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

Samuel Dodini†, Tucker Smith‡, and Yaling Xu§

September 11, 2026

Preliminary. Please do not cite or circulate.

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 nationwide 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 equally by reduced employment rates and lower conditional wages. In response to soft labor demand, graduates from 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 suggests there may be limited returns to such formal "upskilling." 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

*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.

†Dodini: Federal Reserve Bank of Dallas (*samuel.dodini@dal.frb.org*)
‡Smith: Federal Reserve Bank of Dallas (*tucker.smith@dal.frb.org*)
§Xu: FAIRE 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 – tasks central to many knowledge-intensive occupations. While GenAI has the potential to increase productivity growth and aggregate income, the technology's capability to automate large volumes of tasks will likely translate to substantial changes in the skills firms demand from their employees, prompting concerns of widespread labor market disruption that may be unique among technology shocks throughout history.

In this paper, we examine the effects of rapid deployment of GenAI on (1) labor market outcomes for recent college graduates and (2) their educational decisions in response to changes in the demand environment. Individuals aging into key educational decisions and labor market entry may be a particularly insightful population for studying the early effects of GenAI on the workforce for multiple reasons. Concerns over GenAI displacing significant shares of workers in knowledge-intensive occupations have been particularly amplified by the coincident timing of GenAI's rapid deployment with unusually large increases in the unemployment rate for young college graduates (Figure 1). Moreover, initial empirical evidence suggests that GenAI may substitute for less-experienced labor while complementing more-experienced workers (Brynjolfsson et al., 2025; Hosseini and Lichtinger, 2025). Amid uncertainty of AI's capabilities and continued development and adoption, reducing hiring is a cheaper method to decrease headcounts than layoffs as firms adopt the technology in their production processes. Finally, as firms continue to adopt AI to perform tasks typically undertaken by their employees, workers will need to develop skills that complement the new technology's capabilities. Such skill development responses are difficult to observe among entrenched labor market participants but are more easily identifiable among individuals attached to the formal education system.

Despite the potential importance of such analysis, there is relatively little causal evidence of GenAI's effects on new college graduates and entrants. While a plethora of data sources allow for analysis of GenAI's effects on labor demand, a key challenge to estimating effects of changing demand across skills on 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 four-year universities in Texas linked to state earnings records. We observe postsecondary course histories, major choice, and graduation timing along with post-college employment, earnings, and industry of employment in Texas, allowing us to estimate effects

of the shock on earnings of new labor market entrants and on student investments in skills across fields of study. The data also allow us to observe typically unobserved factors like proxies of underlying ability, key background variables like origin institutions, and other factors that non-panel datasets miss. The richness of the data across cohorts also allows us to track the extent to which the effects we observe for new graduates have affected those further up the experience distribution.

Armed with this panel dataset, we project recently developed "task-based" measures of GenAI exposure onto postsecondary fields of study and compare outcomes of individuals from majors that typically feed into more- vs. less-automatable occupations across cohorts observed after vs. before the release of ChatGPT.

We first use job posting data to document how labor demand for AI-exposed occupations changed within the state of Texas, allowing us map the labor market experience of those covered by our rich administrative data to demand changes in the same geographic markets.

We then examine how the GenAI shock to labor demand we document affected earnings and employment outcomes for new graduates from more-exposed majors. We show that earnings in the initial year after graduation fell by 10% for graduates from majors with a 10 percentage-point greater share of tasks that could be automated by AI. Similarly-sized declines in both employment rates and earnings conditional on finding employment comprise the total effect, and these results support that GenAI's estimated impact on *relative* labor demand across occupations (e.g., Brynjolfsson et al., 2025) have harmed new labor market entrants that would typically sort into more-exposed occupations. Moreover, research on the lasting effects of initial labor market conditions for college graduates suggest that such harm may persist (e.g., Rothstein, 2023; Schwandt and von Wachter, 2019; Kahn, 2010).

One way that college graduates may respond to worse labor market prospects is to pursue more education. Students deciding whether to make educational investments weigh expected increases in lifetime earnings against direct, opportunity, and effort costs. We find that in response to facing a relatively soft labor market, graduates from more-exposed majors were significantly more likely to enroll in graduate school, consistent with a decline in opportunity costs. However, we find little evidence that students that continued their education shifted toward studying less-exposed fields. Such investments in education beyond a four-year degree could yield sizable returns if the additional expertise is more complementary to AI than the knowledge gained from undergraduate studies or if the signal value of the advanced degree sets them apart from their peers who have passed lower barriers to entry into the field. Contrary to this notion, we find that first-year earnings for new master's recipients fell by 5% for graduates from fields with a 10 percentage-point automatable task share after the release of ChatGPT. This suggests that the returns to formal "upskilling" within AI-exposed

fields may be limited.

We next explore whether GenAI has affected undergraduate major choice and course enrollment. Fields with greater AI exposure—such as computer science, management, computer engineering, and languages—increased in popularity over the decade leading up to the GenAI shock. Since the release the ChatGPT, the existing enrollment growth trends in more-exposed fields continued for upperclassmen through Fall 2025 – likely due to the high cost of switching majors. Continued growth in high-exposure major enrollment does not appear to mask more nuanced course-level adjustments: students that were already declared for more-exposed majors prior to the release of ChatGPT completed similar coursework in the subsequent semesters as prior cohorts. In contrast, fall freshmen enrollment declined by 6% more between 2022 and 2025 in majors with a 10 percentage-point greater share of automatable tasks. 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 enrollment responses to prior sectoral shocks (Weinstein, 2022).

Our empirical strategy relies on the key assumption 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).[1] 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 analysis contributes to multiple emerging strands of research on educational responses to technology and generative AI. First, a growing body of work finds mixed evidence on the degree to which AI has affected labor market prospects for young workers. Brynjolfsson et al. (2025) document substantial declines in employment (e.g., nearly 20% for software developers) for early-career workers aged 22–25 in occupations with larger automatable task shares since the release of ChatGPT. Hosseini and Lichtinger (2025) reinforce this, finding that firms that adopt GenAI sharply reduce junior hiring while senior employment remains

[1]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).

stable or rises. However, critics of this narrative posit monetary tightening, pandemic-era over-hiring, and remote work as alternative explanations for reduced demand for young workers in recent years (Iscenko and Millet, 2026; Frank et al., 2026; Emanuel et al., 2026; Mask, 2026; Lambert and Schindler, 2026; Liu et al., 2025). We provide evidence of AI-induced declines in labor demand that are robust to addressing proposed alternative explanations. Our administrative panel dataset confirms that these changes in hiring negatively affected early-career earnings and employment for AI-exposed labor market entrants, consistent with the negative wage effects found by Azar et al. (2025). We also show one method in which such entrants are adjusting to the decline in labor demand: continuing their education.

Second, our study contributes to research on the effects of AI and sectoral demand shocks on skill development and career choice. Existing evidence supports that, as predicted by a basic human capital model, students shift field choices upon realization of a sectoral demand shock (Conzelmann et al., 2023; Weinstein, 2022). Adopting a similar approach to Timmerman (2025), we identify fields of study exposed to AI automation. Unlike prior sectoral shocks, GenAI affects both the demand for skills and how skills are produced (Light, 2026; Peterson, 2025), and the availability of GenAI tools likely reduces the effort costs required for academic success. If LLMs are similarly capable of automating schoolwork for the same fields of study in which they automate commercial tasks, marginal students may *increase* their enrollment in such fields. We find delayed shifts in undergraduate enrollment away from AI-exposed fields that resemble responses to prior sectoral shocks (Weinstein, 2022; Conzelmann et al., 2023), suggesting that demand-side considerations are dominating reductions in effort costs for coming-of-age students. This result matches empirical evidence that Swiss students reduced apprenticeship applications for AI-exposed occupations (Goller et al., 2025) and experimental evidence showing students are wary of GenAI automation risk when choosing career paths (Gschwendt et al., n.d.).

More broadly, our results relate to 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. Although exposed graduates may benefit from investing in skills that complement the new technology through additional education, we find that undergraduates from more-exposed majors returned to grad school to largely study the same fields, and that master's recipients from those same fields also exhibited earnings declines.

2 Conceptual framework

Our empirical analysis of educational responses to GenAI is centered around a simple human-capital investment model in the tradition of Becker (1962). We first consider undergraduate field choice. In this framework, current and prospective students make educational decisions by weighing perceived lifetime earnings premia for earning a degree in field m against the monetary, opportunity, and effort costs of doing so. A standard field-specific demand shock would incentivize students to pivot from studying m if it reduced the perceived earnings premia for obtaining a degree in that field. 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, and our empirical analysis of the effects of GenAI on labor demand suggests that the new technology reduces demand for graduates from more-exposed fields.

A decline in the earnings premia 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 (Weinstein, 2022).

Generative AI diverges from a standard technology shock that affects demand for particular skills in the fact that GenAI tools not only can automate tasks in the labor market, but also many components of 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.

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

(a) Switching from an exposed major increases with GenAI whenever effort costs fall by less than the returns (the labor demand channel dominates),

(b) Switching is decreasing in seniority

(c) Reallocation away from exposed majors will be relatively slow in the first few years of large and fast technology adoption and advancement, as large steps forward in AI technology progress faster than more senior students can predict or adapt to given switching costs. 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.

We also apply this framework toward 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. While we explicitly model GenAI as a substitute in the labor market for new entrants with four-year degrees, we allow for ambiguity in whether the students consider additional knowledge in exposed fields to complement or substitute the new technology. This yields the following predictions:

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

(f) Graduate school enrollment increases in AI-exposed fields if students perceive additional education as complementary to GenAI, but decreases if the perceived earnings premia for graduate degrees in AI-exposed fields falls by more than effort costs

3 Empirical Framework

3.1 Data Sources

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 Industrial 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.

Job 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 A1 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 overrepresents postings in financial services and insurance, information services, manufacturing, mining, professional and business services, real estate, and wholesale trade. The Lightcast data underrepresents 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 using the Lightcast data as a proxy for labor demand in a panel setting if such representativeness remains roughly constant over time. Figure A2 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). Evidence of 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, quarter observations from Q2 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 (approximately 873,000 job ads).

University of Houston Education Research Center: Our other primary dataset consists of administrative student-level data from the University of Houston Education Research Center (UHERC), featuring linked individual-level K-12, postsecondary, and workforce records from 1994 to present. K-12 data from the Texas Education Agency include 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), course enrollment, declared majors, and degree and certificate records from all public colleges and universities in the state. 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.[2] The records

[2]Self-employed workers, independent contractors, many federal employees, members of the military, and

report employer industry codes (6-digit NAICS) in all years but do not report firm identifiers, occupation, county of residency, or county of employment.

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 time period. 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 A2 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.

3.2 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 ($OccExp_c$) measures as weighted averages of its O*NET core and supplemental tasks.[3] 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 requirements, and other hurdles yield notable divergences between theoretical exposure and

participants in the informal sector are not covered by the state unemployment insurance system.

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

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. 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 robustnesss of our results to using theoretical exposure measures in Appendix B.

Although many students likely have a career path in mind, in college they choose fields of study to pursue, rather than occupations directly. We define a field of study's (formally, four-digit Classification of Instructional Programming 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 A1 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.

