management	74.94%
Computer Engineering	74.42%
Electrical Engineering Technologies/Technicians	73.43%
Mathematics	73.13%
Genetics	72.63%
Construction Engineering	72.50%
Industrial Production Technologies/Technicians	71.92%
Astronomy and Astrophysics	71.41%
Geological and Earth Sciences/Geosciences	70.35%
Chemistry	68.92%
Science, Technology and Society	67.98%
Finance and Financial Management Services	66.74%
Business/Managerial Economics	66.67%
Architectural Engineering Technologies/Technicians	66.67%
Insurance	66.55%
Engineering, General	65.64%
Management Sciences and Quantitative Methods	65.12%
Economics	63.25%
Biochemistry, Biophysics and Molecular Biology	62.12%
Food Science and Technology	61.88%
Accounting and Related Services	61.82%
Construction Engineering Technologies	61.56%
Agricultural Mechanization	61.55%
General Sales, Merchandising and Related Marketing Operations	59.83%
Computer and Information Sciences, General	59.55%
Computer Science	59.43%
Drafting/Design Engineering Technologies/Technicians	59.20%
Real Estate	58.46%
Cell/Cellular Biology and Anatomical Sciences	58.29%
Plant Sciences	57.44%
Computer Programming	57.14%
Building/Construction Finishing, Management, and Inspection	55.97%
Animal Sciences	55.19%
Agricultural Business and Management	54.63%
Quality Control and Safety Technologies/Technicians	54.37%
Agricultural Production Operations	53.54%
Agricultural and Food Products Processing	52.88%
Information Science/Studies	52.06%
Biotechnology	51.39%
Computer Systems Analysis	51.35%
Ecology, Evolution, Systematics, and Population Biology	51.17%
Microbiological Sciences and Immunology	51.10%
Natural Resources Conservation and Research	50.77%
Clinical/Medical Laboratory Science/Research and Allied Professions	49.53%
Multi Disc: Biological and Physical Sciences	49.20%
Animal studies, General	49.09%
Biology, General	48.16%
Business/Commerce, General	47.27%
Geography and Cartography	47.13%
Air Transportation	46.52%
Management Information Systems and Services	45.52%
City/Urban, Community and Regional Planning	44.51%
Physiology, Pathology and Related Sciences	44.36%
Natural Resources Management and Policy	44.25%
Foods, Nutrition, and Related Services	44.14%
Cognitive Science	43.16%
Nutrition Sciences	42.83%
Entrepreneurial and Small Business Operations	42.57%
International Business	42.49%
Business Administration, Management and Operations	42.20%
Applied Horticulture and Horticultural Business Services	42.02%
Forestry	41.96%
Zoology/Animal Biology	41.28%
Family and Consumer Sciences/Human Sciences Business Services	41.08%
Wildlife and Wildlands Science and Management	40.15%
Audiovisual Communications Technologies/Technicians	40.15%

Quant-Light Majors	Avg. Share of Quant-Skill Courses Taken
Marketing	39.78%
Bioethics/Medical Ethics	38.26%
Computer/Information Technology Administration and Management	38.05%
Specialized Sales, Merchandising and Marketing Operations	36.95%
Family and Consumer Economics and Related Studies	36.17%
Dietetics and Clinical Nutrition Services	35.87%
Neurobiology and Neurosciences	34.98%
Architecture	34.00%
Agriculture, General	33.27%
Environmental Design	31.58%
Slavic, Baltic and Albanian Languages, Literatures, and Linguistics	30.80%
Hospitality Administration/Management	30.79%
Health and Medical Administrative Services	30.09%
Computer Software and Media Applications	29.89%
Urban Studies/Affairs	28.98%
Human Resources Management and Services	28.52%
Public Health	27.80%
Apparel and Textiles	27.64%
Landscape Architecture	27.51%
Public Policy Analysis	27.40%
Allied Health Diagnostic, Intervention, and Treatment Professions	27.04%
Multi-/Interdisciplinary Studies, General	22.56%
Liberal Arts and Sciences, General Studies and Humanities	21.79%
Sustainability Studies	21.78%
Interior Architecture	21.42%
Allied Health and Medical Assisting Services	21.31%
Health professions, General	21.07%
Veterinary/Animal Health Technology/Technician and Veterinary Assistant	20.57%
Public Administration	20.22%
Germanic Languages, Literatures, and Linguistics	19.11%
Multi/Interdisciplinary Studies, Other	18.12%
Agricultural Public Services	17.74%
Mental and Social Health Services and Allied Professions	17.24%
Public Admin and social services, general	16.61%
International Relations and National Security Studies	15.79%
Family and Consumer Sciences/Human Sciences, General	15.61%
Public Relations, Advertising, and Applied Communication	15.49%
Psychology, General	15.00%
East Asian Languages, Literatures, and Linguistics	14.94%
Sociology	14.92%
International/Global Studies	14.29%
Dental Support Services and Allied Professions	14.08%
Human Development, Family Studies, and Related Services	13.74%
Research and Experimental Psychology	13.63%
Anthropology	13.49%
Political Science and Government	13.25%
Romance Languages, Literatures, and Linguistics	12.21%
Philosophy	12.14%
Social Sciences, General	11.42%
Area Studies	11.39%
Classics and Classical Languages, Literatures, and Linguistics	11.36%
Middle/Near Eastern and Semitic Languages, Literatures, and Linguistics	11.32%
Clinical, Counseling and Applied Psychology	11.03%
Social Sciences, Other	10.80%
Business/Corporate Communications	10.09%
Communication Disorders Sciences and Services	10.07%
Linguistic, Comparative, and Related Language Studies and Services	9.94%
Ethnic, Cultural Minority, Gender, and Group Studies	9.72%
Criminology	9.71%
Religion/Religious Studies	9.38%
Communication and Media Studies	9.26%
Rehabilitation and Therapeutic Professions	9.07%
Journalism	8.71%
Rhetoric and Composition/Writing Studies	8.69%
Radio, Television, and Digital Communication	8.41%
History	7.77%
Archeology	6.91%
Social Work	6.36%
English Language and Literature, General	6.11%
Classical and Ancient Studies	5.20%
Registered Nursing, Nursing Administration, Nursing Research and Clinical Nursing	4.74%
Graphic Communications	3.35%
American Sign Language	3.23%
Gerontology	0.00%

