

Federal Reserve
Bank of Dallas

Degrees for GPTs: The Effects of Generative AI on Skill Demand and Supply

Samuel Dodini, Tucker Smith and Yaling Xu

Working Paper 2642 **November 2026**

Research Department

https://doi.org/10.24149/wp2642

Working papers from the Federal Reserve Bank of Dallas are preliminary drafts circulated for professional comment. The views in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Degrees for GPTs: The Effects of Generative AI on Skill Demand and Supply*

Samuel Dodini† Tucker Smith‡ and Yaling Xu§

November 2026

Abstract

We examine how the rapid deployment of generative AI affects labor market outcomes and educational decisions of recent college graduates and current students. Using administrative records that link all students at Texas public four-year universities to state earnings data, along with detailed job postings data, we measure exposure to AI automation by mapping occupational task measures onto college majors. We find that graduates from majors with 10 percentage points higher exposure to AI automation experienced 10% lower first-year earnings following ChatGPT's release in November 2022, driven by both reduced employment rates and lower conditional earnings. In response to soft labor demand, graduates from more-exposed majors were more likely to enroll in graduate school. However, they largely continued studying the same field, and earnings declines for recent master's recipients in more-exposed fields suggest that graduate-level education does not unlock complementarities large enough to offset "automation penalties" in these fields. In contrast, graduates from exposed majors whose coursework tilted toward less-exposed fields experienced smaller earnings losses. Undergraduate enrollment data shows that younger students are pivoting away from AI-exposed fields as of Fall 2025, consistent with the delayed effects of prior sectoral shocks on major choice.

JEL Codes: J24, J23, J31, O33, I26

Keywords: generative AI, college major choice, human capital, labor demand

*We thank Robert Leigh for excellent research assistance. We also wish to thank audience members and participants at the Dallas Fed, University of Texas at Austin, Furman University, the Peterson Institute for International Economics, and participants at the 96th Southern Economics Association Annual Meeting. The views expressed in this paper are those of the author(s) and do not necessarily reflect those of the Federal Reserve Bank of Dallas or the Federal Reserve System. The conclusions of this research also do not necessarily reflect the opinion or official position of the University of Houston Education Research Center, the Texas Education Agency, the Texas Higher Education Coordinating Board, the Texas Workforce Commission, or the State of Texas.

†Samuel Dodini, Federal Reserve Bank of Dallas, samuel.dodini@dal.frb.org.

‡Tucker Smith, Federal Reserve Bank of Dallas, tucker.smith@dal.frb.org.

§Yaling Xu, FAIR Center, Department of Economics, Norwegian School of Economics, yaling.xu@nhh.no.

1 Introduction

The release of ChatGPT in November 2022 rapidly expanded the use of generative artificial intelligence (GenAI) and intensified debate about how these tools may reshape work. GenAI tools can perform complex writing, coding, and information synthesis tasks—activities central to many knowledge-intensive occupations. While GenAI has the potential to raise productivity and aggregate income, its ability to automate tasks previously performed by workers may substantially alter labor demand across occupations and the skills that employers value.

These changes are particularly important for college students and recent graduates. Early evidence suggests that the new technology may complement experienced workers because of their tacit knowledge base gained through on-the-job learning but may substitute for less-experienced labor (Brynjolfsson et al., 2025; Hosseini and Lichtinger, 2025). College graduates typically enter the labor market with formal training but limited professional experience, and much of the human capital investment of recent graduates was determined before the potential capabilities of GenAI became apparent. The coincident timing of GenAI's rapid development with unusually large increases in the unemployment rate for young college graduates (Figure 1) amplifies these concerns. Students and recent graduates also offer a unique opportunity to study how the current and prospective workforce is affected by GenAI and adapts to it: attained human capital and evolving skill investments that are difficult to observe among established workers are directly measurable within the formal education system.

In this paper, we examine how GenAI affects the early-career labor market outcomes of college graduates and students' educational choices. Our analyses shed light on both the labor demand consequences of GenAI and the role of educational investments in shaping young workers' exposure and responses to the technological shock. These questions extend beyond the cohorts and analysis period we observe: previous research documents that graduates who enter the labor market when demand is weak suffer earnings losses that persist for years (Kahn, 2010; Schwandt and von Wachter, 2019; Rothstein, 2023), suggesting that an "automation penalty" at entry may shape the careers of a generation of college-educated workers. In light of this evidence, early disadvantages associated with GenAI exposure may have lasting consequences. Overall, understanding the skill demand and supply responses to GenAI is essential for current and prospective workers, firms, and policymakers to effectively navigate the transition into an AI world.

Despite these stakes, there is relatively little causal evidence of GenAI's effects on new college graduates and entrants (Orr et al., 2026). While many data sources allow for analysis

of GenAI's effects on labor demand, a key challenge to estimating how changing demand across skills affects the earnings and employment outcomes of *individuals* is observing both an individual's trained skill set and their labor market outcomes. We overcome this limitation by using a panel of individual-level administrative records for all students at public two- and four-year universities in Texas that link pre-collegiate academic achievement, post-secondary course histories, major choice, and graduation timing along with post-college employment, earnings, and industry of employment in Texas. We use a difference-in-differences strategy that exploits variation in GenAI automation exposure across college majors and cohorts before and after ChatGPT's release, allowing us to estimate the effects of the shock on earnings of new labor market entrants and on student investments in skills across fields of study. The richness of our data across graduation cohorts, including detailed information on students' academic preparation, demographic backgrounds, fields of study, and institutional characteristics, also allows us to examine who is most affected by GenAI and whether the earnings effects we observe for new graduates extend to workers further up the experience distribution.

We measure exposure to GenAI automation using occupation-level measures derived from observed Claude usage (Handa et al., 2025), which capture the share of occupational tasks associated with automation-oriented uses of the technology. We map these measures to college majors using the occupational distribution of pre-ChatGPT job postings listing each field of study as a desired qualification. A college major is considered to have higher automation exposure if more-exposed occupations list the major as a target degree in their job ads. This approach treats each major as a bundle of potential occupational destinations and allows us to compare graduates from fields that differ in their exposure to GenAI automation. We also construct individual-level measures using students' detailed course histories in which those taking more courses in exposed majors are considered more exposed on an individual basis. This captures variation in exposure across graduates who earned degrees in the same major but pursued different combinations of coursework.

We begin by examining changes in skill demand in Texas using job posting data, providing complementary evidence on the labor-demand channel through which GenAI may affect recent graduates. Consistent with emerging literature (e.g., Brynjolfsson et al., 2025), we find that job postings fell disproportionately for occupations with greater exposure to GenAI automation following ChatGPT's release. Within industries, occupations with a 10 percentage-point higher share of automatable tasks experienced approximately 5% fewer postings within a year of the release, with the relative decline approaching 8% by mid-2026. We also document adjustments within incumbent firms: those whose pre-ChatGPT recruitment was more concentrated in exposed occupations subsequently reduced their total postings and shifted

their remaining recruitment toward less-exposed occupations. Together, these results establish the labor-market relevance of our exposure measure and provide context for examining how graduates from more-exposed fields have fared following the introduction of GenAI.[1]

We then examine how the GenAI shock to labor demand affected earnings and employment outcomes for new graduates. We show that earnings in the initial year after graduation fell by 10% for 2024 graduates from majors with a 10 percentage-point greater share of tasks that could be automated by AI. Declines in both employment rates and earnings conditional on employment contribute to the total effect, and earnings have yet to recover for graduates from more-exposed majors for whom we observe multiple years of labor market outcomes. The effects are remarkably similar across race/ethnicity, family income, and gender. These results indicate that GenAI's estimated impact on relative labor demand across occupations (e.g., Brynjolfsson et al., 2025) has harmed new labor market entrants who would typically sort into more-exposed occupations (Orr et al., 2026). Our results for job postings and early-career earnings are robust to several alternative explanations from the recent literature, including monetary policy sensitivity, industry-specific shocks, and the dynamics of remote work/work-from-home.

How should students respond when a new technology reduces demand for the skills taught in their field? A simple human capital framework posits two broad strategies for investing in skills that complement GenAI. First, students can "upskill," or deepen their expertise within exposed fields. This strategy pays off if automation removes less-expert tasks in an occupation and raises the value of expertise in the tasks that remain (Autor and Thompson, 2025; Hosseini and Lichtinger, 2025). Alternatively, students can "pivot," or shift their skill investments toward fields that GenAI is less able to automate. The unique richness of our data allows us to closely track these choices.

We provide evidence that some recent graduates upskilled in response to the shock. Facing a relatively soft labor market, graduates from more-exposed majors were 6% more likely to enroll in graduate school, consistent with a decline in opportunity costs. However, we find little evidence that students who continued to graduate school pivoted toward studying less-exposed fields. For example, approximately two-thirds of 2024 computer science bachelor's recipients who continued their education enrolled in a computer science program, essentially the same share as among pre-GenAI cohorts.

We also explore undergraduate major and coursework selections for evidence of pivoting or upskilling. Fields with greater AI exposure – driven by computer science – increased in popularity over the decade leading up to the GenAI shock. Since the release of ChatGPT, existing enrollment growth trends in more-exposed fields continued for upperclassmen

[1]These results also hold for the nation as a whole when we expand the scope beyond Texas.

through Fall 2025, consistent with the high cost of switching majors. This continued growth in high-exposure major enrollment does not appear to mask course-level pivots toward less-exposed fields nor upskilling efforts via increased specialization among students who declared for more-exposed majors prior to the release of ChatGPT: such students completed similar coursework in the subsequent semesters as prior cohorts. In contrast, in majors with a 10 percentage-point greater share of automatable tasks, Fall semester Freshman enrollment declined by 7% more between 2022 and 2025. This suggests that students that entered college in the wake of the shock had greater leeway to shift their field choices, a result consistent with pivoting enrollment responses to prior sectoral shocks (Weinstein, 2022). Adjustment also became evident as these later entrants progressed through college: by Fall 2025, students who entered in Fall 2023 with more-exposed majors shifted toward less-exposed coursework relative to prior trends and became more likely to switch majors.

Given evidence of skill investment responses by recent graduates and current undergraduates, we assess early signs of whether upskilling or pivoting strategies yield labor market benefits. In particular, we provide two pieces of evidence that upskilling investments have not fully compensated for the new technology's direct substitution effects and that pivoting has generated protective benefits. First, we define an individual component of a graduate's exposure to GenAI as a weighted average of the automation exposure of the individual *courses* they took in school minus their own major's automation exposure. Positive values indicate that a student's coursework was more exposed to AI automation than their major's typical occupational destinations. Estimates of specifications that interact major-level and individual components of GenAI exposure indicate that earnings losses were even larger for graduates from more-exposed majors with stronger concentrations of courses in exposed fields. On the other hand, diversification of course portfolios to include less-exposed fields alleviated approximately 25% of the would-be earnings losses for graduates from high-exposure majors. These results suggest that less-exposed coursework might provide a potential buffer against changes in labor demand. Second, we show that new master's recipients from more-exposed fields experienced similar relative earnings declines to those of four-year graduates after the release of ChatGPT, suggesting that graduate-level upskilling does not unlock complementarities large enough to offset the automation penalty.

Our empirical approach compares outcomes across graduation cohorts from majors with different degrees of AI automation exposure, controlling for persistent differences across fields and university-specific cohort shocks. Interpreting these differences as causal responses to GenAI exposure requires that outcomes of students from more- and less-exposed majors would trend in parallel if they experienced the same treatment "dose" (i.e., a similarly sized

decline in labor demand) (Callaway et al., 2025).[2] We address four key concerns that would threaten causal interpretation of our estimated effects on labor market outcomes of new graduates. First, we show the robustness of our earnings results to controlling for time-varying industry-specific shocks, such as a general boom and bust in tech hiring or the differential incidence of monetary tightening (Zens et al., 2020). Similarly, we show robustness to controlling for time-varying effects of remote work uptake (Lambert and Schindler, 2026; Emanuel et al., 2026; Mask, 2026). We also control for school-wide shocks common across cohorts, such as remote learning policies adopted during the pandemic. Finally, we provide evidence that our estimated effects are not spurious relationships appearing from selection on ability into more-exposed majors or withdrawal from the labor market by graduate school enrollees.

Our paper contributes to several strands of research. First, we add to the emerging literature on GenAI's effects on the labor market, particularly for young and inexperienced workers. Recent papers document employment declines for early-career workers in exposed occupations (Brynjolfsson et al., 2025), reduced junior hiring at adopting firms (Hosseini and Lichtinger, 2025), and slowing wage growth in AI-exposed occupations (Azar et al., 2025). We show that AI-induced declines in labor demand reduced early-career earnings and employment, and that these declines survive the alternative explanations proposed in the literature (monetary tightening, pandemic-era over-hiring, and remote work) (Iscenko and Millet, 2026; Frank et al., 2026; Emanuel et al., 2026; Mask, 2026; Lambert and Schindler, 2026; Liu et al., 2025). The closest paper to ours is Orr et al. (2026), which also examines early-career earnings and employment across college majors using linked graduation and earnings records.[3] Our administrative data allow us to use pre-collegiate test scores to proxy for student ability, so we can more precisely demonstrate that our results are not driven by time-varying selection on ability into more-exposed majors.

Our broader contribution relative to Orr et al. (2026) is to provide evidence of how recent graduates and current students are adjusting to the GenAI shock, enabled by our dataset's rich information on course history and subsequent enrollment. Consistent with a basic human capital model, prior work shows that students shift field choices after sectoral

[2]Formally, we identify a convex combination of "slope" effects (causal responses) of marginal increases in AI exposure on outcomes under this "strong" parallel trends assumption (Callaway et al., 2025).

[3]Orr et al. (2026) use an event-study design to compare changes in employment rates and initial earnings for graduates from majors with greater versus lesser exposure to GenAI. They use linked graduation and earnings records from the Census Bureau's Post-Secondary Employment Outcomes (PSEO), and also find that early-career earnings and employment fell for more-exposed relative to less-exposed graduates after the release of ChatGPT. Their dataset's earnings records for 48 states make it particularly robust to cross-state migratory responses to the shock. We provide evidence suggesting there was no such response to the GenAI shock among Texas college graduates.

demand shocks, though often with a delay (Conzelmann et al., 2023; Weinstein, 2022). GenAI complicates this prediction because it affects not only the demand for skills but also how they are produced (Light, 2026; Peterson, 2025): if the technology lowers the effort cost of coursework in the same fields in which it automates work, marginal students may *increase* their enrollment in those fields. We study both strategies available to students – upskilling within exposed fields and pivoting away from them – across four margins: the major chosen at entry, switching majors, coursework within a major, and graduate school enrollment and field. We find that the perceived decline in demand for automatable skills dominates lower effort costs for students entering college after observing the GenAI shock, while enrolled students far into their collegiate studies and returning graduates largely stayed in exposed fields. The pivot response of younger cohorts is consistent with evidence that European students reduced apprenticeship applications for AI-exposed occupations (Goller et al., 2025) and weigh automation risk when choosing career paths (Gschwendt et al., 2026).

By providing early evidence on the returns to each strategy, we also contribute to work on whether expertise complements AI (Autor and Thompson, 2025; Hosseini and Lichtinger, 2025). Using each student's own course history, we find that course-level pivoting, not upskilling, protected early-career earnings, and that master's recipients in more-exposed fields experienced similar declines in earnings to bachelor's recipients. Our findings align with those of Hyman et al. (2025), who show that workers displaced from AI-exposed occupations in recent years gained only small returns from retraining programs that reflected deeper investment in AI-exposed work. Instead, retraining programs that enabled pivots to less-exposed occupations garnered substantial returns. We add evidence on the cohorts and margins that training programs do not reach: students making initial field choices, for whom pivoting is cheapest, and those already committed to exposed fields, who largely did not adjust if they were far into their studies. These patterns suggest that for those graduating in the middle of this AI transition, the expertise in exposed fields that complements GenAI may need to be acquired through experience rather than additional coursework or formal training. For those early in their educational paths during this transition, pivoting within the formal education system appears to be a viable and beneficial option.

More broadly, our results relate to the literature on the effects of labor demand shocks on new labor market entrants. While many studies find evidence of lasting harm from graduating college during a recession (e.g., Kahn, 2010), relatively little is known about the effects of field-specific demand shocks on college graduates (Berge, 2018). We show that negative earnings effects of the GenAI shock have lasted at least two years for graduates from more-exposed majors. Results from the "scarring" effects literature and early research on AI suggest multiple reasons why this labor market harm may persist. Liu et al. (2016)

and Brunner and Kuhn (2014) provide evidence that skill mismatch and hindered accumulation of early-career occupation- and industry-specific human capital are mechanisms for the persistence of the negative effects of entering the labor market during cyclical downturns. If GenAI complements expertise (Autor and Thompson, 2025; Hosseini and Lichtinger, 2025), then exposed college graduates without opportunities for on-the-job skill-building are doubly disadvantaged.

The paper proceeds in three steps. First, we estimate the effects of emerging GenAI adoption on labor demand as measured by job postings in Texas collected by Lightcast (formerly Burning Glass Technologies). Second, we use administrative data from Texas to directly estimate the effects of GenAI automation exposure on recent graduates' earnings and employment prospects, including heterogeneity by student characteristics and institutional quality. This examines whether there is an "automation penalty" for exposed majors and who is most affected by this penalty. Third, we examine student responses to this change in labor demand and whether initial labor market outcomes suggest such adjustments pay off. We frame this analysis around a human capital model that lays out the tradeoffs and constraints faced by students in choosing how to respond to this shock and how GenAI might differ from prior technological changes via its effect on the effort cost of attaining particular majors.

2 Data and Empirical Framework

2.1 Data Sources

Our empirical analysis uses two main data sources. We describe each below.

Lightcast: We quantify labor demand using job postings data collected since 2010 from over 40,000 job boards and company websites by Lightcast (formerly known as "Burning Glass"). These data represent the near universe of online U.S. job postings, are processed to remove duplicate postings, and are parsed into standardized fields including job title, six-digit Standard Occupational Classification (SOC) code, North American Industry Classification System (NAICS) industry code, location, and desired education levels and work experience. Moreover, firm IDs allow for analysis of within-firm changes in posting behavior.

Jobs posted online are a non-representative subset of all U.S. vacancies. Following Hershbein and Kahn (2018), we assess the representativeness across sectors of job openings observed in Lightcast by comparing them to those observed in the BLS Job Openings and Labor Turnover Survey (JOLTS), a monthly nationally representative survey of approximately 16,000 business establishments sampled from unemployment insurance records. Figure B1

shows the share of openings from 2010 through 2024 across sectors (2-digit NAICS codes) in our Lightcast postings data and JOLTS, respectively. Across the entire sample, the Lightcast data overrepresent postings in financial services and insurance, information services, manufacturing, mining, professional and business services, real estate, and wholesale trade. The Lightcast data underrepresent accommodation and food services, arts, construction, health services, other services, public administration/government, and transportation and warehousing.

The differences in shares across industries, most of which are small in magnitude, do not threaten the internal validity of empirical analysis that uses the Lightcast data as a proxy for labor demand in a panel setting if such representativeness remains roughly constant over time. Figure B2 presents the relationship between gaps in shares of openings in the Lightcast data vs. JOLTS in a baseline period (2018 - 2022) and the respective gaps in the years following the release of ChatGPT (2023 - 2024). Consistency in the representativeness of Lightcast openings supports using trends in postings as evidence of changing labor demand, rather than changing composition of the data.

Our analysis of the effects of generative AI on openings across occupations uses Lightcast postings structured in cells defined by 6-digit SOC code, 4-digit NAICS code, and quarter with observations from Q3 2020 to Q2 2026. We also structure the data in firm-industry-quarter cells to examine job posting behavior within firms. To match the setting of our educational analysis, we limit the sample to postings from Texas. Table 1 presents descriptive statistics of the occupation-industry level dataset. We construct this from approximately 28.3 million job postings in our analysis sample, and we drop postings missing occupation, industry, or company identifiers (roughly 873,000 job ads).

University of Houston Education Research Center: Our primary data source consists of statewide administrative records available through the University of Houston Education Research Center (UHERC), linking individual-level K–12 education, postsecondary education, and workforce records from 1994 to the present. K-12 data from the Texas Education Agency include statewide standardized test scores, course schedules, graduation records, and demographic information for all students who attended public schools in Texas. Postsecondary data from the Texas Higher Education Coordinating Board include college entrance exam scores (e.g., SAT scores), semester-level information on course enrollment, declared majors, academic standing, and degree attainment for all public colleges and universities in the state. Beginning in 2012, these records also contain detailed course histories. Finally, wage reports from the Texas Workforce Commission include quarterly earnings and employment records by employer for all Texas workers earning wages in positions covered by

state unemployment insurance.[4] The wage records report employer industry codes (6-digit NAICS) in all years but do not report firm identifiers or occupation.

We use the UHERC data to construct three analysis samples. For analysis of the effects of GenAI on the labor market outcomes of new college graduates, we construct a panel of cohorts of May graduates from public four-year universities in Texas from 2016 to 2024. We observe each graduate's major, GPA, course history, college entrance exam score, demographic information, and subsequent enrollment in graduate programs at Texas public universities. Using linkages to earnings records, we also define a graduate's cumulative earnings during their first four quarters since graduating and the industry in which they earned the most during that period.[5] Table 2 presents summary statistics of the panel of May graduates receiving bachelor's degrees. We similarly define a panel of new master's degree recipients from Texas public universities from 2016 to 2024. Table B1 presents descriptive statistics of this panel.

We also construct a panel of four-year college *entrants* to examine the effects of GenAI on choices in fields of study, observing each student's initial major choice, college entrance exam scores, and demographic information, and subsequent course selections. Table 3 presents descriptive statistics of this panel of students that enter a public four-year college with a declared major.

2.2 Defining AI Exposure

We follow recent literature (e.g., Eloundou et al., 2023) in measuring an occupation's (or college major's) "exposure" to automation by GenAI as the share of its typical tasks that AI could perform more efficiently. Eloundou et al. (2023) define task-level exposure indices according to a rubric applied by human annotators and GPT-4 that scores O*NET tasks as exposed if an LLM could complete the task with similar or improved quality in less than 50% of the time of the average human worker. We follow the authors in defining occupation-level exposure ($OccExp_o$) as the weighted average of an occupation's O*NET core and supplemental tasks.[6] The resulting "theoretical" exposure measure roughly reflects the share of an occupation's tasks that could be automated by GenAI tools.

Both survey data (Bick et al., 2025) and internal records of GenAI usage (Massenkoff and McCrory, 2026) support that the Eloundou et al. (2023) measure of occupation-level AI exposure correlates with observed usage. However, legal constraints, human verification

[4]Self-employed workers, independent contractors, many federal employees, members of the military, and participants in the informal sector are not covered by the state unemployment insurance system.

[5]We consider May graduates and measure earnings starting in Q3 of their graduation year.

[6]In particular, we replicate the authors' "beta" measure, which measures efficiency gains from either an LLM alone or an LLM in combination with complementary software.

requirements, and other hurdles yield notable divergences between theoretical exposure and realized usage (Bick et al., 2025; Massenkoff and McCrory, 2026). In practice, theoretical exposure measures also vary dramatically by which LLM is used to rate GenAI's ability to perform a given task (Yin et al., 2026).

Given our goals of estimating the real-time effects of GenAI on labor demand and education choices, our primary specifications utilize a task-based exposure measure developed by Anthropic based on observed Claude usage (Handa et al., 2025). The Anthropic exposure measure for a given occupation reflects the share of its O*NET tasks that Claude users used the tool for automation as of Fall 2024. Although the Anthropic measure has a stronger implicit relationship with GenAI usage than the measures from Eloundou et al. (2023), endogenous uptake of Claude for particular tasks makes the Anthropic measure more susceptible to bias. Our approach balances the need to adequately capture adoption while avoiding strong endogeneity of adoption over time by fixing adoption and usage early in the adoption cycle. However, a confounding labor demand shock could lead workers to use Claude for tasks that their firm can no longer afford to hire a human to do; a confounding labor supply shock could cause workers to employ Claude for tasks that there are no longer quality candidates available to fill. Therefore, we test the robustness of our results to using theoretical exposure measures in Appendix C.

Although many students likely have a career path in mind, in college they choose fields of study to pursue, rather than occupations directly (Arcidiacono et al., 2012). We define a field of study's (formally, four-digit Classification of Instructional Programs codes) exposure to AI ($MajExp_m$) as the weighted average exposure of occupations with pre-AI job postings listing the field as a desired candidate major:

$$MajExp_m = \sum_o \omega_{mo} \, OccExp_o$$

where ω_{mo} is the share of job postings listing major m that were for occupation o. This index treats education fields as bundles of occupational destinations and maps an external technology measure into the education setting using pre-period demand and transition patterns (from Q1 2018 through Q2 2022). The resulting major-level exposure measure reflects the average share of tasks in the pre-AI job set for that graduate at risk of automation.

Table 4 lists majors across bins of $MajExp_m$. Major-level exposure levels are fairly intuitive: the least-exposed majors include psychology, nursing, and education, while the most-exposed majors include computer science, computer engineering, and language studies.

Students graduating with the same major may nevertheless make different investments in skills through their course selections. To capture this within-major variation, we define

each graduate's "idiosyncratic" coursework exposure as the difference between the course-weighted average of GenAI automation exposure across the fields in which they completed courses and the automation exposure of their graduating major:

$$CourseExp_i = \sum_m \omega_{im} \, MajExp_m$$
$$IdExp_{im} = CourseExp_i - MajExp_m$$

where ω_{im} is the share of student i's completed coursework that took place in subject m. $IdExp_{im}$ captures whether a student's course portfolio is tilted toward fields associated with more- or less-AI-exposed occupations relative to their major. Positive values indicate greater exposure through coursework, while negative values indicate lower exposure. For students in exposed majors, the measure maps onto two adjustment strategies that we propose in Section 4.1: concentrating coursework in the major or in similarly exposed fields ("upskilling") raises idiosyncratic exposure, while taking courses in less-exposed fields ("pivoting") lowers it.[7] We use this measure to test for heterogeneous labor market effects of the GenAI shock across students with course bundles consistent with either of these strategies.

2.3 Identification

Our research design seeks to identify the effects of GenAI on labor market outcomes and educational decisions. To identify these effects, we leverage the release of ChatGPT by OpenAI in November 2022 as a shock to both information and technology that began rapid GenAI proliferation to commercial uses, carrying differential incidence across occupations and fields of study. In particular, we estimate the following event-study specification:

$$y_{imct} = \sum_{\tau \neq t^*-1}^{\tau} \pi_\tau^{ES} \mathbf{1}\{t = \tau\} \, MajExp_m + \Phi S_i + \alpha_m + \alpha_{ct} + \epsilon_{imct} \tag{1}$$

where y_{imct} denotes the outcome for individual i, who graduated from (or entered) school c in year t with major m. We regress this outcome on the major's share of automatable tasks, $MajExp_m$, interacted with cohort indicators, $\mathbf{1}\{t = \tau\}$. Major α_m and school-by-cohort α_{ct} fixed effects absorb time-invariant differences in student outcomes across majors (e.g., if Economics majors always earn more on average than English majors) and time-varying shocks common to students across all majors at the same school (e.g., if labor demand

[7] Idiosyncratic exposure and specialization within a student's major coincide only for exposed majors. For a student in a less-exposed major, specializing further lowers idiosyncratic exposure, and a student can raise their exposure without specializing by taking exposed courses outside their major.

fell for all college graduates from university X during the pandemic). Thus, the operative comparison is students in more exposed majors compared to their peers graduating the same year from the same institution. This approach rules out local macroeconomic shocks affecting the local labor market and overall job prospects of graduates from the school, school-wide policy changes, school reputation, or any other institutional features that affect all students.

We additionally control for student-level demographic characteristics S_i, including race, ethnicity, gender, and college entry-exam scores.[8] The event study coefficients of interest $\hat{\pi}_\tau^{ES}$ reflect the differences in outcomes associated with a 10 percentage-point difference in a major's share of automatable tasks, relative to this difference among the final cohort whose outcomes we observe before the release of ChatGPT. For our estimations of the effects of exposure to AI automation on earnings and employment in a graduate's first year in the labor market, graduates from the 2021 cohort serve as the reference group – their earnings from 2021:Q3 through 2022:Q2 were realized before ChatGPT's release.

$MajExp_m$ is an inherently continuous measure of a major's treatment "dose" of automation risk, and no major is completely unexposed to AI. Interpreting π_τ^{ES} for post-ChatGPT cohorts as a *causal* response to greater AI exposure requires assuming that outcomes of students from more- and less-exposed majors would trend in parallel if they experienced the same treatment dose (i.e., a similarly sized decline in labor demand) (Callaway et al., 2025).[9] Throughout, we report $\hat{\pi}_\tau^{ES}$ scaled to a 10 percentage-point difference in exposure. Even under the required assumption, $\hat{\pi}_\tau^{ES}$ reflects a weighted average of the "slope" effects of marginal increases in GenAI exposure on outcomes across the exposure distribution, rather than a discrete 10 percentage-point change for any particular major (Callaway et al., 2025).[10]

Any time-varying shock correlated with $MajExp_m$ would threaten identification of π_τ^{ES}. We address four forms of such threats. First, one may be concerned that graduates from more-exposed majors – particularly computer science students – disproportionately find work in the tech sector, and that a reversal of a post-pandemic hiring boom in tech could be mistakenly interpreted as effects of AI. More generally, monetary policy can have differential

[8] We observe either an SAT score or ACT score from the admissions records of 72% of four-year graduates between 2012 and 2024. We convert composite ACT scores to SAT scores using a concordance table published by ACT Education Corp (ACT and College Board, 2018).

[9] Because no major is a pure control group that does not experience the AI shock, comparisons of student outcomes across more- vs. less-exposed majors cannot identify the effects of AI on student outcomes under solely the standard difference-in-differences assumption of parallel trends in the *absence* of the GenAI shock. Callaway et al. (2025) denote the above assumption "strong" parallel trends. Under strong parallel trends, π_τ^{ES} is equivalent to a convex combination of the "slope" effects of marginal increases in AI exposure on outcomes and to a convex combination of two-by-two comparisons of changes in outcomes of more- vs. less-exposed majors.

[10] We acknowledge that the underlying causal-response functions of student outcomes with respect to major-level AI exposure may not be linear. Therefore, we present estimates from specifications including higher-order polynomials of $MajExp_m$ in Section C.1.

bite across industries (Zens et al., 2020), and sharp hikes in interest rates during 2022 and early 2023 may have caused coincident reductions in hiring in sectors that typically hire more graduates from AI-exposed majors. We alleviate such concerns by testing the robustness of estimated effects on earnings to controlling for industry-by-year fixed effects, which limits identification of π_{τ}^{ES} to comparisons of students from different majors that find employment in the same industry.