3.3 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 an information and technology shock and the start of rapid

GenAI proliferation to commercial uses, carrying differential incidence across occupations and fields of study. In particular, we estimate the following event-study specification:

$$y_{imt} = \sum_{\substack{\tau \\ \tau \neq t^* - 1}} \pi_{\tau}^{ES} \mathbf{1}\{t = \tau\} * MajExp_m + \alpha_m + \alpha_t + \Phi S_i + \epsilon_{imt} \tag{1}$$

an outcome y_{imt} for individual i who graduated from (or entered) with major m in year t is regressed on their major's share of automatable tasks $MajExp_m$ interacted with indicator variables for belonging to a particular cohort τ. Major α_m and cohort α_t fixed effects absorb time-invariant differences in student outcomes across majors (i.e., if Economics majors always earn more on average than English majors) and time-varying shocks common to students across all majors (i.e., if labor demand fell for all college graduates during the pandemic). We additionally control for student-level demographic characteristics S_i, including race, ethnicity, gender, and college entry-exam scores.[4] 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 GPT-3's release.

$MajorExp_m$ 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-GPT cohorts as the *causal* effect of 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).[5]

Any time-varying shock correlated with $MajExp_m$ would threaten identification of π_{τ}^{ES}. We test the robustness of our results to three forms of such threats. First, one may be concerned that 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

[4]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).

[5]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). 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.

tech could be mistakenly interpreted as effects of AI. More generally, monetary policy can have differential 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 address 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-year, limiting identification to comparisons of students with different majors from the same school. Controlling for college entry-exam scores, a proxy for student ability, also alleviates concerns of bias from changes in the composition of students studying more- vs. less-exposed majors (i.e., if computer science became more popular with lower-ability students 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-GPT Lightcast job postings that offered remote work and test robustness to controlling for WFH exposure interacted with cohort fixed effects.

3.4 Design Validation: AI and Labor Demand

Before estimating the effects of GenAI on outcomes of college graduates and entrants, we first characterize the AI shock to labor demand using job postings data from Lightcast. To estimate the effects of the emergence of GenAI tools on labor demand, we estimate Poisson event study estimates from a quarterly, occupation-industry level version of Equation 1 with posting counts on the left-hand side. 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. 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 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: $log(E(y_{okt}|\mathbf{x_{okt}}) = \alpha_{ok} + \alpha_{kt} + \sum_{\tau \neq t^*-1}^{\tau} \pi_{\tau}^{ES} * \mathbf{1}\{t = \tau\} * OccExp_o + \epsilon_{okt}$ where k denotes the industry. In the second model, we structure the data as firm-quarter observations on a balanced panel of firms who were present in the posting data through the entire panel. 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. 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 model clarifies demand mechanisms. If the exposure effects for incumbent firms is smaller than the occupation-industry estimates, then new firms requiring fewer human workers (or a reduction of surviving firms) explains 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 for 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.[6]

Figure 2 presents these event study estimates for our occupation-industry 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 GPT-3. 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 GPT's release and continued to fall after the beginning of 2025. Postings for more-exposed occupations remained nearly 10 percent lower through the mid point of 2026.[7] 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.

[6]To smooth volatility in firms' posting behavior from quarter to quarter, we use a four-quarter moving average of job postings and automation score as the dependent variables.

[7]In complementary analyses, we also structure the data spatially at the occupation-commuting zone level and find comparable results (Figure B1). We estimate $log(E(y_{oct}|\mathbf{x_{oct}}) = \alpha_o + \alpha_{ct} + \sum_{\tau \neq t^*-1}^{\tau} \pi_{\tau}^{ES} * \mathbf{1}\{t = \tau\} * OccExp_o + \epsilon_{oct}$ where o, c, and t index occupation, commuting zone, and quarter, respectively.

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). The coefficients in Panel A suggest that firms whose pre-ChatGPT workforce was more concentrated in occupations whose tasks became more AI-automatable after ChatGPT reduced their job posting rate by approximately 8 percent by the beginning of 2024 for every 10 percentage point increase in exposure. Panel B shows that the composition of the firms' post-ChatGPT job ads shifted away from more automatable occupations as well. For every 10 percentage point increase in the firm's exposure, the average automation exposure score of the firm's job ads fell by approximately two percent by early 2024.

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 position 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 posting 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 estimate the functions with 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. As a 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, and we interact this sensitivity measure with quarter fixed effects as an additional regressor in our Poisson model of Equation 1.

Figure B2 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.

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 the early incidence of this shock may be particularly painful for new college graduates. In response to these shifts in labor demand, new graduates and current students may adjust their educational choices.[8]

4 Results

4.1 Effects of GenAI on New Labor Market Entrants

We next examine how the GenAI labor demand shock affected new graduates who studied fields aligned with AI-exposed occupations. Figure 4 presents estimates of Equation 1 with total earnings during a graduate's first four quarters after graduation as the outcome. Prior to the release of ChatGPT, earnings trended similarly across cohorts for more- relative to less-exposed majors. However, earnings trajectories across majors diverged after the release of the new technology. May 2023 graduates holding a major with a 10 percentage-point higher share of automatable tasks experienced a relative earnings decline of approximately $2,000 (2025$) in their first year in the labor market compared to graduates that entered the labor market before ChatGPT's release. The estimated effect size grew to over $3,000 for May 2024 graduates, representing nearly a 10% decline relative to the raw average earnings across 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 posting shown in Section 3.4. 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.

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

[8]The labor demand effects we find are particularly salient for this group because jobs in the Lightcast data tend to orient toward those with less job experience. Figure A3 shows that firms' typical job postings require, on average, only a minimum of 1-3 years of experience.

demand shock, as well, particularly if GenAI complements "tacit" knowledge, gained through experience, while substituting for "codified" knowledge (Brynjolfsson et al., 2025).

Although we only observe earnings through the first half of 2025, we explore early indicators of whether the initial effects of the GenAI demand shock 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 significant relative declines of \$2,351 in second-year earnings, 0.045 quarters worked, and \$2,102 in earnings, conditional on working all four quarters. Only the estimated effect on second-year employment is an economically significant improvement relative to the first-year effects.

Relative declines in earnings suggest GenAI substitutes for new four-year graduates from fields with higher automatable task shares. We next explore whether the additional expertise gained from graduate programs in more-exposed fields is complementary or also substitutable. We estimate Equation 1 with the same field-level exposure levels and a cohort-level panel of new master's degree recipients. Figure 6 shows that first-year earnings grew by \$2,500 more across cohorts for master's recipients in fields with a 10 percentage-points 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-GPT cohort. 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.

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 7 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.[9] 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.[10]

[9]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.

[10]Figure B3 presents estimates that control for each confounder, separately.

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-GPT graduating cohorts relative to previous cohorts could cause average earnings and employment rates for those majors to fall, regardless of an AI shock. However, Figure 8 shows a relative *increase* in SAT scores for graduates of more- vs. less-exposed majors in this cohort relative to previous ones, suggesting that the estimated declines in earnings for more-exposed graduates come in spite of potentially positive ability bias.

Our estimates could 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 4.2. Figure B5 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 B6 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. Tables B1 and B2 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 9 out of 27 specifications using an exposure measure grounded in an LLM's own definition of what tasks it can perform are not statistically distinguishable from zero. Given evidence from Yin et al. (2026) on the fragility of LLM-based rubric scores across different models, we place greater weight on results derived from other exposure measures.

4.2 GenAI and Educational Decisions

Canonical human capital models (Becker, 1962) imply multiple mechanisms through which GenAI automation could affect educational decisions by students still enrolled in college and recent graduates. By reducing the opportunity cost of education, the contemporaneous decline in labor demand for more-exposed majors may incentivize graduates to re-enroll in post-baccalaureate programs. To the extent that students believe the observed

decline in labor demand for AI-exposed majors represents a *persistent* shift in demand for particular skills, both undergraduate and graduate students may pivot toward studying less-exposed fields.[11] However, the new technology represents more than just a pure shock to labor demand: just as firms utilize LLMs to replicate human intelligence, so too can students. The availability of GenAI tools likely reduces the effort costs required for academic success. If LLMs are similarly capable of automating schoolwork for the same fields of study in which they automate commercial tasks, marginal students may *increase* their enrollment in such fields.[12]

How did four-year graduates from more-exposed majors respond to soft labor demand? Figure 9 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). However, if GenAI causes a more persistent shift in the types of skills rewarded in the labor market, the returns to graduate school investments could hinge on newly developed skills being more complementary to the technology than those acquired through undergraduate studies. Evidence of declines in earnings for new master's recipients from AI-exposed fields in Section 4.1 suggest limited returns to formal upskilling in such fields. Returning graduates from more-exposed fields thus would likely benefit from pivoting toward studying less-exposed fields.

To test whether students from more-exposed majors shifted toward less-exposed fields in graduate school, 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 10 shows little discernible evidence of such shifts. A 10 percentage-point greater undergraduate major automatable task share is associated with only a 0.5% decline

[11] GenAI may also affect undergraduate and graduate enrollment decisions by affecting demand for *all* workers of a given education level, regardless of their field of study. Because our event study design leverage comparisons across majors of differing exposure to GenAI automation, we do not identify effects of this channel.

[12] GenAI tools likely also affect the cost of instruction. Peterson (2025) posits that if GenAI reduces the costs of instruction for the same fields in which it automates labor, institutions may increase offerings in AI-exposed fields, resulting in sup-optimal counts of young adults trained in those fields. GenAI may increase instruction costs, though, by enabling students to complete assignments while acquiring less human capital and reducing the signal value of out-of-class assignments in measuring learning (Light, 2026).

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.[13]

We next explore whether undergraduates shifted their fields of study in response to GenAI. Figure 11 presents estimates of the relationship between major-level GenAI exposure and enrollment counts, using a Poisson analog of Equation 1. Panel 11a 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 the 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 significant heterogeneity across classifications (Panel 11b). Switching majors as an upperclassmen often delays graduation, yielding substantial financial and effort costs. Moreover, 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 ten percentage-point difference in major exposure correlated with a 6% decline in freshmen enrollment between Fall 2022 and Fall 2025, consistent with their greater flexibility to adjust to a demand shock. 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 gradual adjustments as cohorts aged.

Null effects on major choice for students already enrolled prior to the GenAI shock may mask subtler adjustments to skill investments. We examine whether course-taking behavior of students that entered college declared for more- vs. less-exposed majors across cohorts that made course enrollment decisions after vs. before the GenAI shock. If students in more-exposed majors internalized declining labor demand for the skills taught in their declared fields, we would expect them to diversify their course selections beyond their initial major. The event study in Figure 12 suggests that, at least in the initial two years following the release of ChatGPT, they made little, if any, adjustments in course choices. The statistically insignificant estimates indicate that students that entered college as a major with 10 p.p. higher exposure reduced their courseload in the initial major by only 0.5% through their junior fall semester. For example: the average student that entered college as a computer science major in 2022 took 5.1 computer science courses through the fall of their junior year, only 0.1 fewer courses than did computer science majors that entered in 2019.