Writing-Heavy Majors	Avg. Share of Writing-Skill Courses Taken	Writing-Light Majors	Avg. Share of Writing-Skill Courses Taken
Archeology	48.21%	Computer and Information Sciences, General	16.08%
East Asian Languages, Literatures, and Linguistics	48.16%	Mathematics	16.07%
History	47.65%	Information Science/Studies	16.05%
Germanic Languages, Literatures, and Linguistics	46.41%	Computer Science	15.98%
Rhetoric and Composition/Writing Studies	45.74%	Family and Consumer Sciences/Human Sciences, General	15.93%
Middle/Near Eastern and Semitic Languages, Literatures, and Linguistics	45.21%	Urban Studies/Affairs	15.75%
International Relations and National Security Studies	45.02%	Entrepreneurial and Small Business Operations	15.44%
Communication and Media Studies	44.04%	Economics	15.13%
Slavic, Baltic and Albanian Languages, Literatures, and Linguistics	43.69%	Insurance	15.07%
Ethnic, Cultural Minority, Gender, and Group Studies	43.41%	Physiology, Pathology and Related Sciences	14.72%
Social Sciences, General	41.24%	Cognitive Science	14.42%
Anthropology	41.06%	Family and Consumer Sciences/Human Sciences Business Services	14.42%
Religion/Religious Studies	41.05%	Drafting/Design Engineering Technologies/Technicians	14.28%
Area Studies	40.46%	Mental and Social Health Services and Allied Professions	14.24%
Multi/Interdisciplinary Studies, Other	39.50%	Multi Disc: Biological and Physical Sciences	14.19%
Social Sciences, Other	38.64%	Management Sciences and Quantitative Methods	13.85%
Business/Corporate Communications	35.76%	Animal studies, General	13.79%
Public Admin and social services, general	35.34%	Statistics	13.71%
Public Administration	35.09%	Gerontology	13.64%
Graphic Communications	31.76%	Building/Construction Finishing, Management, and Inspection	13.33%
Public Policy Analysis	31.59%	Real Estate	13.19%
Sociology	29.82%	Hospitality Administration/Management	13.11%
Criminology	27.47%	Zoology/Animal Biology	13.10%
Liberal Arts and Sciences, General Studies and Humanities	27.16%	Agricultural Business and Management	12.97%
Bioethics/Medical Ethics	27.04%	Nutrition Sciences	12.95%
Landscape Architecture	26.61%	Business/Managerial Economics	12.82%
Construction Management	26.12%	Natural Resources Conservation and Research	12.78%
Multi-/Interdisciplinary Studies, General	25.51%	Forestry	12.50%
Computer Software and Media Applications	25.46%	Biotechnology	12.47%
Dental Support Services and Allied Professions	25.25%	Foods, Nutrition, and Related Services	12.45%
Architecture	25.18%	Microbiological Sciences and Immunology	12.35%
Psychology, General	24.88%	Architectural Engineering	12.06%
American Sign Language	23.97%	Biology, General	12.03%
Agriculture, General	23.23%	Computer Engineering	11.98%
Social Work	23.01%	Finance and Financial Management Services	11.89%
Agricultural Mechanization	22.45%	Family and Consumer Economics and Related Studies	11.86%
Health and Medical Administrative Services	22.44%	Neurobiology and Neurosciences	11.83%
Sustainability Studies	22.41%	Chemistry	11.78%
City/Urban, Community and Regional Planning	22.23%	Aerospace, Aeronautical and Astronautical Engineering	11.70%
Architectural Engineering Technologies/Technicians	22.22%	Cell/Cellular Biology and Anatomical Sciences	11.54%
Research and Experimental Psychology	22.21%	Applied Mathematics	11.46%
Human Resources Management and Services	21.47%	Apparel and Textiles	10.99%
Computer Programming	21.43%	Electrical Engineering Technologies/Technicians	10.83%
Construction Engineering Technologies	20.89%	Ecology, Evolution, Systematics, and Population Biology	10.77%
International Business	20.02%	Electrical, Electronics and Communications Engineering	10.65%
Marketing	19.67%	Biochemistry, Biophysics and Molecular Biology	10.55%
Accounting and Related Services	18.99%	Wildlife and Wildlands Science and Management	10.49%
Interior Architecture	18.39%	Agricultural Engineering	10.43%
Business/Commerce, General	18.28%	Electromechanical Instrumentation and Maintenance Technologies/Technicians	10.36%
Environmental Design	17.81%	Construction Engineering	10.21%
Rehabilitation and Therapeutic Professions	17.56%	Animal Sciences	10.21%
Dietetics and Clinical Nutrition Services	17.22%	Geological and Earth Sciences/Geosciences	10.03%
Registered Nursing, Nursing Administration, Nursing Research and Clinical Nursing	17.18%	Astronomy and Astrophysics	9.76%
Health professions, General	17.02%	Plant Sciences	9.73%
Communication Disorders Sciences and Services	16.94%	Mechanical Engineering	9.65%
Computer/Information Technology Administration and Management	16.78%	Manufacturing Engineering	9.54%
Business Administration, Management and Operations	16.62%	Engineering-Related Fields	9.32%
Human Development, Family Studies, and Related Services	16.60%	Natural Resources Management and Policy	9.30%
Audiovisual Communications Technologies/Technicians	16.58%	Agricultural Production Operations	9.29%
Public Health	16.58%	Biomedical/Medical Engineering	9.28%
Geography and Cartography	16.40%	Industrial Production Technologies/Technicians	9.28%
Clinical, Counseling and Applied Psychology	16.31%	Mechanical Engineering Related Technologies/Technicians	9.17%
Management Information Systems and Services	16.17%	Computer Systems Analysis	9.16%
Specialized Sales, Merchandising and Marketing Operations	16.15%	Engineering, General	9.08%
		Geological/Geophysical Engineering	8.96%
		Engineering Technology, General	8.94%
		Systems Engineering	8.85%
		Air Transportation	8.75%
		Allied Health Diagnostic, Intervention, and Treatment Professions	8.67%
		Civil Engineering	8.64%
		Computer Engineering Technologies/Technicians	8.55%
		Petroleum Engineering	8.31%
		Physics	7.95%
		Quality Control and Safety Technologies/Technicians	7.54%
		Genetics	7.35%
		Environmental/Environmental Health Engineering	7.31%
		Agricultural and Food Products Processing	7.25%
		Clinical/Medical Laboratory Science/Research and Allied Professions	7.25%
		Metallurgical Engineering	7.13%
		Applied Horticulture and Horticultural Business Services	7.01%
		Chemical Engineering	6.96%
		Science, Technology and Society	6.89%
		Nuclear Engineering	6.75%
		Engineering-Related Technologies	6.74%
		Computational Science	6.69%
		Food Science and Technology	6.67%
		Industrial Engineering	6.54%
		Engineering Physics	6.18%
		General Sales, Merchandising and Related Marketing Operations	6.08%
		Atmospheric Sciences and Meteorology	5.15%
		Ocean Engineering	5.11%
		Allied Health and Medical Assisting Services	3.87%
		Civil Engineering Technologies/Technicians	3.60%
		Materials Engineering	3.17%
		Veterinary/Animal Health Technology/Technician and Veterinary Assistant	2.06%

Notes: This table classifies CIP-4 majors based on the average share of quantitative or writing skill courses taken by students who earned a degree in each major. To construct these categories, I pooled majors across all institutions and calculated the median skill share across majors. Majors with average skill shares above the median are classified as "Heavy"; those below are "Light."

Table A.3: OLS Estimates of the Association Between Coursework-Based Skills and Earnings

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Quantitative Skill						
Quant Skill	0.020***	0.021***	0.018***	0.018***	0.015***	0.015***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Individual controls	No	No	No	No	Yes	Yes
Cohort FE	No	Yes	Yes	No	Yes	Yes
School × Major FE	No	No	Yes	No	Yes	Yes
School × Major × Cohort FE	No	No	No	Yes	Yes	No
High School FE	No	No	No	No	No	Yes
Adj. R^2	0.001	0.010	0.280	0.284	0.290	0.296
Within R^2	0.001	0.001	0.001	0.001	0.016	0.017
N	239577	239577	239577	239223	239223	175372
Panel B: Writing Skill						
Writing Skill	-0.005**	-0.006***	-0.008***	-0.008***	-0.007***	-0.006***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Individual controls	No	No	No	No	Yes	Yes
Cohort FE	No	Yes	Yes	No	Yes	Yes
School × Major FE	No	No	Yes	No	Yes	Yes
School × Major × Cohort FE	No	No	No	Yes	Yes	No
High School FE	No	No	No	No	No	Yes
Adj. R^2	0.000	0.009	0.279	0.283	0.290	0.296
Within R^2	0.000	0.000	0.000	0.000	0.015	0.016
N	239577	239577	239577	239223	239223	175372

Notes: Each column reports coefficients from an OLS regression of log quarterly earnings on the indicated skill measure. Quant Skill (Panel A) and Writing Skill (Panel B) are the standardized shares of a student's upper-division courses that teach quantitative and writing skills, respectively. Columns (2)–(6) add controls sequentially as indicated: cohort fixed effects, school–major fixed effects, school–major–cohort fixed effects, individual pre-collegiate covariates, and high school fixed effects. Sample sizes differ across columns because some controls are not available for all students. Standard errors, in parentheses, are clustered at the school–major level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.4: Heterogeneous Effects of Coursework-Based Quantitative Skills on Log Quarterly Earnings: OLS, First Stage, Reduced Form, and 2SLS Estimates

Group	(1) OLS	(2) First Stage	(3) First Stage F-stats	(4) Reduced Form	(5) 2SLS	(6) N	(7) Share of Courses w/ Skill (mean, [SD])
Panel A. Gender							
Male	0.019***	0.100***	59.337	0.004	0.035	98892	49.1%
	(0.002)	(0.013)		(0.003)	(0.028)		[28.1%]
Female	0.011***	0.097***	66.014	0.006**	0.065**	140639	28.4%
	(0.002)	(0.012)		(0.003)	(0.029)		[24.4%]
Panel B. Race							
White	0.018***	0.115***	79.851	0.005	0.040	124450	37.6%
	(0.002)	(0.013)		(0.003)	(0.025)		[28.0%]
Black/Hispanic	0.010***	0.085***	39.231	0.008***	0.098***	91989	34.2%
	(0.002)	(0.014)		(0.003)	(0.033)		[27.8%]
Asian	0.018***	0.065***	11.168	-0.000	-0.003	18917	46.2%
	(0.006)	(0.019)		(0.009)	(0.134)		[25.6%]
Panel C. Family Income							
Non-Econ-Disadvantaged	0.017***	0.103***	79.383	0.005**	0.053**	182063	37.7%
	(0.002)	(0.012)		(0.002)	(0.023)		[27.9%]
Econ-Disadvantaged	0.010***	0.087***	38.907	0.006	0.067	57467	34.6%
	(0.002)	(0.014)		(0.004)	(0.043)		[27.8%]
Panel D. Academic Preparation							
Gifted	0.015***	0.100***	49.861	-0.000	-0.001	52805	41.8%
	(0.003)	(0.014)		(0.005)	(0.046)		[28.9%]
Non-Gifted	0.015***	0.099***	78.936	0.007***	0.070***	186719	35.6%
	(0.002)	(0.011)		(0.002)	(0.021)		[27.5%]
HS math above median	0.014***	0.105***	86.105	0.004	0.035	119318	42.8%
	(0.002)	(0.011)		(0.003)	(0.028)		[28.9%]
HS math below median	0.016***	0.094***	59.102	0.007***	0.078***	120214	31.1%
	(0.002)	(0.012)		(0.003)	(0.027)		[25.6%]
Panel E. Major Type							
Quant-Heavy Major	0.013***	0.108***	45.269	0.002	0.019	109065	60.8%
	(0.003)	(0.016)		(0.003)	(0.030)		[19.8%]
Quant-Light Major	0.017***	0.087***	38.524	0.009***	0.106***	130512	17.0%
	(0.002)	(0.014)		(0.003)	(0.034)		[15.1%]
Panel F. School Selectivity							
Tier 1/2 School	0.020***	0.093***	42.400	0.003	0.036	155814	39.6%
	(0.003)	(0.014)		(0.003)	(0.030)		[28.1%]
Tier 3 School	0.007***	0.109***	42.928	0.008**	0.069**	83763	31.9%
	(0.002)	(0.017)		(0.003)	(0.028)		[26.8%]
Other Individual Controls	Yes	Yes		Yes	Yes		
Cohort FE	Yes	Yes		Yes	Yes		
School × Major FE	Yes	Yes		Yes	Yes		