Second, one may be concerned that universities that typically produce more graduates in majors with higher AI exposure experienced school-wide shocks that harmed outcomes of their graduates (e.g., remote learning policies adopted during the COVID-19 pandemic). To address this concern, our main specification controls for school-by-cohort fixed effects, limiting identification to comparisons of students with different majors from the same school. Controlling for college entry-exam scores, a proxy for student ability, addresses a third concern: bias from changes in the pre-collegiate academic achievement of students studying more- vs. less-exposed majors (e.g., if computer science became more popular among students with weaker pre-collegiate achievement in recent years).

Finally, estimates of π_{τ}^{ES} could be confounded by occupation-specific demand shocks that correlate with GenAI exposure. One particular shock of note is the rapid expansion of remote work opportunities in the 2020s. By raising the cost of monitoring workers and slowing on-the-job learning, work-from-home (WFH) proliferation could also reduce labor demand for young workers (Lambert and Schindler, 2026; Emanuel et al., 2026; Mask, 2026). We define major-level shares of pre-ChatGPT Lightcast job postings that offered remote work and test robustness to controlling for WFH exposure interacted with cohort fixed effects.

3 Labor Demand and New College Graduates

3.1 AI and Labor Demand

Interpreting changes in outcomes of graduates from AI-exposed majors as evidence of automation requires that task-based exposure measures capture a real shift in employers' demand for the skills associated with each field of study. We test this in the same labor market that our graduates enter using job postings data from Lightcast. If GenAI reduced demand for occupations comprised of automatable tasks, job postings for those occupations should fall relative to postings for less-exposed occupations after ChatGPT's release, the decline should begin or accelerate with the release rather than continue an earlier trend, and the decline should not be explained by other shocks that hit similar occupations around the same time, such as monetary tightening or the spread of remote work. Because job postings

measure firms' hiring demand directly and skew toward the early-career positions that new graduates fill (Figure B3), they are well suited for this test.

To measure the effects of the emergence of GenAI tools on labor demand, we estimate Poisson event studies from a quarterly, occupation-industry level version of Equation 1 with posting counts on the left-hand side.[11] In the model, we measure occupations using O*NET occupation codes and industry using four-digit NAICS codes.

We structure the data in two ways so as to leverage both firm-level and occupation-industry level variation of susceptibility to automation within the state of Texas. In the first model, we leverage occupation-industry-quarter observations and include occupation-industry cell fixed effects and industry-quarter fixed effects. Formally, we estimate the following Poisson regression:

$$E(y_{okt} \mid \mathbf{x}_{okt}) = \exp\left(\alpha_{ok} + \alpha_{kt} + \sum_{\substack{\tau \\ \tau \neq t^*-1}} \pi_{\tau}^{ES} \, \mathbf{1}\{t = \tau\} \, OccExp_o\right) \tag{2}$$

where k denotes the industry and $t^* - 1$ is the last quarter before ChatGPT's release. In the second model, we structure the data as firm-quarter observations (y_{fkt}) on a balanced panel of firms who were present in the posting data through the entire panel and estimate the model with firm and industry-quarter fixed effects (i.e., replace α_{ok} with α_f). We define exposure at the firm level as the mean exposure of the occupations listed in the firm's job postings before the launch of ChatGPT: if a firm posted more intensely for exposed occupations before Q3 2022, that firm is considered more treated in the post period. This balanced panel allows us to estimate the effects of AI exposure on the behavior of incumbent firms that were recruiting before the launch of ChatGPT, and thus to distinguish whether the decline comes from existing employers cutting exposed hiring or from changes in the mix of firms. If the exposure effects for incumbent firms are smaller than the occupation-industry estimates, then new firms requiring fewer human workers (or the exit of existing firms) explain a large share of the decline at the occupation-industry level.

As an additional examination of firm-level behavior, we also estimate the average automation exposure value of the occupations that are posted by firms using a simple OLS model with the same fixed effects. Firms continuing to post job openings while reducing the average automation score of the jobs they do post is evidence that firms are shifting their labor demand away from more GenAI-automatable positions relative to less automatable positions even with the same company.[12]

[11] The Poisson estimator allows us to correctly handle the count nature of the data as well as model zeroes during quarters in which there were no job postings in a particular cell in any given quarter.

[12] To smooth volatility in firms' posting behavior from quarter to quarter, we use a four-quarter moving

Figure 2 presents the occupation-industry event study estimates, with Panel A representing estimates that are unweighted and Panel B representing estimates that are weighted by the total number of job postings in the occupation-industry cell before Q3 2022. In Panel A, the event study shows that job postings trended fairly similarly across more- vs. less exposed occupations in the 8 quarters preceding the release of ChatGPT. However, the coefficients indicate that job postings for occupations with a 10 p.p. greater share of automatable tasks fell by 5 percent in relation to less exposed occupations within a year of ChatGPT's release and continued to fall after the beginning of 2025. Postings for more-exposed occupations remained nearly 8 percent lower through the midpoint of 2026.[13] When weighting by the pre-ChatGPT importance of the occupation-industry cell, the decline is quantitatively similar and is not statistically distinguishable from the unweighted estimates. The sharp and persistent decline in postings is consistent with the results of Brynjolfsson et al. (2025) and other recent literature.

We next examine the posting behavior of incumbent firms. Figure 3 presents the results of estimating our Poisson model for job ads within the firm (Panel A) and the OLS model for the composition of the firm's job ads as represented by their ads' average automation exposure score (Panel B). By the beginning of 2024, postings had fallen more at firms whose pre-ChatGPT postings were concentrated in more-exposed occupations: relative to the quarter before ChatGPT's release, a 10 percentage-point difference in firm-level exposure is associated with approximately 8 percent fewer postings (Panel A). Panel B shows that the composition of the firms' post-ChatGPT job ads shifted away from more automatable occupations as well. The average automation exposure score of the firm's job ads fell by approximately two percentage points by early 2024 per 10 percentage-point difference in firm-level exposure.

We test the robustness of these results against other candidate explanations from the literature for the decline in labor demand for AI-exposed jobs, namely occupational sensitivity to monetary policy (Zens et al., 2020), which tightened significantly in 2022, and changes in labor demand for positions more likely to be hired into remote or hybrid work (Lambert and Schindler, 2026; Emanuel et al., 2026; Mask, 2026). To test this, we calculate the average share of each occupation's job ads before Q3 2022 that are listed as remote or hybrid option positions and interact this occupation-level mean with quarter fixed effects to control for any differential demand changes in more remote occupations. For monetary policy responsiveness, we structure the Lightcast data into monthly counts of occupation-specific job postings

average of job postings and automation score as the dependent variables.

[13] In complementary analyses, we also structure the data spatially at the occupation-commuting zone level and find comparable results (Figure C1). We estimate $E(y_{ozt} \mid \mathbf{x}_{ozt}) = \exp\left(\alpha_o + \alpha_{zt} + \sum_{\tau \neq t^*-1}^{\tau} \pi_\tau^{ES} \mathbf{1}\{t = \tau\} OccExp_o\right)$ where o, z, and t index occupation, commuting zone, and quarter, respectively.

from 2010 to Q1 2020 (to avoid the COVID-19 recession) and estimate occupation-specific impulse response functions to monetary policy changes during this period, following a similar approach to Liu et al. (2025).

We then interact the occupation-specific sensitivity measure with quarter fixed effects as an additional regressor in our Poisson model of Equation 1.[14] Figure C2 shows the result of these exercises. Neither these controls separately or collectively materially change the point estimates or time path of the results in Figure 2, with one small exception: controlling for remote work shares flattens what might be perceived as a small pre-trend in the estimates from Q2 2022 to Q3 2022 and shifts the timing of the beginning of the drop in postings back one quarter.

Interestingly, Figure C3 shows that even when controlling for sensitivity to monetary policy and GenAI exposure, occupation-industry cells with more remote work prior to ChatGPT experienced lower demand. The coefficients for remote work shares indicate that, *independent* of GenAI effects, demand for more remote occupations fell quickly in the middle of 2022, stopped falling within one year at approximately 3-4%, and entered into a new steady state. The correlation between our measure of Claude task automation and the occupation-industry level remote work shares in 2022 as defined by Hansen et al. (2026) and extracted from the full text of Lightcast job postings is surprisingly low at 0.21 (0.40 when weighting by total postings). These results suggest that remote work shares might be responsible for the effects of GenAI that show up "too early" in 2022, that remote work cannot directly explain the decline in demand attributable to GenAI in 2023-2026, and that remote work has an independent, measurable relationship to labor demand separable from GenAI effects.[15]

We interpret the totality of these results as strong evidence that GenAI reduced labor demand for occupations with more tasks that could be automated by the new technology. If the technology is particularly adept at replicating the "codified" knowledge new labor market entrants may possess, but less so the "tacit" knowledge gained from experience (Hosseini and Lichtinger, 2025; Brynjolfsson et al., 2025), then new college graduates in particular may face an automation penalty. In response to these shifts in labor demand, new graduates and current students may adjust their educational choices.

[14] We estimate the impulse response functions from six months before to twelve months after changes in monetary policy, and estimate responses to the inverse hyperbolic sine transformation of postings to account for (infrequent) zeroes. To form the measure of responsiveness, we extract the largest coefficient (in absolute value terms) for each occupation across the post-change coefficients as our summary measure of sensitivity.

[15] Instead of examining Texas alone, we also estimate the model using job postings from the entire United States in Figure C4b, which shows a similar effect in direction, magnitude, and time path.

3.2 Effects of GenAI on New Labor Market Entrants

Section 3.1 shows that GenAI reduced firms' demand for workers in AI-exposed occupations. Lower demand for workers in exposed occupations should reach the graduates trained to fill them. If GenAI imposes an automation penalty on a standard degree in exposed fields, it should appear first among new graduates, who enter the labor market without the experience that may protect incumbent workers and whose skill investments were made before the technology existed.

Figure 4 presents estimates of Equation 1 with total earnings during a graduate's first four quarters after graduation as the outcome. Before the release of ChatGPT, earnings trended similarly across cohorts for graduates of more- and less-exposed majors. Earnings trajectories diverged after the release. For May 2023 graduates, a 10 percentage-point difference in major-level exposure is associated with approximately $2,000 (2025$) lower first-year earnings, relative to the same difference for graduates who entered the labor market before ChatGPT's release. The estimate grows to over $3,000 for May 2024 graduates, nearly 10% of average first-year earnings among pre-GenAI cohorts.

The decline in raw earnings may reflect both extensive-margin declines in employment and intensive-margin declines in wages, both of which would be consistent with the decline in postings shown in Section 3.1. Figure 5 provides evidence of both mechanisms: 2024 graduates from more-exposed majors worked in Texas for 0.08 fewer quarters in their first year (2.8%, relative to the pre-GenAI mean) and earned $2,427 (5%) less, conditional on working all four quarters.[16]

Observed earnings declines for graduates from more-exposed majors would be particularly harmful if they persist. A robust literature finds that entering the labor market during recessions causes declines in earnings for at least a decade (Kahn, 2010) and even into middle age (Schwandt and Von Wachter, 2020). Proposed mechanisms for this sustained harm include that the mismatch between skills demanded and supplied by new entrants rises during recessions, and that extended non-employment and lower-quality job matching hinder the development of occupation- and industry-specific human capital (Liu et al., 2016; Brunner and Kuhn, 2014). These dynamics could plausibly hold in the wake of a field-specific demand shock, as well, particularly if GenAI complements tacit knowledge, gained through

[16]Because we can only track the earnings and employment of those employed in the state of Texas, it is important to ask if differential exit from the state could explain some of the employment effects we find. Figure B4 uses microdata from the American Community Survey annual extracts to examine outmigration rates from Texas by age, educational attainment, college enrollment, and GenAI exposure. No strong differential changes in outmigration rates are detectable after 2022 by education, enrollment, or by AI exposure in workers' degree field. Differential exit rates cannot explain our findings. Moreover, prior research shows that Texas has among the lowest out-migration rates of any state, and in-state earnings records produce little migration bias there (Mountjoy and Hickman, 2021; Foote and Stange, 2022).

experience, while substituting for codified knowledge (Brynjolfsson et al., 2025).

Relative declines in earnings suggest GenAI substitutes for new four-year graduates from fields with higher automatable task shares, consistent with the automation penalty proposed in our conceptual framework. Although we only observe earnings through the first half of 2025, we explore early indicators of whether these initial effects on exposed graduates are persisting. Table 5 presents estimates of π^{ES}_{2023} with labor market outcomes measured across graduates' first (Panel A; the estimates from Figures 4 and 5) and second (Panel B) years in the labor market. Estimates suggest the shock's negative effects persist: 2023 graduates from more-exposed majors experienced a significant relative decline of $2,351 in second-year earnings. The estimated decline in earnings conditional on working all four quarters increases from an insignificant $934 in the first year to a statistically significant $2,102 in the second year. Meanwhile, the estimated employment decline shrank from 0.083 to 0.045 quarters. These patterns suggest that earnings disadvantages among employed graduates may persist even as employment differences narrow.

We assess the plausibility of multiple threats to interpreting the above estimates as the causal effects of GenAI automation. Our main specification already controls for school-specific shocks that vary across cohorts. Just as with the job posting estimations, we test the robustness of estimated effects on earnings to controlling for (1) time-varying industry-specific shocks, such as the differential impact of higher interest rates, along with (2) the differential adoption of remote work across occupations. Figure 6 presents estimated effects of AI exposure on earnings of employed graduates, controlling for industry-by-cohort fixed effects and cohort-specific effects of major-level remote work shares.[17] Controlling for both proposed confounders still yields an estimated earnings decline of $2,599 for 2024 graduates from more-exposed majors, nearly identical to the estimate in Figure 5b.[18]

Although our main specification controls for observable student characteristics, including a proxy for academic ability (ACT/SAT scores), changing in selection on unobservable student characteristics into more-exposed majors across cohorts could bias our estimates. In particular, increased enrollment in more-exposed majors by lower-ability students in the post-ChatGPT graduating cohorts relative to previous cohorts could cause average earnings and employment rates for those majors to fall, regardless of an AI shock. However, we find no evidence that recent graduates from more-exposed majors were negatively selected on observable academic preparation. Figure C7 shows a relative increase in SAT scores for graduates of more- vs. less-exposed majors in this cohort relative to previous ones, sug-

[17]We define major-specific remote work adoption rates as the share of job ads from before Q3 2022 for major m that list remote or hybrid options for the position. We assign industries to new graduates as the industry in which they earned the most in their first four quarters after graduating.

[18]Figure C5 presents estimates that control for each confounder, separately.

gesting that the estimated declines in earnings for more-exposed graduates come in spite of potentially positive ability bias.

Our estimates could also exaggerate the true effects of GenAI on potential earnings if new graduates respond to the technological shock by returning to graduate school, as examined in Section 4. Figure C8 presents estimates of the relationship between GenAI exposure and first-year earnings for new graduates, excluding students that enroll in graduate school. Once again, the estimated decline in earnings associated with a 10 percentage-point difference in AI exposure hardly changes and maintains its statistical significance. Figure C9 also shows that observable proxies for ability are still positively correlated with major-level AI exposure in the leftover sample. Our results do not appear to reflect the withdrawal of higher-ability graduates from the labor market and into graduate school.

Table 6 presents an exhaustive set of estimates of $\hat{\pi}^{2024}$ across all combinations of controls for the confounders discussed above.[19] Tables C1 and C2 present estimates from the same set of specifications, using Eloundou et al.'s (2023) rubric-based exposure measures derived by human annotators and GPT-4, respectively. Our estimates are extremely robust across specifications using exposure measures based on Claude usage or rubrics scored by human annotators. However, estimates from 10 out of 30 specifications using an exposure measure grounded in an LLM's own definition of what tasks it can perform are not statistically distinguishable from zero at the 10% level. Given evidence from Yin et al. (2026) on the fragility of LLM-based rubric scores across different models and the weak correlation between this measure and reported AI usage (Bick et al., 2025), we place greater weight on results derived from other exposure measures.

We also examine whether the earnings losses we document are driven by any single field of study or industry. We re-estimate π^{ES}_{2024} for earnings of graduates that worked all four quarters of their first year, dropping one 2-digit CIP code (31 estimates) or one 2-digit NAICS industry (25 estimates) at a time. Figure C10 plots the distribution of these estimates. Most are tightly clustered: 27 of the 31 estimates that drop a field and 17 of the 25 estimates that drop an industry fall between -$2,800 and -$2,600, with medians of -$2,670 and -$2,667, respectively. Every estimate remains negative and statistically significant at the 5% level. The clear exception is Computer and Information Sciences (CIP 11). Dropping computer science graduates shrinks the estimated decline to -$1,586, roughly 40% smaller than the median estimate, though it remains statistically significant. A meaningful share of the estimated earnings losses for employed graduates can therefore be explained by worsening outcomes for computer science graduates, but the decline is not solely a computer science phenomenon. In contrast, dropping graduates employed in Professional, Scientific, and Technical Services

[19] Figure C6 shows event studies for our earnings regressions for our main specification.

(NAICS 54) leaves the estimate essentially unchanged (-$2,627). Relative to the median, no single industry moves the estimate by more than approximately $500. The largest change comes from dropping Health Care and Social Assistance (NAICS 62), which increases the estimated decline to -$3,169.

3.2.1 Heterogeneous Effects on Recent Graduates

We examine whether the earnings declines associated with GenAI exposure differ across student backgrounds by fully interacting Equation 1 with subgroup indicators for sex, race and ethnicity, family income, college entrance exam scores, and institutional selectivity. Figures D1 and D5 (a)–(b) report the estimates. The estimated effects for the 2024 graduating cohort are negative across these groups, including graduates with higher entrance exam scores and those attending more-selective institutions. The earnings declines are therefore broad-based rather than confined to students with weaker prior academic preparation or those attending less-selective institutions. This breadth is consistent with a large, field-specific decline in labor demand that reached most graduates who would typically sort into AI-exposed occupations.

3.2.2 Effects Over the Experience Distribution

We next examine whether GenAI automation is affecting workers further up the experience distribution. Figure D7 plots the earnings growth over time of graduates in more-exposed majors by graduation cohort, with cohorts aged approximately 45 (2000 graduates), 40 (2005), 35 (2010), 30 (2015), and 25 (2019) at ChatGPT's release. Under a typical Mincerian life-cycle with a concave wage profile, earnings growth is steepest early in a worker's career, so we would expect more recent cohorts to experience faster growth relative to older cohorts. This held before the pandemic: relative growth was steeper for the 2015 cohort than for the 2010 cohort, and steeper still for the 2010 cohort than for earlier graduates. However, after the release of ChatGPT, the 2019 cohort's relative earnings growth fell below that of the 2015 cohort to barely outpacing the 2010 cohort, which is about ten years further along the experience curve. Earlier cohorts continued on their pre-pandemic trends while the 2019 cohort of exposed graduates had their trends bend downward noticeably. These results constitute suggestive evidence that GenAI has not strongly affected the earnings trajectories of workers with at least six years of experience before its proliferation, but may be slowing early-career earnings growth for junior workers with exposed skillsets.

4 Skill Adjustments to GenAI: Pivoting and Upskilling

4.1 Conceptual Framework

Sections 3.1 and 3.2 demonstrate that there is an apparent "automation penalty" affecting students completing their education in more exposed fields compared to their peers in other majors coming from the same graduating cohorts at the same institutions. These changes may alter students' incentives to invest in different fields and levels of education. To understand how current and prospective students respond, we draw on a simple human-capital investment framework in the tradition of Becker (1962), summarized here and presented as a full model in Appendix A.

In this framework, current and prospective students choose educational levels of attainment and fields of study by weighing the perceived lifetime earnings premium for earning a particular credential against the monetary, opportunity, and effort costs of doing so. We apply this framework to students' choices of college majors, course enrollment, and the joint decision of graduate school enrollment and program choice.

An automating technology reduces the perceived earnings premium for field m if it substitutes for tasks in occupations that graduates of m typically enter. We denote this reduction the "automation penalty," which the prior sections demonstrate is present among these cohorts. Students have two broad ways to respond. First, they can "upskill," deepening their expertise in m in the hope that more advanced knowledge complements GenAI (Autor and Thompson, 2025). Alternatively, they can "pivot," reallocating their skill investments toward fields that GenAI is less able to automate.[20] Students can upskill by specializing further in their major's coursework or by pursuing a graduate degree in the same field. They can pivot along four margins: the major they declare when entering college, the courses they take given a declared major, switching to another major, and the field of any graduate program they pursue. We model a major's "exposure" to GenAI as the share of tasks that the new technology can automate in occupations typically held by graduates from that major. Our empirical analysis of the effects of GenAI on labor demand suggests that the new technology imposes a substantial automation penalty for graduates from more-exposed fields.

A large automation penalty for major m incentivizes students already pursuing a degree in major m at the time of the demand shock to switch fields. However, the time, monetary, and effort costs to gain class credits toward a particular major make reallocation to another field costly. We assume that these costs increase the further along a student is in their education

[20] Hyman et al. (2025) similarly frame occupational training programs for workers displaced from AI-exposed occupations as "upskilling" if the program leads to placement in an even more-exposed occupation and "downskilling" if the program leads to placement in a less-exposed occupation.

(Weinstein, 2022). These costs also differ across margins of adjustment. Pivoting majors is cheapest for students entering college, who have not yet accumulated major-specific credits, and most expensive for upperclassmen. Adjusting coursework without switching majors – for example, through electives – is an intermediate option, and one that can move in either direction. Students already attached to an exposed major can upskill by concentrating more of their coursework in that major, or pivot by taking more courses in less-exposed fields. This is less costly because they do not forfeit credits earned toward their final degree. If students perceive a lasting decline in demand for the skills taught in exposed majors, we therefore expect those already declared in these majors to pivot at the course level even if switching majors is too costly. Conversely, if they perceive that deeper expertise in their field complements GenAI, or if GenAI lowers the effort cost of advanced coursework in exposed fields, they may instead specialize further.

Generative AI diverges from a standard technology shock that affects demand for particular skills in that GenAI tools not only can automate tasks in the labor market but also many components of *training* to do the tasks of the job via coursework. This enters our framework as a reduction in the effort costs required to earn a degree in field m. If GenAI's capabilities to automate coursework are correlated with its ability to automate tasks in the labor market, then declines in effort costs may be largest for the same fields that experience large declines in labor demand. This is plausible given that GenAI simplifies access to technically advanced concepts and allows non-experts to apply them immediately. This channel works in the opposite direction as pivoting: lower effort costs in exposed fields make staying in or upskilling within exposed majors relatively more attractive.

Although the opposing earnings premia and effort cost channels make GenAI's overall effect on field choices ambiguous, our framework predicts:

(a) Pivoting away from exposed majors increases with GenAI whenever effort costs fall by less than the perceived returns.

(b) The cost of pivoting increases with seniority, so students already declared in exposed majors are more likely to adjust their coursework – specializing further ("upskilling") or taking less-exposed courses ("pivoting") – than to switch majors.

(c) Reallocation away from exposed majors will be relatively slow in the first few years of large and fast technology adoption and advancement, because AI capabilities progress faster than students further along in their degrees can anticipate or, given switching costs, adapt to. Recent graduates in exposed fields are most vulnerable to the falling labor demand to which they had no or limited scope to adjust prior to completing their education, which could result in short- and medium-term scarring.

(d) If GenAI purely substitutes for the skills taught in exposed fields, then within exposed majors, graduates whose coursework is more exposed than that of their major suffer larger earnings losses. In other words, course-level pivoting pays off relative to course-level upskilling. If instead deeper expertise complements GenAI, the reverse holds.

We also apply this framework to the joint decision of whether to enroll in graduate school and what field to study. GenAI reduces the opportunity cost of additional education for new graduates from more-exposed fields. Just as with bachelor's degrees, the availability of GenAI tools also reduces the effort costs of earning graduate degrees, and reductions in effort cost may correlate with changes in field-specific labor demand. Graduates who return to school face the same choice between upskilling and pivoting. They upskill by enrolling in graduate programs in exposed fields, which pays off if advanced expertise complements GenAI, and they pivot by enrolling in less-exposed programs. This yields the following predictions:

(e) GenAI increases graduate school enrollment for students from more-exposed undergraduate majors.

(f) Graduates of exposed majors upskill, remaining in exposed fields, if they perceive advanced education as complementary to GenAI, but pivot if the perceived earnings premium for graduate degrees in exposed fields falls by more than effort costs.

(g) Post-shock earnings of new master's recipients in more-exposed fields fall if the automation penalty outweighs complementarity from upskilling. If the expertise acquired through graduate programs in exposed fields complements GenAI, earnings of new master's recipients in those fields should continue along trends or rise.

These predictions motivate our empirical analysis of student responses to the decline in labor demand resulting from GenAI automation exposure.

4.2 How Did Students Adjust?

How did four-year graduates from more-exposed majors respond to soft labor demand? Figure 7 presents estimates of Equation 1 with an indicator for enrolling in a graduate program within a year of a cohort's May graduation as the outcome. Consistent with the decline in opportunity costs, graduate enrollment increased by 0.9 percentage-points (6%) for 2024 four-year graduates from a major with a 10 percentage-point higher share of automatable tasks. If the decline in demand for automatable skillsets is merely a temporary phenomenon while firms adopt AI, after which productivity gains increase demand for workers with complementary skills, then graduate programs may serve as a safe harbor to buffer against the

"scarring" effects of entering the labor market during a period of depressed demand (Rothstein, 2023; Schwandt and von Wachter, 2019).

Conditional on enrolling, graduates of exposed majors could either upskill by staying in exposed fields or pivot toward less-exposed ones. We estimate our event study specification on the sample of graduate school enrollees with the automatable task share of the *graduate program* field as the outcome variable. Figure 8 shows little discernible evidence of pivoting toward less-exposed majors. A 10 percentage-point greater undergraduate major automatable task share is associated with only a 0.5 percentage-point decline in graduate program exposure to GenAI. In simpler terms, nearly two-thirds of computer science majors from the 2024 graduating cohort that enrolled in graduate school chose a computer science program, roughly the same share as for pre-ChatGPT graduate enrollees.[21] This evidence of continued investments in exposed fields is consistent with an upskilling response.

If students perceive a lasting automation penalty, undergraduate enrollment may shift away from exposed majors. Figure 9 presents estimates of the relationship between major-level GenAI exposure and enrollment counts, using a Poisson analog of Equation 1 estimated on counts aggregated to the major-by-year level, with major and year fixed effects in place of the student-level controls. Panel 9a shows that, driven by the increasing popularity of computer science, undergraduate enrollment grew by 10% more in more-exposed majors relative to less-exposed majors between Fall 2016 and Fall 2022. At face value, the GenAI shock appears to have hardly affected major choices: enrollment continued to grow along its underlying trend through Fall 2024, before slightly falling in Fall 2025.

However, the aggregate trend masks substantial heterogeneity across classifications that is consistent with switching costs that rise with seniority (Panel 9b). Students already enrolled in college at the time of the shock pivoted away from exposed majors at lower rates. Switching majors as an upperclassman often delays graduation, yielding substantial financial and effort costs. Senior enrollment continued along its existing trend – even in Fall 2025, all on-time Seniors entered college before the release of ChatGPT – and the estimates for Sophomores and Juniors suggest only gradual adjustments as cohorts aged.

In contrast, students entering college after the shock pivoted toward less-exposed fields. These students had not yet accumulated major-specific credits, and the GenAI shock to labor demand likely became more salient over time, as both adoption and declines in demand for graduates of more-exposed fields grew. A 10 percentage-point difference in major exposure correlated with a 7% decline in Freshman enrollment between Fall 2022 and Fall 2025 (Panel

[21] 66% of computer science majors from our sample's 2024 graduating cohort that enrolled in a graduate program at a Texas public university within a year of receiving their undergraduate diploma enrolled in a computer science graduate program. The average across the 2016 - 2022 graduating cohorts is 68%.

9b), consistent with their greater flexibility to adjust to a demand shock. By the end of the third year after the release of ChatGPT, enrollment in exposed majors among Freshman had reversed the enrollment trends going back to 2016.

There are two other margins at which students may respond: shifting their course load within majors or by changing their declared major, which constitutes a costlier adjustment than shifting their initial major declaration. We estimate both the probability of switching majors as a function of the exposure in their initial major and the average exposure of their semester-specific course load for Juniors in the data over time. Figure 10 shows that the average GenAI exposure of the courses taken by college Juniors in Texas was trending upward from 2016 through 2024 and then abruptly shifted downward from the prior trend in Fall 2025, reflecting overall adjustments in course-taking both within and across declared majors. This lower-cost alternative to major switching appears strongest in the Fall semester course enrollment of 2025 Juniors (Panel (a)), and to a lesser extent in the cumulative credits they attempted through their Junior fall (Panel (b)). College underclassmen appear to be pivoting at the course level away from more AI-exposed subjects.

At the same time, Figure 11 shows that, by their Junior year, students who were initially declared for more exposed majors began changing their declared majors at higher rates in Fall 2025, consistent with the other changes in course enrollment above. For a 10 percentage point increase in initial major exposure, students were approximately 2.5 percentage points (5%, relative to the mean likelihood) more likely to change majors. This represents costlier action taken by students in response to perceived changes in labor demand. These students were initially more persistent in choosing an exposed major even after some of their fellow students made pivots away from exposed majors as Freshmen and Sophomores (Figure 9b), making them a slightly differently selected group.[22] As these results show, a notable fraction changed majors by the end of their Junior year.

Despite GenAI's unique implications on both the demand and supply side of higher education, sticky field choices among enrolled students and a delayed pivoting response are consistent with the effects of prior sectoral demand shocks on major choice (Weinstein, 2022; Conzelmann et al., 2023).

4.2.1 Heterogeneity in Educational Responses

We next examine whether graduate enrollment, graduate-field choice, and Freshman enrollment responses differ across socioeconomic backgrounds, academic preparation, and

[22]This time-varying selection on initial major could bias major-switching estimates. However, the increased likelihood of marginal would-be enrollees in more-exposed majors to have selected less-exposed majors upon initial college enrollment would likely bias against finding a positive relationship between AI exposure and major switching.

institutional selectivity. We find suggestive evidence of heterogeneity in graduate-school enrollment responses in Figures D2 and D5. For the 2024 cohort, a 10 percentage-point difference in major-level exposure is associated with graduate school enrollment increases of 1.5 percentage points for Hispanic graduates and 1.8 percentage points for non-Hispanic Asian graduates, compared with estimates near zero for non-Hispanic white and Black graduates. The estimated response is also larger for graduates from families earning less than $80,000 than for those from higher-income families. Differences across SAT groups are less pronounced. Several of these subgroup differences are imprecisely estimated and also appear among pre-ChatGPT cohorts, so we interpret them cautiously.[23]

Among graduate-school enrollees, the point estimates suggest a modest shift toward less-exposed programs for men and little change for women (Figure D3). Non-Hispanic white graduates, who did not return to graduate school at higher rates, exhibited the largest shift toward less-exposed programs among those that did enroll.