[13] 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%.

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

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, we add to a growing literature on how GenAI is affecting labor market prospects for young adults. In particular, our findings suggest the GenAI transition is causing varying degrees of harm across three groups. First, we provide direct evidence that recent graduates from majors that develop automatable skillsets exhibit deteriorating labor market outcomes. 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 and 2026 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 shifts toward studying other fields may buffer them against the labor demand shock.

Our results are consistent with the GenAI transition operating as a normal technological demand shock. However, the technology could carry much broader implications for both the labor market and education system. 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

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

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.

Autor, David and Neil Thompson, "EXPERTISE," *NBER Working Paper 33941*, 2025.

Azar, José, Mireia Gine, and Javier Sanz-Espín, "AI Is Already Eroding Wages: Quasi-Experimental Evidence From Occupational Exposure," 2025.

Becker, Gary S., "Investment in Human Capital: A Theoretical Analysis," *Journal of Political Economy*, 1962, *70* (5), 9–49.

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.

Bick, Alexander, Adam Blandin, and David Deming, "THE RAPID ADOPTION OF GENERATIVE AI," 2025.

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.

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.

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

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. arXiv:2303.10130 [econ].

Emanuel, Natalia, Emma Harrington, and Amanda Pallais, "THE POWER OF PROXIMITY TO COWORKERS," *NBER Working Paper*, June 2026.

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

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.

Gschwendt, Christian, Martina Viarengo, and Thea S Zollner, "Generative AI and Career Choices," *IZA discussion paper*, *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. arXiv:2503.04761 [cs].

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.

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.

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.

Lambert, Peter John and Yannick Schindler, "The Broken Ladder: AI, Remote Work, and Early-Career Hiring," 2026.

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.

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.

Mask, Joshua, "Same Occupations, Different Clocks: Separating Remote Work from Generative AI for Labor-Market Entrants," August 2026.

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

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

Rothstein, Jesse, "The Lost Generation?," *Journal of Human Resources*, 2023, *58* (5), 1452–1479.

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.

_ **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.

Timmerman, Jean Xia, "Educational Exposure to Generative Artificial Intelligence," *FEDS Notes*, February 2025, (2025-02-26), None–None.

Weinstein, Russell, "Local Labor Markets and Human Capital Investments," *Journal of Human Resources*, September 2022, *57* (5), 1498–1525.

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.

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.

Tables and Figures

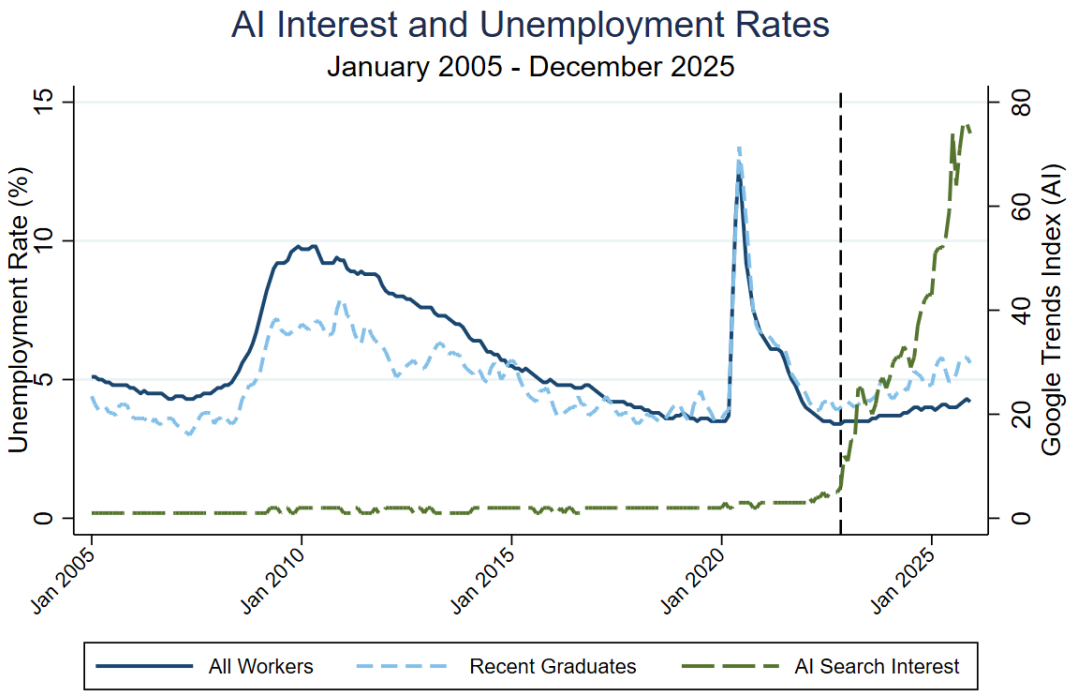
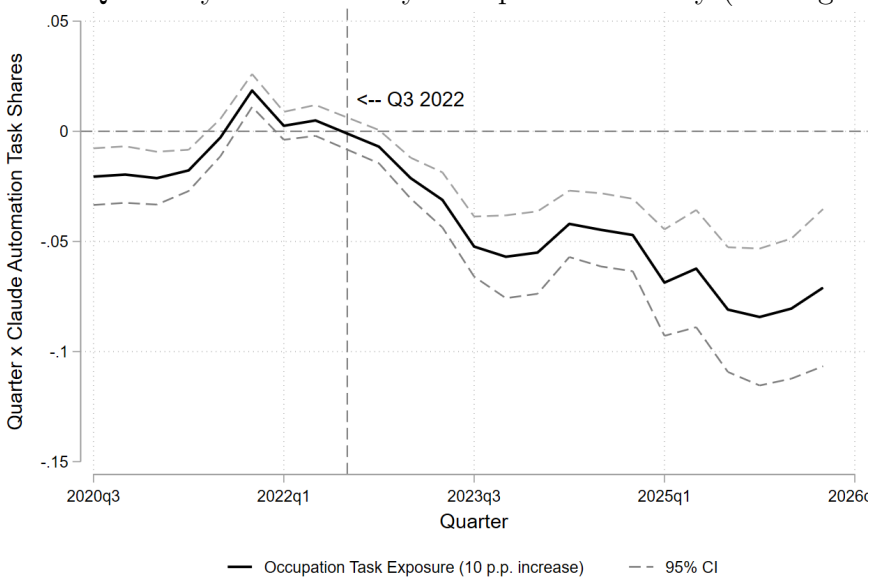


Figure 1: AI Boom Coincided with Increase in Unemployment for Young College Graduates

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.

Panel A: Quarterly Ad Counts by Occupation-Industry (Unweighted)



Panel B: Quarterly Ad Counts by Occupation-Industry (Weighted)

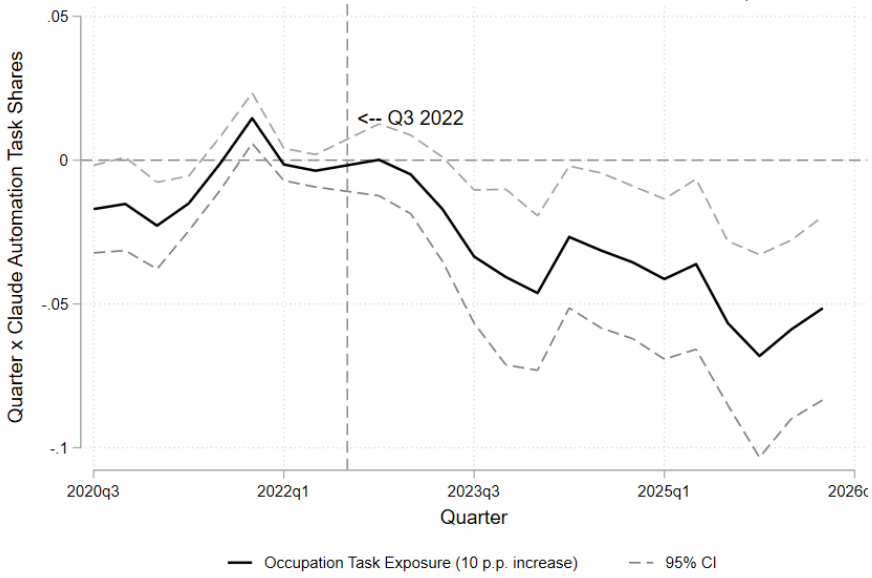


Figure 2: More Exposed Occupations Experienced a Decline in Job Posting Rates

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.

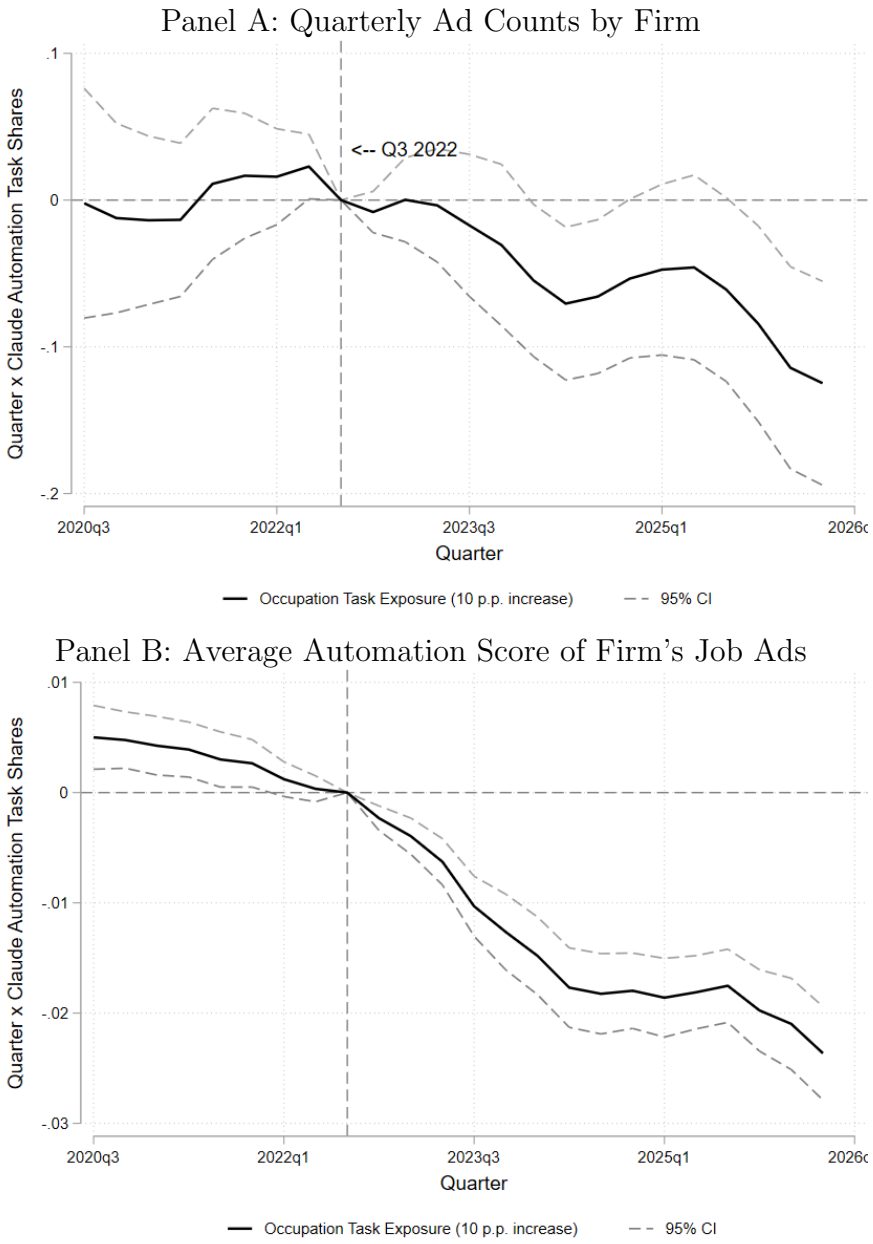


Figure 3: More Exposed Firms Reduced ther Job Posting Rates and Pivoted Away from AI-Automatable Occupations

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.