Notes: Each row reports estimates for the indicated subsample (by panel heading), with log quarterly earnings as the dependent variable. Column (6) reports the subsample size N. Column (7) reports the raw mean (first line) and standard deviation (second line, in brackets) of the share of courses taken by students with the relevant skill (quantitative) in each subsample. "HS math above median" refers to students whose standardized state math test scores in high school are above the median within their degree college. "Major type" in Panel E is defined based on whether the average quantitative skill share of a CIP-4 major is above or below the median across all CIP-4 majors pooled across schools. School selectivity tiers follow the categorization presented in Table 1. All regressions include individual-level controls (excluding the group variable used for heterogeneity), school-by-major fixed effects, and cohort fixed effects. Standard errors (in parentheses) are clustered at the school-by-major level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.5: Heterogeneous Effects of Coursework-Based Writing Skills on Log Quarterly Earnings: OLS, First Stage, Reduced Form, and 2SLS Estimates

Group	(1) OLS	(2) First Stage	(3) First Stage F-stat	(4) Reduced Form	(5) 2SLS	(6) N	(7) Share of Courses w/ Skill (mean, [SD])
Panel A. Gender							
Male	-0.012***	0.101***	83.844	0.000	0.001	98892	20.3%
	(0.002)	(0.011)		(0.003)	(0.031)		[17.9%]
Female	-0.001	0.116***	64.395	-0.001	-0.013	140639	27.2%
	(0.002)	(0.014)		(0.002)	(0.021)		[20.5%]
Panel B. Race							
White	-0.009***	0.117***	69.711	-0.000	-0.001	124450	24.6%
	(0.002)	(0.014)		(0.003)	(0.022)		[20.1%]
Black/Hispanic	-0.002	0.110***	94.046	-0.001	-0.008	91989	25.4%
	(0.002)	(0.011)		(0.003)	(0.027)		[19.8%]
Asian	-0.005	0.066***	14.742	-0.008	-0.117	18917	18.0%
	(0.005)	(0.017)		(0.009)	(0.141)		[15.4%]
Panel C. Family Income							
Non-Econ-Disadvantaged	-0.008***	0.111***	89.730	-0.002	-0.019	182063	24.2%
	(0.002)	(0.012)		(0.002)	(0.021)		[19.9%]
Econ-Disadvantaged	-0.001	0.110***	74.281	0.003	0.030	57467	24.7%
	(0.003)	(0.013)		(0.003)	(0.031)		[19.4%]
Panel D. Academic Preparation							
Gifted	-0.008***	0.101***	61.861	-0.002	-0.018	52805	22.8%
	(0.003)	(0.013)		(0.005)	(0.045)		[19.5%]
Non-Gifted	-0.006***	0.113***	94.051	0.000	0.000	186719	24.8%
	(0.002)	(0.012)		(0.002)	(0.018)		[19.8%]
HS reading above median	-0.005***	0.101***	83.378	-0.004	-0.037	117206	24.5%
	(0.002)	(0.011)		(0.003)	(0.028)		[20.2%]
HS reading below median	-0.007***	0.120***	92.250	0.003	0.024	122347	24.2%
	(0.002)	(0.012)		(0.002)	(0.021)		[19.3%]
Panel E. Major Type							
Writing-Heavy Major	-0.008***	0.112***	44.657	-0.001	-0.006	119695	35.5%
	(0.002)	(0.017)		(0.002)	(0.021)		[20.8%]
Writing-Light Major	-0.004**	0.110***	61.190	0.000	0.002	119882	13.2%
	(0.002)	(0.014)		(0.003)	(0.031)		[9.9%]
Panel F. School Selectivity							
Tier 1/2 School	-0.010***	0.081***	40.829	-0.001	-0.009	155814	22.7%
	(0.002)	(0.013)		(0.003)	(0.036)		[19.2%]
Tier 3 School	0.000	0.166***	70.464	-0.001	-0.008	83763	27.5%
	(0.002)	(0.020)		(0.003)	(0.017)		[20.3%]
Individual controls	Yes	Yes		Yes	Yes		
Cohort FE	Yes	Yes		Yes	Yes		
School x Major FE	Yes	Yes		Yes	Yes		

Notes: Each row reports estimates for the indicated subsample (by panel heading), with log quarterly earnings as the dependent variable. Column (6) reports the subsample size N. Column (7) reports the raw mean (first line) and standard deviation (second line, in brackets) of the share of courses taken by students with the relevant skill (writing) in each subsample. "HS reading above median" refers to students whose standardized state reading test scores in high school are above the median within their degree college. "Major type" in Panel E is defined based on whether the average writing skill share of a CIP-4 major is above or below the median across all CIP-4 majors pooled across schools. School selectivity tiers follow the categorization presented in Table 1. All regressions include individual-level controls (excluding the group variable used for heterogeneity), school-by-major fixed effects, and cohort fixed effects. Standard errors (in parentheses) are clustered at the school-by-major level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.6: Balance Check for Number of High School Courses by Category

	(1)	(2)	(3)	(4)
Dep. var. *(number of HS courses)*	Actual Quant Skill	Offered Quant Skill	Actual Writing Skill	Offered Writing Skill
English lang arts	-0.088***	0.004	0.075***	0.043
	(0.008)	(0.027)	(0.008)	(0.026)
Mathematics	-0.060***	-0.018	0.050***	0.021
	(0.008)	(0.023)	(0.008)	(0.022)
Science	-0.060***	0.000	0.064***	0.014
	(0.007)	(0.023)	(0.007)	(0.022)
Social studies	-0.049***	0.001	0.048***	0.019
	(0.005)	(0.020)	(0.006)	(0.020)
Foreign lang	-0.041***	-0.010	0.056***	0.001
	(0.006)	(0.014)	(0.006)	(0.013)
Fine arts	-0.138***	-0.007	0.126***	0.008
	(0.011)	(0.019)	(0.015)	(0.020)
Business edu	-0.000	-0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
CTE + Tech App	-0.026*	-0.034	-0.058***	-0.022
	(0.012)	(0.021)	(0.013)	(0.020)
Others	-0.020*	0.039	-0.050***	0.012
	(0.009)	(0.020)	(0.010)	(0.019)
English lang arts (AP)	-0.047***	0.011	0.076***	-0.001
	(0.005)	(0.008)	(0.006)	(0.007)
Math + Tech App (AP)	-0.002	0.002	0.031***	-0.009*
	(0.007)	(0.005)	(0.006)	(0.004)
Science (AP)	-0.009	-0.002	0.041***	0.001
	(0.006)	(0.008)	(0.006)	(0.008)
Social studies (AP)	-0.031***	0.020	0.091***	0.005
	(0.008)	(0.012)	(0.008)	(0.011)
Foreign lang (AP)	-0.001	0.000	0.012***	0.003
	(0.002)	(0.003)	(0.002)	(0.003)
Fine arts (AP)	-0.001	0.002	0.006***	0.004*
	(0.001)	(0.002)	(0.001)	(0.002)
Cohort FE	Yes	Yes	Yes	Yes
School × Major FE	Yes	Yes	Yes	Yes
Observations	239577	239577	239577	239577

Notes: Each row reports a separate regression in which the dependent variable is the number of high-school courses completed in the indicated category (student-level counts; AP rows denote Advanced Placement). Columns (1) and (3) show associations with the treatments (T_i)—actual quantitative and writing skills; columns (2) and (4) show associations with the instruments (Z_i)—offered quantitative and writing skills. Parentheses contain standard errors clustered at the school-by-major level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.7: Heterogeneous Effects of Coursework-Based Skills on Log Quarterly Earnings: Fully Saturated 2SLS Model