Freshman enrollment responses were broad-based across sex and family income (Figure D4). Between Fall 2022 and Fall 2025, enrollment fell by roughly 11% for women and 7% for men in majors with a 10 percentage-point greater share of automatable tasks. The corresponding declines were 7% and 8% for students from families earning above and below $80,000, respectively. The point estimates suggest a larger shift toward less-exposed fields at the most-selective institutions (Figure D6), although differences in pre-period trends make it difficult to determine whether responses differ by institutional selectivity.[24]

4.3 Do Adjustments Pay Off?

Two pieces of evidence speak to whether upskilling or pivoting insulated recent graduates from GenAI's automation penalty: the earnings of new master's recipients in exposed fields, and the earnings returns to coursework concentrated in exposed fields among four-year graduates.

Similar declines in earnings for graduates from master's programs in exposed fields to those documented for bachelor's recipients would suggest that GenAI's automation penalty outweighs potential complementarities of the expertise derived from graduate programs. We perform two analyses aimed at understanding these dynamics.

[23] We observe graduate enrollment only at Texas public universities. It is possible that differences across groups in enrollment responses at private or out-of-state institutions therefore affect these comparisons.

[24] We do not estimate Freshman enrollment responses by SAT score due to the stark drop-off in the share of students reporting SAT scores after many institutions adopted test-blind admissions policies. Between 2017 and 2020, less than 10% of four-year entrants in the dataset are observed without an SAT or ACT score. This jumped to over one-third of entrants between 2021 and 2024, and the resulting changes in sample composition make such estimates difficult to interpret.

First, we estimate Equation 1 with the same field-level exposure levels and a cohort-level panel of new master's degree recipients. Figure 12 shows that first-year earnings grew by $2,500 more across cohorts for master's recipients in fields with a 10 percentage-point greater share of automatable tasks between 2016 and 2021. Since the release of ChatGPT, these relative earnings gains have more than fully reversed. Employed 2024 graduates from fields with 10 percentage-point higher field-level AI exposure earned nearly $4,000 (5%) less compared to the last pre-ChatGPT cohort.

Second, we construct the data to include both master's and bachelor's degree earners and estimate Equation 1 but include master's degree-by-cohort fixed effects and master's degree-by-year interactions with automation exposure. We then present this latter interaction, which tests whether the returns to a more-exposed master's degree *relative to a bachelor's degree* have declined after 2022. In the context of our conceptual framework, this coefficient reflects the potential complementarity of graduate-level education in AI-exposed fields after netting out a common automation penalty imposed for bachelor's and master's degree recipients. Figure 13 presents the coefficients on the interaction between master's degree by year and automation exposure. Before the release of ChatGPT, the returns to a more exposed master's degree relative to less-exposed master's degrees were growing. After the release of ChatGPT, that relative trend halted and may have even reversed. The relative gains for the 2024 cohort were nearly $3,000 lower than what would have been implied by the pre-GenAI trend.

To summarize, the initial stage of GenAI automation appears to bite just as hard for new labor market entrants with additional education as it does for four-year graduates. This suggests that the additional skills obtained through graduate programs in exposed fields do not complement the new technology, after accounting for a common automation penalty at both degree levels.

Upskilling and pivoting can also happen within a degree through the courses a student takes alongside their major. We define a graduate's idiosyncratic course-based exposure as the course-weighted automatable task share across all fields in which they completed courses, minus their major's automatable task share. The measure maps onto the two strategies in Section 4.1 for students in exposed majors: concentrating coursework in the major or in similarly exposed fields (upskilling) raises idiosyncratic exposure, while taking courses in less-exposed fields (pivoting) lowers it. To assess whether course mix magnified or buffered against major-level shock exposure, we estimate a modified two-period version of equation 1 that further includes idiosyncratic exposure and its interaction with major-level exposure.

Figure 14 plots marginal effects of a one-standard deviation increase in idiosyncratic exposure on earnings for graduates that worked all four quarters of their first year, across

values of major-level exposure.[25] Positive estimates near the top of the major-level exposure distribution would suggest that specialized expertise in AI-exposed fields offsets some of the automation penalty, supporting an upskilling strategy. Positive estimates at lower levels of major exposure would also suggest that familiarity with the technical skills GenAI automates complements skills that the technology cannot automate. We find no evidence that additional coursework in more-exposed fields raised earnings at any point in the distribution. The marginal effect is near zero for graduates from the least-exposed majors ($38) and declines steadily with major-level exposure, becoming statistically significant for majors with more than roughly 15% of automatable tasks.

Framed in reverse, a less-exposed course bundle is associated with a meaningful share of the losses being buffered for graduates from highly exposed majors. At a major-level exposure of 30%, the estimated marginal effect implies that a one-SD less-exposed course bundle is associated with $1,910 higher first-year earnings. Extrapolating the main estimate linearly, this equals roughly a quarter of the implied earnings gap for such a major. This suggests that there are substantial benefits to pivoting toward developing skills in less-exposed fields, to the extent that sorting into coursework on unobserved earnings potential did not change differentially in exposed majors after ChatGPT's release.[26] Earnings benefits to pivoting investments directionally align with Hyman et al.'s (2025) finding that displaced workers from AI-exposed occupations in recent years gained sizable earnings returns from retraining programs that prepared them for less-exposed jobs.

Taken together, our results suggest that for young workers, automation penalties dominated potential complementarities from specialized expertise in exposed fields during the early stages of the AI transition. We note two caveats to this interpretation, though. Initial earnings are only a small portion of the lifetime earnings streams that individuals consider when making human capital decisions. Still, evidence of rapid skill obsolescence for workers in technology-intensive fields (Deming and Noray, 2020) suggests that early-career earnings may be particularly informative in assessing the benefits of adjustment strategies to GenAI. Moreover, these estimates reflect labor market outcomes of students largely completing "pre-GenAI" curriculum. As universities monitor what skills employers require from workers in AI-integrated production processes, coursework designed to develop newly demanded skills may unlock complementarities that are not present in our data.

[25] Estimates come from the following specification: $y_{imct} = \gamma_0 IdExp_{im} + \gamma_1 IdExp_{im} \times MajExp_m + \gamma_2 Post_t \times IdExp_{im} + \gamma_3 Post_t \times MajExp_m + \gamma_4 Post_t \times IdExp_{im} \times MajExp_m + \alpha_m + \alpha_{ct} + \Phi S_i + \epsilon_{imct}$.

[26] The implied loss scales the $2,427 decline in first-year earnings for employed 2024 graduates per 10 percentage points of major-level exposure (Figure 5b) to a major with 30% of tasks automatable, relative to a hypothetical major with none. Treatment of average causal response parameters from continuous difference-in-differences specifications as linearly scalable merits caution, though (Callaway et al., 2025).

5 Conclusion

Economists vary widely both in their predictions of how AI will ultimately affect the labor market (Chen and Lahart, 2026) and in their assessments of its effects to date (Brynjolfsson et al., 2025; Lambert and Schindler, 2026). Although widespread labor market disruption has yet to appear, our findings suggest that the GenAI transition is already causing varying degrees of harm depending on the degree level, degree field, and individual course load. We provide direct evidence that recent graduates from majors that develop automatable skillsets exhibit deteriorating labor market outcomes. If these workers are unable to gain the "tacit" knowledge necessary to complement GenAI, the automation penalty apparent in their earnings may persist. Because students already in college at the time of the AI shock made only small adjustments, if any, to their field choices, similar harm may fall on 2025 college graduates in the aggregate. However, students that entered college more recently appear to be internalizing the decline in demand for more-exposed fields, and their pivots toward studying other fields may buffer them against the labor demand shock.

Our results suggest that the net effects of the GenAI transition on early-career earnings and human capital investments appear similar to those of a normal sectoral demand shock, thus far. However, the technology could carry much broader implications for both the labor market and education system. Unlike conventional demand shocks, GenAI may affect not only the labor market value of skills but also the costs and methods of acquiring them. Given AI's capability to automate many out-of-classroom assignments, researchers should continue to study the technology's effects on learning. Whether individuals view GenAI's function of democratizing access to expertise as a substitute for investing in formal education also deserves attention. Regardless of how the AI transition continues to unfold, the robust literature on the persistent effects of initial wage conditions (e.g., Arellano-Bover, 2024) highlights the importance of our findings of its effects on cohorts caught in the transition.

References

Acemoglu, Daron and Pascual Restrepo, "Tasks, Automation, and the Rise in U.S. Wage Inequality," *Econometrica*, 2022, *90* (5), 1973–2016, DOI: https://doi.org/10.3982/ECTA19815.

ACT and College Board, "Guide to the ACT/SAT Concordance," Technical Report, ACT and the College Board 2018.

Arcidiacono, Peter, V. Joseph Hotz, and Songman Kang, "Modeling college major choices using elicited measures of expectations and counterfactuals," *Journal of Econometrics*, 2012, *166* (1), 3–16, DOI: https://doi.org/https://doi.org/10.1016/j.jeconom.2011.06.002. Annals Issue on "Identification and Decisions", in Honor of Chuck Manski's 60th Birthday.

Arellano-Bover, Jaime, "Career Consequences of Firm Heterogeneity for Young Workers: First Job and Firm Size," *Journal of Labor Economics*, April 2024, *42* (2), 549–589, DOI: https://doi.org/10.1086/723500.

Autor, David and Neil Thompson, "EXPERTISE," *NBER Working Paper 33941*, 2025, DOI: https://doi.org/10.3386/w33941.

Azar, José, Mireia Gine, and Javier Sanz-Espín, "AI Is Already Eroding Wages: Quasi-Experimental Evidence From Occupational Exposure," 2025, DOI: https://doi.org/10.2139/ssrn.5842084.

Becker, Gary S., "Investment in Human Capital: A Theoretical Analysis," *Journal of Political Economy*, 1962, *70* (5), 9–49, DOI: https://doi.org/10.1086/258724.

Berge, Wiljan Van Den, "Bad start, bad match? The early career effects of graduating in a recession for vocational and academic graduates," *Labour Economics*, August 2018, *53*, 75–96, DOI: https://doi.org/10.1016/j.labeco.2018.05.011.

Bick, Alexander, Adam Blandin, and David Deming, "THE RAPID ADOPTION OF GENERATIVE AI," 2025, DOI: https://doi.org/10.3386/w32966.

Brunner, Beatrice and Andreas Kuhn, "The impact of labor market entry conditions on initial job assignment and wages," *Journal of Population Economics*, July 2014, *27* (3), 705–738, DOI: https://doi.org/10.1007/s00148-013-0494-4.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen, "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence," 2025.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna, "Difference-in-Differences with a Continuous Treatment," 2025, DOI: https://doi.org/10.48550/ARXIV.2107.02637.

Chen, Te-Ping and Justin Lahart, "Economists Weigh In on the Future of Work and AI," *The Wall Street Journal*, June 2026.

Chen, Wilbur Xinyuan, Suraj Srinivasan, and Saleh Zakerinia, "Displacement or Complementarity? The Labor Market Impact of Generative AI," 2026.

Conzelmann, Johnathan G., Steven W. Hemelt, Brad Hershbein, Shawn M. Martin, Andrew Simon, and Kevin M. Stange, "Skills, Majors, and Jobs: Does Higher Education Respond?," August 2023, DOI: https://doi.org/10.3386/w31572.

Deming, David J and Kadeem Noray, "Earnings Dynamics, Changing Job Skills, and STEM Careers*," *The Quarterly Journal of Economics*, November 2020, *135* (4), 1965–2005, DOI: https://doi.org/10.1093/qje/qjaa021.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, "GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models," August 2023, DOI: https://doi.org/10.48550/arXiv.2303.10130. arXiv:2303.10130 [econ].

Emanuel, Natalia, Emma Harrington, and Amanda Pallais, "THE POWER OF PROXIMITY TO COWORKERS," *NBER Working Paper*, June 2026, DOI: https://doi.org/10.3386/w31880.

Foote, Andrew and Kevin M. Stange, "Attrition from Administrative Data: Problems and Solutions with an Application to Postsecondary Education," NBER Working Paper 30232, National Bureau of Economic Research July 2022, DOI: https://doi.org/10.3386/w30232.

Frank, Morgan R., Alireza Javadian Sabet, Lisa Simon, Sarah H. Bana, and Renzhe Yu, "AI-exposed jobs deteriorated before ChatGPT," January 2026, DOI: https://doi.org/10.48550/arXiv.2601.02554. arXiv:2601.02554 [econ.GN].

Garicano, Luis, Jin Li, and Yanhui Wu, "Weak Bundle, Strong Bundle: How AI Redraws Job Boundaries," CEPR Discussion Papers 21453, Centre for Economic Policy Research May 2026, DOI: https://doi.org/None.

Goller, Daniel, Christian Gschwendt, and Stefan C. Wolter, "This time it's different – Generative artificial intelligence and occupational choice," *Labour Economics*, August 2025, *95*, 102746, DOI: https://doi.org/10.1016/j.labeco.2025.102746.

Gschwendt, Christian, Martina Viarengo, and Thea S Zollner, "Generative AI and Career Choices," *IZA discussion paper*, 2026, *18456.*

Handa, Kunal, Alex Tamkin, Miles McCain, Saffron Huang, Esin Durmus, Sarah Heck, Jared Mueller, Jerry Hong, Stuart Ritchie, Tim Belonax, Kevin K. Troy, Dario Amodei, Jared Kaplan, Jack Clark, and Deep Ganguli, "Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations," February 2025, DOI: https://doi.org/10.48550/arXiv.2503.04761. arXiv:2503.04761 [cs].

Hansen, Stephen, Peter John Lambert, Nicholas Bloom, Steven J Davis, Raffaella Sadun, and Bledi Taska, "Remote Work Across Jobs, Companies, and Space," *NBER working paper 31007*, 2026, DOI: https://doi.org/10.3386/w31007.

Hershbein, Brad and Lisa B. Kahn, "Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings," *American Economic Review*, July 2018, *108* (7), 1737–72, DOI: https://doi.org/10.1257/aer.20161570.

Hosseini, Seyed M and Guy Lichtinger, "Generative AI as Seniority-Biased Technological Change: Evidence from U.S. R´esum´e and Job Posting Data," *Working Paper*, 2025, DOI: https://doi.org/10.2139/ssrn.5425555.

Hyman, Benjamin G, Benjamin Lahey, Karen Ni, and Laura Pilossoph, "How Retrainable Are AI Exposed Workers?," *NBER Working Paper*, 2025, *34174*, DOI: https://doi.org/10.3386/w34174.

Iscenko, Zanna and Fabien Curto Millet, "Looking for the Ladder: Is AI Impacting Entry-Level Jobs?," January 2026.

Kahn, Lisa B., "The long-term labor market consequences of graduating from college in a bad economy," *Labour Economics*, April 2010, *17* (2), 303–316, DOI: https://doi.org/10.1016/j.labeco.2009.09.002.

Lambert, Peter John and Yannick Schindler, "The Broken Ladder: AI, Remote Work, and Early-Career Hiring," 2026, DOI: https://doi.org/10.2139/ssrn.6787638.

Light, Jacob, "How exposed is higher education to Artificial Intelligence?," 2026.

Liu, Kai, Kjell G. Salvanes, and Erik Ø. Sørensen, "Good skills in bad times: Cyclical skill mismatch and the long-term effects of graduating in a recession," *European Economic Review*, May 2016, *84*, 3–17, DOI: https://doi.org/10.1016/j.euroecorev.2015.08.015.

Liu, Yan, He Wang, and Shu Yu, "Labor Demand in the Age of Generative Ai: Early Evidence from the US Job Posting Data," *Available at SSRN 5504741*, 2025, DOI: https://doi.org/10.2139/ssrn.5504741.

Mask, Joshua, "Same Occupations, Different Clocks: Separating Remote Work from Generative AI for Labor-Market Entrants," August 2026, DOI: https://doi.org/10.2139/ssrn.7075479.

Massenkoff, Maxim and Peter McCrory, "Labor market impacts of AI: A new measure and early evidence," Technical Report, Anthropic March 2026.

Mountjoy, Jack and Brent R. Hickman, "The Returns to College(s): Relative Value-Added and Match Effects in Higher Education," NBER Working Paper 29276, National Bureau of Economic Research September 2021, DOI: https://doi.org/10.3386/w29276.

Orr, Cody, Lee C Tucker, and Lawrence Warren, "Graduating Into Disruption: Labor Market Outcomes for AI-Exposed College Majors," 2026.

Peterson, Andrew J., "Training for Obsolescence? The AI-Driven Education Trap," November 2025, DOI: https://doi.org/10.48550/arXiv.2508.19625. arXiv:2508.19625 [econ.GN].

Rothstein, Jesse, "The Lost Generation?," *Journal of Human Resources*, 2023, *58* (5), 1452–1479, DOI: https://doi.org/10.3368/jhr.58.5.0920-11206R1.

Schwandt, Hannes and Till von Wachter, "Unlucky Cohorts: Estimating the Long-Term Effects of Entering the Labor Market in a Recession in Large Cross-Sectional Data Sets," *Journal of Labor Economics*, 2019, *37* (S1), S161–S198, DOI: https://doi.org/10.1086/701046.

_ **and Till Von Wachter**, "Socio-Economic Decline and Death: The Life-Cycle Impacts of Recessions for Labor Market Entrants," Technical Report w26638, National Bureau of Economic Research, Cambridge, MA January 2020, DOI: https://doi.org/10.3386/w26638.

Weinstein, Russell, "Local Labor Markets and Human Capital Investments," *Journal of Human Resources*, September 2022, *57* (5), 1498–1525, DOI: https://doi.org/10.3368/jhr.58.1.1119-10566r2.

Yin, Michelle, Hoa Vu, and Claudia Persico, "HOW (UN)STABLE ARE LLM OCCUPATIONAL EXPOSURE SCORES? EVIDENCE FROM MULTI-MODEL REPLICATION," *NBER Working Paper 35110*, 2026, DOI: https://doi.org/10.3386/w35110.

Zens, Gregor, Maximilian Böck, and Thomas O. Zörner, "The heterogeneous impact of monetary policy on the US labor market," *Journal of Economic Dynamics and Control*, 2020, *119*, 103989, DOI: https://doi.org/https://doi.org/10.1016/j.jedc.2020.103989.

Tables and Figures

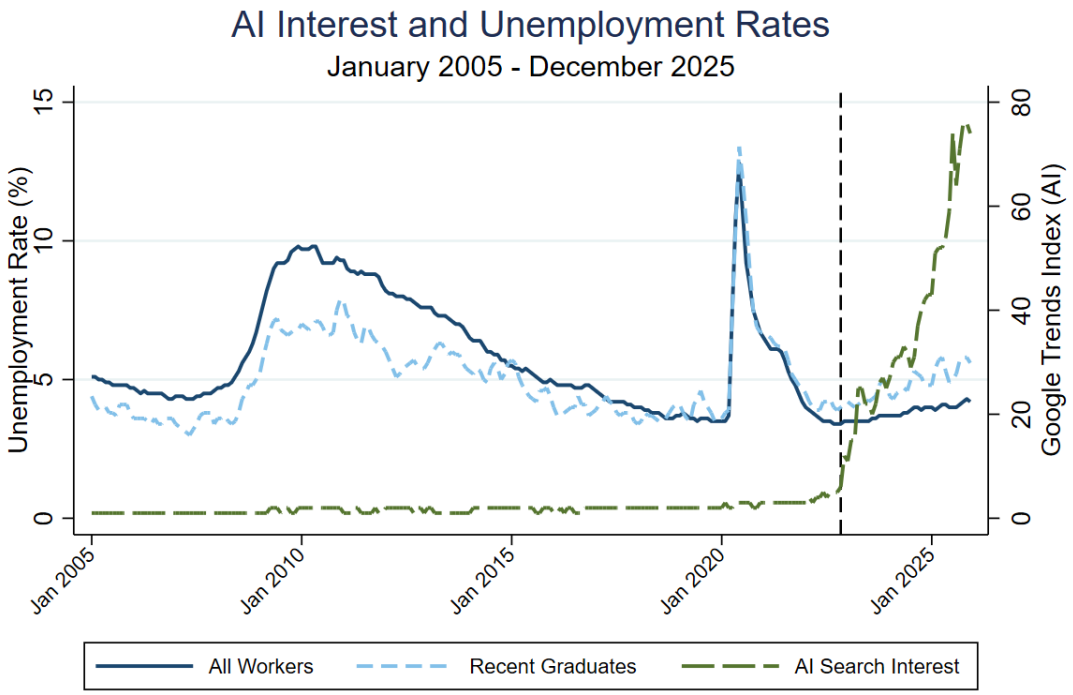
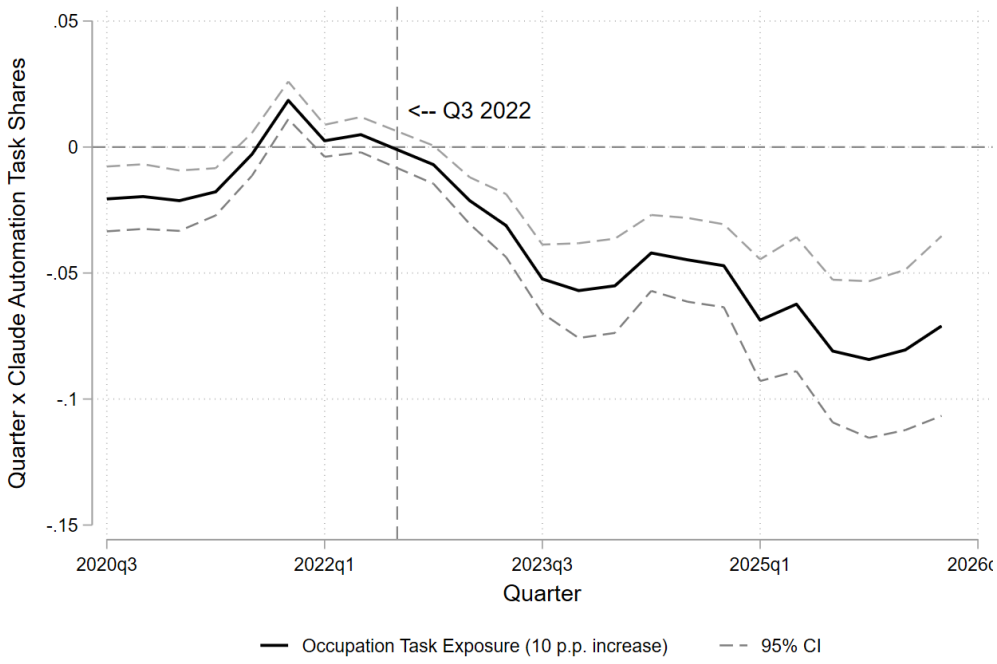
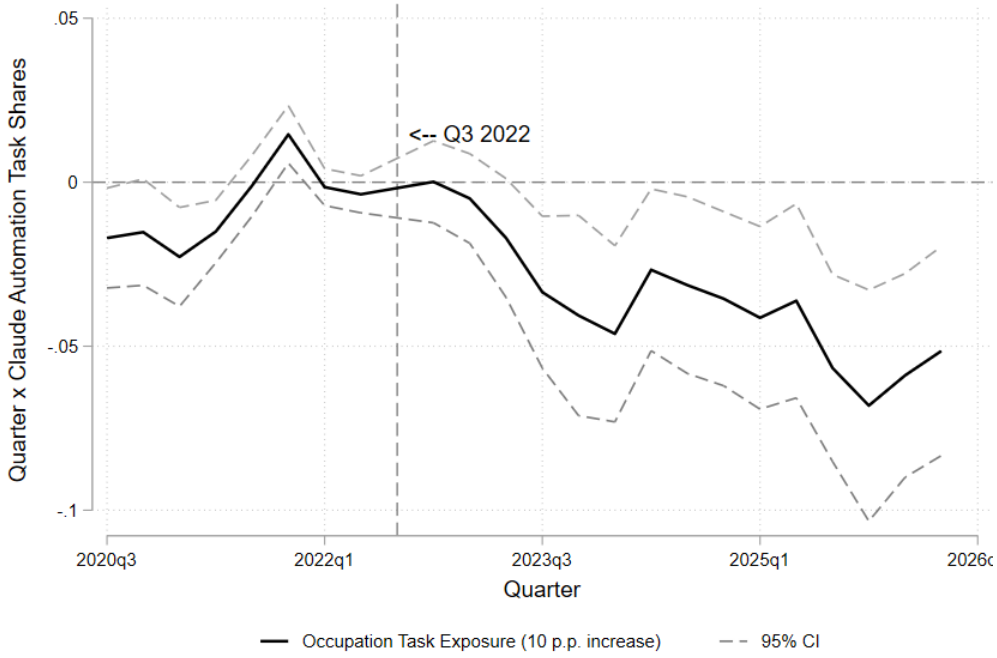


Figure 1: AI Boom Coincided with Increase in Unemployment for Young College Graduates

Notes: This figure presents trends in unemployment and AI interest. Unemployment rates for all workers (ages 16 - 64) and recent college graduates (ages 22 - 27) are seasonally adjusted three-month moving averages series constructed by the Federal Reserve Bank of New York using data from the Current Population Survey.



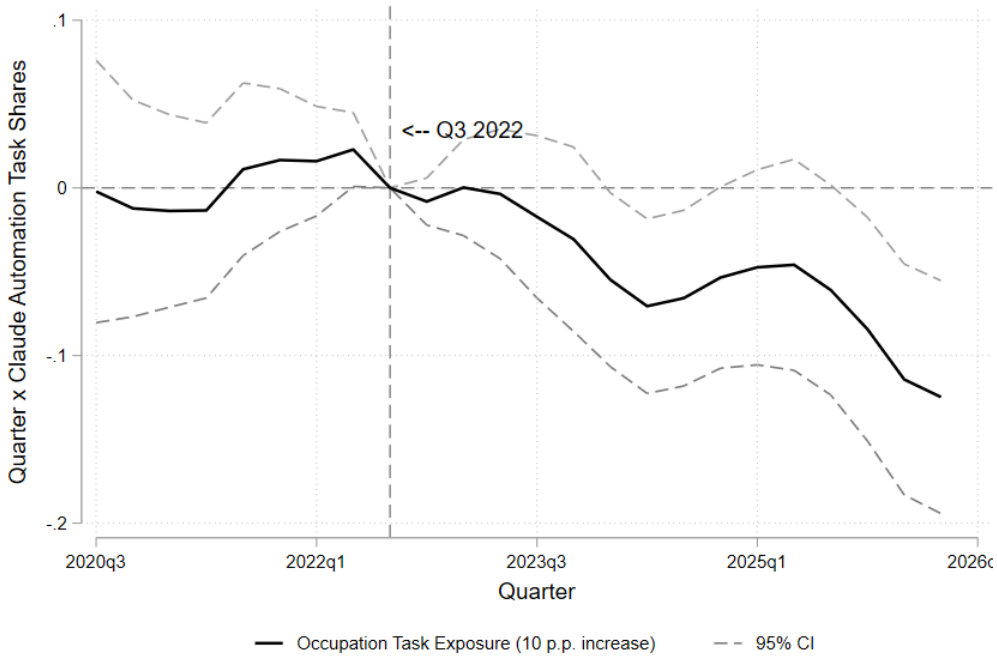
(a) Quarterly Ad Counts by Occupation-Industry (Unweighted)



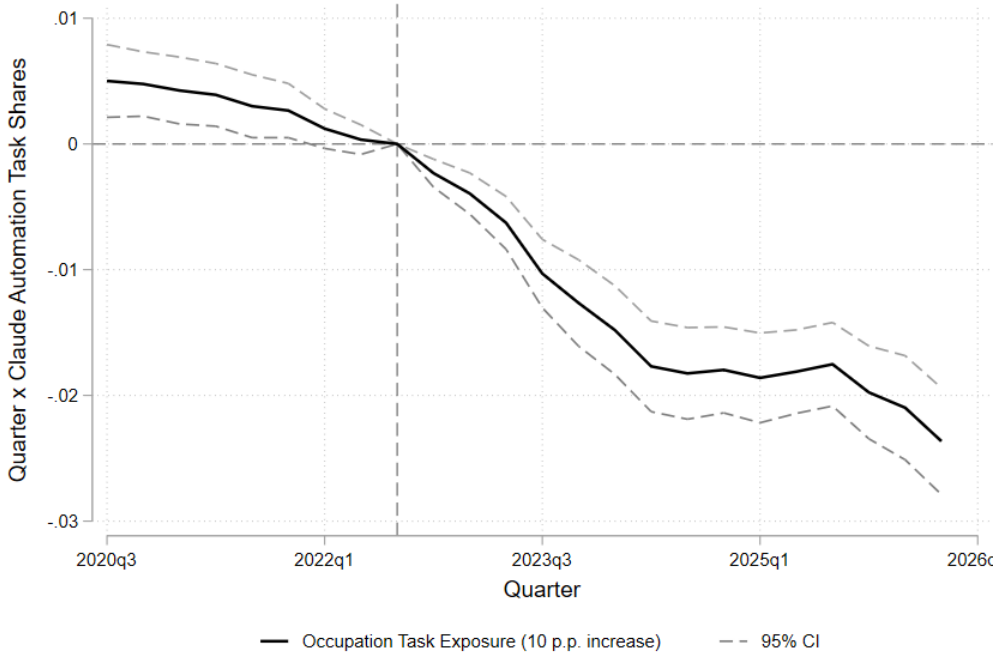
(b) Quarterly Ad Counts by Occupation-Industry (Weighted)

Figure 2: More Exposed Occupations Experienced a Decline in Job Posting Rates

Notes: This figure presents estimates of the relationship between occupation-level exposure to GenAI automation and the rate of job postings in the Lightcast data at the occupation-industry levels, where the estimates are unweighted (Panel A) or weighted (Panel B) by pre-ChatGPT total posts in the cell. Estimates reflect the difference in job postings associated with an occupation having a 10 p.p. greater share of automatable tasks relative to Quarter 3 of 2022. Both panels include fixed effects for occupation-industry cell and industry-by-quarter. Standard errors are clustered at the occupation level. Occupations are defined by the 2019 O*NET classifications, and industry is defined by the four-digit NAICS classification.



(a) Quarterly Ad Counts by Firm



(b) Average Automation Score of Firm's Job Ads

Figure 3: More Exposed Firms Reduced their Job Posting Rates and Pivoted Away from AI-Automatable Occupations

Notes: This figure presents estimates of the relationship between firm-level exposure to GenAI automation and the rate of job postings (Panel A) and average automation exposure in postings (Panel B) in the Lightcast data at the firm level. Outcomes are smoothed using a four-quarter moving average to account for quarterly volatility and are weighted by the total number of postings by the firm prior to the release of ChatGPT. Estimates reflect the difference in job postings associated with an occupation having a 10 p.p. greater baseline share automatable tasks in the firm's job postings before treatment, and the coefficients are measured relative to Quarter 3 of 2022. Both panels include fixed effects for firm and industry-by-quarter. Standard errors are clustered at the firm level. Industry is defined by the four-digit NAICS classification.