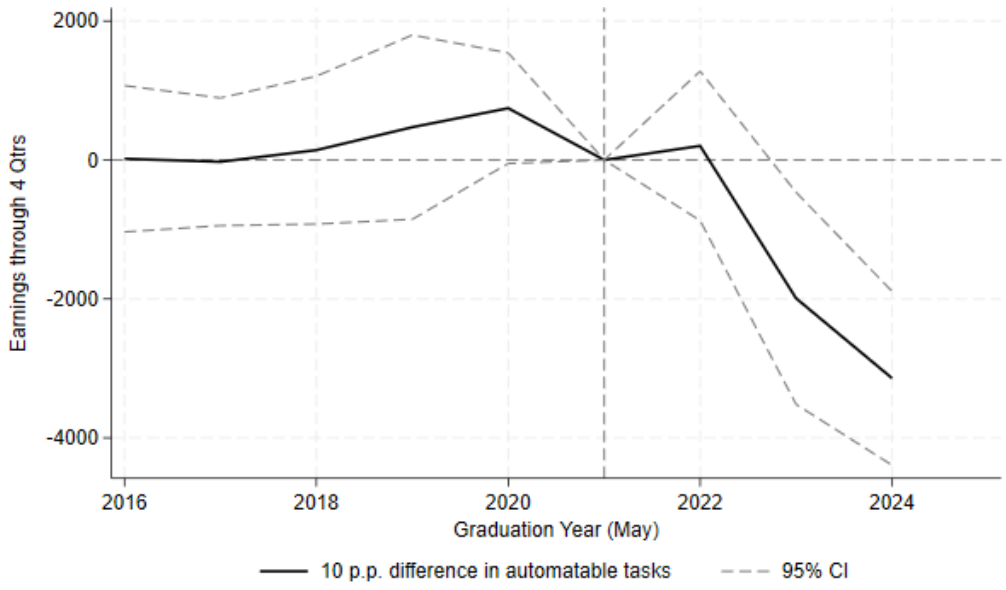
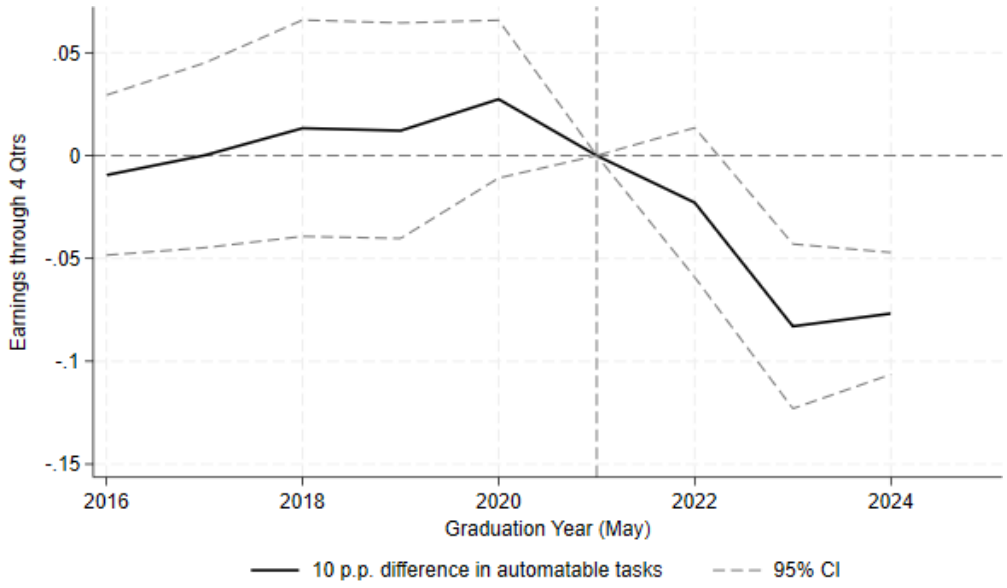
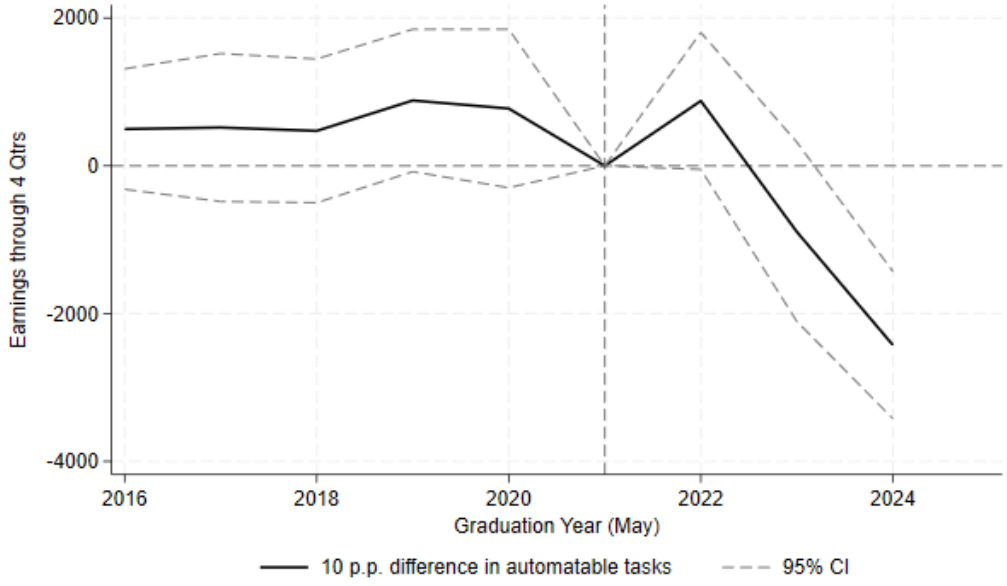


Figure 4: Earnings Fell for New Graduates from AI-Exposed Majors

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 entrant 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 entrant exam scores. Standard errors are clustered by major.

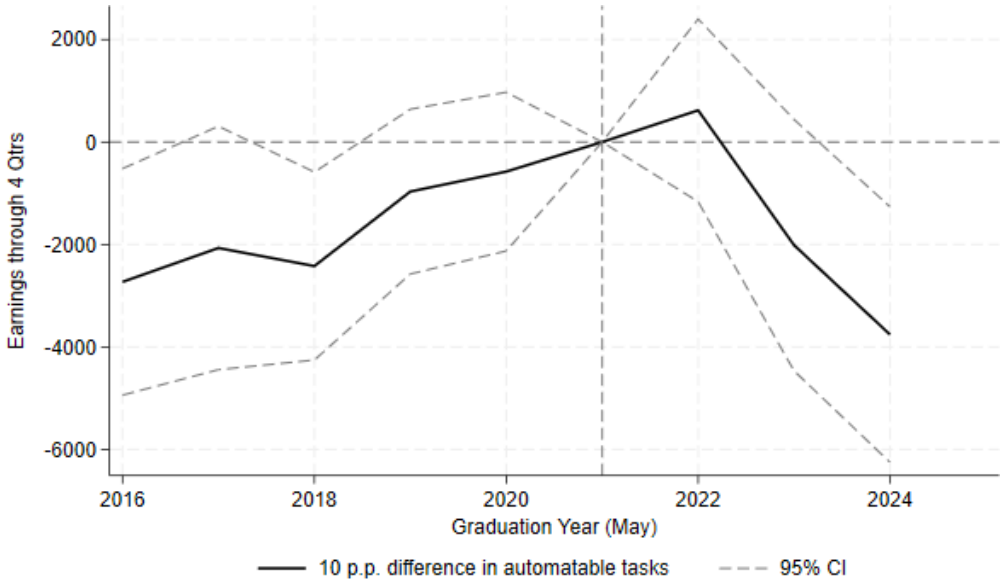


Figure 6: Earnings Fell for New Master's Recipients in AI-exposed Fields

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 2021 graduates. The specification controls for major (4-digit CIP code), school-by-year fixed effects, and student demographics. Standard errors are clustered by major.

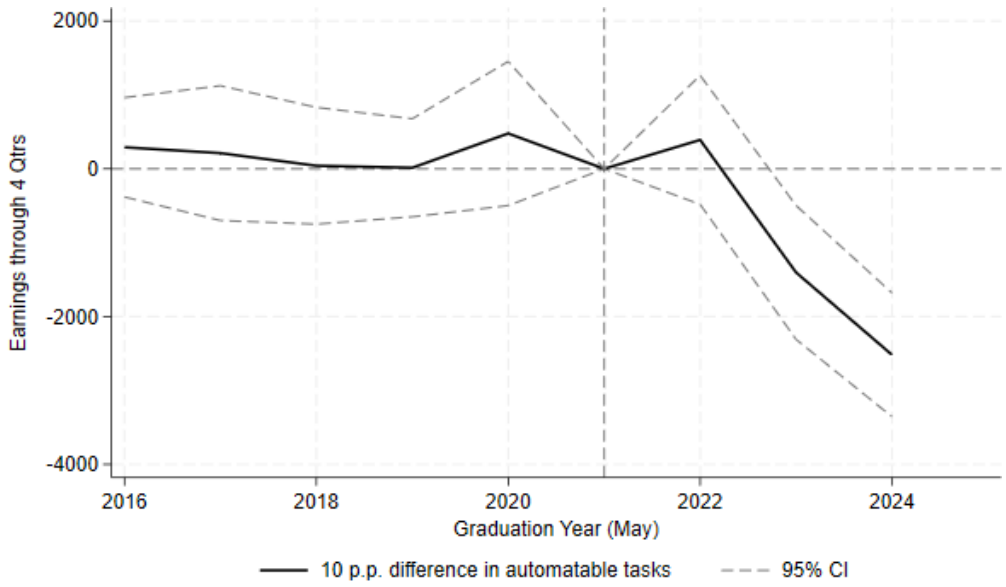


Figure 7: Estimated Declines in Earnings are Robust to Controlling for Time-Varying Industry Shocks & Remote Work Shares

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.