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. Quant Skill							
Quant	0.030	0.037	0.052**	0.003	0.042	0.035	0.014
	(0.026)	(0.024)	(0.023)	(0.042)	(0.030)	(0.027)	(0.029)
Quant × Female	0.043						
	(0.035)						
Quant × Black/Hispanic		0.064*					
		(0.038)					
Quant × Econ. Disadv.			0.017				
			(0.043)				
Quant × Non-Gifted				0.066*			
				(0.040)			
Quant × Tier 3 School					0.034		
					(0.041)		
Quant × HS Math Below Median						0.041	
						(0.033)	
Quant × Quant-Light Major							0.091**
							(0.045)
Individual Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
School × Major FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	239577	216491	239577	239577	239577	239577	239577
Panel B. Writing Skill							
Writing	−0.006	−0.005	−0.017	−0.014	−0.001	−0.033	−0.018
	(0.031)	(0.022)	(0.020)	(0.045)	(0.035)	(0.028)	(0.021)
Writing × Female	0.000						
	(0.033)						
Writing × Black/Hispanic		0.009					
		(0.031)					
Writing × Econ. Disadv.			0.049				
			(0.031)				
Writing × Non-Gifted				0.011			
				(0.041)			
Writing × Tier 3 School					−0.009		
					(0.039)		
Writing × HS Reading Below Median						0.049*	
						(0.028)	
Writing × Writing-Light Major							0.028
							(0.036)
Individual Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
School × Major FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	239577	216491	239577	239577	239577	239577	239577

Notes: This table reports the main effect of the skill variable (first row) and its interactions with each subgroup indicator from a single, fully saturated 2SLS model. For each subgroup G (e.g., Female, Black/Hispanic, Econ. Disadv.), the endogenous skill measure, its instrument, and all control variables are interacted with G. The coefficient on the main skill row represents the effect for the baseline group (e.g., male, White, non-Econ. Disadv.); the interaction terms represent differences in effects relative to that baseline. For race, the contrast is Black/Hispanic versus White. Standard errors (in parentheses) are clustered at the school-by-major level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.8: Heterogeneous Effects of Coursework-Based Quantitative Skills on Log Quarterly Earnings by Field

	(1)	(2)	(3)	(4)	(5)	(6)
Field	OLS	First Stage	Reduced Form	2SLS	Share of Courses w/ Quant Skills (mean, [SD])	N
Agriculture + Natural Resources	0.015	0.202***	0.016	0.079	48.1%	12153
	(0.010)	(0.038)	(0.011)	(0.053)	[19.5%]	
Communications	0.024***	0.073**	0.013	0.176	15.1%	23386
	(0.005)	(0.026)	(0.007)	(0.112)	[14.8%]	
CS and Information Sciences	0.014	0.116**	0.002	0.017	55.7%	7619
	(0.010)	(0.042)	(0.009)	(0.073)	[17.9%]	
Vocational	0.004	0.095*	-0.002	-0.019	67.5%	4564
	(0.006)	(0.042)	(0.008)	(0.090)	[15.7%]	
Engineering + Architecture	0.003	0.091**	-0.005	-0.051	79.8%	21128
	(0.004)	(0.028)	(0.008)	(0.086)	[18.6%]	
Liberal Arts + Humanities	0.022***	0.031	0.001	0.033	15.7%	43206
	(0.004)	(0.021)	(0.005)	(0.143)	[15.1%]	
Biology + Health	0.002	0.080*	0.010	0.127	29.9%	35632
	(0.004)	(0.032)	(0.008)	(0.086)	[22.9%]	
Physical Sciences + Math	0.015	0.056*	0.000	0.001	73.3%	6812
	(0.008)	(0.026)	(0.015)	(0.267)	[15.2%]	
Social Sciences	0.022***	0.084**	0.008	0.089	14.9%	28930
	(0.004)	(0.026)	(0.006)	(0.082)	[13.5%]	
Business + Economics	0.015***	0.124***	0.003	0.024	50.3%	56147
	(0.004)	(0.022)	(0.003)	(0.025)	[17.9%]	
Individual Controls	Yes	Yes	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
School by Major FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Estimates are reported separately by broad academic field. Columns (1)–(4) report OLS, first-stage, reduced-form, and 2SLS estimates (standard errors in parentheses). Column (5) reports the mean (and standard deviation in brackets) of the share of courses with quantitative skills in each field. Column (6) reports the number of students in each given field. Specifications are identical to those in the main table. See Appendix Table A.1 for details on the classification of fields. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.9: Heterogeneous Effects of Coursework-Based Writing Skills on Log Quarterly Earnings by Field

Field	(1) OLS	(2) First Stage	(3) Reduced Form	(4) 2SLS	(5) Share of Courses w/ Writing Skills (mean, [SD])	(6) N
Agriculture + Natural Resources	0.001	0.104**	-0.015	-0.143	15.5%	12153
	(0.007)	(0.035)	(0.010)	(0.114)	[14.6%]	
Communications	-0.011*	0.113***	-0.001	-0.008	41.7%	23386
	(0.004)	(0.030)	(0.006)	(0.053)	[23.6%]	
CS and Information Sciences	-0.015	0.024	0.008	0.353	16.2%	7619
	(0.010)	(0.053)	(0.014)	(1.033)	[10.2%]	
Vocational	0.010*	0.232***	0.002	0.009	14.7%	4564
	(0.005)	(0.046)	(0.008)	(0.035)	[10.0%]	
Engineering + Architecture	-0.006	0.123***	0.011	0.092	10.6%	21128
	(0.004)	(0.036)	(0.008)	(0.075)	[8.3%]	
Liberal Arts + Humanities	-0.001	0.101**	0.002	0.024	43.3%	43206
	(0.005)	(0.039)	(0.004)	(0.043)	[21.2%]	
Biology + Health	0.002	0.157***	-0.009	-0.056	14.6%	35632
	(0.003)	(0.028)	(0.007)	(0.047)	[11.0%]	
Physical Sciences + Math	-0.011	0.079*	-0.006	-0.081	12.8%	6812
	(0.006)	(0.033)	(0.011)	(0.150)	[10.2%]	
Social Sciences	-0.015***	0.134***	-0.000	-0.004	29.4%	28930
	(0.003)	(0.025)	(0.005)	(0.038)	[16.8%]	
Business + Economics	-0.009***	0.095***	0.004	0.043	16.5%	56147
	(0.002)	(0.021)	(0.003)	(0.039)	[9.9%]	
Individual Controls	Yes	Yes	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
School × Major FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Estimates are reported separately by broad academic field. Columns (1)–(4) report OLS, first-stage, reduced-form, and 2SLS estimates (standard errors in parentheses). Column (5) reports the mean (and standard deviation in brackets) of the share of courses with writing skills in each field. Column (6) reports the number of students in each given field. Specifications are identical to those in the main table. See Appendix Table A.1 for details on the classification of fields. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.10: Event Study: Two-Way Fixed Effects Model

	(1) Log of Quarterly Earnings Post-College
Event Time = –4	–0.0113
	(0.0188)
Event Time = –3	–0.0007
	(0.0240)
Event Time = –2	–0.0093
	(0.0156)
Event Time = –1	0.0000
	(0.0000)
Event Time = 0	0.0014
	(0.0182)
Event Time = 1	0.0418
	(0.0277)
Event Time = 2	0.0535**
	(0.0235)
Event Time = 3	0.0777***
	(0.0224)
Event Time = 4	0.0572***
	(0.0204)
Individual Controls	Yes
School × Major FE	Yes
Cohort FE	Yes
Observations	239577

Notes: This table reports event-time estimates from a two-way fixed effects regression of log earnings on event-time indicators. Event time is defined relative to the year in which a program experiences a ≥50-percentage-point increase in the share of courses teaching quantitative skills. Ninety of 1,190 programs are treated. Event time = 0 refers to the partially treated cohort (seniors in the year of the curricular change). The omitted category is event time = –1. Standard errors are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.11: Group Differences in HS Math and Reading Test Scores

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. HS Math Test Score						
Tier 3 School	-0.292***					
	(0.006)					
Black/Hispanic		-0.169***				
		(0.004)				
Female			-0.023***			
			(0.004)			
Gifted				0.352***		
				(0.008)		
Econ. Disadv.					-0.091***	
					(0.004)	
HS Reading Above Median						0.336***
						(0.006)
School × Major × Cohort FE		Yes	Yes	Yes	Yes	Yes
Major × Cohort FE	Yes					
N	239555	216491	239577	239577	239577	239577
Panel B. HS Reading Test Score						
Tier 3 School	-0.150***					
	(0.005)					
Black/Hispanic		-0.096***				
		(0.003)				
Female			0.117***			
			(0.003)			
Gifted				0.191***		
				(0.004)		
Econ. Disadv.					-0.085***	
					(0.003)	
HS Math Above Median						0.277***
						(0.004)
School × Major × Cohort FE		Yes	Yes	Yes	Yes	Yes
Major × Cohort FE	Yes					
N	239555	216491	239577	239577	239577	239577