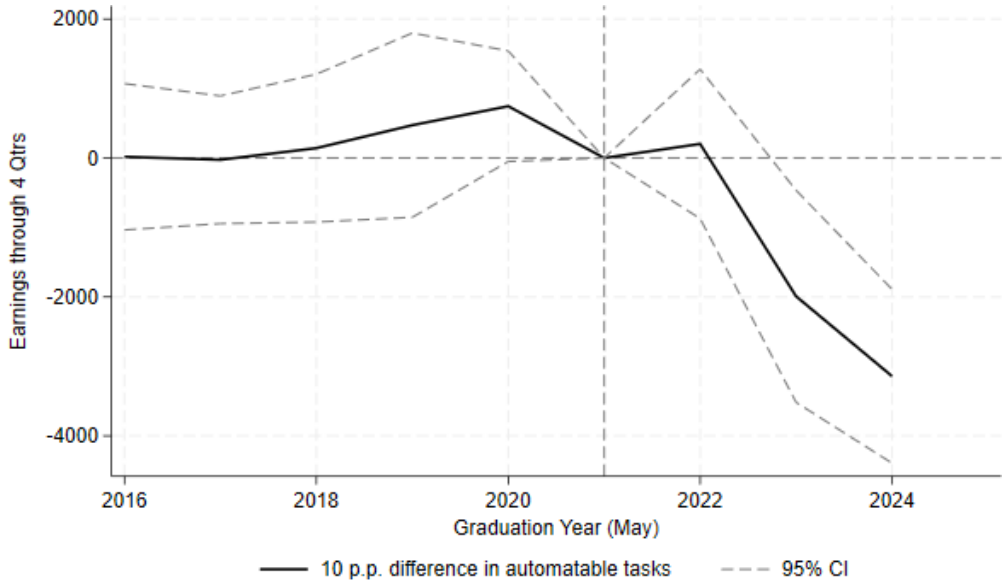
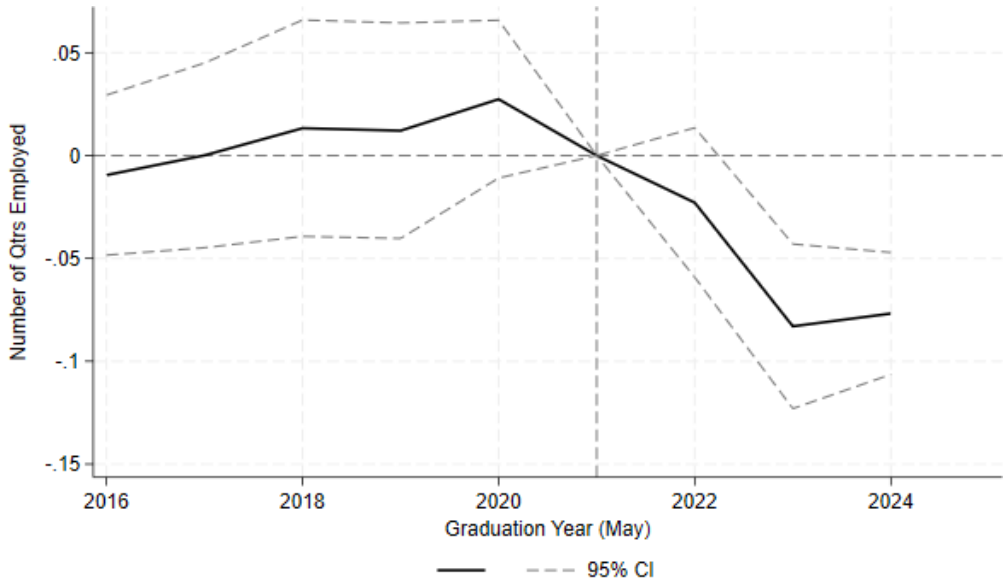
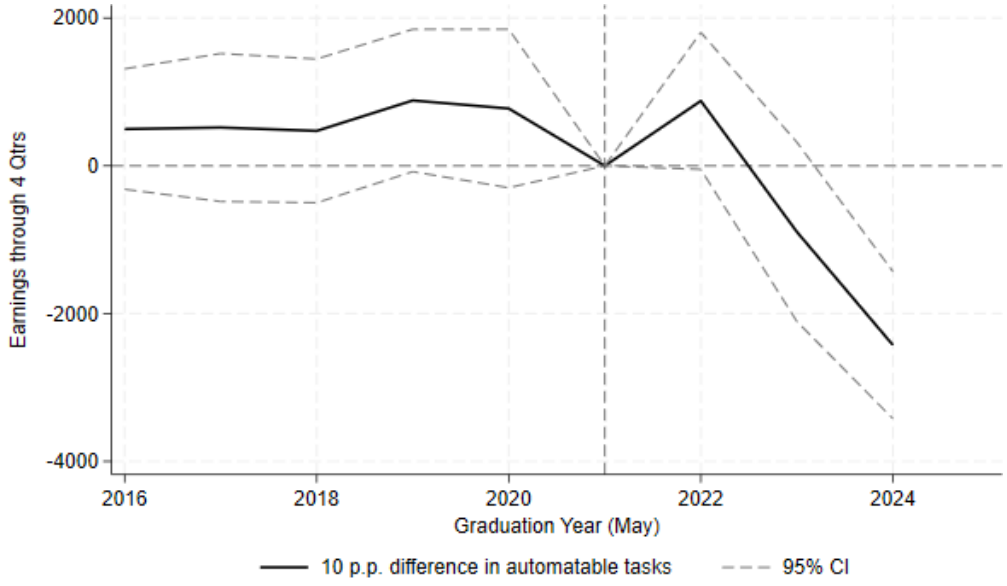


Figure 4: Earnings Fell for New Graduates from AI-Exposed Majors

Notes: This figure presents estimates of the relationship between major-level exposure to GenAI automation and earnings across cohorts of May graduates from four-year universities. The outcome is total earnings in Texas during an individual's first four quarters after completing a four-year degree, including zeroes. Estimates reflect the difference in earnings (2025$, including zeroes) associated with a graduate's major having a 10 p.p. greater share of automatable tasks, relative to this difference for May 2021 graduates. The specification controls for major (4-digit CIP code) and school-by-year fixed effects, student demographics, and college entrance exam scores. Standard errors are clustered by major.



(a) # of Qtrs Employed



(b) Earnings, Employed Grads

Figure 5: GenAI Exposure Associated with Declines in Both Extensive and Intensive Margins of Earnings

Notes: This figure presents estimates of the relationship between major-level exposure to GenAI automation and labor market outcomes across cohorts of May graduates from four-year universities. The outcome in (a) is the number of quarters a graduate was employed in Texas during their first four quarters after completing a four-year degree. The outcome in (b) is earnings (2025$) for graduates that worked all four quarters of their first year after graduation. Earnings are winsorized at the 1st and 99th percentiles. Estimates reflect the difference in outcomes associated with a graduate's major having a 10 p.p. greater share of automatable tasks, relative to this difference for May 2021 graduates. The specification controls for major (4-digit CIP code) and school-by-year fixed effects, student demographics and college entrance exam scores. Standard errors are clustered by major.

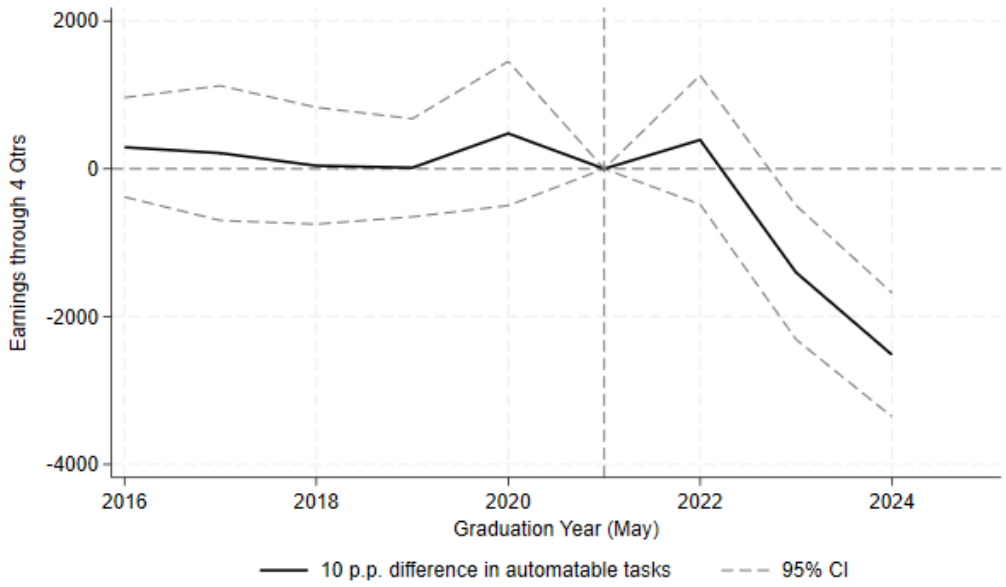


Figure 6: Estimated Declines in Earnings are Robust to Controlling for Time-Varying Industry Shocks & Remote Work Shares

Notes: This figure presents estimates of the relationship between major-level exposure to GenAI automation and earnings across cohorts of May graduates from four-year universities. The outcome is earnings (2025$) for graduates that worked all four quarters of their first year after graduation. Earnings are winsorized at the 1st and 99th percentiles. Estimates reflect the difference in earnings associated with a graduate's major having a 10 p.p. greater share of automatable tasks, relative to this difference for May 2021 graduates. The specification controls for major (4-digit CIP code), school-by-year, and industry-by-year (2-digit NAICS) fixed effects, and student demographics. Standard errors are clustered by major.

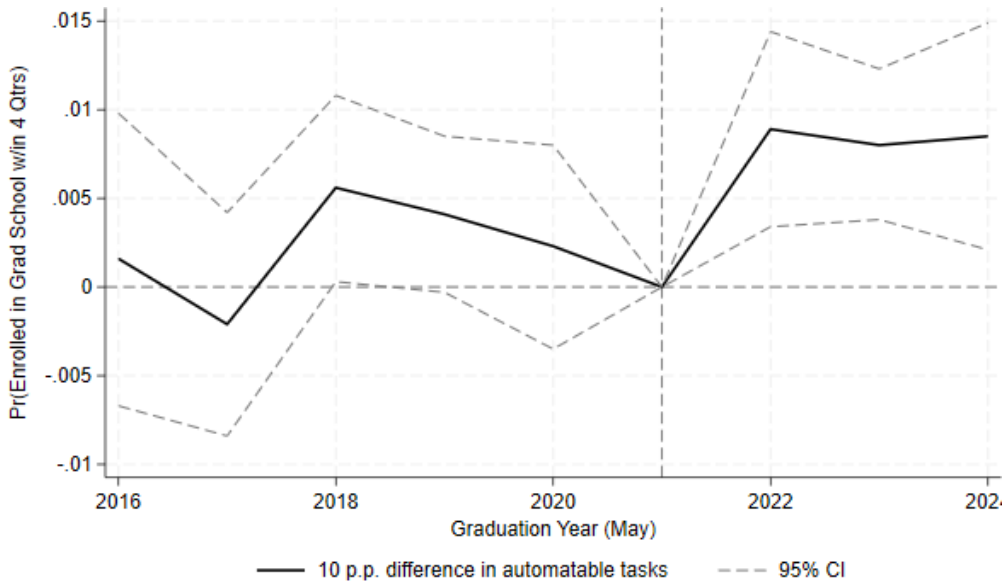


Figure 7: Graduate-School Enrollment Increased for Students from More-Exposed Undergraduate Majors

Notes: This figure presents estimates of the relationship between major-level exposure to GenAI automation and subsequent enrollment in a graduate program. The outcome is an indicator for enrolling in graduate studies within a year of completing their undergraduate degree. Estimates reflect the difference in outcomes associated with a graduate's major having a 10 p.p. greater share of automatable tasks, relative to this difference for May 2021 graduates. The specification controls for major (4-digit CIP code) and school-by-year fixed effects, student demographics, and college entrance exam scores. Standard errors are clustered by major.

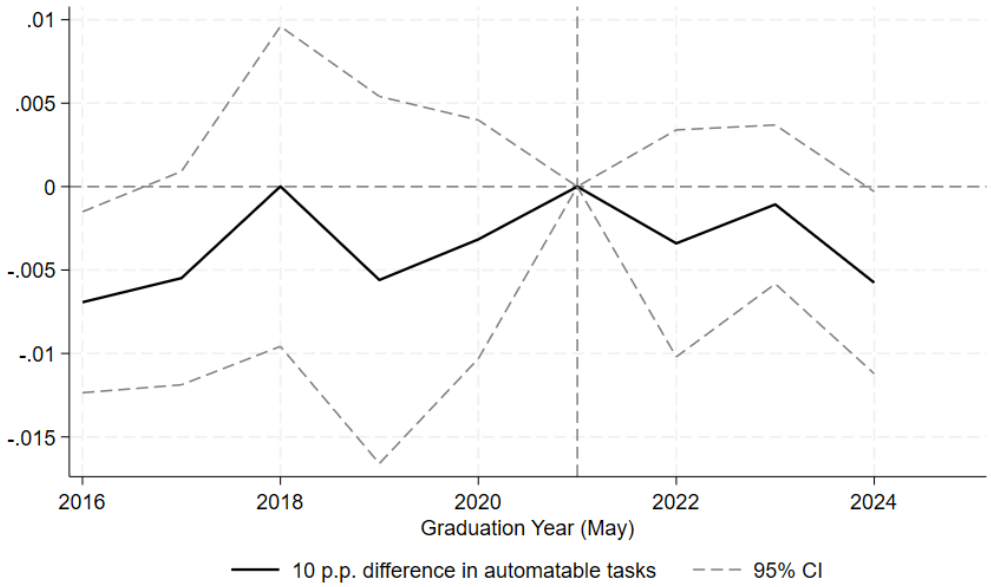
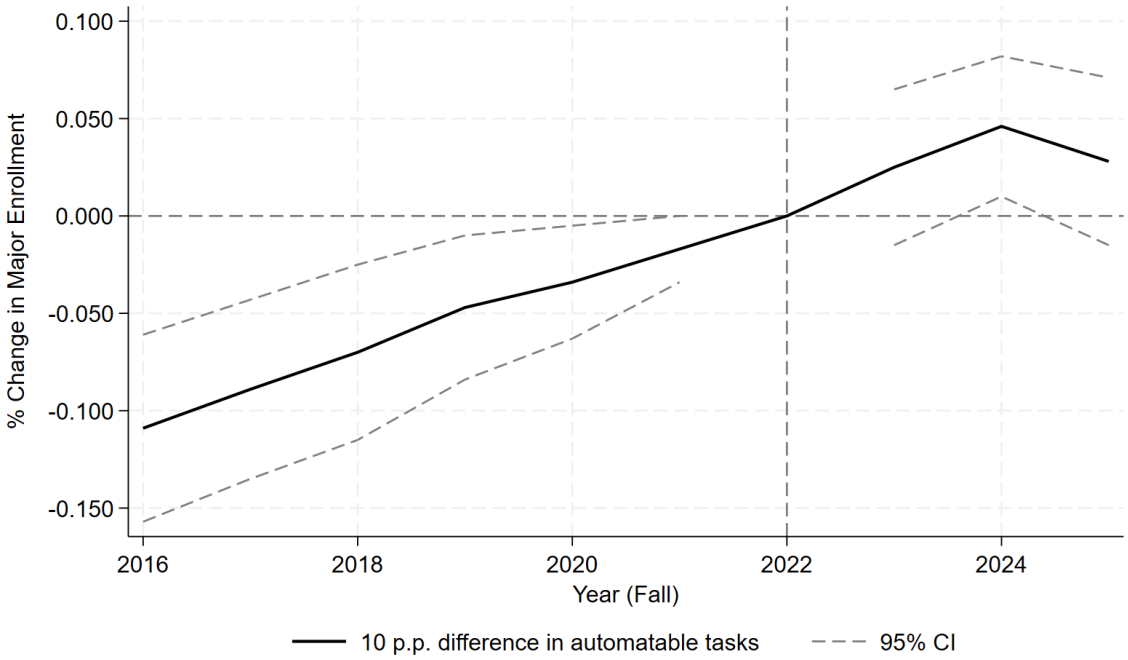
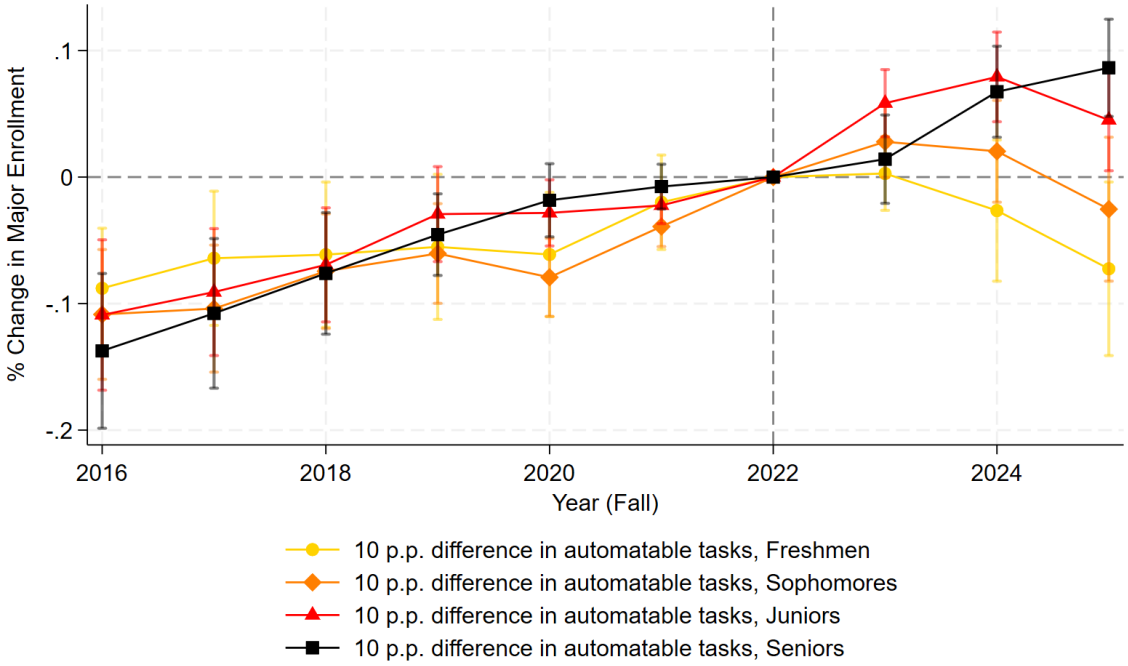


Figure 8: Little Evidence that Graduates From More-Exposed Majors Shifted to Less-Exposed Graduate Programs

Notes: This figure presents estimates of the relationship between major-level exposure to GenAI automation and their graduate program's field-level exposure to GenAI. The outcome is $MajExp_m$ for the *graduate* program in which they enrolled. Estimates reflect the difference in outcomes associated with a graduate's major having a 10 p.p. greater share of automatable tasks, relative to this difference for May 2021 graduates. The specification controls for major (4-digit CIP code) and school-by-year fixed effects, student demographics, and college entrance exam scores. Standard errors are clustered by major.



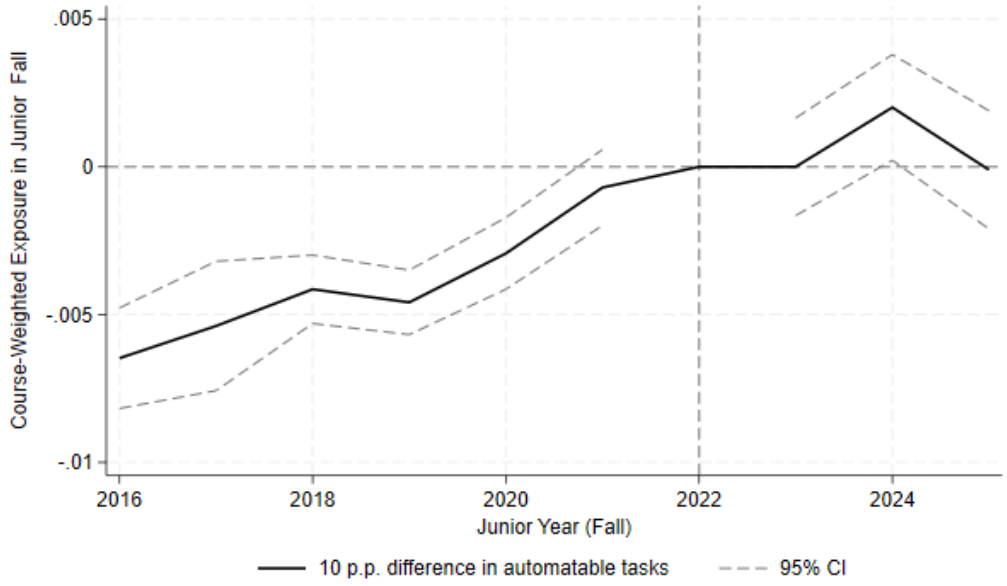
(a) Total Enrollment



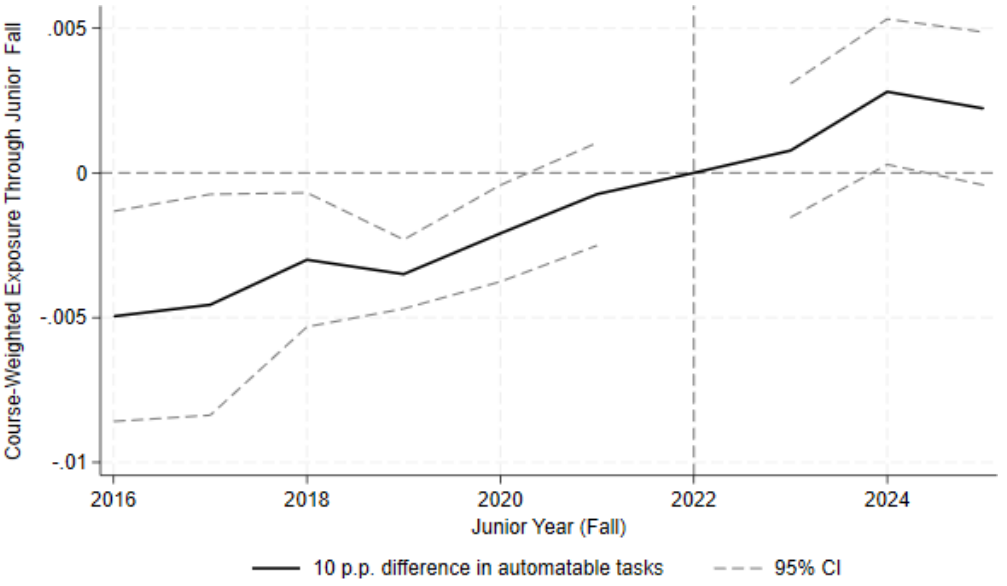
(b) Enrollment by Classification

Figure 9: Underclassmen Shifted Away From High-Exposure Fields

Notes: This figure presents estimates of the relationship between major-level exposure to GenAI automation and enrollment growth by classification. Estimates reflect a Poisson analog of Equation 1 with total and classification-level enrollment counts as the outcome. Coefficients are interpretable as the differences in percent change of Fall enrollment counts between 2021 and a given year associated with a 10 percentage-point difference in major-level GenAI exposure. The specification controls for major and year fixed effects. Standard errors are clustered by major.



(a) Fall Junior Semester Course-Level AI Exposure



(b) Fall Junior Cumulative Course-Level AI Exposure

Figure 10: Course Selections Deviated from Underlying Trends in Fall 2025 for Students that Initially Declared for AI-Exposed Majors

Notes: This figure presents estimates of the relationship between the exposure of a student's Freshman declared major to GenAI automation and their subsequent course-taking behavior. The outcome is the average field-level AI exposure of courses taken in the fall of their Junior year (a) and the cumulative average field-level AI exposure of courses taken through the fall of their Junior year (b). Estimates reflect the difference in outcomes associated with a student's college-entrance major having a 10 p.p. greater share of automatable tasks, relative to this difference for May 2022 graduates. The specification controls for major (4-digit CIP code) and school-by-year fixed effects, student demographics, and college entrance exam scores. Standard errors are clustered by major.

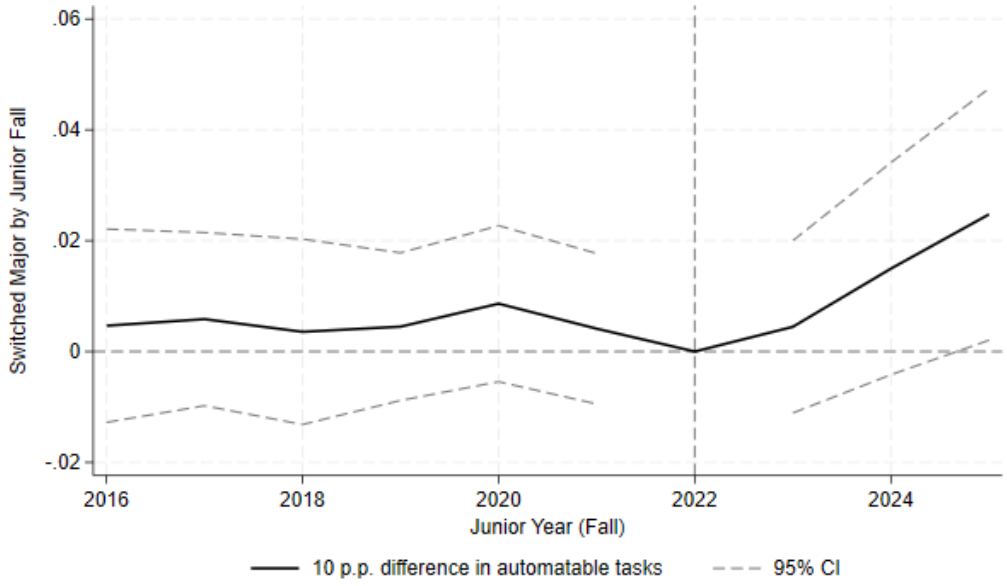


Figure 11: Major-Switching Increased Among Students that Initially Declared for AI-Exposed Majors

Notes: This figure presents estimates of the relationship between the exposure of a student's Freshman declared major to GenAI automation and the likelihood they switched majors by their Junior year. The outcome is an indicator for having a different four-digit CIP code major two falls after initial enrollment. Estimates reflect the difference in outcomes associated with a student's college-entrance major having a 10 p.p. greater share of automatable tasks, relative to this difference for May 2021 graduates. The specification controls for major (4-digit CIP code) and school-by-year fixed effects, student demographics, and college entrance exam scores. Standard errors are clustered by major.

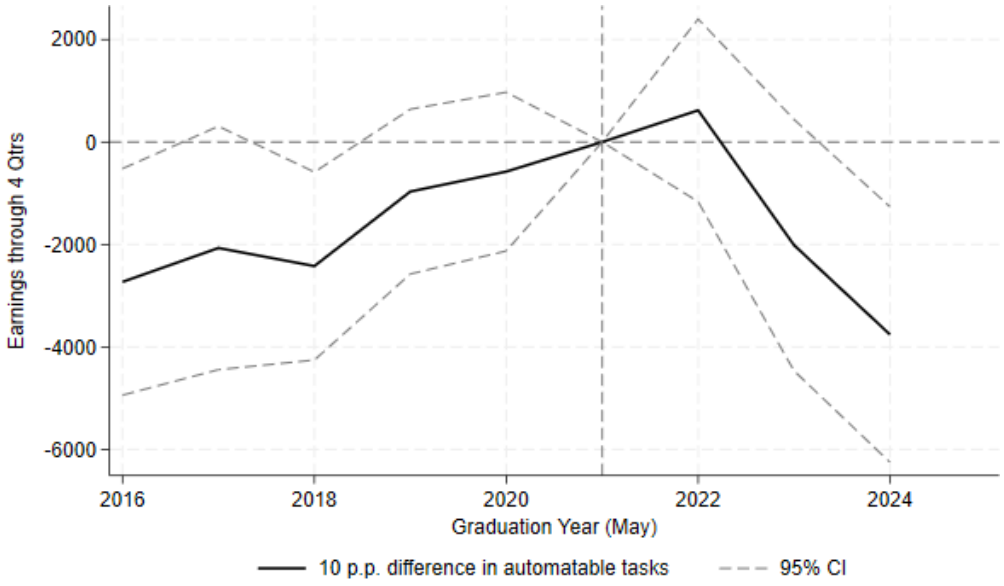


Figure 12: Earnings Fell for New Master's Recipients in AI-exposed Fields

Notes: This figure presents estimates of the relationship between field-level exposure to GenAI automation and earnings across cohorts of May graduates from Master's programs at Texas public four-year universities. Graduates are assigned the same field-level exposure measure used in estimations with four-year graduates. The outcome is earnings (2025$) for graduates that worked all four quarters of their first year after graduation. Earnings are winsorized at the 1st and 99th percentiles. Estimates reflect the difference in earnings associated with a graduate's field having a 10 p.p. greater share of automatable tasks, relative to this difference for May 2022 graduates. The specification controls for major (4-digit CIP code), school-by-year fixed effects, and student demographics. Standard errors are clustered by major.