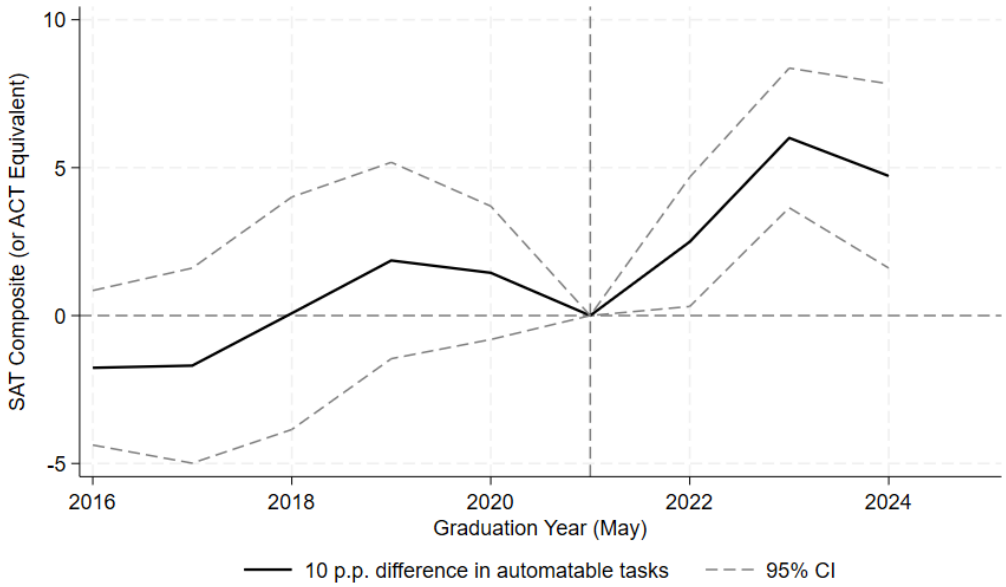


Figure 8: Recent Grads From More-Exposed Majors were Positively Selected on Observable Ability

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.

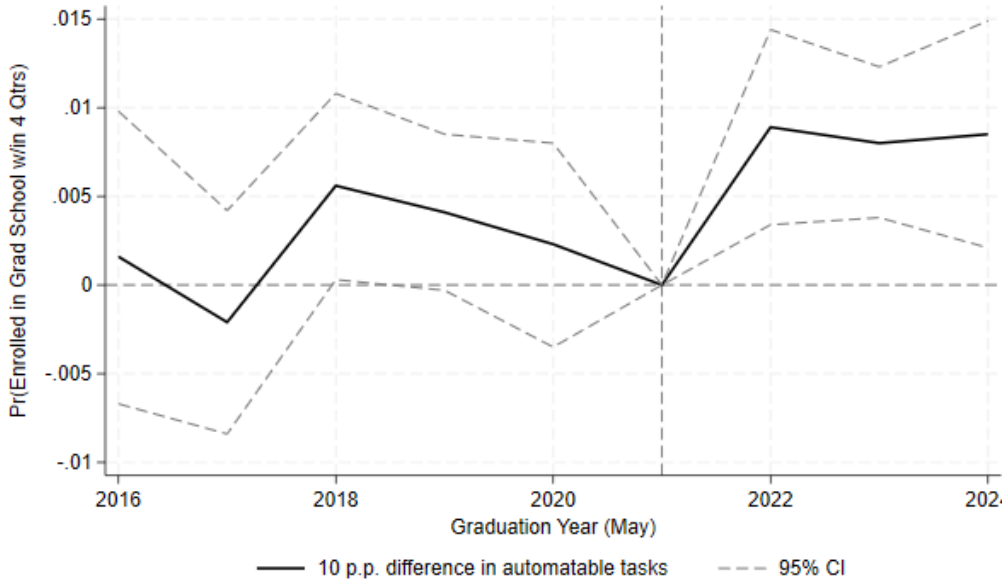


Figure 9: Graduate-School Enrollment Increased for Students from More-Exposed Undergraduate Majors

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.

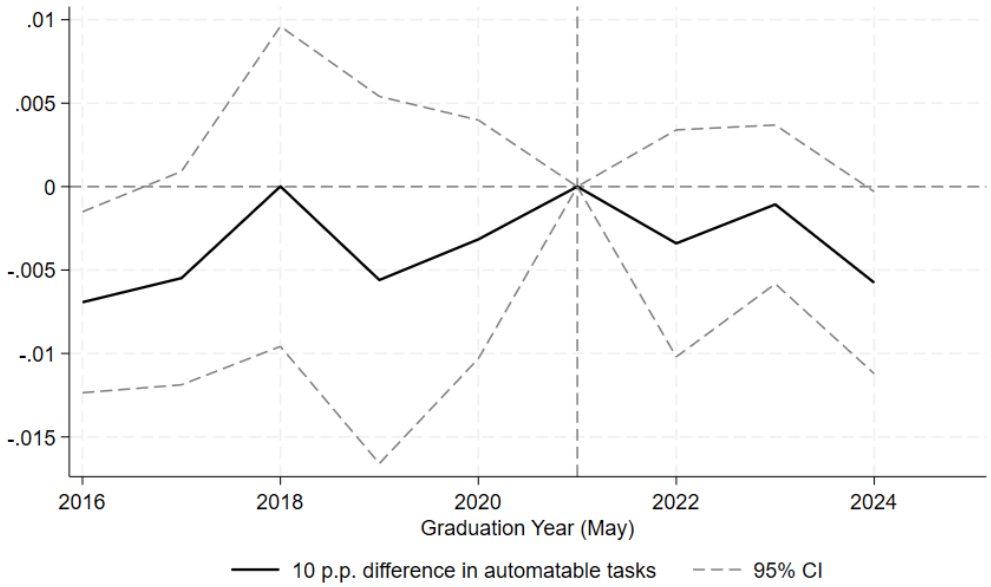
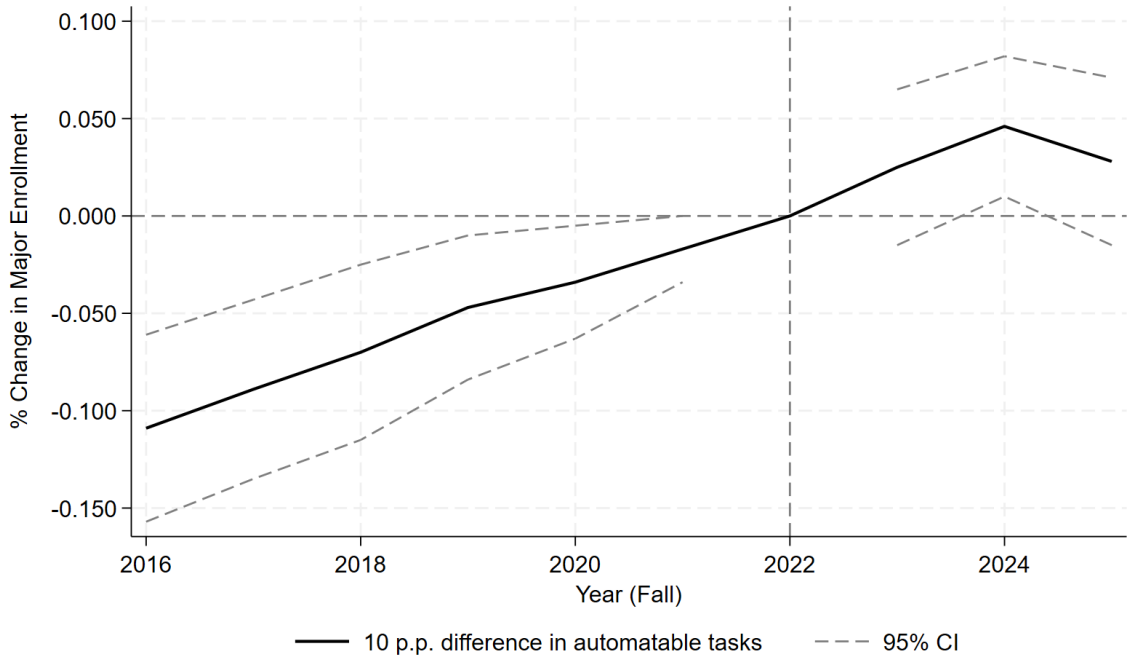
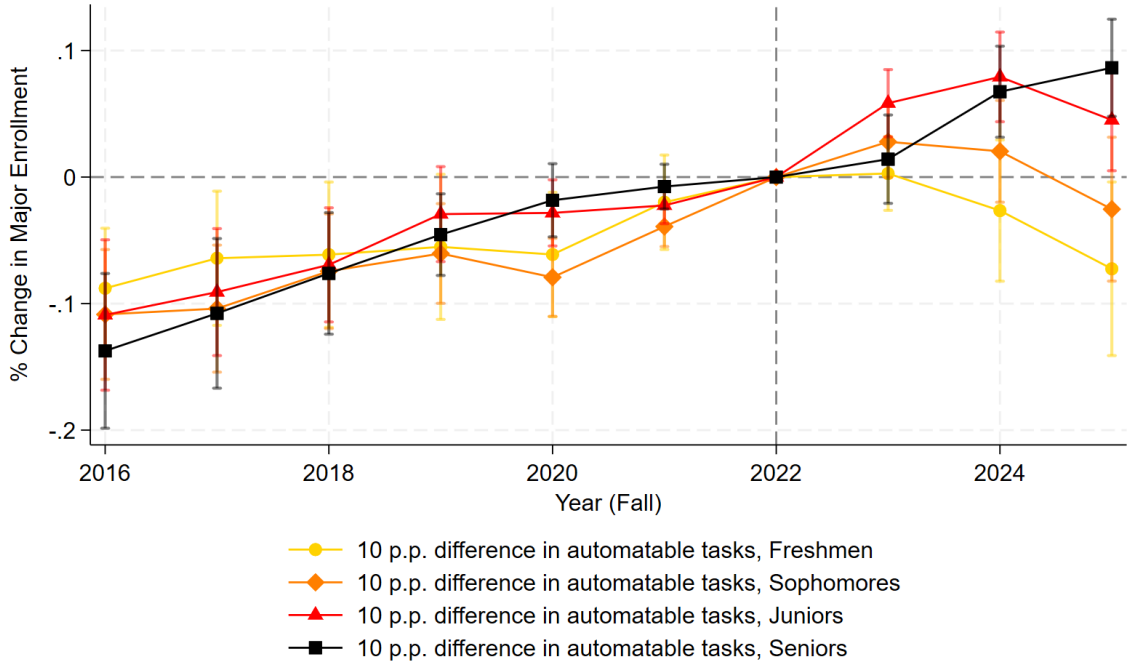


Figure 10: Little Evidence that Graduates From More-Exposed Majors Shifted to Less-Exposed Graduate Programs

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 $MajorExp_m$ 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 11: 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.