Notes: Each column reports a separate OLS regression where the dependent variable is the student's standardized HS math test score (Panel A) or HS reading test score (Panel B). Column (1) includes Major × Cohort fixed effects; the coefficient on Tier 3 School is the within–Major × Cohort difference between Tier-3 institutions and Tiers 1–2. Columns (2)–(6) include School × Major × Cohort fixed effects; the coefficient on the indicator shown in each row equals the mean difference in the outcome between students with indicator= 1 and indicator= 0 within the same School × Major × Cohort cell. In column (6) of Panel A, the regressor captures variation in HS reading performance; in column (6) of Panel B, it captures variation in HS math performance. Standard errors are clustered at the fixed-effect level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.12: Complier Characterization

	All		White		Black/Hispanic	
Var.	$\Pr_\omega(X=1 \mid g)$	$\Pr(X=1 \mid g)$	$\Pr_\omega(X=1 \mid g)$	$\Pr(X=1 \mid g)$	$\Pr_\omega(X=1 \mid g)$	$\Pr(X=1 \mid g)$
Panel A. Quant Skill						
White	0.622	0.520				
Black/Hispanic	0.336	0.385				
Asian	0.053	0.080				
Gifted Program	0.217	0.221	0.219	0.232	0.204	0.198
Female	0.574	0.587	0.553	0.567	0.598	0.628
Econ. Disadvantaged	0.201	0.240	0.061	0.063	0.494	0.478
Tier 3 School	0.386	0.350	0.423	0.359	0.362	0.391
Quant-Light Major	0.455	0.545	0.404	0.533	0.511	0.601
Within-Program HS Math Q1	0.263	0.255	0.241	0.224	0.323	0.305
Within-Program HS Math Q2	0.222	0.252	0.206	0.248	0.260	0.260
Within-Program HS Math Q3	0.240	0.251	0.263	0.262	0.224	0.234
Within-Program HS Math Q4	0.275	0.242	0.291	0.266	0.193	0.202
Within-Program HS Reading Q1	0.254	0.261	0.224	0.234	0.308	0.293
Within-Program HS Reading Q2	0.227	0.253	0.244	0.253	0.204	0.255
Within-Program HS Reading Q3	0.242	0.248	0.236	0.260	0.237	0.236
Within-Program HS Reading Q4	0.277	0.237	0.296	0.254	0.251	0.215
Panel B. Writing Skill						
White	0.579	0.520				
Black/Hispanic	0.388	0.385				
Asian	0.042	0.080				
Gifted Program	0.190	0.221	0.187	0.232	0.182	0.198
Female	0.642	0.587	0.646	0.567	0.642	0.628
Econ. Disadvantaged	0.220	0.240	0.062	0.063	0.473	0.478
Tier 3 School	0.520	0.350	0.550	0.359	0.503	0.391
Writing-Light Major	0.469	0.500	0.492	0.501	0.412	0.461
Within-Program HS Math Q1	0.258	0.255	0.239	0.224	0.312	0.305
Within-Program HS Math Q2	0.228	0.252	0.207	0.248	0.240	0.260
Within-Program HS Math Q3	0.261	0.251	0.266	0.262	0.252	0.234
Within-Program HS Math Q4	0.253	0.242	0.287	0.266	0.196	0.202
Within-Program HS Reading Q1	0.258	0.261	0.229	0.234	0.314	0.293
Within-Program HS Reading Q2	0.251	0.253	0.240	0.253	0.264	0.255
Within-Program HS Reading Q3	0.237	0.248	0.253	0.260	0.204	0.236
Within-Program HS Reading Q4	0.254	0.237	0.278	0.254	0.218	0.215

Notes: This table reports complier characteristics for the quantitative skill (Panel A) and writing skill (Panel B) first stages, following Frandsen et al. (2023). For each predetermined characteristic X_i and group $g \in \{\text{All}, \text{White}, \text{Black/Hispanic}\}$, I estimate a 2SLS regression in which the dependent variable is the interaction $X_i D_i$, instrumenting the acquired skill D_i with the offered skill Z_i, and controlling for school–major fixed effects and cohort fixed effects. The resulting estimand is $\hat{\delta} = \text{Cov}(Z_i, X_i D_i)/\text{Cov}(Z_i, D_i) = \mathbb{E}[\omega_i X_i \mid g]/\mathbb{E}[\omega_i \mid g]$, where $\omega_i = \partial D_i/\partial Z_i$ is each student's first-stage responsiveness. For a binary characteristic X_i, this equals the responsiveness-weighted share $\Pr_\omega(X = 1 \mid g)$, reported in odd-numbered columns. The column $\Pr(X = 1 \mid g)$ reports the raw population share of each characteristic in group g. If $\Pr_\omega(X = 1 \mid g) > \Pr(X = 1 \mid g)$, that characteristic is over-represented among students whose skill acquisition responds most strongly to changes in course availability; if below, it is under-represented. Within-program HS math (reading) quartiles are defined by ranking students' 9th-grade math (reading) test scores within each school–major cell.

Table A.13: Reweighting Complier URM Students to Match Characteristics of Complier White Students

	(1) All	(2) White	(3) Black/Hispanic	(4) Asian	(5) Black/Hisp. Fully Saturated
Panel A. Quant Skill					
Quant Skill	0.048**	0.042*	0.062*	−0.029	0.039*
	(0.019)	(0.022)	(0.033)	(0.096)	(0.022)
Black/Hispanic × Quant Skill					0.034
					(0.037)
Controls	Yes	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes	Yes
School×Major FE	Yes	Yes	Yes	Yes	Yes
Panel B. Writing Skill					
Writing Skill	0.001	0.004	0.006	−0.162	0.006
	(0.022)	(0.023)	(0.038)	(0.165)	(0.024)
Black/Hispanic × Writing Skill					0.004
					(0.039)
Controls	Yes	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes	Yes
School×Major FE	Yes	Yes	Yes	Yes	Yes

Notes: This table reports 2SLS estimates of the returns to coursework-based quantitative skill (Panel A) and writing skill (Panel B) after reweighting the complier distribution of URM students to match that of White students. To construct the weights, I estimate the first-stage regression separately for URM and White students, including binarized individual characteristics as covariates, including at-risk status, gifted status, limited-English-proficiency status, gender, economic disadvantage, and quartiles (Q1–Q4) of within program high-school math and reading scores; the math and reading variables enter as quartiles because they are continuous. I then compute average propensity scores within each subgroup by race and construct inverse probability weights that reweight the distribution of complier characteristics among non-white students to match the complier distribution among White students. Columns (1)–(4) report estimates for the indicated subsamples; column (5) reports estimates from a pooled specification interacting all regressors with an indicator for Black/Hispanic, so that the main coefficient captures the effect for White students and the interaction term captures the URM–White difference. All specifications include individual-level baseline controls, school×major fixed effects, and cohort fixed effects. Standard errors, reported in parentheses, are clustered at the school×major level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.14: Effects of Coursework-based Skills on Industry Employment

Industry	(1) Quant Skill	(2) Writing Skill	(3) Observations
Agriculture, Forestry, Fishing and Hunting	0.000	-0.000	1388
	(0.003)	(0.002)	
Mining, Quarrying, and Oil and Gas Extraction	-0.017**	0.003	5428
	(0.008)	(0.009)	
Utilities	-0.002	0.004	2328
	(0.004)	(0.003)	
Construction	-0.007	-0.008	13206
	(0.009)	(0.008)	
Manufacturing	0.003	-0.006	19126
	(0.011)	(0.009)	
Wholesale Trade	-0.008	0.003	19219
	(0.012)	(0.009)	
Retail Trade	-0.010	-0.013	28689
	(0.013)	(0.010)	
Transportation and Warehousing	0.006	0.005	7644
	(0.008)	(0.006)	
Information	0.008	0.023***	10938
	(0.010)	(0.008)	
Finance and Insurance	0.024*	-0.029**	25146
	(0.013)	(0.011)	
Real Estate and Rental and Leasing	0.013*	-0.000	8868
	(0.007)	(0.006)	
Professional, Scientific, and Technical Services	0.034**	0.013	55011
	(0.017)	(0.015)	
Management of Companies and Enterprises	-0.002	0.010*	4542
	(0.006)	(0.005)	
Administrative and Support and Waste Mgmt. Services	-0.011	0.011	41636
	(0.018)	(0.016)	
Educational Services	-0.010	-0.021	58650
	(0.018)	(0.015)	
Health Care and Social Assistance	-0.006	0.004	45827
	(0.014)	(0.014)	
Arts, Entertainment, and Recreation	0.012**	0.004	4968
	(0.005)	(0.004)	
Accommodation and Food Services	-0.009	-0.004	16496
	(0.010)	(0.008)	
Other Services (except Public Administration)	-0.003	0.005	7511
	(0.007)	(0.006)	
Public Administration	0.009	0.005	11446
	(0.008)	(0.007)	
Individual Controls	Yes	Yes	
School × Major FE	Yes	Yes	
Cohort FE	Yes	Yes	

Notes: Each row reports a separate 2SLS estimate from a linear probability model in which the dependent variable is an indicator for ever working in the listed industry during Years 1–5 after college. The Quant Skill and Writing Skill columns are estimated in separate regressions. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure A.1: Example Course Catalog Entries

Texas A&M University–Commerce

CSCI 303 - Technical Communication for Computing Professionals

The course will consist of a study of formal and informal communications for computing professionals. Types of communications that will be examined will include academic conference and journal publications; powerpoint presentations for technical and non-technical audience; writing clean code with comments; collaborative software development; soft skills for IT job interviews; in-house technical reports, progress reports, and email messages; writing blog posts and wiki articles. Some of these communications/documents will be created as an individual requirement and will be completed as a team project.