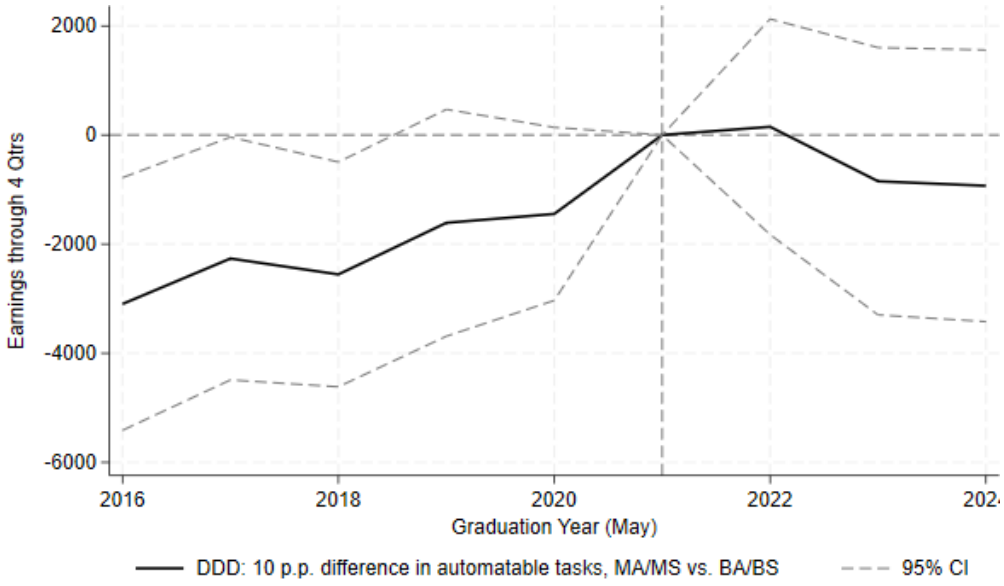


Figure 13: The Earnings Premium for an Exposed Master's Degree Relative to Bachelor's Degree Stopped Growing

Notes: This figure presents estimates of the relationship between field-level exposure to GenAI automation and earnings across cohorts of May graduates from programs at Texas public four-year universities for both Bachelor's and Master's degrees. The dataset is stacked for both degree types, and the coefficients plotted here are for Master's degree by year interacted with major exposure. The outcome is earnings (2025$) for graduates that worked all four quarters of their first year after graduation. Earnings are winsorized at the 1st and 99th percentiles. Estimates reflect the difference in earnings associated with a graduate's Master's degree field having a 10 p.p. greater share of automatable tasks, relative to this difference for May 2021 graduates. The specification controls for major (4-digit CIP code), school-by-year fixed effects, master's degree-by-year fixed effects, and student demographics. Standard errors are clustered by major.

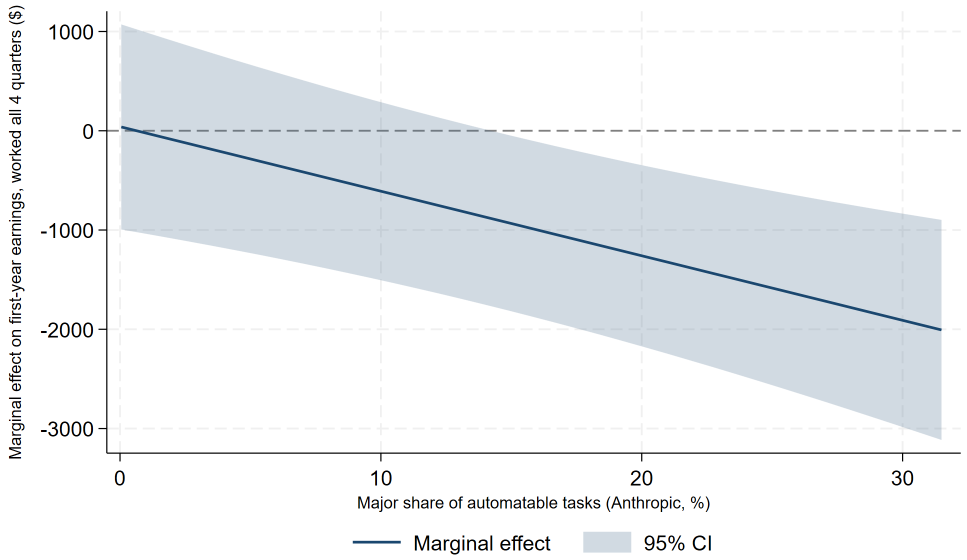


Figure 14: Earnings Losses were Worse for Students from High-Exposure Majors that Took More High-Exposure Courses

Notes: This figure presents marginal effects from a specification that interacts major-level GenAI exposure with a graduate's idiosyncratic course-based exposure, evaluated across values of major-level exposure (the share of the major's tasks that could be automated, in percent). Idiosyncratic exposure is defined as an individual's course-weighted automatable task share across all fields in which they completed courses minus their major's automatable task share. The marginal effects are scaled to reflect a one-standard deviation increase in idiosyncratic exposure. The outcome is earnings (2025$) for graduates that worked all four quarters of their first year after graduation. Shaded areas are 95% confidence intervals.

Table 1: Descriptive Statistics: Lightcast Sample

Panel A: Unweighted

	(1) Mean	(2) SD	(3) N
Total Ads in Quarter	6.08	31.89	2,706,723
OpenAI GPT-4 Beta Rating (10 ppt units)	3.955	2.203	2,253,030
Anthropic Overall Usage Exposure (10 ppt units)	1.020	1.757	2,253,030
Anthropic Automation Exposure (10 ppt units)	0.411	1.038	2,706,723
Contraction from Monetary Policy (Peak)	-3.786	1.026	2,706,304
Contraction from Monetary Policy (Mean)	-1.221	0.195	2,706,304
Share of Job Ads Hybrid or Remote	0.059	0.159	2,211,026

Panel B: Weighted by Pre-ChatGPT Total Ads

	(1) Mean	(2) SD	(3) N
Total Ads in Quarter	153.80	163.60	2,108,803
OpenAI GPT-4 Beta Rating (10 ppt units)	4.127	2.146	2,097,238
Anthropic Overall Usage Exposure (10 ppt units)	1.177	1.776	2,097,238
Anthropic Automation Exposure (10 ppt units)	0.576	1.217	2,108,803
Contraction from Monetary Policy (Peak)	-3.817	0.633	2,108,602
Contraction from Monetary Policy (Mean)	-1.183	0.158	2,108,602
Share of Job Ads Hybrid or Remote	0.067	0.101	2,107,658

Job Ads	Unique Occ-Industry Cells	Unique Occupations	Unique Industries
28,322,499	128,776	963	309

Notes: This table presents descriptive statistics for Lightcast job posting data in Texas by quarter from Q3 2020 through Q2 2026. Observations are structured at the occupation-industry-quarter level. Occupation-industry cells with at least one job posting are present in the data.

Table 2: Descriptive Statistics: New Graduates Sample

	(1) Mean
Male	0.433
	(0.446)
White	0.676
	(0.468)
Hispanic	0.333
	(0.471)
Black	0.102
	(0.303)
Asian	0.143
	(0.350)
Composite SAT Score	1163
	(181)
Major share of automatable tasks (Observed)	0.076
	(0.071)
Major share of automatable tasks (Theoretical)	0.474
	(0.087)
# of Qtrs Employed in First Year after Graduation	2.73
	(1.64)
Earnings in First Year after Graduation (2025$)	31,273
	(31,378)
Enrolled in Graduate School	0.150
	(0.353)
N	320,301
Earnings, Conditional on Being Employed Every Qtr	49,306
	(27,340)
N	178,973

Notes: This table presents descriptive statistics for students from the 2016 to 2024 cohorts of May graduates from public four-year universities in Texas. SAT scores are converted from ACT scores when only ACT is observed. Employment variables reflect employment in a position covered by unemployment insurance in Texas in the first four quarters after graduation. Earnings include zeros. Conditional earnings include only individuals with positive wages in all four quarters their first year after graduation and are winsorized at the 1st and 99th percentiles. Both earnings outcomes are deflated to 2025 dollars.

Table 3: Descriptive Statistics: College Entrants Sample

Age, Fall Freshman Sem.	18.2
	(1.64)
Major share of automatable tasks (Observed), Fall Freshman Sem.	0.083
	(0.073)
Major share of automatable tasks (Theoretical), Fall Freshman Sem.	0.475
	(0.094)
Major share of automatable tasks (Observed), Fall Junior Sem.	0.082
	(0.078)
Swtiched majors from Fall Fresh. to Fall Junior Sem.	0.459
	(0.498)
Course Share in Freshan Major through Fall Jun. Sem	0.256
	(0.199)
N	334,552

Notes: This table presents descriptive statistics for students from the 2014 to 2023 cohorts of first-time undergraduate enrollees at public four-year universities in Texas that enter college with a declared major. SAT scores are converted from ACT scores when only ACT is observed.

Table 4: Example Undergraduate Fields of Study by Exposure Level

Exposure Level	**Example Fields of Study**
$MajExp_m < 0.1$	Agricultural/Biological Engineering, Animal Sciences, Astronomy, Business Administration, Ecology, Economics, Elementary Education, History, Hospitality Administration, Human Development, Mental and Social Health Services, Mining and Mineral Engineering, Nutrition Sciences, Philosophy, Plant Sciences, Psychology, Real Estate, Registered Nursing, Urban Studies
$0.1 \leq MajExp_m < 0.2$	Engineering, Journalism, Mathematics, Political Science, Statistics
$MajExp_m \geq 0.2$	Computer Science, Computer Engineering, Geography and Cartography, Linguistic, Comparative, and Related Language Studies and Services, Systems Science and Theory

Notes: Fields of study represent four-digit Classification of Instructional Programs codes from the National Center for Education Statistics. $MajExp_m$ is the average share of tasks among pre-AI job postings listing major m that can be automated based on Claude usage reported by Anthropic.

Table 5: GenAI Exposure is Associated with Persistent Declines in Earnings and Employment for New Graduates

	(1) Earnings	(2) # of Qtrs Employed	(3) Earnings, Employed All 4 Qtrs
A. 1- 4 Qtrs After Graduation			
$MajExp_m$	-2,012**	-0.083***	-934
	(776)	(0.020)	(622)
N	280,532	280,532	156,339
B. 5 - 8 Qtrs After Graduation			
$MajExp_m$	-2,351***	-0.045**	-2,102***
	(726)	(0.018)	(481)
N	280,532	280,532	133,365

Standard errors clustered by major in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table presents estimated effects of GenAI exposure on first-year (Panel A) and second-year (Panel B) labor market outcomes for May 2023 graduates from Texas public four-year universities. Estimates reflect the difference in earnings associated with a graduate's major having a 10 p.p. greater share of automatable tasks, relative to this difference for May 2021 graduates. The outcomes are (1) total earnings (2025$) in Texas across the panel's designated four quarters, (2) number of quarters with non-zero earnings, and (3) earnings, conditional on working all four quarters. The specification controls for major (4-digit CIP code) and school-by-year fixed effects, student demographics, and college entrance exam scores. Standard errors are clustered by major.

Table 6: Estimated Effects of GenAI Exposure on First-Year Labor Market Outcomes Are Robust Across Specifications: Share of Automatable Tasks (Anthropic)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
A. First-Year Earnings														
$MajExp_m$	-2,853***	-3,141***	-2,239***	-2,360***	-2,817***	-3,070***	-1,962***	-2,029***	–	–	–	–	–	–
	(422)	(488)	(525)	(488)	(834)	(741)	(560)	(580)						
B. Number of Quarters Employed in First Year														
$MajExp_m$	-0.057***	-0.077***	-0.050***	-0.067***	-0.057***	-0.070***	-0.037***	-0.052***	–	–	–	–	–	–
	(0.017)	(0.015)	(0.015)	(0.015)	(0.016)	(0.016)	(0.014)	(0.016)						
C. First-Year Earnings, Conditional on Working All Four Quarters														
$MajExp_m$	-2,557***	-2,427***	-2,019***	-1,708***	-2,396***	-2,265***	-1,701***	-1,401***	-2,790***	-2,518***	-1,921***	-2,702***	-2,427***	-1,672***
	(658)	(508)	(446)	(394)	(724)	(597)	(535)	(491)	(477)	(428)	(373)	(531)	(514)	(435)
Major FEs	X	X	X	X	X	X	X	X	X	X	X	X	X	X
School-by-Year FEs	X	X	X	X	X	X	X	X	X	X	X	X	X	X
SAT Scores		X		X		X		X		X	X		X	X
Industry-by-Year FEs									X	X	X	X	X	X
Major Remote Shares			X	X			X	X			X			X
No Graduate Students					X	X	X	X				X	X	X

Standard errors clustered by major in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table presents estimated effects of GenAI exposure on first-year labor market outcomes for May 2024 graduates from Texas public four-year universities, where exposure is measured by the share of a major's tasks that are automatable by GenAI, measured from Anthropic's Economic Index. Each cell reports the coefficient on the 2024 cohort event study indicator from a separate regression; estimates reflect the difference in outcomes associated with a graduate's major having a 10 p.p. greater exposure to GenAI, relative to this difference for May 2021 graduates. The outcome is total first-year earnings (2025$) in Panel A, the number of quarters with non-zero earnings in the first year in Panel B, and first-year earnings conditional on working all four quarters in Panel C. Columns (9)–(14) are estimated only for the conditional-earnings outcome, so the corresponding cells in Panels A and B are left blank. Rows at the bottom of the table indicate the fixed effects, controls, and sample restrictions included in each specification; all specifications control for student demographics. Standard errors are clustered by major.

Online Appendix to:

"Degrees for GPTs: The Effects of Generative AI on Skill Demand and Supply"

Samuel Dodini **Tucker Smith** **Yaling Xu**

November 19, 2026

A Conceptual Framework: Formal Model

This appendix formalizes the human capital investment framework in Section 4.1, in the tradition of Becker (1962). GenAI enters students' investment problem through three forces: an automation penalty that lowers the earnings premium for a standard degree in fields whose typical occupations contain more automatable tasks, a complementarity premium for additional expertise beyond a field's standard curriculum, which may partly or wholly offset the automation penalty, and a reduction in effort costs for fields with automatable coursework. Students and recent graduates respond on four margins – the major chosen at college entry, switching majors, the allocation of coursework within a declared major, and the decision to enroll in graduate school and in which field. In response to automation, students can *upskill*, deepening their expertise in an exposed field, or *pivot* toward less-exposed fields. Switching costs that rise with seniority determine which margins are available to whom.

A.1 Setup

Each student i has a declared major $m_0(i)$, idiosyncratic preferences over fields $\{\xi_{im}\}$, and seniority $s \in \{0, 1, \ldots, \bar{s}\}$, which indexes accumulated credits. A new student has $s = 0$ and no attachment to a major. Major m's exposure to GenAI is:

$$a_m \;\equiv\; MajExp_m \;\in\; [0, 1],$$

where $MajExp_m$ is the major's automatable task share in our empirical models. $S \in \{0, 1\}$ is an indicator for before ($S = 0$) versus after ($S = 1$) the release of ChatGPT. We assume students form beliefs about how long the GenAI shock to labor demand will last. The parameter $\rho \in [0, 1]$ denotes the expected persistence of the shock into the student's future career.

A.2 Undergraduate payoffs

The present discounted value of earnings from an undergraduate degree in major m is

$$B(m \mid S) \;=\; \beta\big(y_m - \rho\,\delta\,a_m\,S\big), \qquad \beta > 0, \;\; \delta \geq 0, \tag{1}$$

where y_m is the field's pre-AI log earnings level and β discounts post-graduation earnings to the time of the decision. The parameter δ is the **automation penalty channel**: after the release of ChatGPT, earnings in field m fall in proportion to the major's automatable task share a_m.

GenAI could differ from many skill-biased demand shocks in the literature because it also affects the cost of acquiring the skill or credential to the occupation. The time, tuition, and effort cost of completing major m is

$$C(m \mid S) \;=\; c_m - \gamma_m \, a_m \, S, \qquad \gamma_m \geq 0, \tag{2}$$

where c_m is the pre-GenAI completion cost of the major. The parameter γ_m is the **effort cost channel** and characterizes how much GenAI lowers the effort required to complete coursework, allowing that reduction to be correlated with the field's exposure to automation in the labor market.

The net value of an undergraduate degree in major m given these dynamics is

$$V(m \mid S) \;=\; B(m \mid S) - C(m \mid S) \;=\; \underbrace{(\beta y_m - c_m)}_{\equiv v_m} \;-\; S \, a_m \, \underbrace{(\beta \rho \delta - \gamma_m)}_{\equiv \psi_m}, \tag{3}$$

where v_m is the net value of the major before GenAI and ψ_m is the net penalty to exposure after the release of ChatGPT. A more-exposed major becomes less attractive after the technology is adopted if and only if $\psi_m > 0$, i.e., that the perceived discounted earnings penalty $\beta\rho\delta$ exceeds the reduction in effort costs γ_m.

A.3 Choosing and Switching Majors

Switching costs. For a current student, leaving a declared major forfeits the credits they have earned towards their prior major and often delays graduation. The cost of switching is

$$K_i(s) \;=\; \kappa(s) + \eta_i, \qquad \kappa'(s) > 0, \;\; \kappa(0) = 0, \;\; \eta_i \sim G(\cdot) \text{ iid}, \tag{4}$$

where η_i is an idiosyncratic cost of switching. A new or undeclared student ($s = 0$) has no major attachment and faces only η_i.

The switching rule. A student attached to an exposed major m_0 weighs switching to the next-best less-exposed alternative $m^\star$, with $a_{m^\star} < a_{m_0}$. Let $\Delta a \equiv a_{m_0} - a_{m^\star} > 0$ and $\Delta v_i \equiv (v_{m^\star} - v_{m_0}) + (\xi_{im^\star} - \xi_{im_0})$, where the preference term captures tastes and the skill distance between the two fields. For simplicity, let $\psi_m = \psi$ for both majors. The student pivots by switching majors if

$$\underbrace{\Delta v_i}_{\text{pre-AI returns + tastes}} \;+\; \underbrace{S \, \psi \, \Delta a}_{\text{AI-induced gain from leaving}} \;>\; \kappa(s) + \eta_i, \tag{5}$$

so the share of students at seniority s who switch is

$$\text{Switch}(s \mid S) \;=\; G\big(\Delta v + S\,\psi\,\Delta a - \kappa(s)\big). \tag{6}$$

For entering students, the same expression with $\kappa(0) = 0$ gives the share choosing the less-exposed major.

A.4 Coursework within a declared major

A student who stays in major m still chooses how to allocate their remaining coursework. Let $x \in [0,1]$ denote the share of a student's coursework in major m, with the remainder in a less-exposed field o ($a_o < a_m$), and let $\bar{x}$ denote the share in the major's standard curriculum. The exposure of the student's course bundle is $e(x) = x\,a_m + (1-x)\,a_o$, so the student's *idiosyncratic* course exposure – the deviation of their course bundle from their major's typical exposure as measured in Section 4.3 – is

$$e(x) - e(\bar{x}) \;=\; (x - \bar{x})\,(a_m - a_o).$$

Raising x above $\bar{x}$ is course-level *upskilling*; lowering it is course-level *pivoting*.

Deeper coursework in an exposed field has two possible effects on earnings in the labor market. First, it raises exposure to substitution risk, which lowers earnings at rate δ per unit of exposure as in Equation 1. Second, it also builds advanced expertise that may complement GenAI (Autor and Thompson, 2025), which raises earnings at rate $\lambda \geq 0$. Together, effort costs fall with exposure at rate γ_m, but a student at seniority s can only reallocate the share $r(s) \equiv 1 - s/\bar{s}$ of their remaining coursework. Writing $x - \bar{x} = r(s)\,d$, where d is the deviation within remaining coursework, and letting departures from the standard curriculum carry a convex cost $\frac{\phi}{2} r(s) d^2$, the student chooses d to maximize

$$-\,S\,r(s)\,d\,(a_m - a_o)\,\underbrace{\big[\beta\rho(\delta - \lambda) - \gamma_m\big]}_{\equiv\,\psi^c_m} \;-\; \frac{\phi}{2}\,r(s)\,d^2,$$

which gives

$$d^* \;=\; -\,\frac{S\,(a_m - a_o)\,\psi^c_m}{\phi}, \qquad x^* - \bar{x} \;=\; r(s)\,d^*. \tag{7}$$

At the course level, students can pivot ($d^* < 0$) if $\psi^c_m > 0$ and upskill ($d^* > 0$) if $\psi^c_m < 0$. This occurs if the ways in which expertise complements the technology are strong (λ close to or above δ) or effort costs fall sharply (γ_m large). Unlike switching majors, reallocating coursework does not forfeit credits toward their major and carries no large fixed cost, so

any student with coursework before graduation adjusts at the intensive margin, while only students whose gains exceed $\kappa(s) + \eta_i$ will switch majors.

A.5 Graduate school: enrollment and field choice

A new bachelor's degree recipient from major m has the option to choose between entering the labor market and enrolling in a graduate program in field $g \in \{m, o\}$, where $g = m$ is graduate-level *upskilling* and $g = o$ ($a_o < a_m$) is *pivoting*. Entering the labor market gives the graduate earnings $y_m - \delta a_m S$, which reflect the realized shock today, and continuation earnings $y_m - \rho\delta a_m S$ in the years after. A graduate degree in field g yields earnings

$$y_g^G \;-\; \rho\,(\delta - \lambda_G)\,a_g\,S,$$

where $\lambda_G \geq 0$ measures the complementarity between GenAI and the advanced expertise and skills acquired in graduate school. We allow λ_G to differ from the undergraduate course-level complementarity λ because graduate training may build different skills. For example, graduate programs may train more in supervising, integrating, and validating GenAI output (Acemoglu and Restrepo, 2022; Chen et al., 2026; Garicano et al., 2026) than additional undergraduate courses. The cost of the graduate program is characterized by $c_g^G - \gamma_G a_g S$, which, as at the undergraduate level, falls with the field's exposure because of reduced effort costs.

The net value of enrolling in field g, relative to entering the labor market, is

$$\Delta_{ig}^G(S) \;=\; \beta\big[y_g^G - \rho(\delta - \lambda_G)a_g S\big] \;-\; \big(c_g^G - \gamma_G a_g S\big) \;-\; \underbrace{\big(y_m - \delta a_m S\big)}_{\text{foregone earnings}} \;-\; \beta\big(y_m - \rho\delta a_m S\big) \;+\; \zeta_{ig}, \tag{8}$$

where ζ_{ig} is an idiosyncratic taste shock for graduate study in field g. The student enrolls if $\max_g \Delta_{ig}^G > 0$.

Enrollment: The shock raises the value of enrolling in either field:

$$\frac{\partial \Delta_{im}^G}{\partial S} \;=\; a_m\big[\delta + \beta\rho\lambda_G + \gamma_G\big] \;>\; 0, \qquad \frac{\partial \Delta_{io}^G}{\partial S} \;=\; \delta\big[a_m + \beta\rho(a_m - a_o)\big] + a_o\big(\beta\rho\lambda_G + \gamma_G\big) \;>\; 0,$$

and both derivatives increase in a_m. The first term in each is the fall in the opportunity cost of further schooling: the demand shock lowers the earnings a graduate from an exposed major gives up by returning to school. This force operates even if $\rho = 0$, i.e., even if students expect the shock to be temporary, in which case graduate school serves as a safe harbor from

the shock.

Field choice. Conditional on enrolling, the graduate upskills rather than pivots if $\Delta_{im}^{G} > \Delta_{io}^{G}$, i.e., if

$$\left[\beta(y_m^G - y_o^G) - (c_m^G - c_o^G)\right] + (\zeta_{im} - \zeta_{io}) \; - \; S\,(a_m - a_o)\,\underbrace{\left[\beta\rho(\delta - \lambda_G) - \gamma_G\right]}_{\equiv \psi^G} \; > \; 0. \qquad (9)$$

The shock shifts graduate enrollees toward less-exposed programs if $\psi^G > 0$ and toward exposed programs if $\psi^G < 0$. Pivoting is more likely when the shock is perceived as persistent (high ρ) and when graduate expertise is a weak complement to GenAI (low λ_G).

Signaling: Graduate training in an exposed field may also serve as a signal of ability. If GenAI lowers the effort required to earn strong grades in exposed fields, undergraduate credentials in those fields may become less informative about ability, raising the signaling value of an advanced degree in the same field. This force adds to the payoff from graduate-level upskilling in Equation 9 and is observationally similar to a larger λ_G in enrollment and field choices. We do not model it formally or separately identify it.

A.6 Predictions

The predictions below correspond to predictions (a)–(g) in Section 4.1. Predictions (a)–(c), (e), and (f) concern choices and depend on students' *perceived* parameters (ρ, δ, λ, λ_G); predictions (d) and (g) concern realized earnings and therefore reveal the *actual* degree of substitution and complementarity.

(a) **Pivoting away from exposed majors:** From Equation 6, $\partial\,\text{Switch}/\partial S = g(\cdot)\,\psi\,\Delta a > 0$ if and only if $\psi = \beta\rho\delta - \gamma_m > 0$: students pivot away from exposed majors when the labor demand channel dominates the effort cost channel. Because $\kappa(0) = 0$, the response is largest for entering students. (Figures 9 and 9b.)

(b) **Coursework adjusts before majors do:** Switching falls with seniority, $\partial\,\text{Switch}/\partial s = -g(\cdot)\,\kappa'(s) < 0$, and there is a threshold $s^\star$ defined by $\kappa(s^\star) = \Delta v + \psi\,\Delta a$ above which the marginal student ($\eta_i = 0$) completes their original major. Course reallocation (Equation 7) has no fixed cost, so students above $s^\star$ can still adjust – pivoting at the course level if $\psi_m^c > 0$ or specializing further if $\psi_m^c < 0$ – although the size of the adjustment, $r(s)\,d^*$, also shrinks as less coursework remains. (Figures 9b, 10, and 11.)

(c) **Slow reallocation and exposed recent graduates:** Students who had passed $s^\star$ when GenAI arrived, and recent graduates with $r(s) = 0$, cannot pivot at the undergraduate level at all. Reallocation across majors, therefore, occurs mainly through entering cohorts of new students, and recent graduates from exposed fields bear the shock with little scope to adjust. If students update ρ gradually as the shock becomes visible in labor market outcomes, reallocation is slower still. (Figures 4 and 9b.)

(d) **Returns to course-level upskilling:** Among graduates of the same major, the realized effect of a unit increase in idiosyncratic course exposure on post-shock log earnings is $-(\delta - \lambda)$. If $\delta > \lambda$, graduates whose coursework is more exposed than their major's earn less, so course-level pivoting pays; if $\lambda > \delta$, course-level upskilling pays. (Figure 14.)

(e) **Graduate enrollment:** Graduate enrollment rises for graduates of more-exposed majors, $\partial^2 \Delta_{ig}^G / \partial S \, \partial a_m > 0$ for both $g = m$ and $g = o$, driven at minimum by the fall in the opportunity cost of schooling. (Figure 7.)

(f) **Graduate field choice:** From Equation 9, graduate enrollees from exposed majors upskill if $\psi^G = \beta\rho(\delta - \lambda_G) - \gamma_G < 0$ – because they perceive graduate expertise as complementary to GenAI, expect the shock to be temporary, or face sharply lower effort costs – and pivot toward less-exposed programs if $\psi^G > 0$. (Figure 8.)

(g) **Returns to graduate-level upskilling:** Post-shock earnings of new master's recipients fall with the exposure of their field at rate $\delta - \lambda_G$. Earnings hold up if $\lambda_G \geq \delta$; declines indicate that graduate-level upskilling offers limited protection. (Figure 12.)

B Data Validation and Additional Context

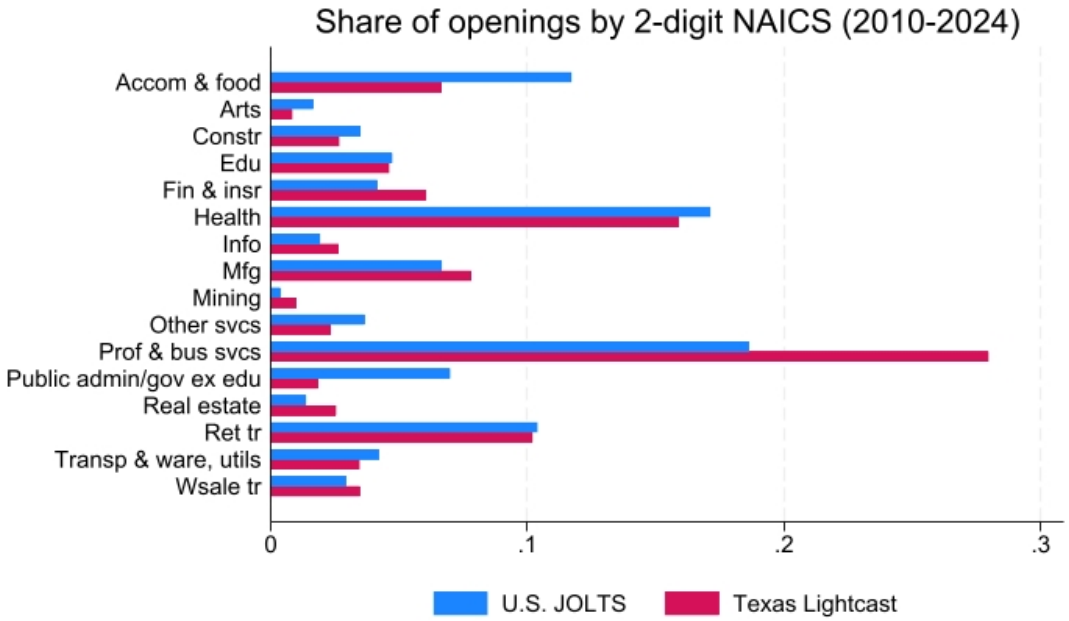
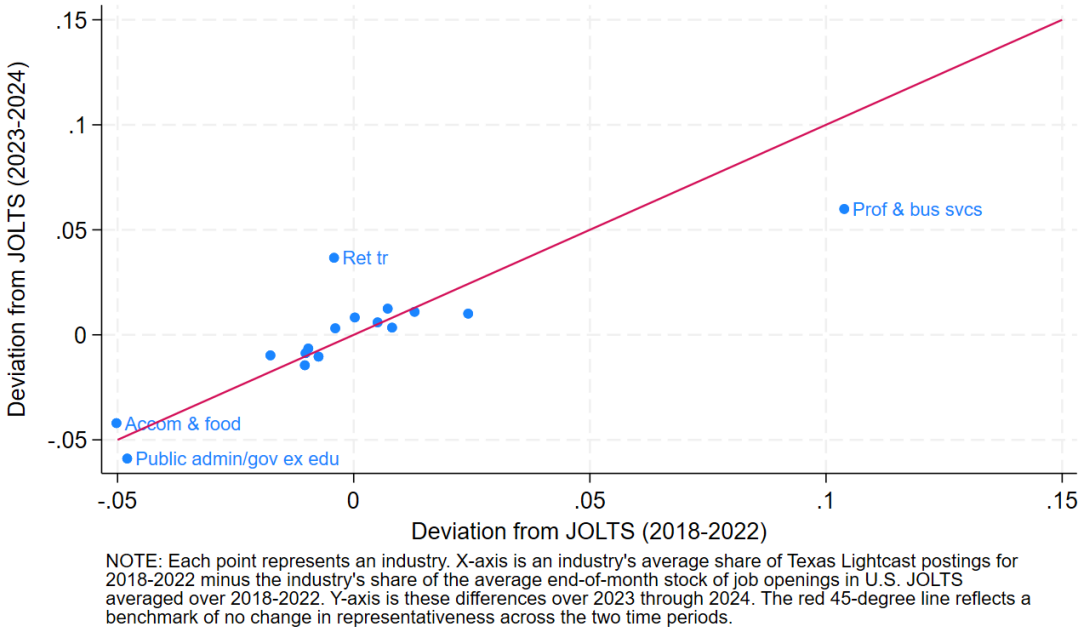


Figure B1: Posting Shares Across Industries in Lightcast vs. JOLTS

Notes: This figure presents the shares of job openings across industries (2-digit NAICS codes) as measured by Lightcast and JOLTS. The shares are averaged across monthly data from each source for 2010 through 2024. JOLTS is nationally representative, while the Lightcast sample is only for openings in Texas.