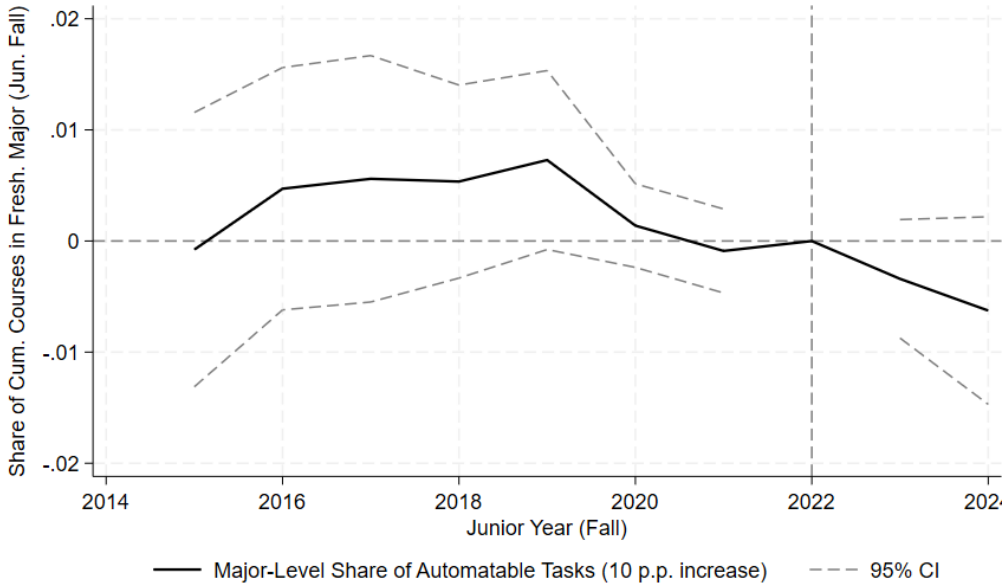


Figure 12: Students that Initially Declared for AI-Exposed Majors Did Not Shift to Safer Fields of Study

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 the share of courses that they took in their initially declared major through the fall of their junior year. 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.

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

	(1) Mean
Age, Fall Freshman Sem.	18.2
	(1.68)
Major share of automatable tasks (Observed), Fall Freshman Sem.	0.078
	(0.067)
Major share of automatable tasks (Theoretical), Fall Freshman Sem.	0.469
	(0.086)
Major share of automatable tasks (Observed), Fall Soph. Sem.	0.077
	(0.071)
Major share of automatable tasks (Observed), Fall Junior Sem.	0.078
	(0.074)
Switched majors from Fall Fresh. to Fall Soph., Sem.	0.214
	(0.410)
Switched majors from Fall Fresh. to Fall Junior, Sem.	0.352
	(0.478)
Course Share in Freshman Major through Fall Soph. Sem	0.206
	(0.198)
Course Share in Freshman Major through Fall Jun. Sem	0.231
	(0.192)
N	451,372

Notes: This table presents descriptive statistics for students from the 2013 to 2022 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 Programming 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 entrant 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.

A Data Validation and Additional Context

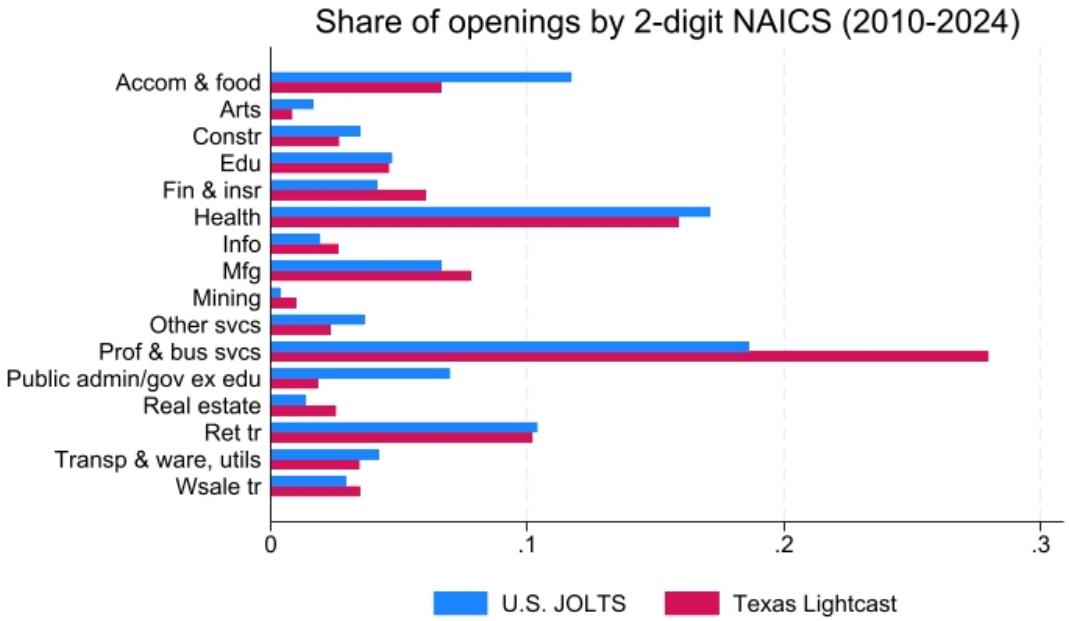
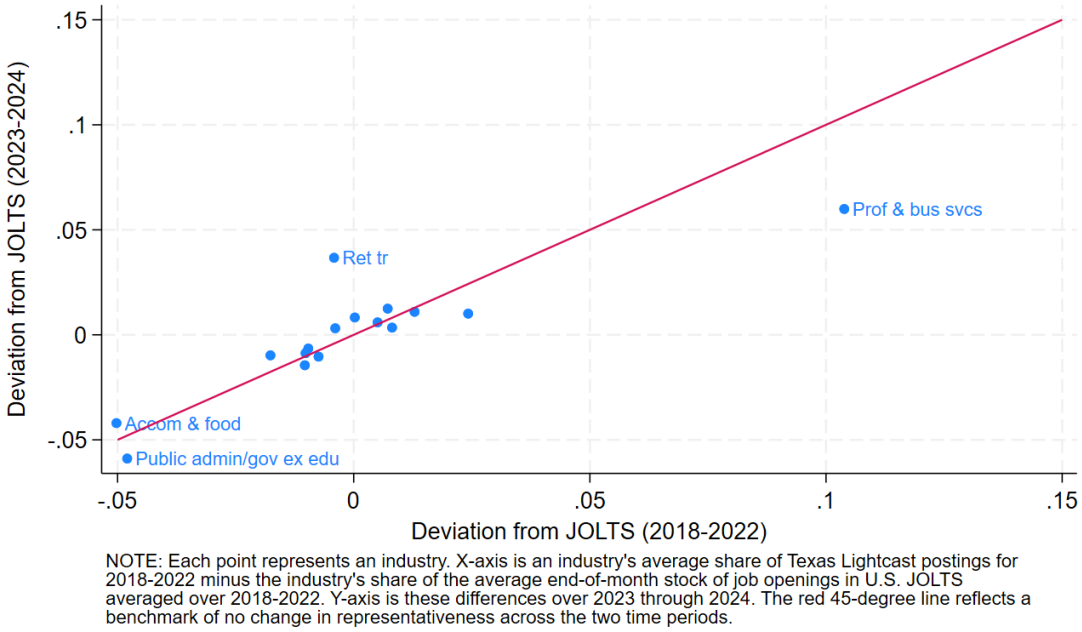


Figure A1: Posting Shares Across Industries in Lightcast vs. JOLTS

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 A2: Posting Shares Across Industries in Lightcast vs. JOLTS

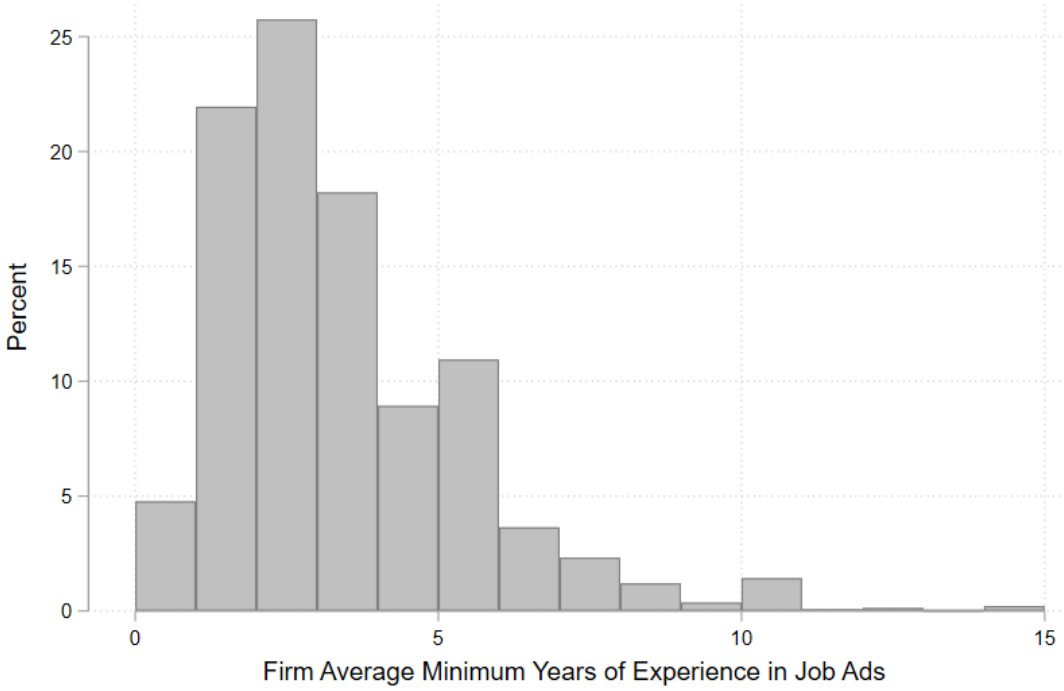


Figure A3: Average Required Minimum Experience for Firms' Job Ads in Lightcast

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.

Table A1: 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 Programming 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 A2: Descriptive Statistics: New Graduates 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.

B Robustness to Alternate Specifications and AI Exposure Measures

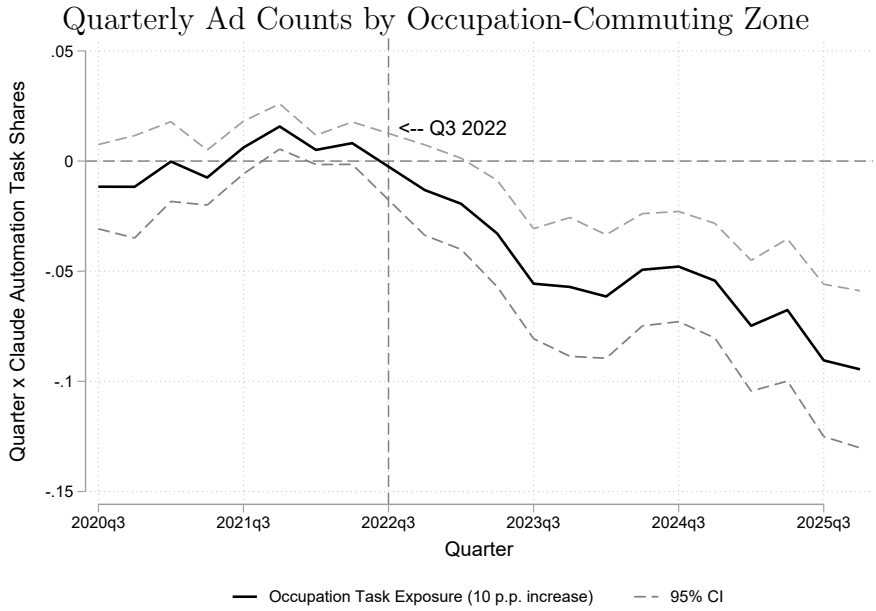


Figure B1: Exposure Results Are Robust to Spatial Controls

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.