The University of Texas at Dallas

EPPS 2302 - Methods of Quantitative Analysis in the Social and Policy Sciences

This course introduces basic concepts and methods of statistical analysis used in different fields of social and policy science research to better understand human relationships and the impacts of government action on them. Topics include data description, using probability to assess the reasonableness of claims about the world based on sample data, exploring cause-effect interactions through regression models, and application of software to ease visualization and calculation. Students completing this course will be good consumers of statistical information and have a solid foundation for pursuing further study of quantitative analysis.

The University of Texas at Austin

GOV 385N - Introduction to Formal Political Analysis

Critical, comparative survey of important formal theories of political processes, stressing general approaches rather than mathematical results. Presupposes no technical background. Field core course.

Texas Tech University

POLS 3302/3314 - Introduction to Political Analysis

The purpose of this course is to introduce students to the use of statistics in political science research. Statistical topics to be covered begin with descriptive statistics where students are introduced to measures of central tendency, measures of dispersion, and the normal curve. We then proceed to inferential statistics which include such things as probability and hypothesis testing, measures of association, and multivariate modeling.

Figure A.2: Skill Distribution Across Majors

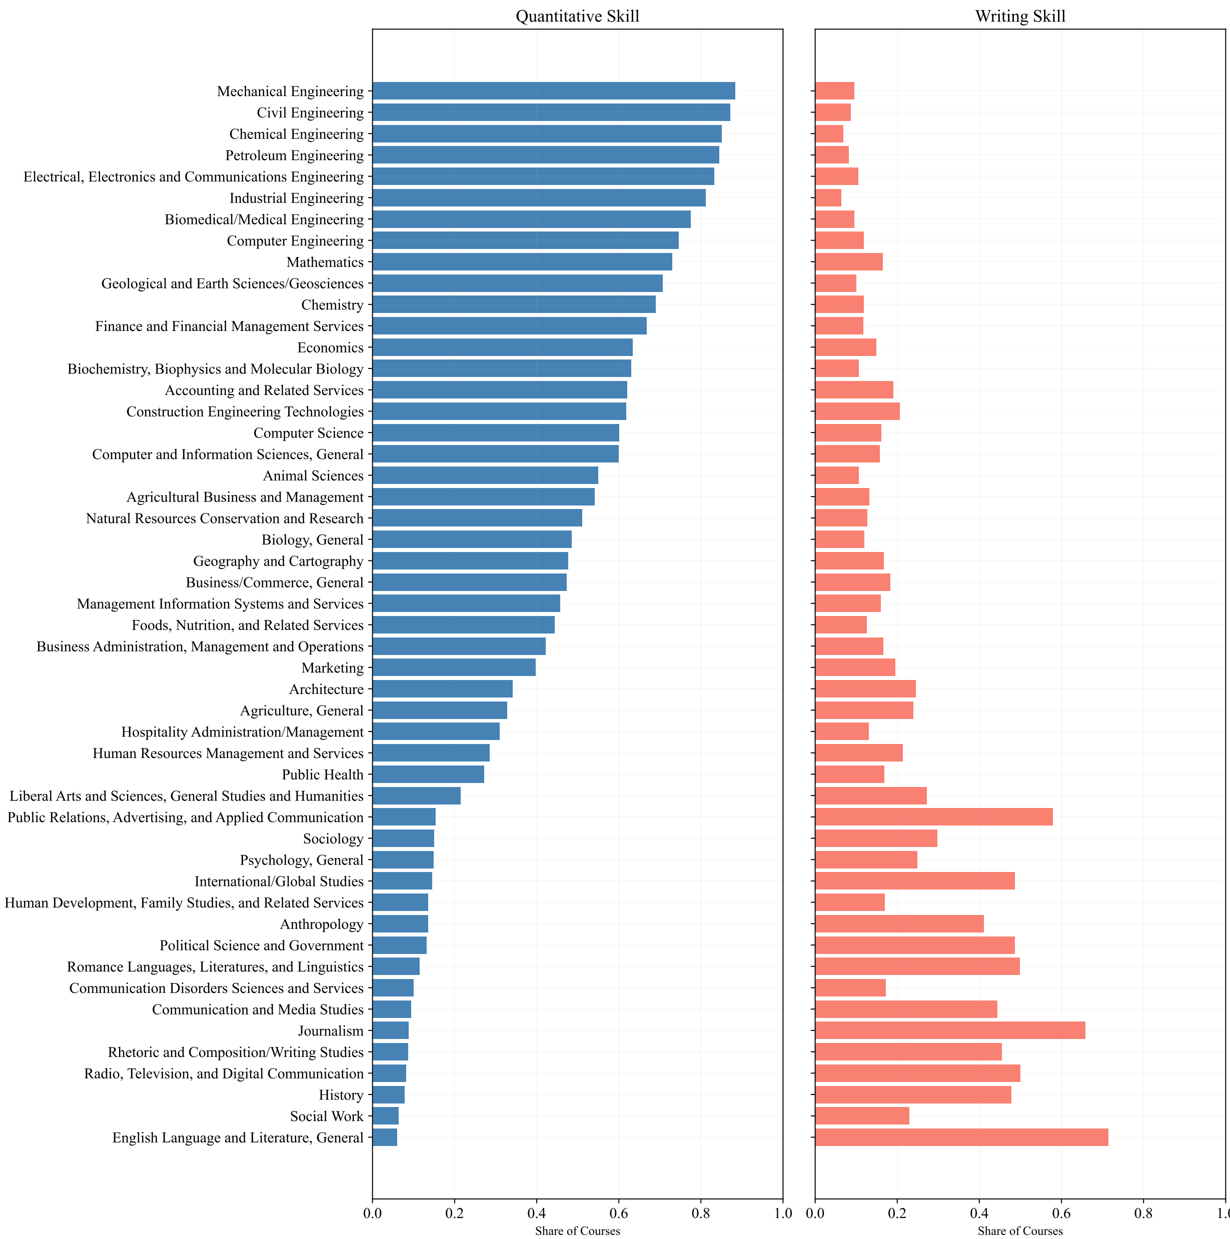

Notes: This figure shows how skill content varies across majors. For each four-digit CIP major, I plot the mean share of courses that are classified as teaching quantitative skills or writing skills. Shares are computed within programs and then averaged across programs for each CIP major.

Figure A.3: Top Departments Offering Courses to Students in Selected Majors at UT Austin

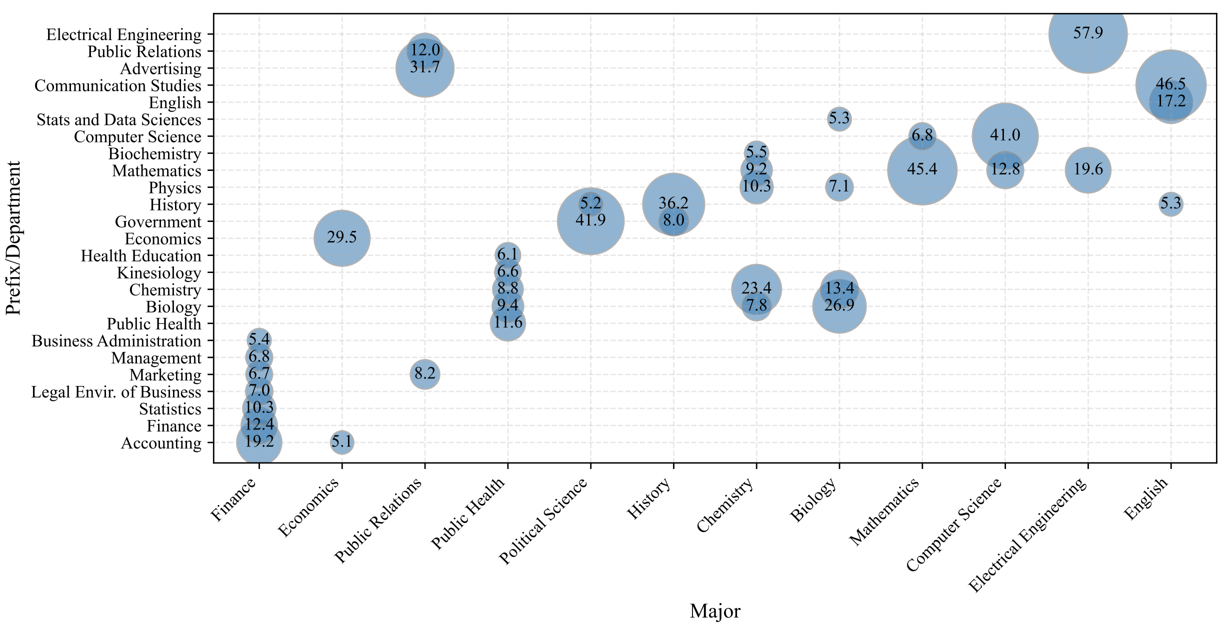

Notes: For each selected major at UT Austin, this figure displays the top departments (course prefixes) from which junior students in that major most frequently take upper-division courses. For a given major, the "share" shown for a department equals the fraction of all upper-division course enrollments by students in that major that are offered by that department. Shares do not sum to one because only departments offering major-relevant courses are displayed. Bubble size reflects the share magnitude. This mapping of where students actually take courses is used to define the set of "major-relevant" courses when constructing the offered-skill instrument.