NOTE: Each point represents an industry. X-axis is an industry's average share of Texas Lightcast postings for 2018-2022 minus the industry's share of the average end-of-month stock of job openings in U.S. JOLTS averaged over 2018-2022. Y-axis is these differences over 2023 through 2024. The red 45-degree line reflects a benchmark of no change in representativeness across the two time periods.

Figure B2: Posting Shares Across Industries in Lightcast vs. JOLTS

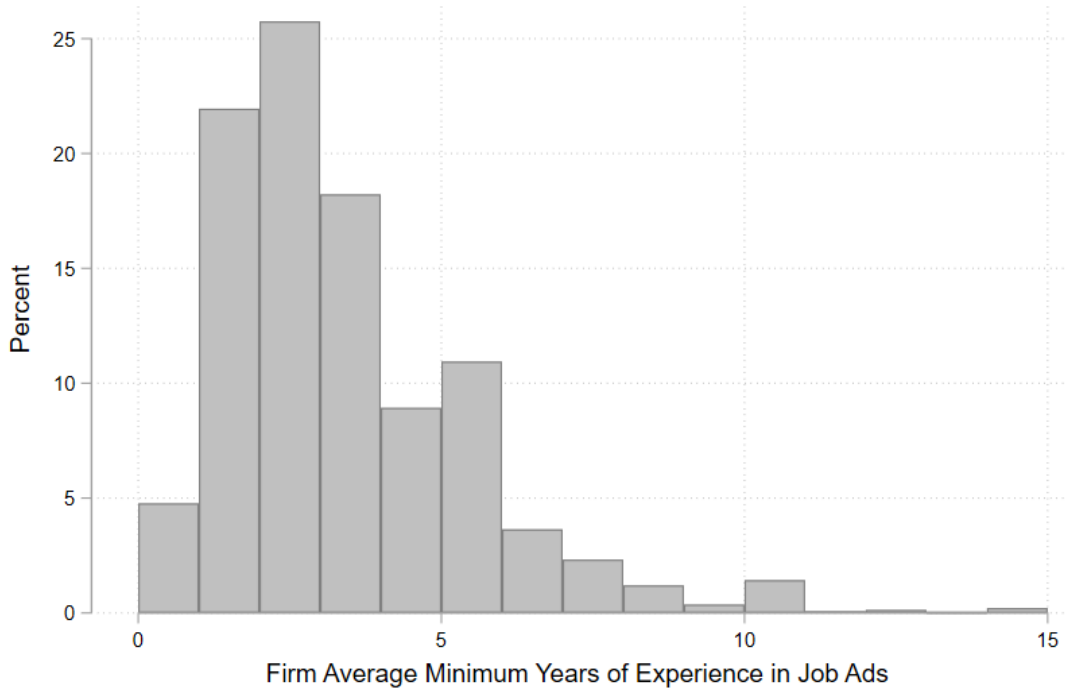
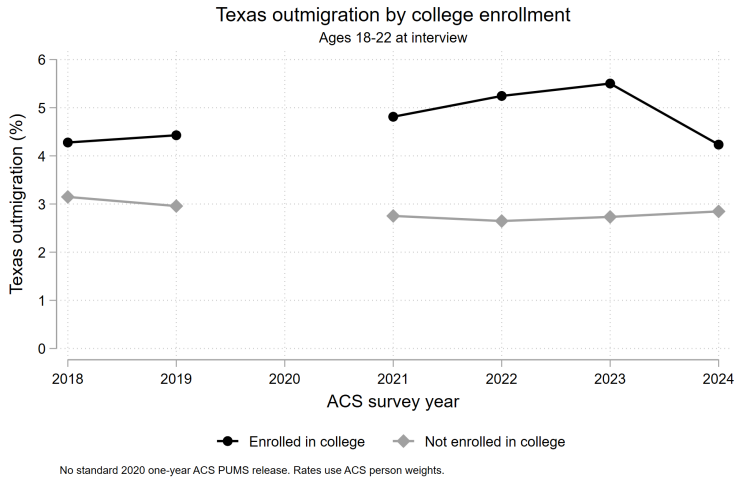
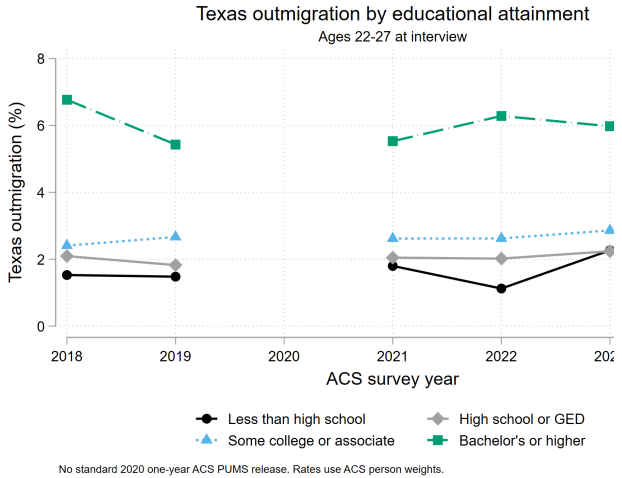


Figure B3: Average Required Minimum Experience for Firms' Job Ads in Lightcast

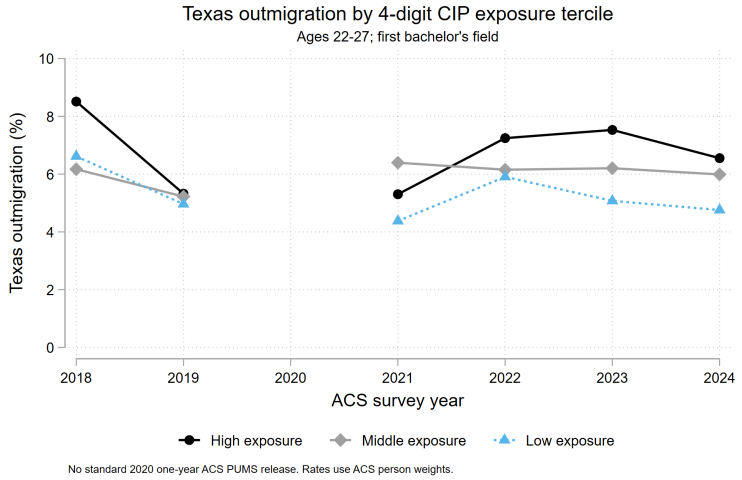
Notes: This figure presents a histogram of the firm-level average of the minimum required years of work experience listed in Lightcast job ads throughout our analysis period.



(a) Age 18-22 by Educational Enrollment



(b) Age 22-27 By Educational Attainment



(c) Age 22-27 By College Major AI Automation Exposure

Figure B4: Outmigration Rates from Texas by Age Group, Educational Attainment, Enrollment, and GenAI Exposure

Notes: Figure presents outmigration rates from the state of Texas in the American Community Survey (ACS) annual microdata for those ages 18-22 or 22-27 by enrollment, educational attainment, and major exposure to GenAI automation based on Anthropic's AI automation measures as applied to college majors in the text. GenAI exposure is classified into terciles of the exposure distribution and crosswalked from 4-digit CIP codes to ACS educational program codes. The 2020 data collection and weighting was interrupted by the COVID-19 pandemic, so the weights for that survey year are experimental.

Table B1: Descriptive Statistics: New Masters Recipients Sample

	(1) Mean
Male	0.435
	(0.496)
White	0.515
	(0.500)
Hispanic	0.192
	(0.394)
Black	0.096
	(0.294)
Asian	0.099
	(0.299)
Major share of automatable tasks (Observed)	0.097
	(0.092)
Major share of automatable tasks (Theoretical)	0.499
	(0.101)
# of Qtrs Employed in First Year after Graduation	2.43
	(1.83)
Earnings in First Year after Graduation (2025$)	48,502
	(62,830)
N	158,955
Earnings, Conditional on Being Employed Every Qtr	84,233
	(44,730)
N	84,233

Notes: This table presents descriptive statistics for students from the 2016 to 2024 cohorts of May graduates from graduate programs at public four-year universities in Texas. Employment variables reflect employment in a position covered by unemployment insurance in Texas in the first four quarters after graduation. Earnings include zeros. Conditional earnings include only individuals with positive wages in all four quarters their first year after graduation and are winsorized at the 1st and 99th percentiles. Both earnings outcomes are deflated to 2025 dollars.

C Robustness to Alternate Specifications and AI Exposure Measures

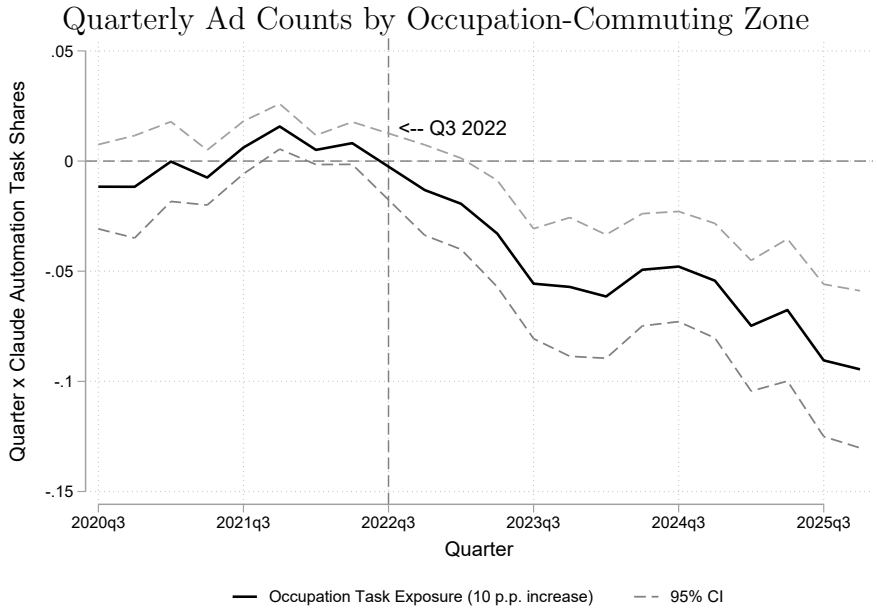
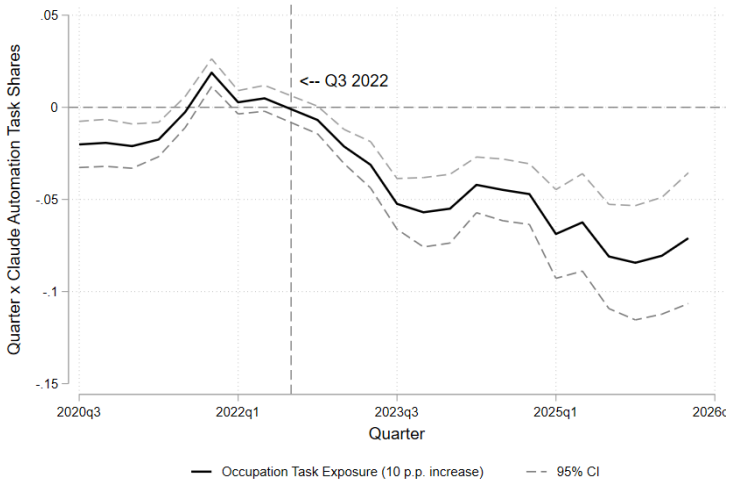
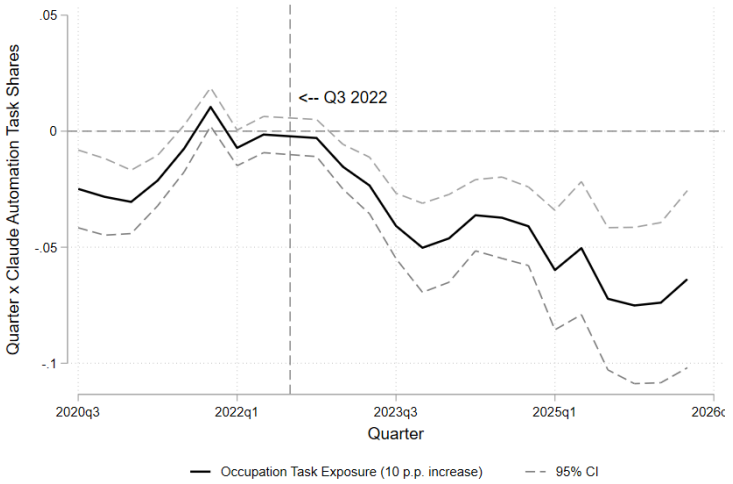


Figure C1: Exposure Results Are Robust to Spatial Controls

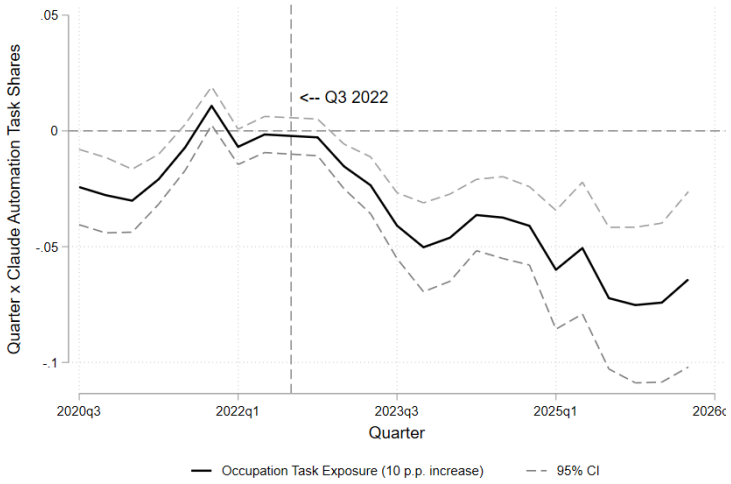
Notes: This figure presents estimates of the relationship between occupation-level exposure to GenAI automation and the rate of job postings in the Lightcast data at the occupation-commuting zone level. Estimates reflect the difference in job postings associated with an occupation having a 10 p.p. greater share of automatable tasks relative to Quarter 3 of 2022. The estimate includes fixed effects for occupation and commuting zone-by-quarter. Standard errors are clustered at the occupation level. Occupations are defined by the 2019 O*NET classifications.



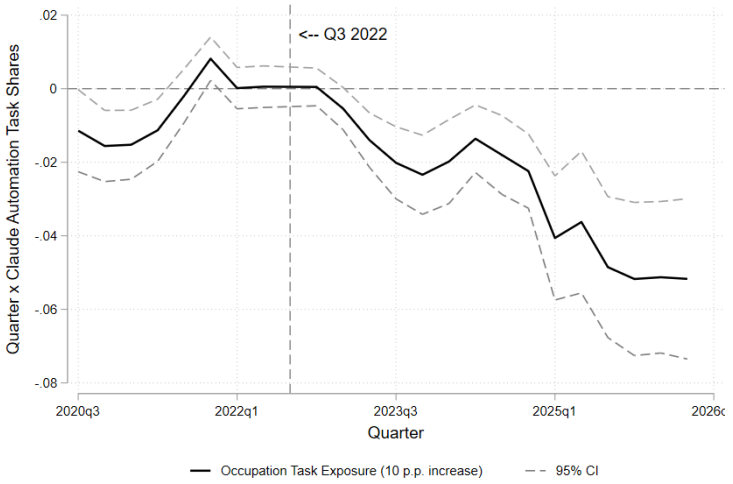
(a) Controls for Monetary Policy Sensitivity



(b) Controls for Remote/Hybrid Work Shares



(c) Controls for Both



(d) Both, Remote Work Defined in Hansen et al. (2026)

Figure C2: Exposure Results Are Robust to Controls for Monetary Policy Sensitivity and Remote/Hybrid Work

Notes: This figure presents the same model as Figure 2 but with controls for occupational sensitivity to monetary policy (Panel a), occupation-industry shares of remote/hybrid work (Panel b), or both (Panel c). Panel d re-estimates panel c, but replaces the share of the occupation-industry with remote/hybrid work with the remote work indicators calculated in Hansen et al. (2026) matched to the job ads in our sample for January 2022-October 2022. These indicators agree with our indicators 96% of the time. Estimates scale to a 10 p.p. greater share of automatable tasks and are relative to Quarter 3 of 2022. Model includes fixed effects for occupation-industry cell and industry-by-quarter. Standard errors are clustered at the occupation level. Occupations = 2019 O*NET classifications, and industry = 4-digit NAICS classification.

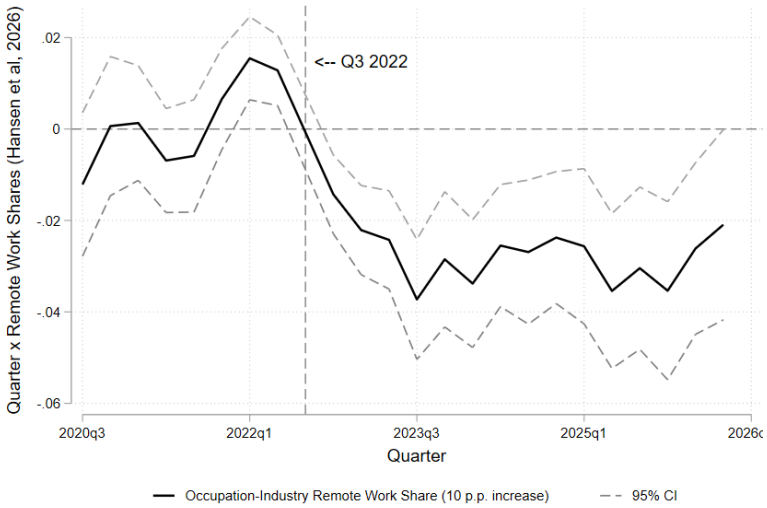
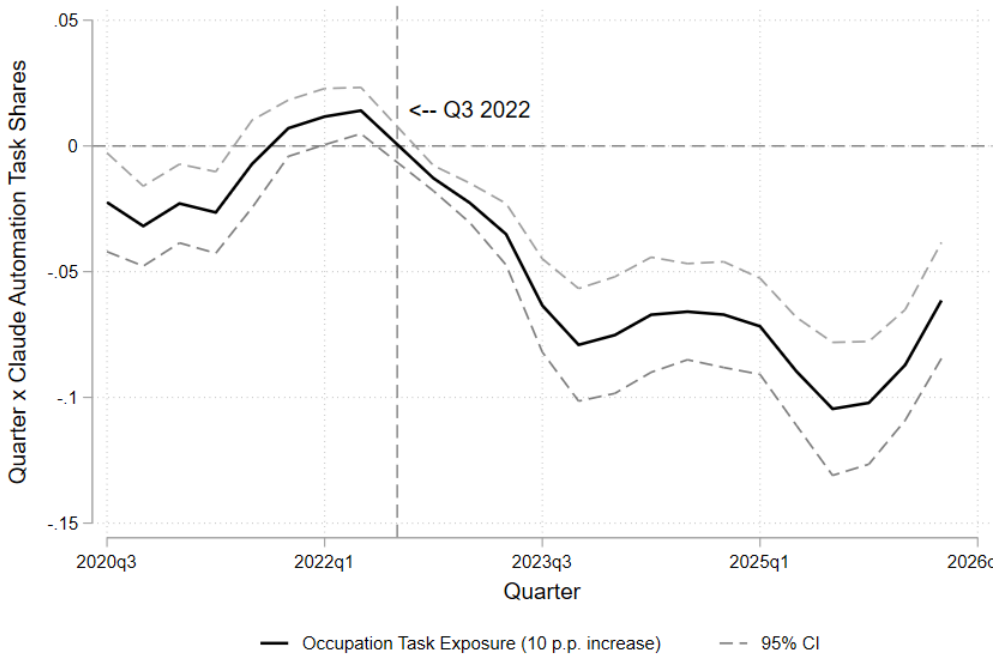
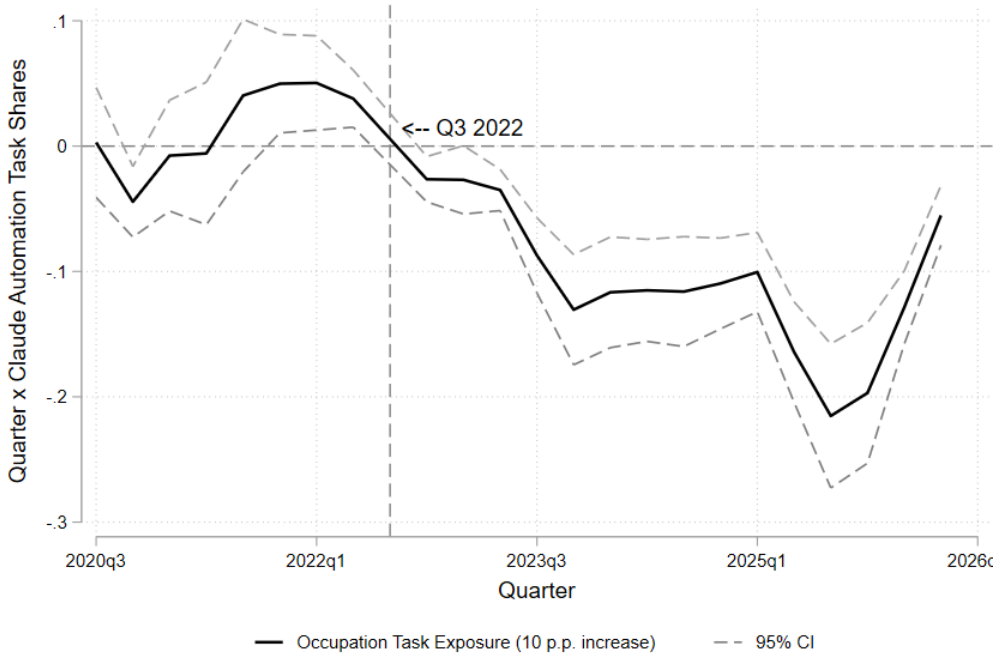


Figure C3: Postings Fell for Remote/Hybrid Occupations Independent of GenAI Effects

Notes: This figure presents estimates of the relationship between occupation-industry level job ads over time and occupation-industry remote work shares calculated in Hansen et al. (2026) matched to the job ads in our sample for January 2022-October 2022. These correspond to the remote work coefficients in the same model as Panel D of Figure C2. Estimates capture the difference in job postings associated with an occupation-industry having a 10 p.p. greater share of 2022 pre-ChatGPT remote work shares in their job ads interacted with quarter fixed effects relative to Quarter 3 of 2022. The estimates control for 1) occupation-level Claude automation task shares x quarter fixed effects, and 2) occupation-level sensitivity to monetary policy x quarter fixed effects. The model includes fixed effects for occupation-industry cell and industry-by-quarter. Standard errors are clustered at the occupation level. Occupations are defined by the 2019 O*NET classifications, and industry is defined by the four-digit NAICS classification.



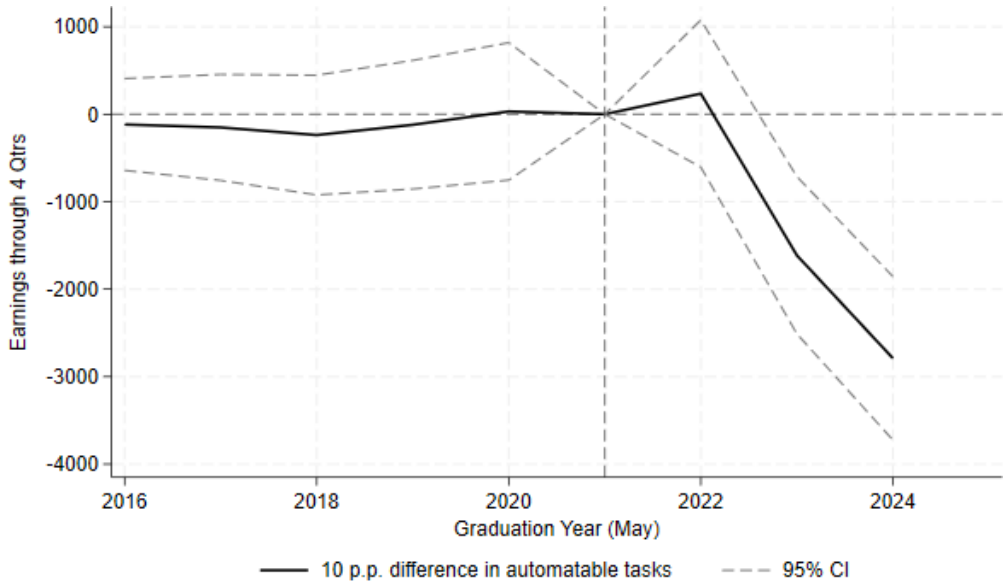
(a) Quarterly Ad Counts by Occupation-Industry (Unweighted)



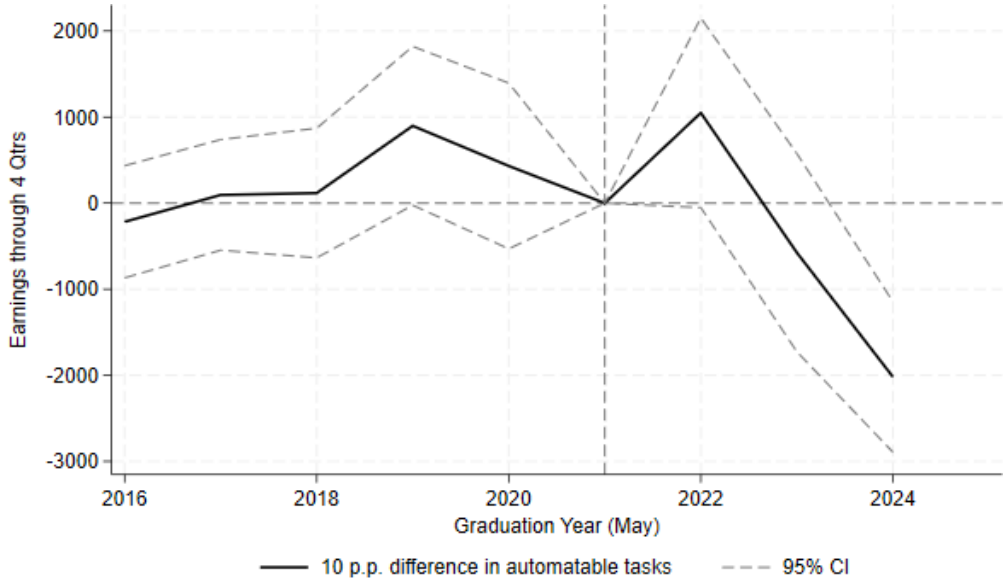
(b) Quarterly Ad Counts by Occupation-Industry (Weighted)

Figure C4: The Job Posting Results Hold Nationally

Notes: This figure presents *national* estimates of the relationship between occupation-level exposure to GenAI automation and the rate of job postings in the Lightcast data at the occupation-industry levels, where the estimates are unweighted (Panel A) or weighted (Panel B) by pre-ChatGPT total posts in the cell. Estimates reflect the difference in job postings associated with an occupation having a 10 p.p. greater share of automatable tasks relative to Quarter 3 of 2022. Both panels include fixed effects for occupation-industry cell and industry-by-quarter. Standard errors are clustered at the occupation level. Occupations are defined by the 2019 O*NET classifications, and industry is defined by the four-digit NAICS classification.



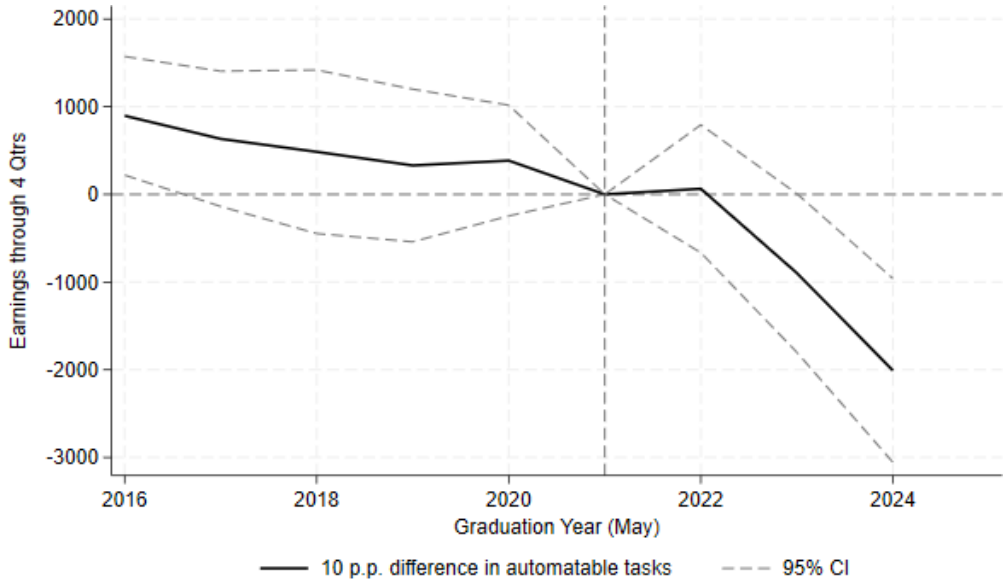
(a) Ind.-by-Year Fixed Effects



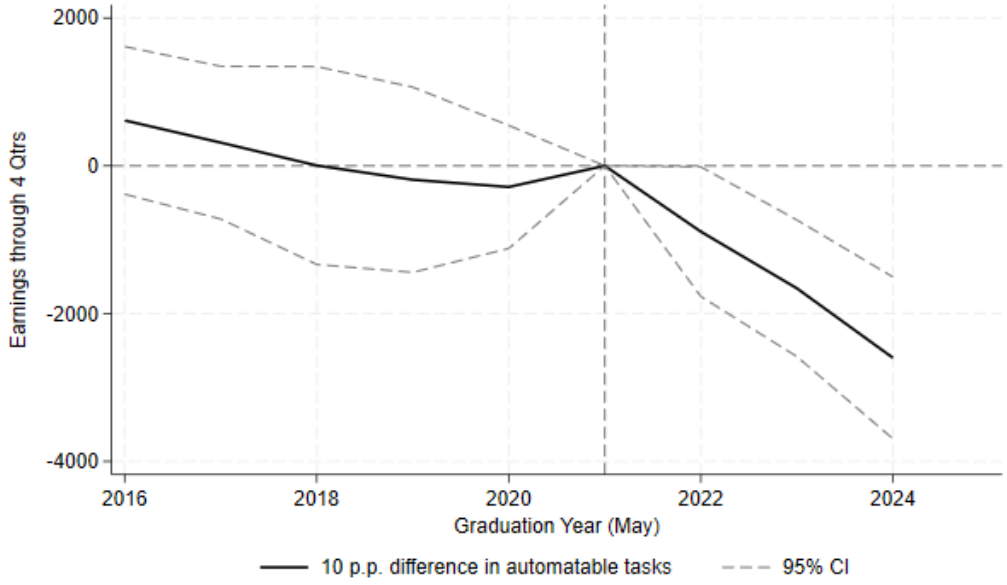
(b) Remote Work Controls

Figure C5: Estimated Declines in Earnings are Robust to Controlling for Time-Varying Industry Shocks & Remote Work Shares

Notes: This figure presents estimates of the relationship between major-level exposure to GenAI automation and earnings across cohorts of May graduates from four-year universities. The outcome is the natural log of earnings for graduates that worked all four quarters of their first year after graduation. Estimates reflect the difference in earnings associated with a graduate's major having a 10 p.p. greater share of automatable tasks, relative to this difference for May 2021 graduates. The specification in (a) controls for major (4-digit CIP code), school-by-year and industry-by-year (2-digit NAICS) fixed effects, and student demographics. The specification in (b) controls for major (4-digit CIP code), school-by-year fixed effects, pre-period remote work shares interacted with year indicators, and student demographics. Standard errors are clustered by major.