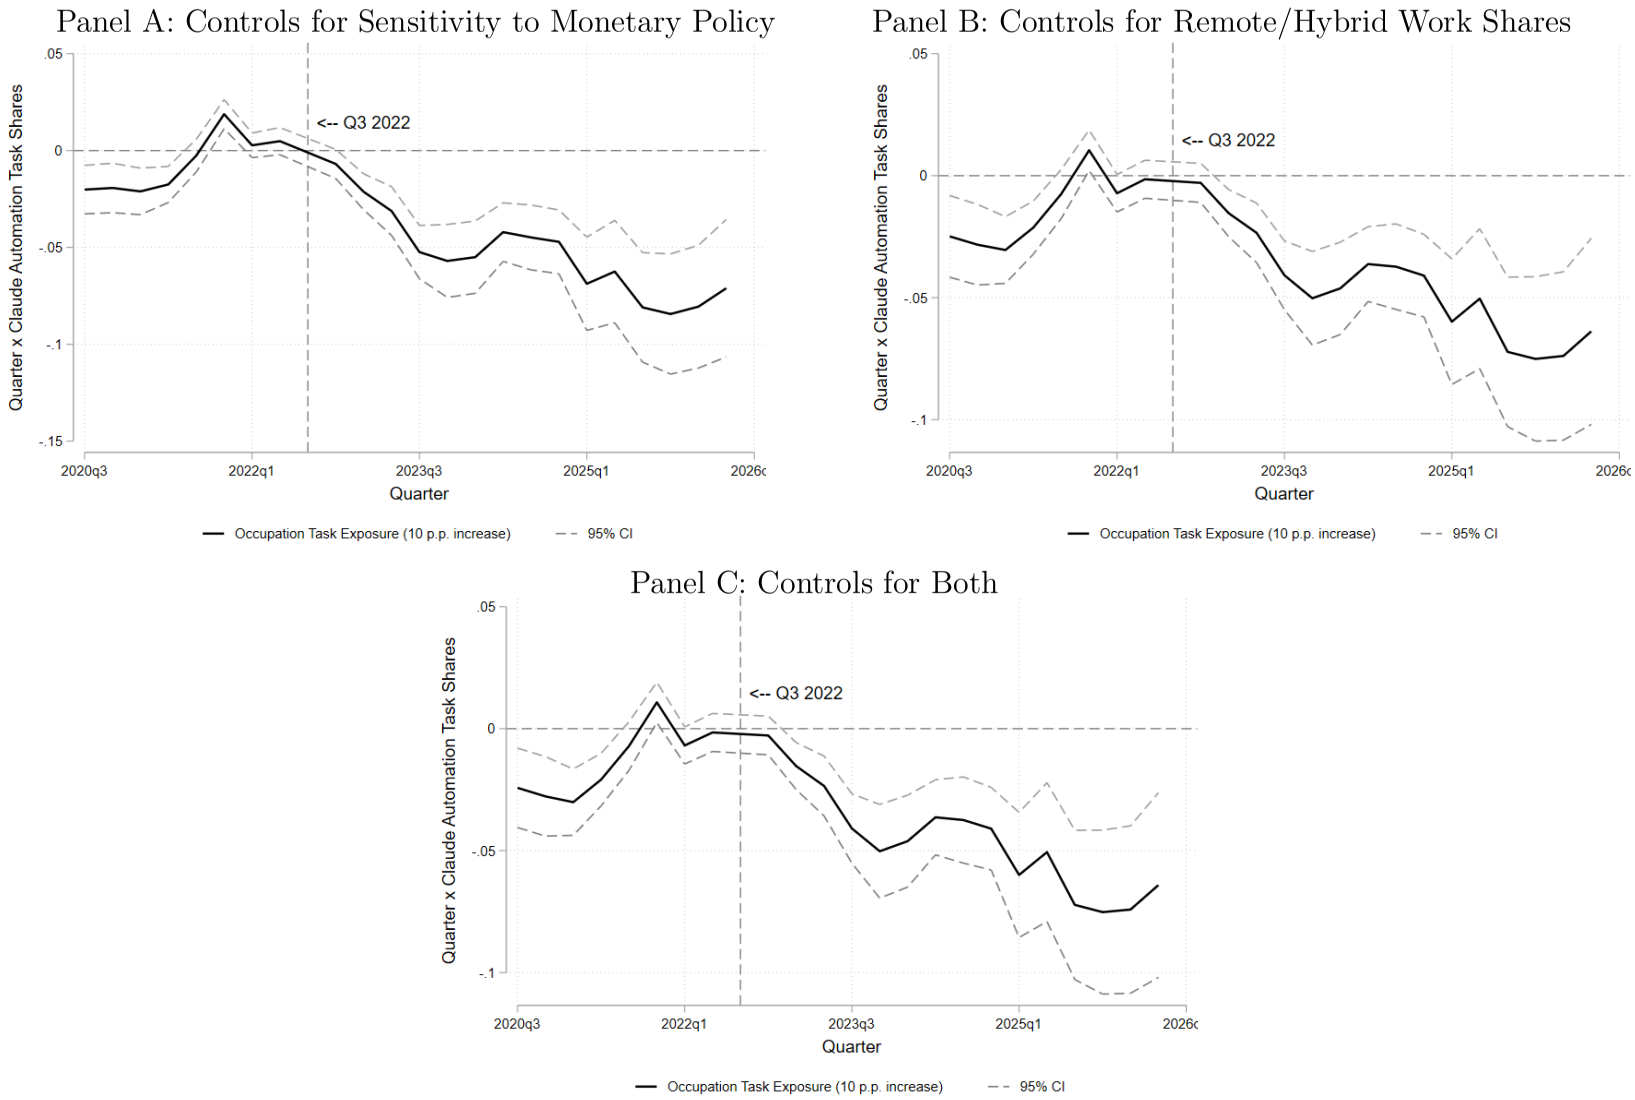
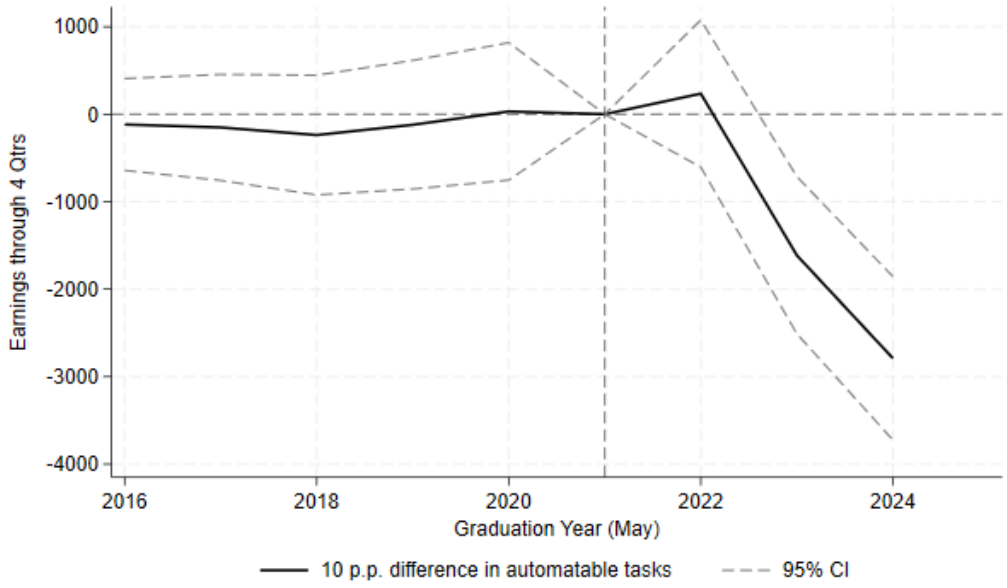
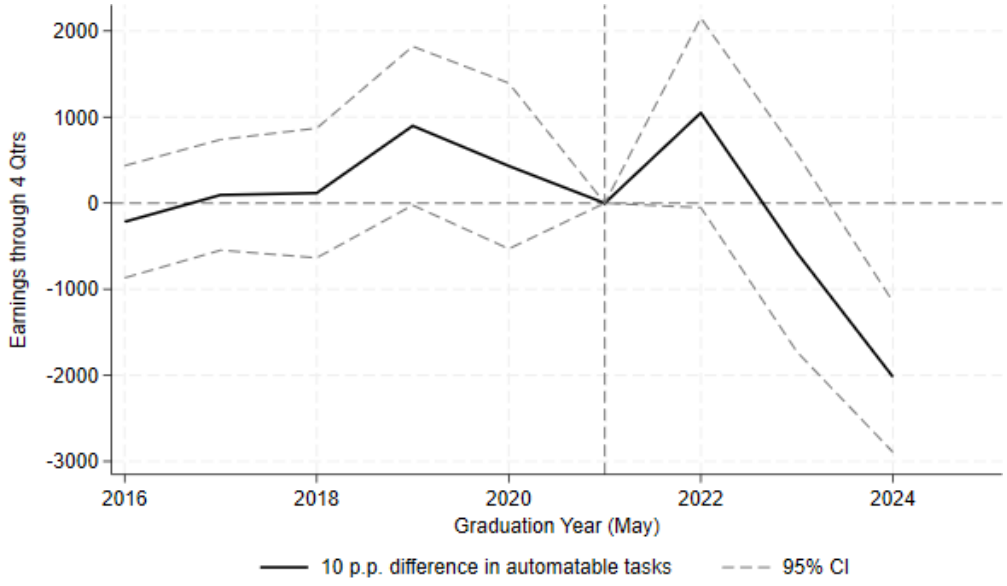


Figure B2: Exposure Results Are Robust to Controls for Monetary Policy Sensitivity and Remote/Hybrid Work

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 control for sensitivity to monetary policy (Panel A), occupation shares of remote/hybrid work (Panel B), or both (Panel C). 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.