Figure A.4: Event Study

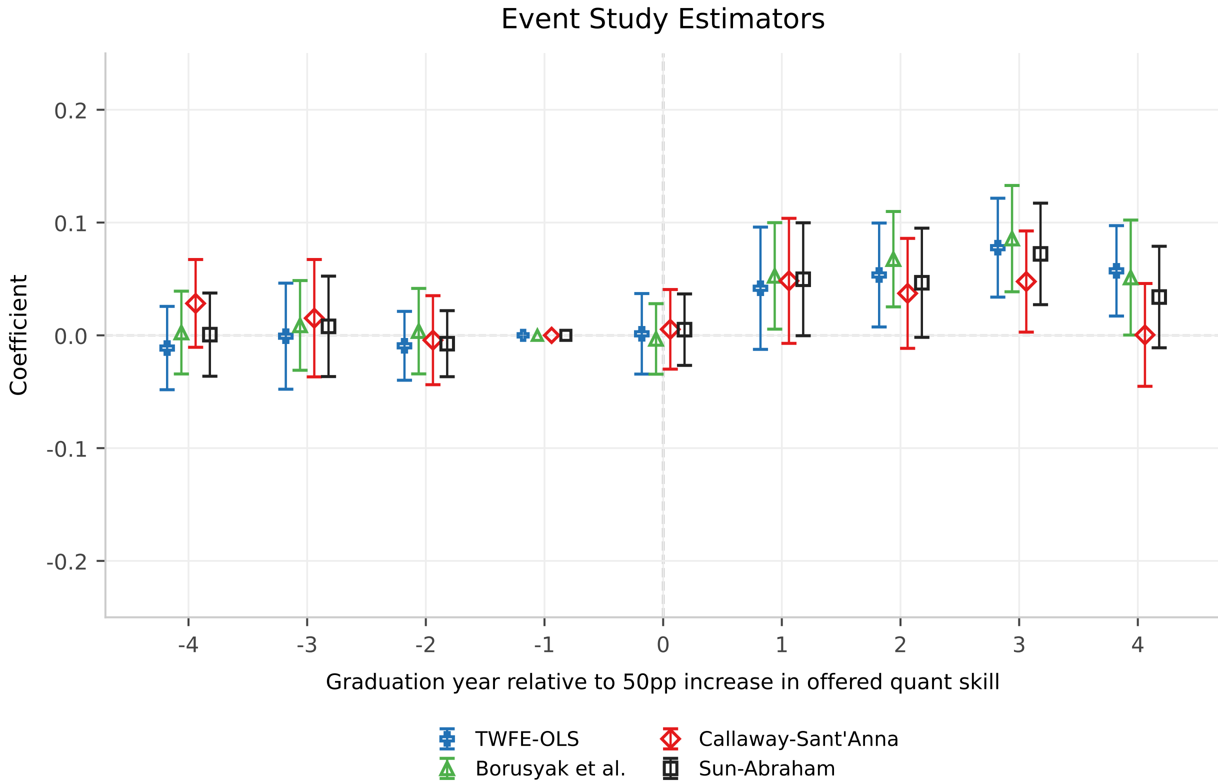

Notes: This figure overlays event–study plots constructed using four estimators: a dynamic version of the TWFE model, Sun and Abraham (2021), Callaway and Sant'Anna (2021) and Borusyak et al. (2024). The outcome is log quarterly earnings. Event time is defined by graduation year relative to the curricular change, where treated programs are those whose share of courses teaching quantitative skills increased by at least 50 percentage points in the event year. Cohorts graduating in the event year are partially treated ($t = 0$); those graduating the following year are fully treated ($t \geq 1$); earlier cohorts are untreated. The omitted category is $t = -1$. Error bars show 95% confidence intervals. Standard errors are clustered at the school×major level.

Figure A.5: Event Study: Peer Quality

Panel A. Avg. Peer HS Math Score

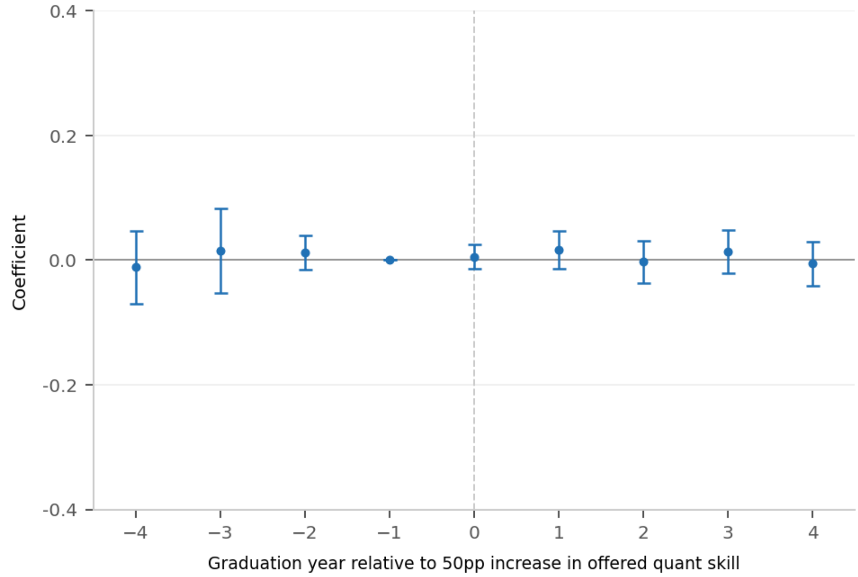

Panel B. Avg. Peer Lower-Division GPA

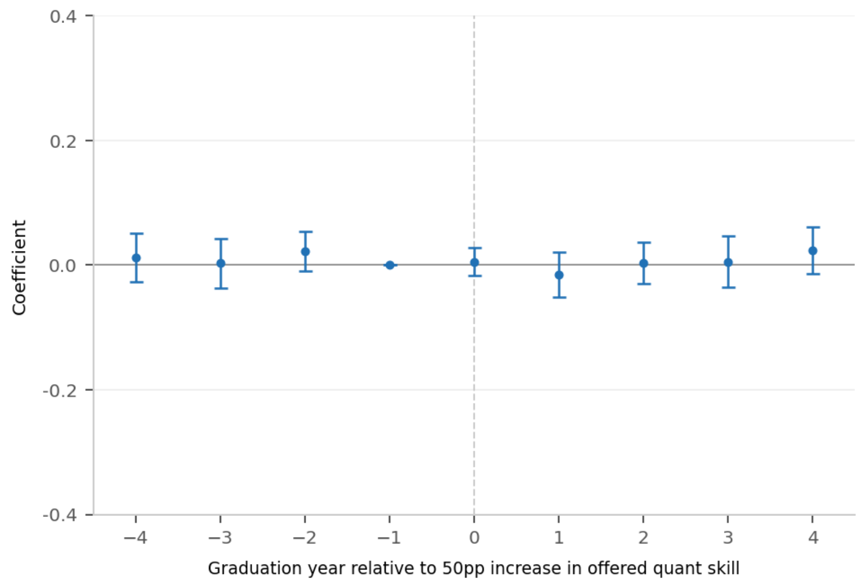

Notes: Each panel plots an event study for a different peer quality outcome: panel (a) shows the average HS math test score of a student's peers, and panel (b) shows the average lower-division GPA of a student's peers. Event time is defined by graduation year relative to the curricular change, where treated programs are those whose share of courses teaching quantitative skills increased by at least 50 percentage points in the event year. Cohorts graduating in the event year are partially treated ($t = 0$); those graduating the following year are fully treated ($t \geq 1$); earlier cohorts are untreated. The omitted category is $t = -1$. Error bars show 95% confidence intervals. Standard errors are clustered at the school×major level.
Variable construction. Both peer quality measures are constructed using a leave-one-out approach. For each upper-division course section, the mean HS math test score (panel a) and the mean lower-division GPA (panel b) of all *other* students enrolled in that section are computed. These section-level peer measures are then averaged across all upper-division courses in each student's schedule to obtain that student's average exposure to peer quality.