(a) GPT-Beta Exposure Measure



(b) Human-Beta Exposure Measure

Figure C6: Estimated Declines in Earnings are Robust to Alternate AI Exposure Measures

Notes: This figure presents estimates of the relationship between major-level exposure to GenAI automation and earnings across cohorts of May graduates from four-year universities. The outcome is the natural log of earnings for graduates that worked all four quarters of their first year after graduation. Estimates reflect the difference in earnings associated with a graduate's major having a 10 p.p. greater share of automatable tasks, relative to this difference for May 2021 graduates. Panel (a) adopts the GPT-Beta exposure measure and panel (b) adopts the Human-Beta measure from Eloundou et al. (2023). The specification controls for major (4-digit CIP code), school-by-year, and industry-by-year (2-digit NAICS) fixed effects, and student demographics. Standard errors are clustered by major.

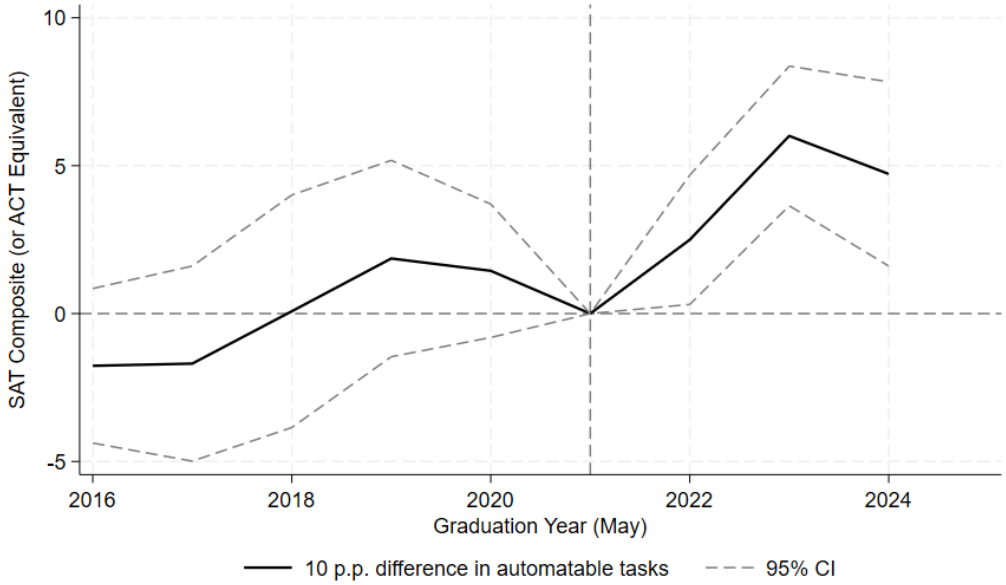
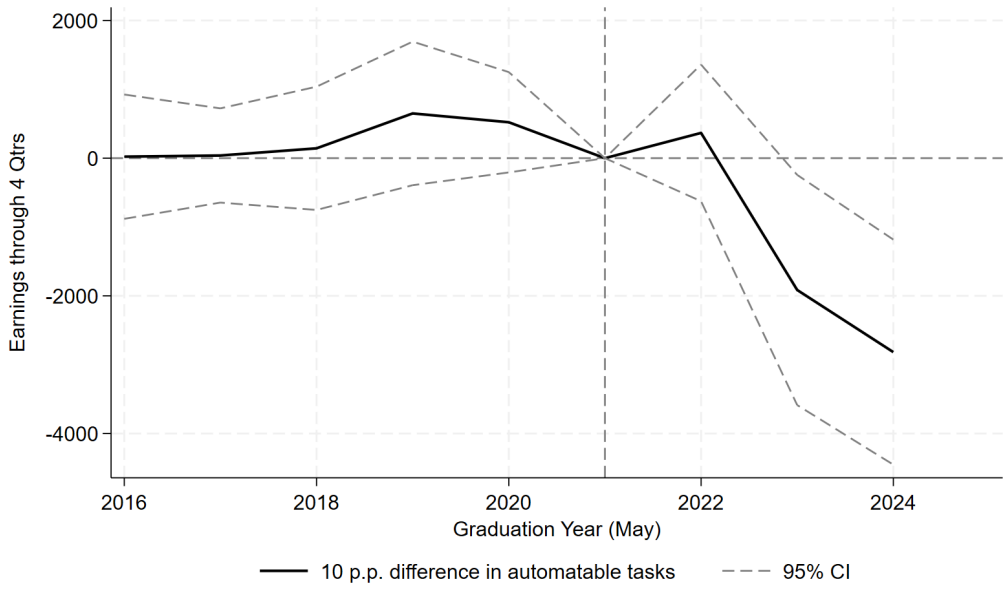
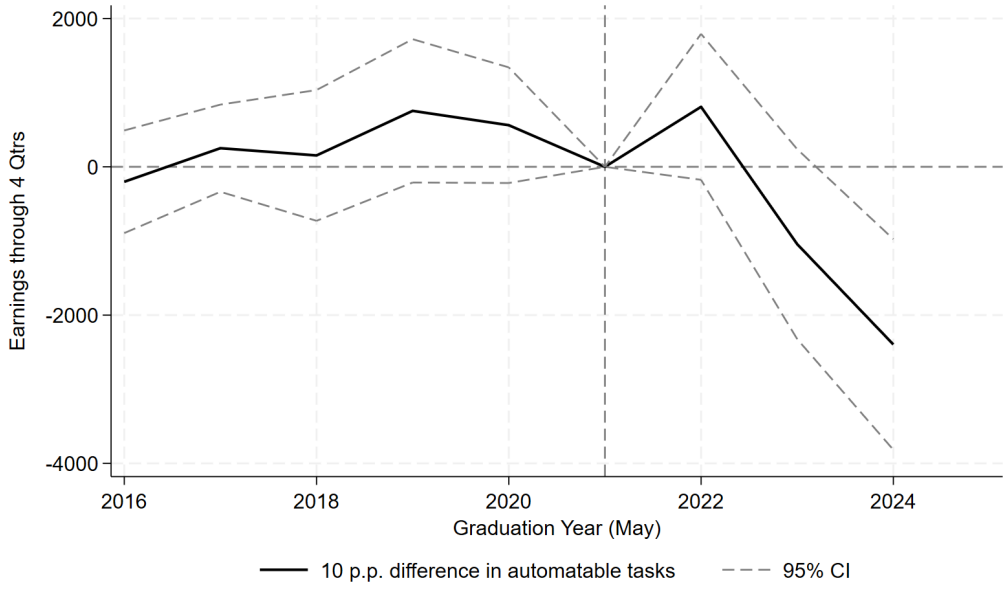


Figure C7: Recent Grads From More-Exposed Majors were Positively Selected on Observable Ability

Notes: This figure presents estimates of the relationship between major-level exposure to GenAI automation and composite SAT scores (or converted ACT equivalents) across cohorts of May graduates from four-year universities. The outcome is the composite SAT score (out of 1600). Estimates reflect the difference in outcomes associated with a graduate's major having a 10 p.p. greater share of automatable tasks, relative to this difference for May 2021 graduates. The specification controls for major (4-digit CIP code) and school-by-year fixed effects, and student demographics. Standard errors are clustered by major.



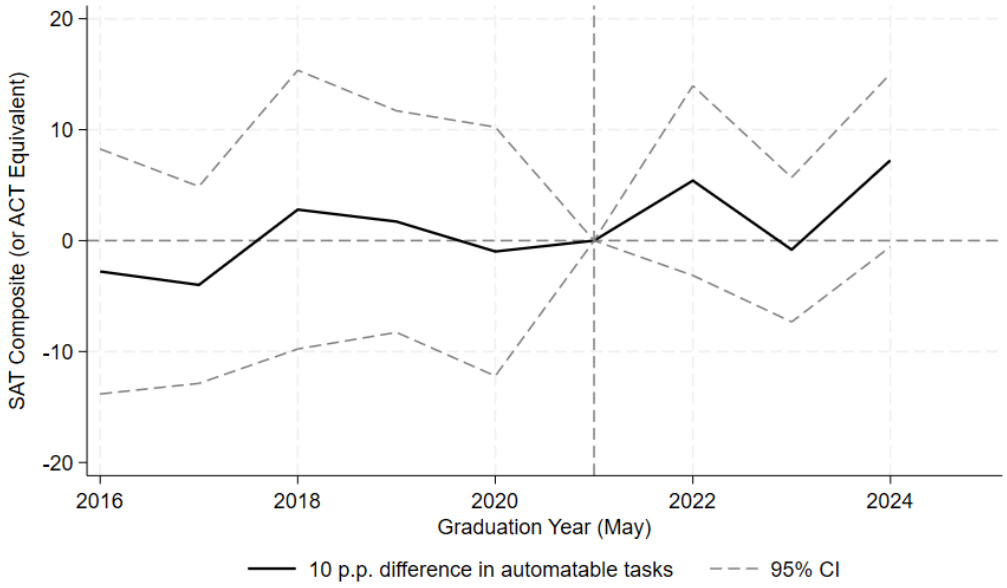
(a) Earnings



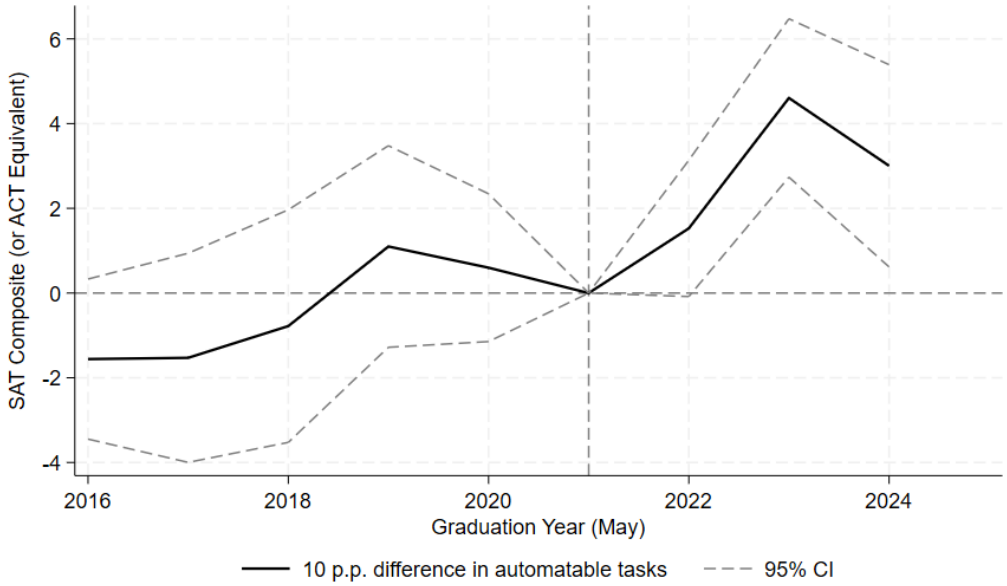
(b) Earnings, Conditional on Working Every Qtr

Figure C8: Estimated Declines in Earnings are Robust to Dropping Graduate School Enrollees

Notes: This figure presents estimates of the relationship between major-level exposure to GenAI automation and earnings across cohorts of May graduates from four-year universities. The outcome is the natural log of earnings for graduates that worked all four quarters of their first year after graduation. Estimates reflect the difference in earnings associated with a graduate's major having a 10 p.p. greater share of automatable tasks, relative to this difference for May 2021 graduates. The specification controls for major (4-digit CIP code), school-by-year fixed effects, and student demographics. Students that enroll in graduate school within a year of completing their undergraduate degree are dropped from the sample. Standard errors are clustered by major.



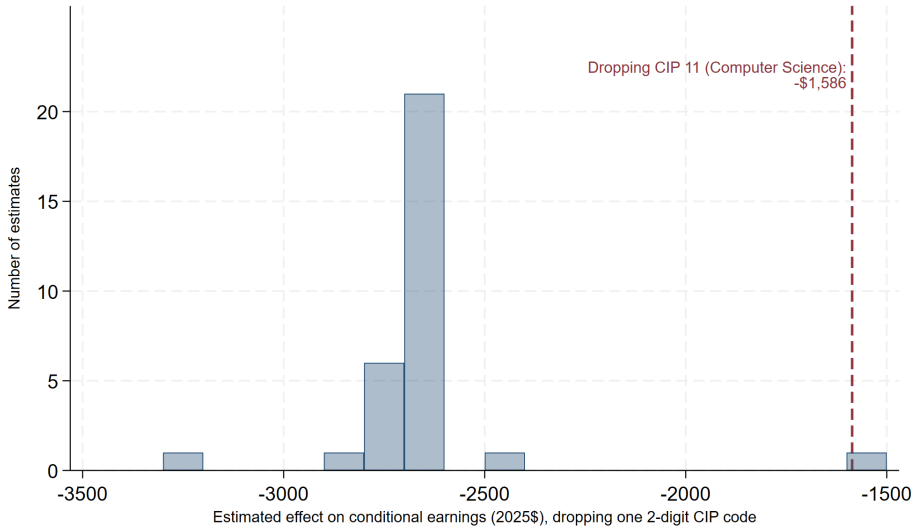
(a) SAT Score — Graduate School Enrollees



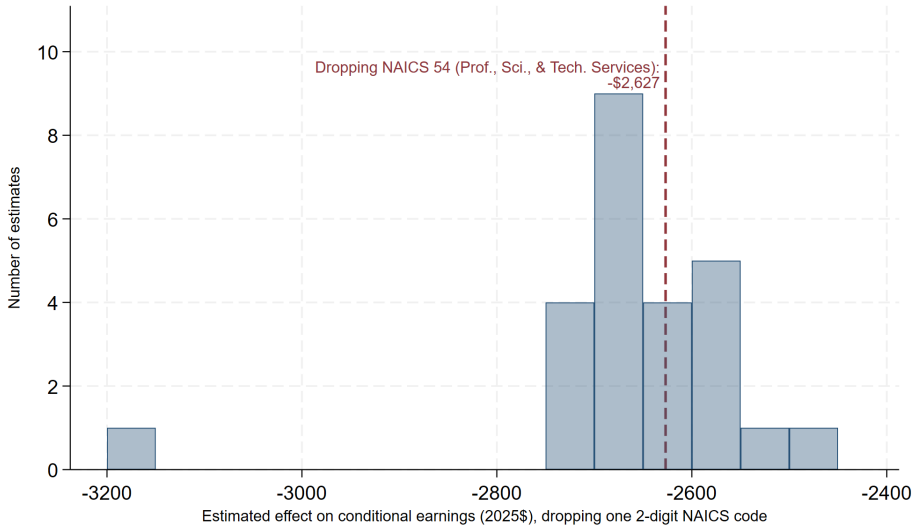
(b) SAT Score — Did Not Enroll in Graduate School

Figure C9: No Evidence of Higher-Ability Graduates from More-Exposed Majors Sorting into Graduate School

Notes: This figure presents estimates of the relationship between major-level exposure to GenAI automation and composite SAT scores (or converted ACT equivalents) across cohorts of May graduates from four-year universities. The outcome is the composite SAT score (out of 1600). Estimates reflect the difference in outcomes associated with a graduate's major having a 10 p.p. greater share of automatable tasks, relative to this difference for May 2021 graduates. The specification controls for major (4-digit CIP code) and school-by-year fixed effects, and student demographics. The sample is limited to only students that enroll in graduate school within a year of completing undergraduate studies in panel (a); in panel (b) graduate school enrollees are dropped. Standard errors are clustered by undergraduate major.



(a) Dropping One 2-Digit CIP Code



(b) Dropping One 2-Digit NAICS Code

Figure C10: Distribution of Estimated Effects on Conditional Earnings, Dropping One Field or Industry at a Time

Notes: This figure presents the distribution of estimates of π^{ES}_{2024} from Equation 1, re-estimated after dropping one group of graduates at a time. The outcome is earnings for graduates that worked all four quarters of their first year after graduation (2025$). Estimates reflect the difference in earnings for May 2024 graduates associated with a graduate's major having a 10 p.p. greater share of automatable tasks, relative to this difference for May 2021 graduates. Panel (a) drops graduates from one 2-digit CIP code per estimate; the dashed line marks the estimate that drops Computer and Information Sciences (CIP 11). Panel (b) drops graduates employed in one 2-digit NAICS industry per estimate; the dashed line marks the estimate that drops Professional, Scientific, and Technical Services (NAICS 54).

Table C1: Estimated Effects of GenAI Exposure on First-Year Labor Market Outcomes Are Robust Across Specifications: Human-Annotator β Exposure

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
A. First-Year Earnings														
$MajExp_m$	-2,097***	-2,779***	-1,278*	-1,581**	-2,633**	-3,126***	-1,431*	-1,784**	–	–	–	–	–	–
	(480)	(841)	(655)	(624)	(1,084)	(925)	(736)	(704)						
B. Number of Quarters Employed in First Year														
$MajExp_m$	-0.055***	-0.071***	-0.053***	-0.062***	-0.083***	-0.083***	-0.062***	-0.072***	–	–	–	–	–	–
	(0.016)	(0.021)	(0.016)	(0.020)	(0.021)	(0.021)	(0.018)	(0.021)						
C. First-Year Earnings, Conditional on Working All Four Quarters														
$MajExp_m$	-1,991**	-2,233***	-1,150*	-1,273**	-2,070**	-2,302***	-1,080	-1,222*	-2,453***	-2,599***	-1,966***	-2,578***	-2,723***	-1,917***
	(880)	(701)	(656)	(585)	(944)	(797)	(738)	(678)	(657)	(539)	(446)	(697)	(633)	(517)
Major FEs	X	X	X	X	X	X	X	X	X	X	X	X	X	X
School-by-Year FEs	X	X	X	X	X	X	X	X	X	X	X	X	X	X
SAT Scores		X		X		X		X		X	X		X	X
Industry-by-Year FEs									X	X	X	X	X	X
Major Remote Shares			X	X			X	X			X			X
No Graduate Students					X	X	X	X				X	X	X

Standard errors clustered by major in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table presents estimated effects of GenAI exposure on first-year labor market outcomes for May 2024 graduates from Texas public four-year universities, where exposure is measured by the human-annotator β exposure measure of Eloundou et al. (2023). Each cell reports the coefficient on the 2024 cohort event study indicator from a separate regression; estimates reflect the difference in outcomes associated with a graduate's major having a 10 p.p. greater exposure to GenAI, relative to this difference for May 2021 graduates. The outcome is total first-year earnings (2025$) in Panel A, the number of quarters with non-zero earnings in the first year in Panel B, and first-year earnings conditional on working all four quarters in Panel C. Columns (9)–(14) are estimated only for the conditional-earnings outcome, so the corresponding cells in Panels A and B are left blank. Rows at the bottom of the table indicate the fixed effects, controls, and sample restrictions included in each specification; all specifications control for student demographics. Standard errors are clustered by major.

Table C2: Estimated Effects of GenAI Exposure on First-Year Labor Market Outcomes Are Robust Across Specifications: GPT-4-Annotator β Exposure

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
A. First-Year Earnings														
$MajExp_m$	-1,464***	-2,031**	-613	-966	-1,741*	-2,184**	-521	-863	–	–	–	–	–	–
	(437)	(811)	(754)	(682)	(1,036)	(884)	(830)	(754)						
B. Number of Quarters Employed in First Year														
$MajExp_m$	-0.041***	-0.058***	-0.037**	-0.049***	-0.054***	-0.064***	-0.039**	-0.051***	–	–	–	–	–	–
	(0.014)	(0.019)	(0.015)	(0.017)	(0.019)	(0.019)	(0.017)	(0.019)						
C. First-Year Earnings, Conditional on Working All Four Quarters														
$MajExp_m$	-1,282	-1,457**	-399	-455	-1,240	-1,371*	-126	-134	-1,931***	-2,008***	-1,356***	-1,930***	-1,983***	-1,110**
	(829)	(652)	(692)	(569)	(884)	(758)	(758)	(629)	(635)	(535)	(417)	(682)	(614)	(484)
Major FEs	X	X	X	X	X	X	X	X	X	X	X	X	X	X
School-by-Year FEs	X	X	X	X	X	X	X	X	X	X	X	X	X	X
SAT Scores		X		X		X		X		X	X		X	X
Industry-by-Year FEs									X	X	X	X	X	X
Major Remote Shares			X	X			X	X			X			X
No Graduate Students					X	X	X	X				X	X	X

Standard errors clustered by major in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table presents estimated effects of GenAI exposure on first-year labor market outcomes for May 2024 graduates from Texas public four-year universities, where exposure is measured by the GPT-4-annotator β exposure measure of Eloundou et al. (2023). Each cell reports the coefficient on the 2024 cohort event study indicator from a separate regression; estimates reflect the difference in outcomes associated with a graduate's major having a 10 p.p. greater exposure to GenAI, relative to this difference for May 2021 graduates. The outcome is total first-year earnings (2025$) in Panel A, the number of quarters with non-zero earnings in the first year in Panel B, and first-year earnings conditional on working all four quarters in Panel C. Columns (9)–(14) are estimated only for the conditional-earnings outcome, so the corresponding cells in Panels A and B are left blank. Rows at the bottom of the table indicate the fixed effects, controls, and sample restrictions included in each specification; all specifications control for student demographics. Standard errors are clustered by major.

C.1 Robustness to Specifying Alternate Functional Forms of Dose-Response Curve

Our primary estimates reflect continuous difference-in-differences specifications with treatment (GenAI exposure) modeled linearly. Under a "strong" parallel trends assumption, these estimates reflect a weighted average of conditional causal response terms with positive weights across the treatment distribution that integrate to 1 (Callaway et al., 2025). Still, the underlying dose-response function may not be linear. Therefore, we test the robustness of our estimates to specifying higher-order polynomials of major-level GenAI exposure.

In particular, we estimate the following specification:

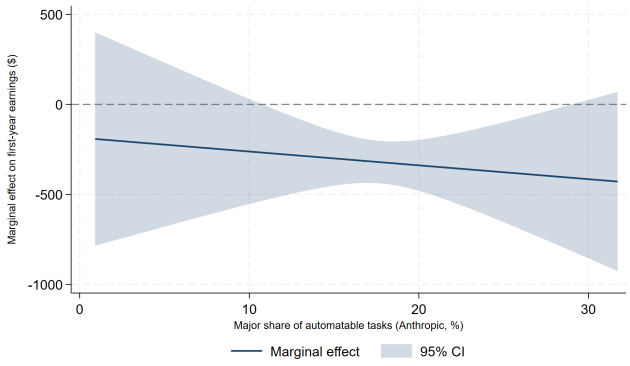
$$y_{imct} = \sum_{k=1}^{k^{max}} \gamma_k \, Post_t \, MajExp_m^k + \alpha_m + \alpha_{ct} + \Phi S_i + \epsilon_{imct}$$

for $k^{max} \in \{2, 3, 4\}$, where $Post_t$ equals 1 for the 2024 graduating cohort and 0 for the 2021 graduating cohort.

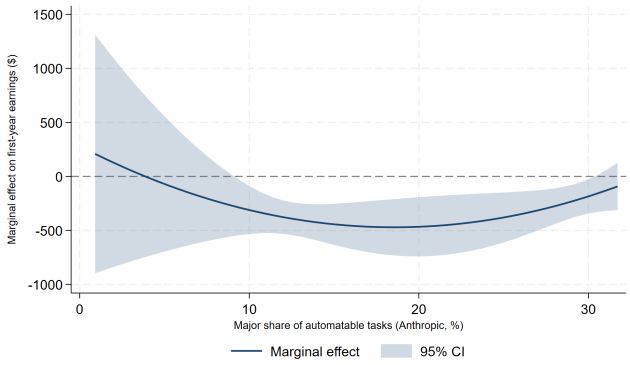
We present estimated marginal effects of major-level AI exposure ($\frac{\partial y_{imct}}{\partial MajExp_m}$) for each polynomial specification with unconditional first-year earnings (Figure C11), number of quarters worked (Figure C12), conditional first-year earnings (Figure C13), and graduate school enrollment (Figure C14) as outcome variables. Our main conclusions are robust to allowing for non-linear dose-response functions. Across all three polynomial specifications, the estimated marginal effect of exposure on unconditional first-year earnings is negative over nearly the entire range of major-level exposure, generally between \$200 and \$500 per percentage point (Figure C11). These magnitudes bracket our linear estimate of a decline of over \$3,000 per 10 percentage points for 2024 graduates (Figure 4), or roughly \$300 per percentage point. The cubic and quartic specifications suggest some curvature: marginal effects of increased exposure are largest in magnitude among majors with automatable task shares between roughly 10 and 20 percent and attenuate toward zero at the tails of the exposure distribution, where estimates are much less precise.

The decomposition into extensive and intensive margins is similarly stable. Marginal effects on the number of quarters worked range from roughly -0.007 to -0.013 per percentage point across most of the exposure distribution (Figure C12), in line with our linear estimate of 0.08 fewer quarters per 10 percentage points, or 0.008 per percentage point (Figure 5a). Marginal effects on earnings conditional on working all four quarters range from roughly -\$200 to -\$400 per percentage point (Figure C13), comparable to our linear estimate of \$2,427 per 10 percentage points, or roughly \$243 per percentage point (Figure 5b). Finally, the estimated marginal effects on graduate school enrollment are positive over most of the

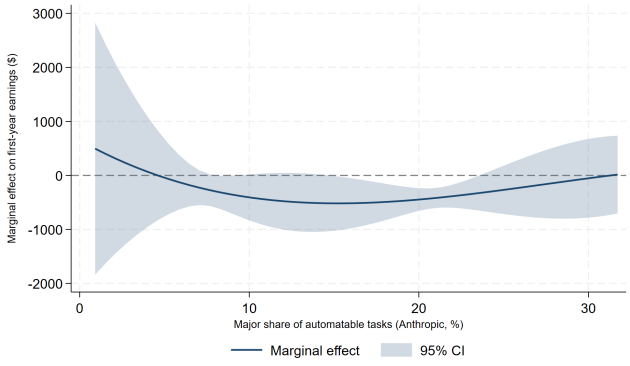
exposure distribution, reaching approximately 0.15 percentage points per percentage point of exposure among majors with automatable task shares near 20 percent (Figure C14), consistent with our linear estimate of a 0.9 percentage-point increase per 10 percentage points, or 0.09 percentage points per percentage point (Figure 7). These estimates are less precise than those for labor market outcomes, and confidence intervals for the cubic and quartic specifications include zero across much of the exposure distribution.



(a) Quadratic



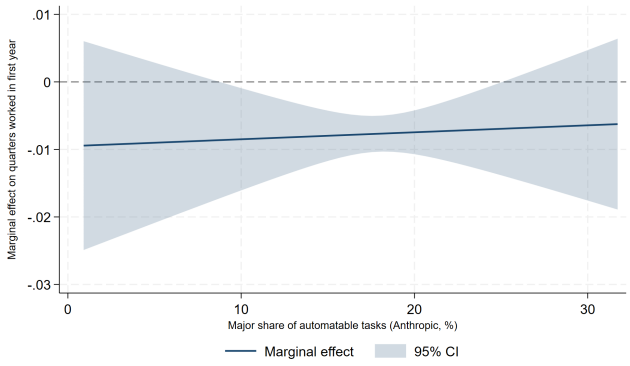
(b) Cubic



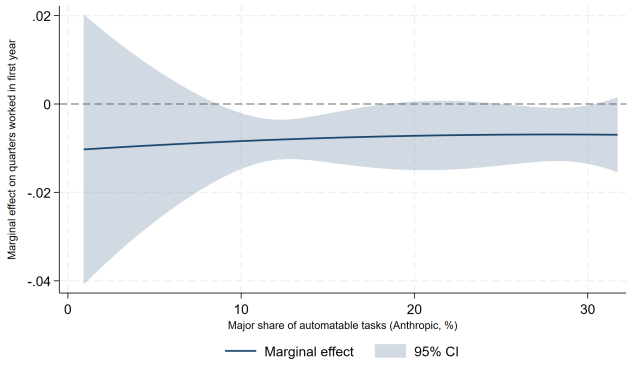
(c) Quartic

Figure C11: Marginal Effects of Major-Level GenAI Exposure on Unconditional First-Year Earnings, Polynomial Specifications

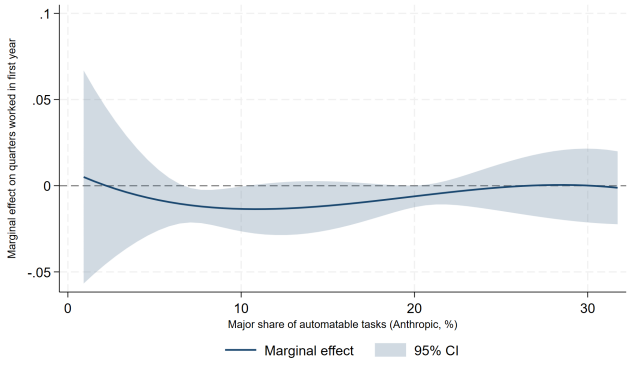
Notes: This figure presents estimated marginal effects of major-level GenAI exposure ($\frac{\partial y_{imct}}{\partial MajExp_m}$), evaluated across the distribution of major-level exposure, from specifications that model exposure as a quadratic (Panel a), cubic (Panel b), and quartic (Panel c) polynomial. The outcome is total earnings in the first year after graduation. Shaded regions denote 95% confidence intervals. The specification controls for major (4-digit CIP code) and school-by-year fixed effects, student demographics, and college entrance exam scores. Standard errors are clustered by major.