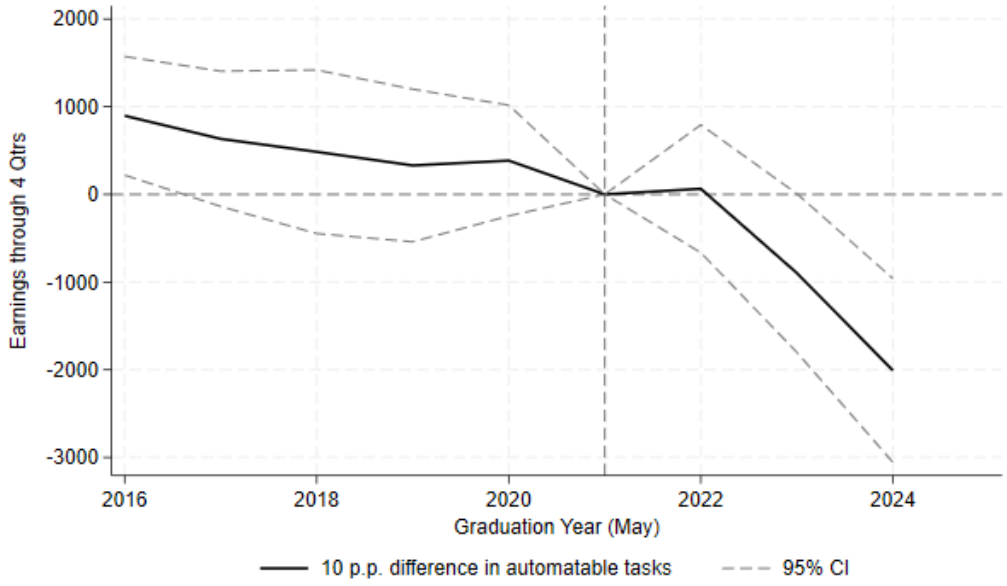
(a) Ind.-by-Year Fixed Effects



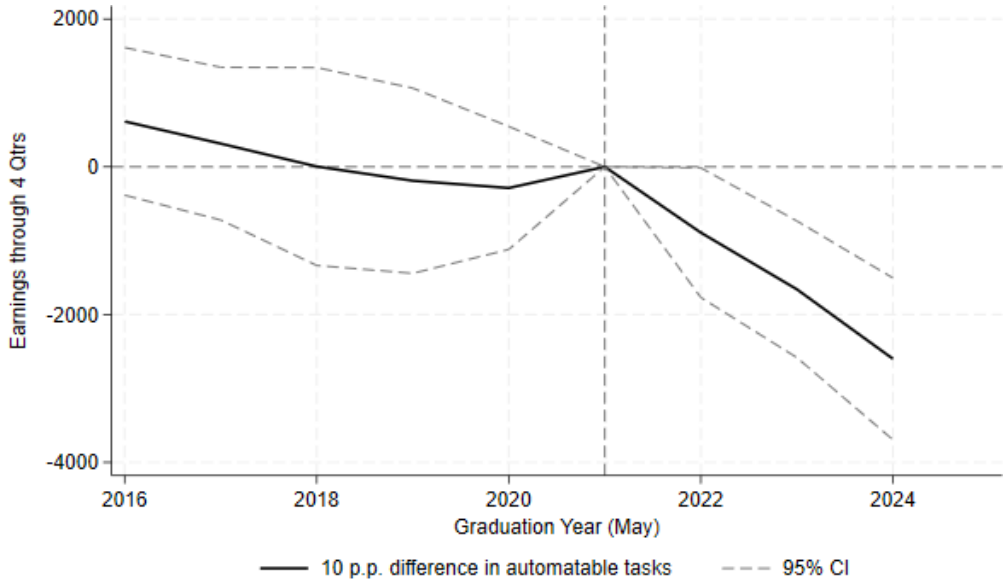
(b) Remote Work Controls

Figure B3: 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.



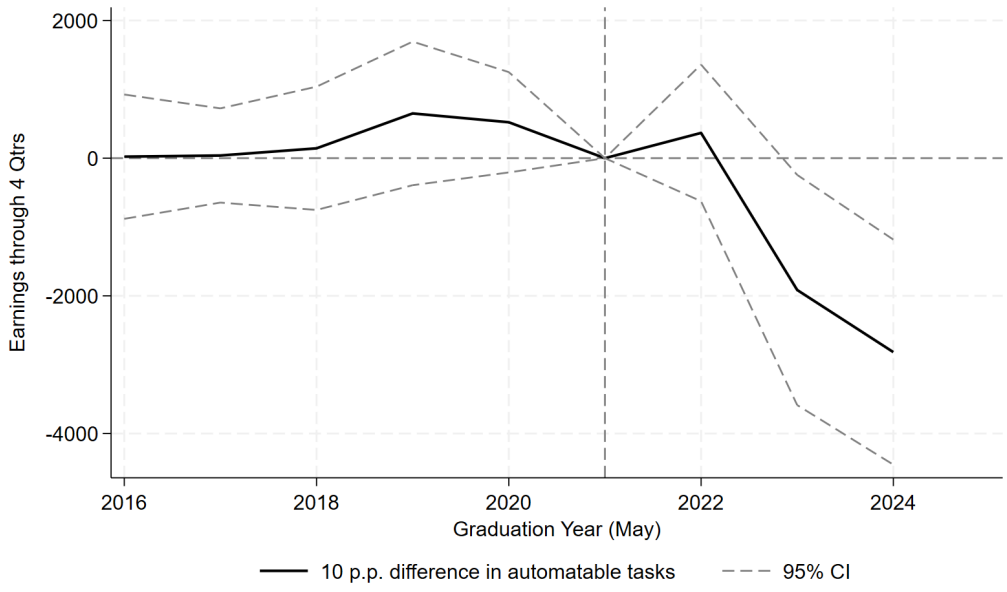
(a) GPT-Beta Exposure Measure



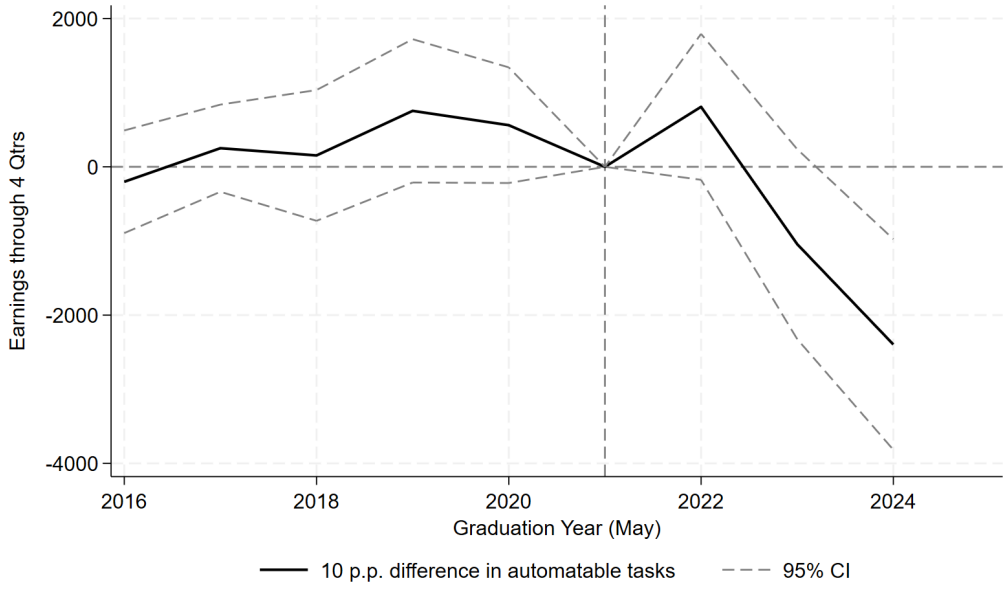
(b) Human-Beta Exposure Measure

Figure B4: 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.



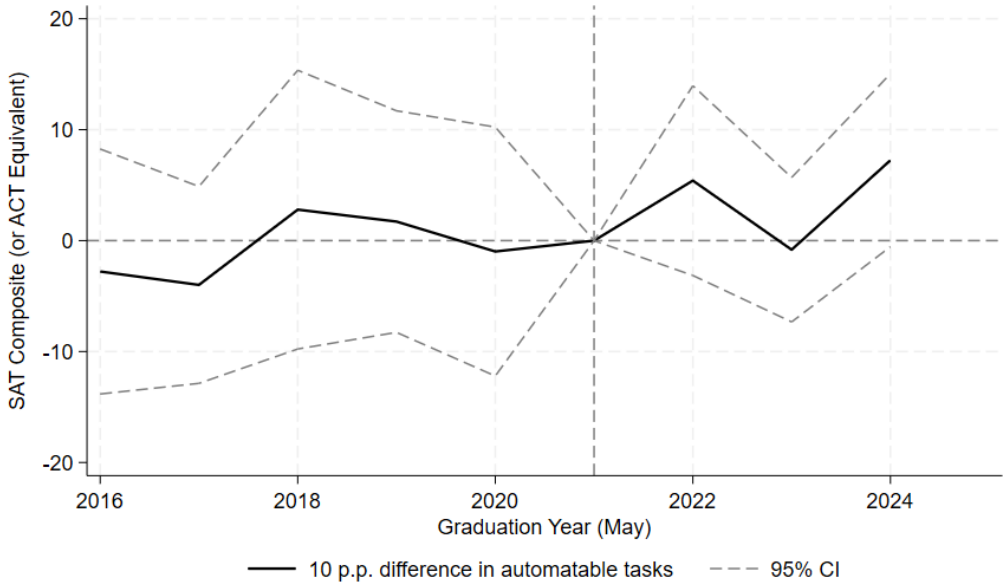
(a) Earnings



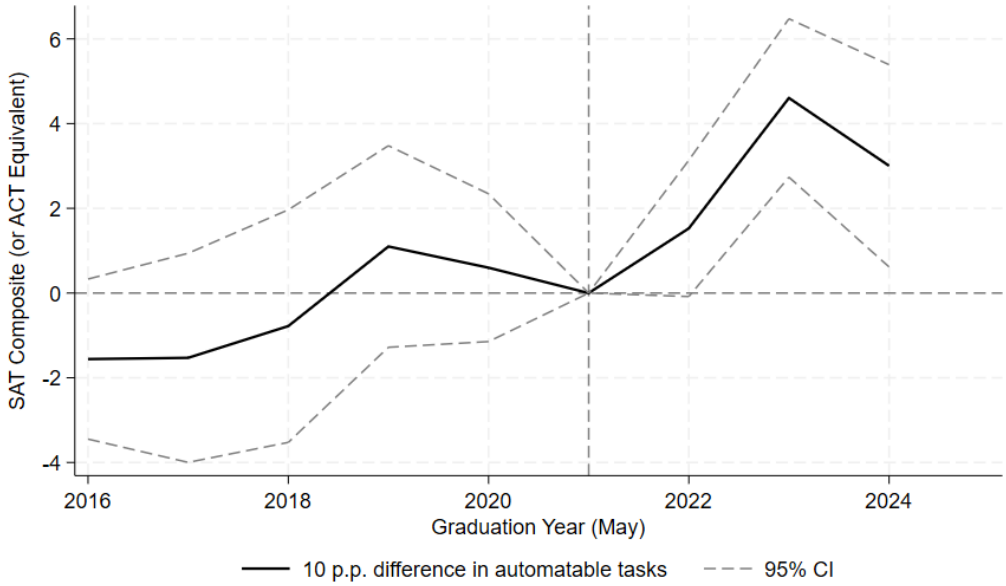
(b) Earnings, Conditional on Working Every Qtr

Figure B5: 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 B6: 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. Standard errors are clustered by major. The sample is limited to only students that enroll in graduate school within a year of completing undergraduate studies in panel (a); in panal (b) graduate school enrollees are dropped. Standard errors are clustered by undergraduate major.

Table B1: 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 B2: 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.