B Data and Variable Construction

B.1 Catalog Data Construction

To construct the course content dataset, I compile course catalogs from 27 public universities in Texas, covering academic years 2012 through 2020. Catalogs published in structured formats (e.g., HTML, XML) are scraped directly, while content from unstructured formats (PDF) is extracted using OCR and parsing tools. I exclude 13 public universities from the final sample because (1) they are newly established and not present in the Texas student administrative data during my sample period, (2) historical catalogs were unavailable or incomplete, or (3) the catalogs exist only as scanned PDFs that do not permit reliable text extraction. The 27 selected institutions include both flagship research universities and most mid-sized regional institutions. Together, they account for over 91% of four-year public university enrollment in Texas.

All text is cleaned to remove formatting artifacts such as HTML tags, line breaks, page headers, and special characters in course descriptions.

I unify cross-listed courses under a shared course ID to avoid duplication. Cross-listing occurs when a course is offered under multiple course codes, either across departments or levels. If unaddressed, this can lead to incorrect matches between course descriptions and administrative course records. I identify two common types of cross-listing in Texas catalogs. One type involves the same department prefix but different course numbers, such as BCOM 3310 and BCOM 3311 at UT Dallas, both titled "Business Communication." Another type involves different department prefixes with the same number and title, such as CLDP 3339 and PSY 3339, both titled "Educational Psychology," offered to Child Learning and Development and Psychology majors, respectively.

To detect these cases, I use a two-step approach. First, some cross-listed courses are explicitly marked in the catalog title (e.g., "BCOM 3310 / BCOM 3311"), which allows for direct identification. Second, when such indicators are absent, I rely on a text-similarity approach. I compare course titles and compute cosine similarity between course descriptions using a TF-IDF representation. Descriptions are preprocessed by removing stop words, lowercasing, and tokenizing the text. I use the `TfidfVectorizer` function from the `scikit-learn` package to transform descriptions into vectors, and I flag pairs with cosine similarity above 0.90 as likely cross-listed. This threshold is calibrated using manual inspection of a validation sample to ensure good precision without excessive false positives.

B.2 Construction of the Instrumental Variable

The IV is constructed to capture plausibly exogenous variation in students' exposure to skills through available coursework. The details are described below.

Step 1: Measuring Offered Skill at the Department Level. A department is the unit that offers courses, and all courses offered by the same department within a given institution share a common course prefix. Let *term* index academic terms, where each term corresponds to a specific (semester, year) pair (e.g., Spring 2015, Fall 2016). Let $sem(term) \in \{\text{Fall}, \text{Spring}\}$ denote the semester component of *term*. For each department p, institution m, academic term *term*, and class level $c \in \{\text{junior}, \text{senior}\}$, let $\mathscr{C}_{p,m,term,c}$ denote the set of courses offered by department p at institution m to students at level c in that term.

I define the department-level offered skill as the class-size-weighted share of courses in $\mathscr{C}_{p,m,term,c}$ that teach skill $s \in \{\text{quantitative}, \text{writing}\}$. Each course is coded as either teaching the skill or not, based on the GPT-4 classification described in Section 2.1. I weight courses by class size, assigning greater weight to larger classes since students are more likely to enroll in them. To avoid endogeneity arising from short-term swings in course popularity, I use each course's average enrollment across all years it is offered in the same semester type (Fall or Spring). For courses offered only once, this weight equals that single observed enrollment. This design ensures that variation in the measure is driven by whether a course is offered, rather than by year-specific demand fluctuations.

$$\text{OfferedSkill}^{(s)}_{p,m,term,c} = \frac{\sum\limits_{k \in \mathscr{C}_{p,m,term,c}} \mathbb{1}\{\text{Course } k \text{ teaches skill } s\} \cdot \overline{\text{Enroll}}_{k,c,sem(term)}}{\sum\limits_{k \in \mathscr{C}_{p,m,term,c}} \overline{\text{Enroll}}_{k,c,sem(term)}} \tag{5}$$

Here, $\overline{\text{Enroll}}_{k,c,sem(term)}$ denotes the average enrollment in course k at class level c, computed across all years the course is offered in the same semester as *term*. For example, this measure captures the proportion of junior-level courses offered by the Economics department at UT Austin in Spring 2015 that emphasize quantitative skills.

Step 2: Aggregating to Major-Level Offered Skill. While departments offer courses, students typically take classes from multiple departments to fulfill the requirements of a given major. For example, economics majors often take coursework from the ECON, MATH, and STAT departments, while public health majors may draw from HLTH, BIOL, and STAT.

Let $\mathscr{S}_{j,m}$ denote the set of students who graduate with major j from institution m. For each class level $c \in \{\text{junior}, \text{senior}\}$, I calculate the share of total course enrollments by students in $\mathscr{S}_{j,m}$ that fall within each department, pooling across all cohorts. I then define the set of *major-relevant departments* as

$$\mathscr{P}_{j,m,c} = \left\{ p : \frac{\text{Enrollments}_{j,m,c,p}}{\sum_{p'} \text{Enrollments}_{j,m,c,p'}} \geq 0.05 \right\},$$

that is, the set of departments that account for at least 5% of total course enrollments among students in $\mathcal{S}_{j,m}$ at class level c. I accordingly define *major-relevant courses* as those offered by departments in $\mathcal{P}_{j,m,c}$. Most majors are associated with two to five such departments. Appendix Figure A.3 shows the departments I classify as major-relevant for selected majors at UT Austin.

Focusing on major-relevant courses captures the realistic set of course options students face. First, students typically complete most of their degree requirements within a small, consistent set of departments, making the instrument more interpretable as variation in feasible skill exposure. Second, students do not have unrestricted access to all departments; many courses have prerequisites or are restricted to specific majors. Third, defining the course choice set at the department level avoids post-treatment bias, which would arise if we conditioned on actual student course-taking. Finally, from a policy perspective, departments—not individual courses—are the units that manage curricula and staffing, so this level of aggregation aligns with how course supply decisions are made.

To construct the *major-level offered skill*, I take a weighted average of the department-level offered skill measures across major-relevant departments. Weights are defined as the share of enrollments within $\mathcal{P}_{j,m,c}$:

$$w_{j,m,c,p} = \frac{\text{Enrollments}_{j,m,c,p}}{\sum_{p' \in \mathcal{P}_{j,m,c}} \text{Enrollments}_{j,m,c,p'}}, \qquad \text{for } p \in \mathcal{P}_{j,m,c},$$

so that $\sum_{p \in \mathcal{P}_{j,m,c}} w_{j,m,c,p} = 1$. The major-level offered skill is then:

$$\text{OfferedSkill}^{(s)}_{j,m,term,c} = \sum_{p \in \mathcal{P}_{j,m,c}} w_{j,m,c,p} \cdot \text{OfferedSkill}^{(s)}_{p,m,term,c}. \tag{6}$$

This measure reflects the proportion of major-relevant courses that teach skill s and are available to a typical student in major j and institution m at class level c during academic term *term*. For example, it captures the share of major-relevant courses available to junior economics majors at UT Austin in Spring 2015 that emphasize quantitative skills.

Step 3: Constructing the Student-Level Instrument. Finally, for each student i, I construct an individual-level instrument that captures exposure to skill s based on the major-relevant courses available during her upper-division years. Let $c(i, term)$ denote student i's class level in term *term*, based on cumulative credit hours. Let $\mathcal{T}_i^{\text{upper}}$ denote the set of academic terms during which student i is enrolled full-time with $c(i, term) \in \{\text{junior, senior}\}$. The raw instrument is defined as the average share of major-relevant courses offered during those terms that teach skill s:

$$\text{RawOfferedSkill}^{(s)}_i = \frac{1}{|\mathcal{T}_i^{\text{upper}}|} \sum_{term \in \mathcal{T}_i^{\text{upper}}} \text{OfferedSkill}^{(s)}_{j(i),m(i),term,c(i,term)}. \tag{7}$$

Each term receives equal weight, so students with longer upper-division sequences are averaged over more terms.

To ensure consistency with the treatment variable, I standardize this raw measure within each institution-major cell (j, m):

$$\text{OfferedSkill}_i^{(s)} = \frac{\text{RawOfferedSkill}_i^{(s)} - \mu_{j,m}^{(s)}}{\sigma_{j,m}^{(s)}}, \tag{8}$$

where $\mu_{j,m}^{(s)}$ and $\sigma_{j,m}^{(s)}$ are the sample mean and standard deviation of $\text{RawOfferedSkill}^{(s)}$ computed over students in the same school–major cell $\mathscr{S}_{j,m}$. This standardization makes the instrument directly comparable to the standardized treatment variable $\text{Skill}_i^{(s)}$ and allows the 2SLS coefficient β in Equation 4 to be interpreted as the effect of a one SD increase in skill exposure on earnings.

For example, for an economics major at UT Austin, this measure reflects how much more (or less) quantitative coursework was available to that student during her junior and senior years relative to the average economics major at UT Austin. Because the instrument is based on course offerings rather than on actual enrollment decisions, it captures plausibly exogenous variation in skill exposure.