(a) Quadratic



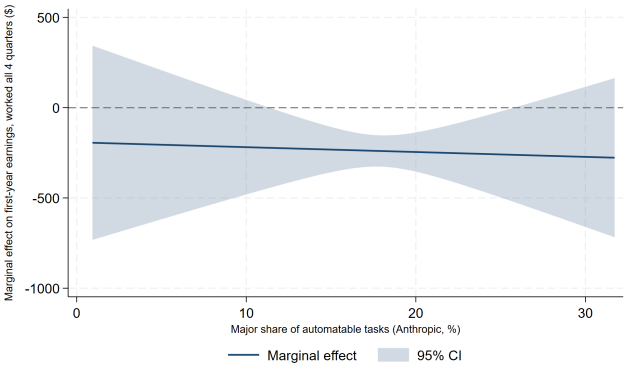
(b) Cubic



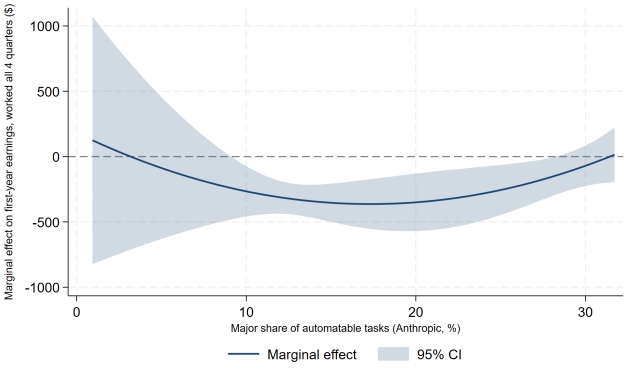
(c) Quartic

Figure C12: Marginal Effects of Major-Level GenAI Exposure on Number of Quarters Worked, Polynomial Specifications

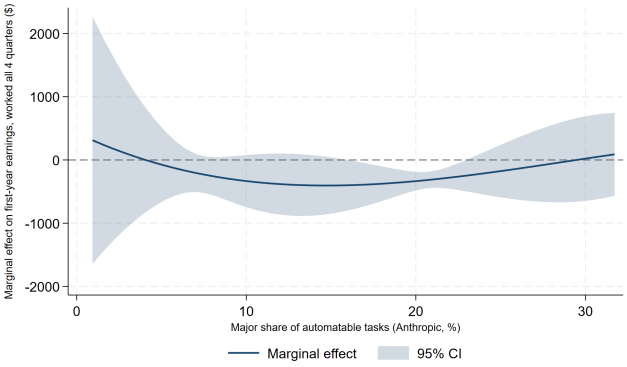
Notes: This figure presents estimated marginal effects of major-level GenAI exposure ($\frac{\partial y_{imct}}{\partial MajExp_m}$), evaluated across the distribution of major-level exposure, from specifications that model exposure as a quadratic (Panel a), cubic (Panel b), and quartic (Panel c) polynomial. The outcome is the number of quarters with positive earnings in the first year after graduation. Shaded regions denote 95% confidence intervals. The specification controls for major (4-digit CIP code) and school-by-year fixed effects, student demographics, and college entrance exam scores. Standard errors are clustered by major.



(a) Quadratic



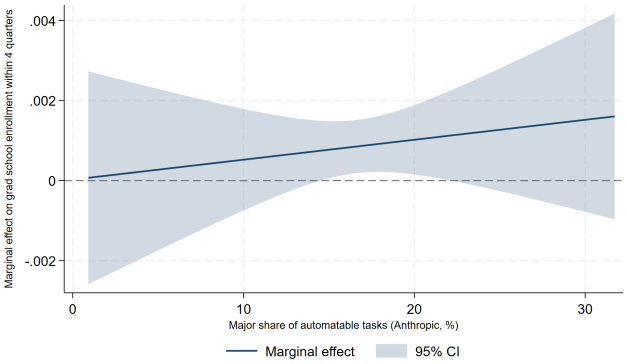
(b) Cubic



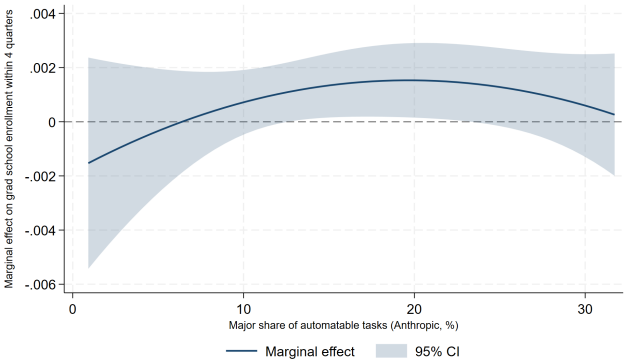
(c) Quartic

Figure C13: Marginal Effects of Major-Level GenAI Exposure on Conditional First-Year Earnings, Polynomial Specifications

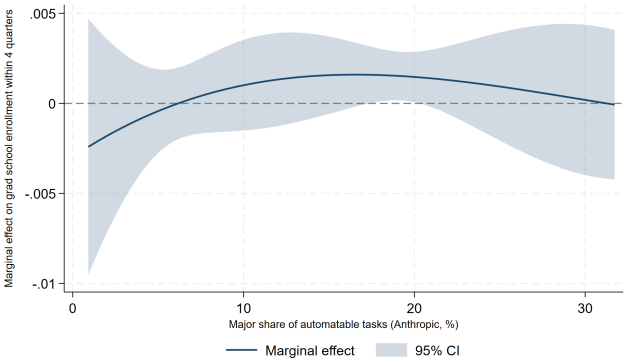
Notes: This figure presents estimated marginal effects of major-level GenAI exposure ($\frac{\partial y_{imct}}{\partial MajExp_m}$), evaluated across the distribution of major-level exposure, from specifications that model exposure as a quadratic (Panel a), cubic (Panel b), and quartic (Panel c) polynomial. The outcome is first-year earnings among graduates employed in all four quarters after graduation. Shaded regions denote 95% confidence intervals. The specification controls for major (4-digit CIP code) and school-by-year fixed effects, student demographics, and college entrance exam scores. Standard errors are clustered by major.



(a) Quadratic



(b) Cubic

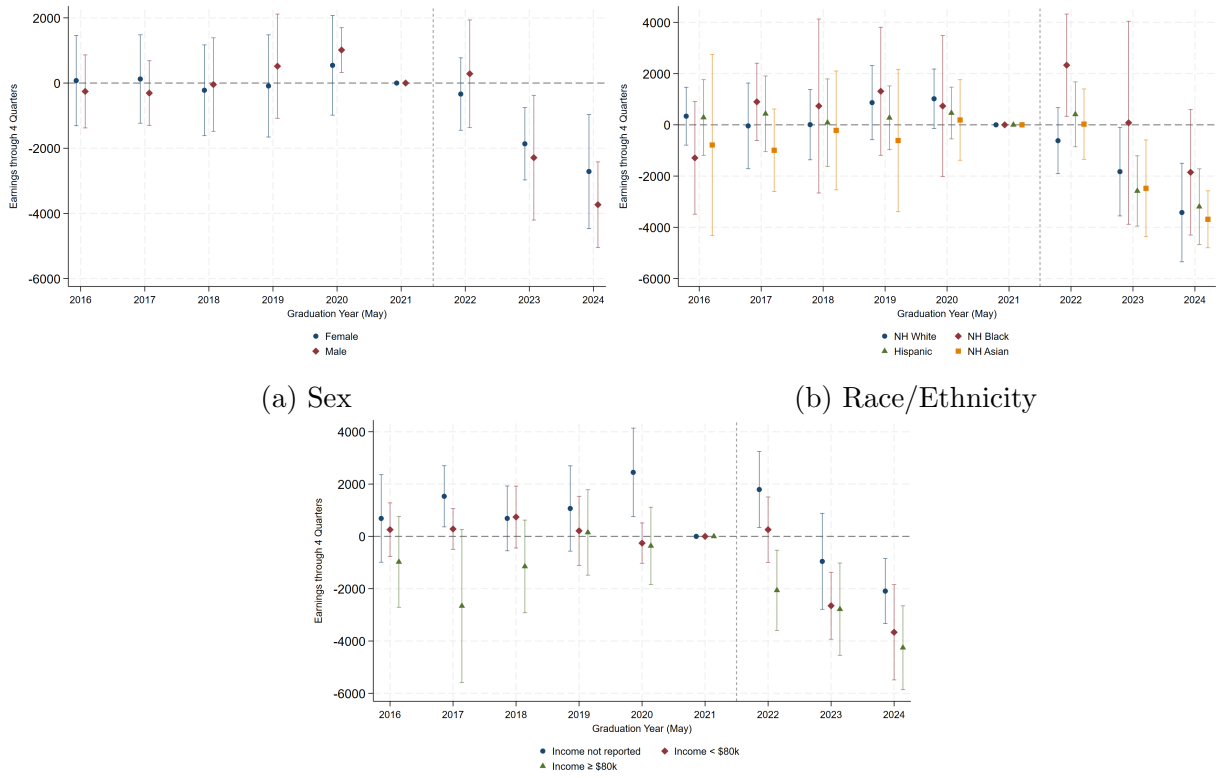


(c) Quartic

Figure C14: Marginal Effects of Major-Level GenAI Exposure on Graduate School Enrollment, Polynomial Specifications

Notes: This figure presents estimated marginal effects of major-level GenAI exposure ($\frac{\partial y_{imct}}{\partial MajExp_m}$), evaluated across the distribution of major-level exposure, from specifications that model exposure as a quadratic (Panel a), cubic (Panel b), and quartic (Panel c) polynomial. The outcome is an indicator for enrolling in graduate studies within a year of completing the undergraduate degree. Shaded regions denote 95% confidence intervals. The specification controls for major (4-digit CIP code) and school-by-year fixed effects, student demographics, and college entrance exam scores. Standard errors are clustered by major.

D Heterogeneous Effects



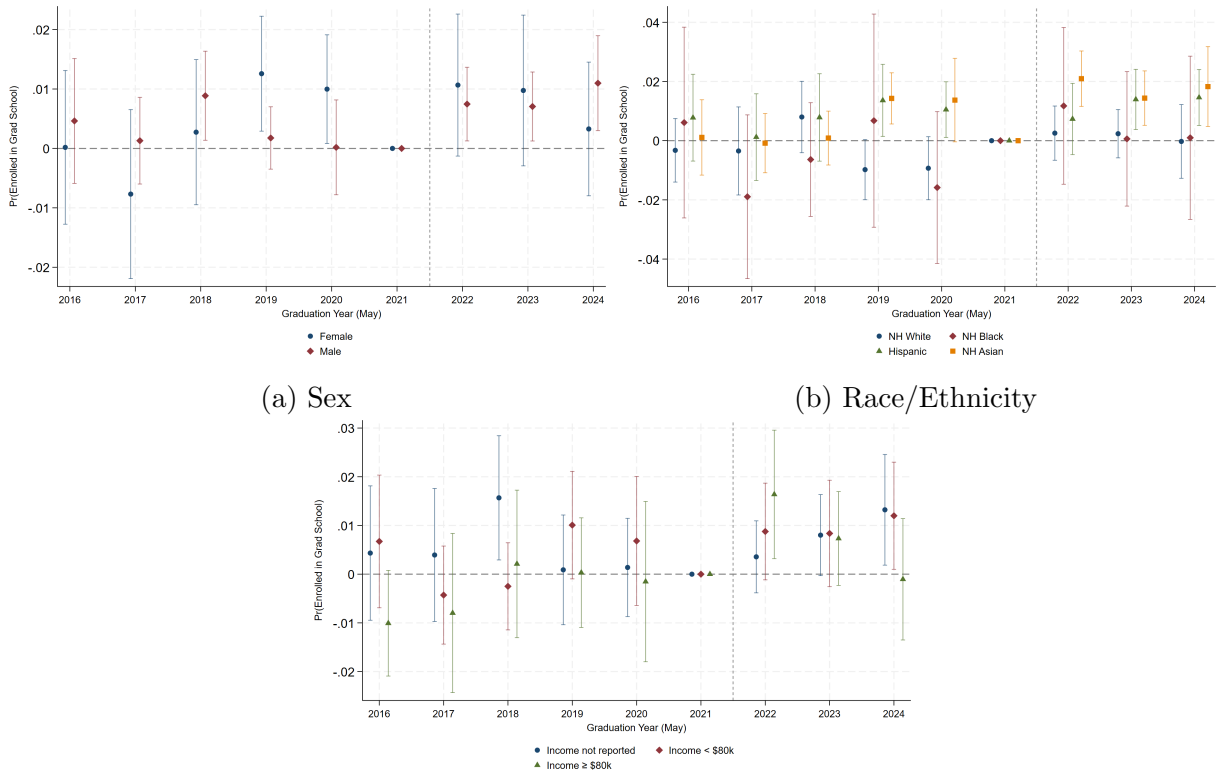
(a) Sex

(b) Race/Ethnicity

(c) Family Income

Figure D1: Earnings Declines Were Broad-Based Across Student Demographic Groups

Notes: This figure presents estimates of the relationship between major-level exposure to GenAI automation and first-year earnings (2025$, including zeroes) across cohorts of May graduates from four-year universities, separately by subgroup. Each panel comes from a single regression in which Equation 1 is fully interacted with subgroup indicators. The first series in each legend is the omitted subgroup; the other series are total effects (omitted-group coefficient plus interaction) for the remaining subgroups. Estimates reflect the difference in earnings associated with a graduate's major having a 10 p.p. greater share of automatable tasks, relative to this difference for May 2021 graduates. The specification controls for major (4-digit CIP code) and school-by-year fixed effects, student demographics, and college entrance exam scores. Standard errors are clustered by major.



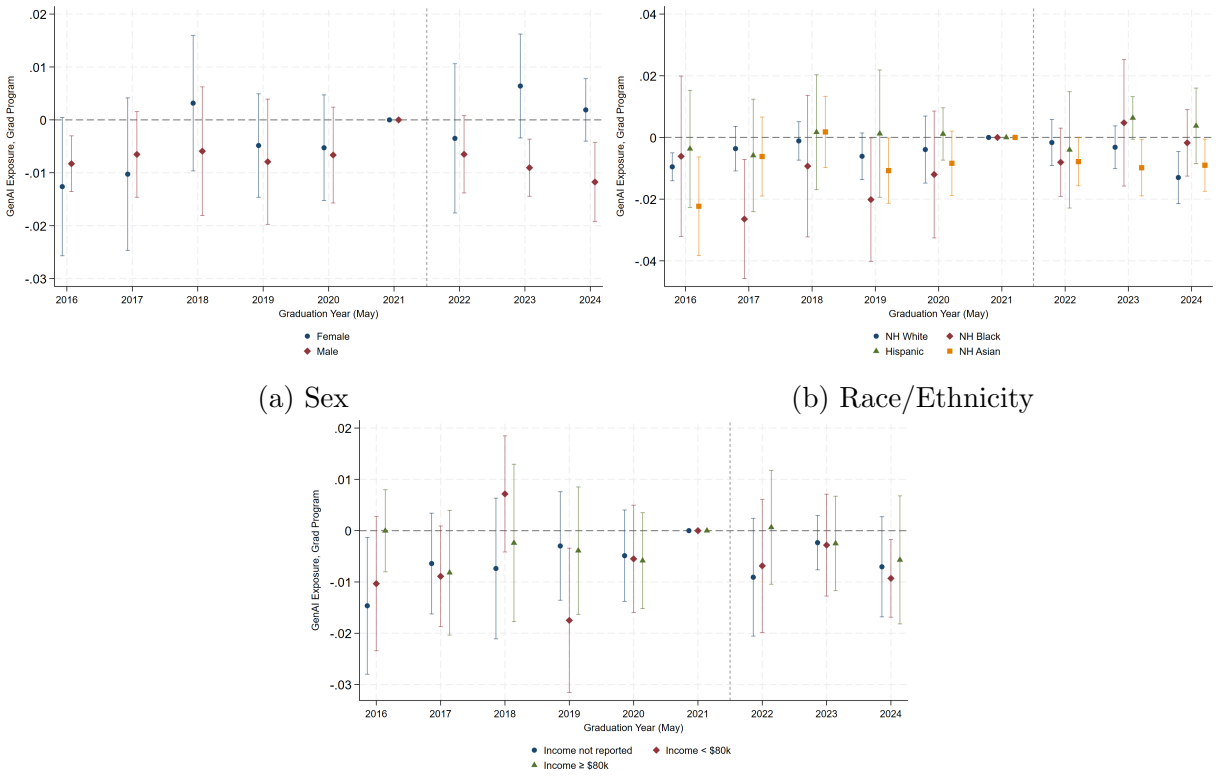
(a) Sex

(b) Race/Ethnicity

(c) Family Income

Figure D2: Graduate-School Enrollment Responses Were Concentrated Among Particular Groups

Notes: This figure presents estimates of the relationship between major-level exposure to GenAI automation and enrollment in a graduate program within a year of completing an undergraduate degree, separately by subgroup. Each panel comes from a single regression in which Equation 1 is fully interacted with subgroup indicators. The first series in each legend is the omitted subgroup; the other series are total effects (omitted-group coefficient plus interaction) for the remaining subgroups. Estimates reflect the difference in outcomes associated with a graduate's major having a 10 p.p. greater share of automatable tasks, relative to this difference for May 2021 graduates. The specification controls for major (4-digit CIP code) and school-by-year fixed effects, student demographics, and college entrance exam scores. Standard errors are clustered by major.



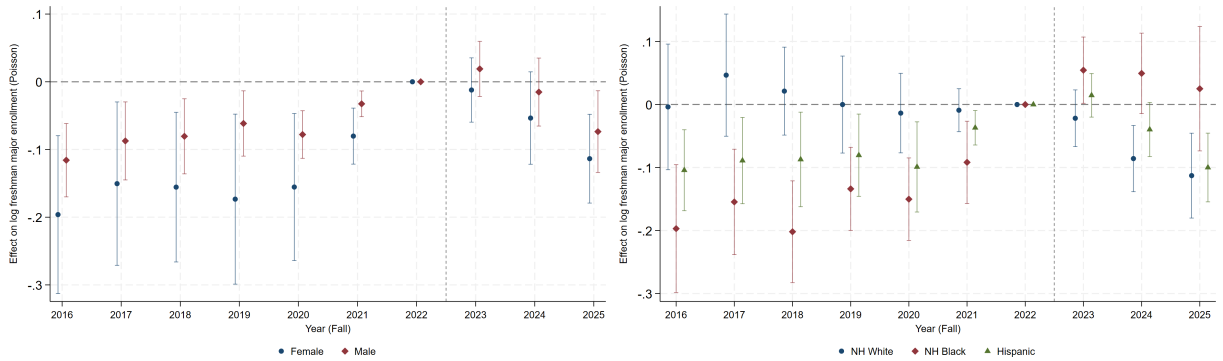
(a) Sex

(b) Race/Ethnicity

(c) Family Income

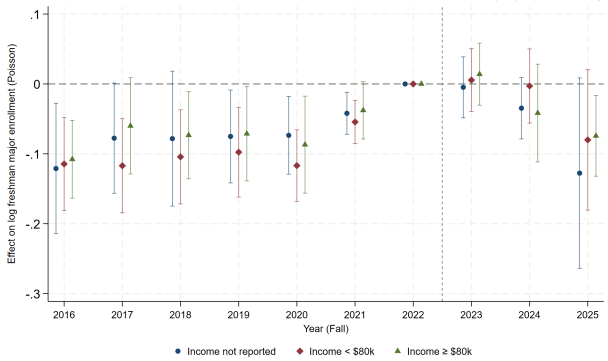
Figure D3: Graduate Program AI Exposure of Graduate-School Enrollees, by Subgroup

Notes: This figure presents estimates of the relationship between major-level exposure to GenAI automation and the GenAI exposure of the graduate program in which a graduate enrolled, separately by subgroup. The outcome is $MajExp_m$ for the *graduate* program. Each panel comes from a single regression in which Equation 1 is fully interacted with subgroup indicators. The first series in each legend is the omitted subgroup; the other series are total effects (omitted-group coefficient plus interaction) for the remaining subgroups. Estimates reflect the difference in outcomes associated with a graduate's major having a 10 p.p. greater share of automatable tasks, relative to this difference for May 2021 graduates. The specification controls for major (4-digit CIP code) and school-by-year fixed effects, student demographics, and college entrance exam scores. Standard errors are clustered by major.



(a) Sex

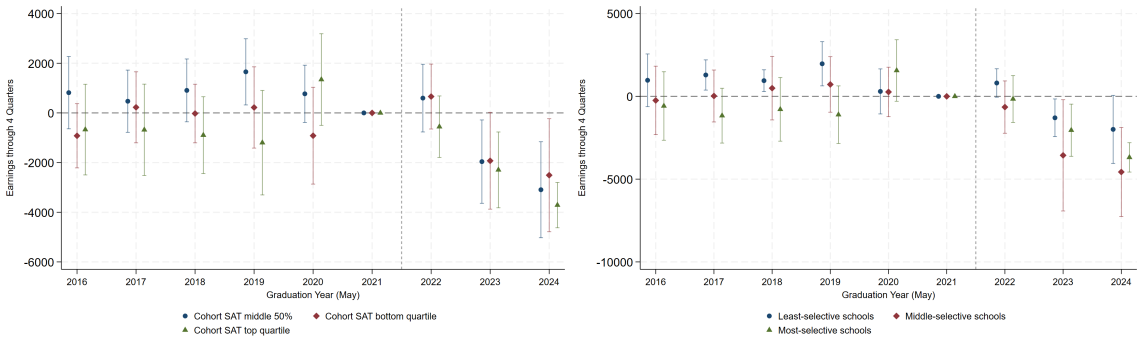
(b) Race/Ethnicity



(c) Family Income

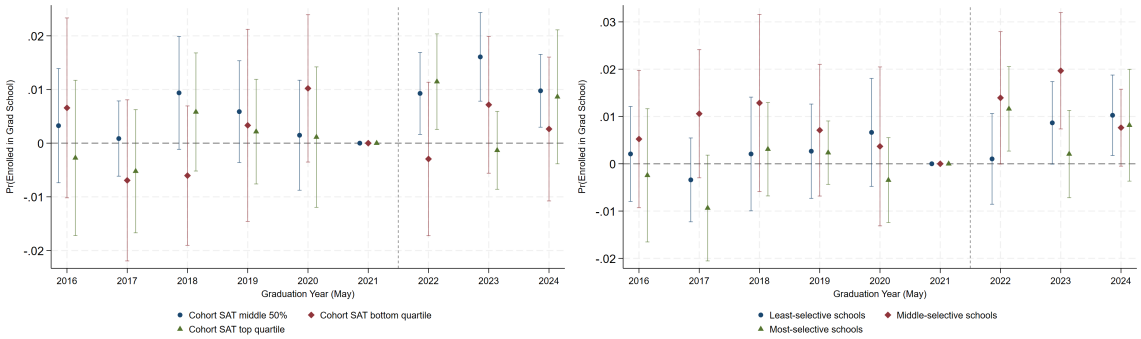
Figure D4: Freshman Enrollment Shifts Away From AI-Exposed Majors Were Broad-Based

Notes: This figure presents Poisson estimates of the relationship between major-level exposure to GenAI automation and fall Freshman enrollment counts, separately by subgroup. Each series is the total effect for a given subgroup. Coefficients are interpretable as differences in log Freshman enrollment between Fall 2022 and a given year associated with a 10 percentage-point difference in major-level GenAI exposure. The specification controls for major and year fixed effects. Standard errors are clustered by major.



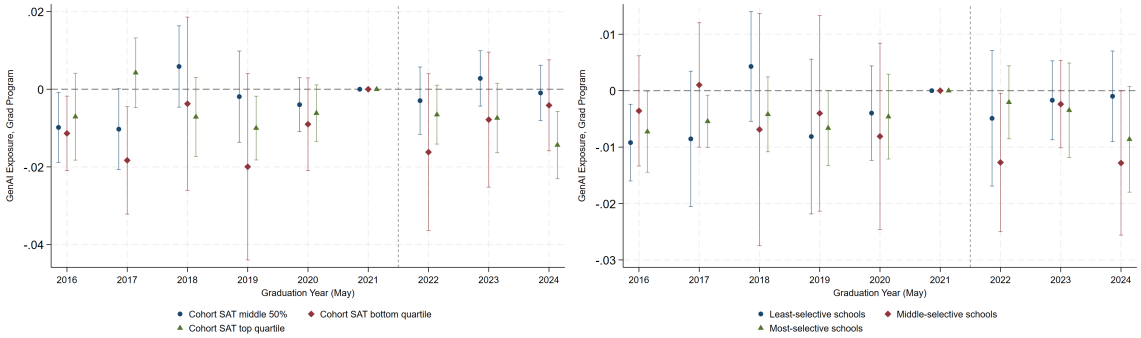
(a) Earnings, by Cohort SAT Quartile

(b) Earnings, by Institutional Selectivity



(c) Grad School Enrollment, by Cohort SAT Quartile

(d) Grad School Enrollment, by Institutional Selectivity



(e) Grad Program Exposure, by Cohort SAT Quartile

(f) Grad Program Exposure, by Institutional Selectivity

Figure D5: Responses to GenAI Exposure by Student Ability

Notes: This figure presents estimates of the relationship between major-level exposure to GenAI automation and first-year earnings (2025$, including zeroes; panels a–b), enrollment in a graduate program within a year of completing an undergraduate degree (panels c–d), and the GenAI exposure of the graduate program in which a graduate enrolled (panels e–f), separately by a graduate's quartile of the SAT distribution within their cohort and by the selectivity of their university. Each panel comes from a single regression in which Equation 1 is fully interacted with subgroup indicators. The first series in each legend is the omitted subgroup; the other series are total effects (omitted-group coefficient plus interaction) for the remaining subgroups. Estimates are relative to May 2021 graduates. The specification controls for major (4-digit CIP code) and school-by-year fixed effects, student demographics, and college entrance exam scores. Standard errors are clustered by major.

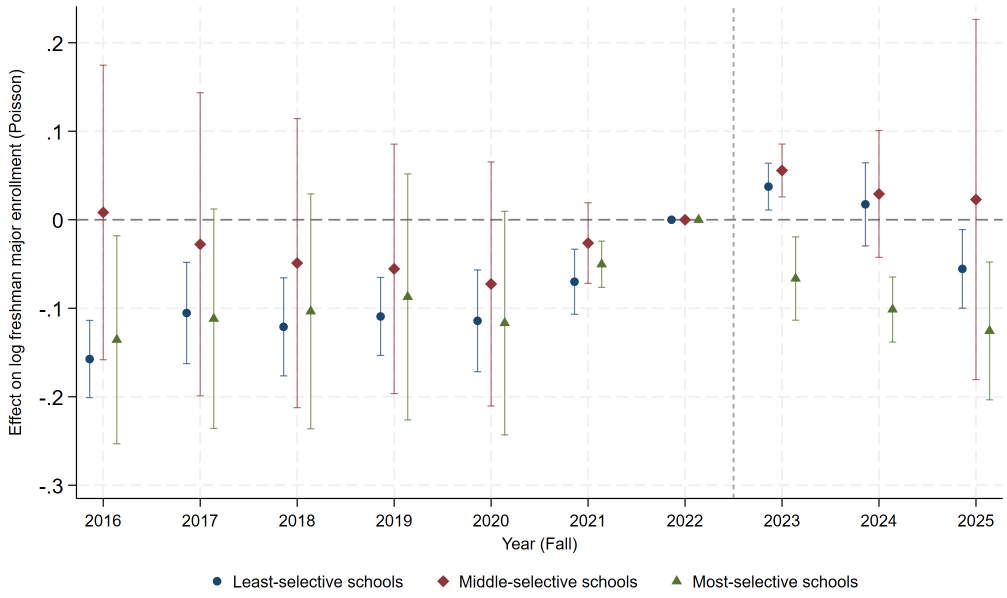


Figure D6: Freshman Enrollment Responses by Institutional Selectivity

Notes: This figure presents Poisson estimates of the relationship between major-level exposure to GenAI automation and fall Freshman enrollment counts, separately by the selectivity of the student's university. Each series is the total effect for a given selectivity group. Coefficients are interpretable as differences in log Freshman enrollment between Fall 2022 and a given year associated with a 10 percentage-point difference in major-level GenAI exposure. The specification controls for major and year fixed effects. Standard errors are clustered by major.

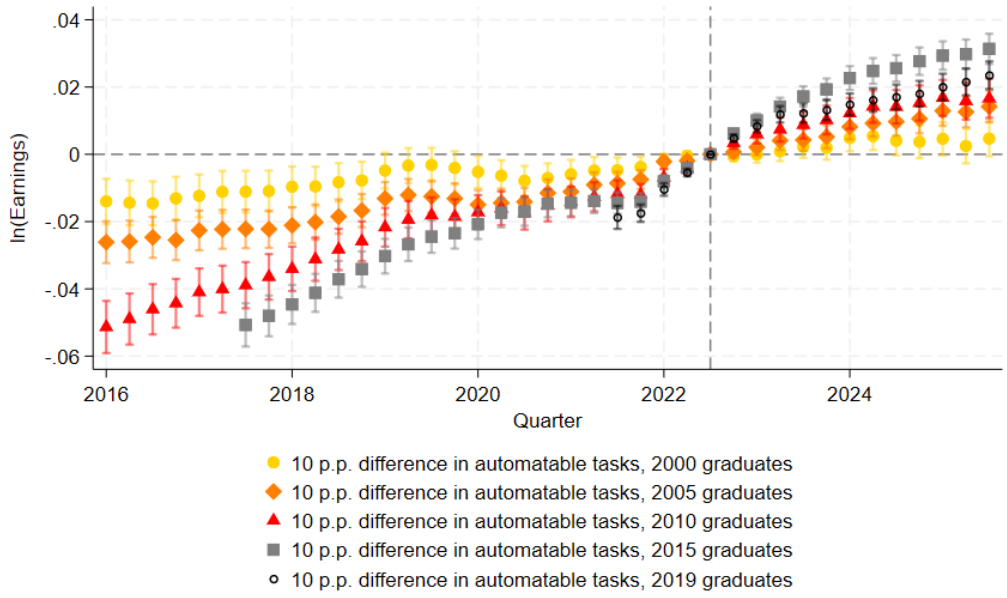


Figure D7: Earnings Growth by Graduation Cohort, Automation Exposure

Notes: This figure presents individual-level regressions of worker earnings over time across graduating cohorts compared to Q3 of 2022 for graduates with a 10 percentage-point difference in major-level GenAI exposure in their major. Each model is estimated separately for each cohort. The specification interacts major automation exposure with quarter fixed effects and controls for individual and quarter fixed effects. With the individual fixed effects, this traces average earnings growth trajectories over time by cohort given a 10 percentage point increase in their majors' GenAI automation. Standard errors are clustered by major.