The impacts of AI exposure on the computing pipeline: an empirical analysis of enrollment and unemployment

by

Christian Item Lutz

B.S., Kansas State University, 2025

A THESIS

submitted in partial fulfillment of the requirements for the degree

MASTER OF ARTS

Department of Economics
College of Arts and Sciences

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2026

Approved by:

Major Professor
Hugh A. Cassidy

Copyright

© Christian Lutz 2026.

Abstract

As artificial intelligence (AI) continues to permeate the labor market, from simplifying tasks to generating code, it could affect the job market for recent U.S. graduates looking for an entry-level occupation. Recent, younger graduates may become more replaceable by AI than workers with more experience, who have accumulated human capital such as firm-specific skills. This thesis therefore examines whether unemployment rose more after 2022 for computing graduates (computer science and computer engineering, CS/CE) than for all graduates of all other majors, and whether the increase is larger for younger graduates. It first documents how enrollment by field of study has changed in the AI era, and then turns to labor market outcomes. Computing degrees and enrollment grew through 2024 and enrollment fell in 2025. Using the American Community Survey (ACS) as a data source from 2010 to 2024 (excluding 2020) for bachelor's degree holders, I estimate with a triple-difference model by age group (22-25, 26-30, and 31-40), with 2023-2024 defined as the post-AI period and all earlier years as the baseline. Unemployment among computing graduates aged 22-25 rose 3.5 points more than among graduates of other majors. That increase is 2.1 points larger than for ages 26-30, and 2.2 points larger than for ages 31-40. The age pattern for non-employment appears to be weaker and does not hold against the oldest group. Among employed graduates, young computing graduates did not shift out of tech occupations, although they were somewhat more likely to be in jobs that do not typically require a bachelor's degree. These results are consistent with AI reducing the labor market prospects for young computing graduates, and they may help explain the decline in computing enrollment in 2025.

Table of Contents

List of Figures vi
List of Tables vii
Chapter 1 - Introduction 1
Chapter 2 - Background and Literature Review 4
Background 4
Literature Review 5
Technology and the demand for tasks 5
Generative AI and early-career employment 6
The technology hiring cycle 7
Entering the labor market under weak conditions 7
College major, enrollment, and labor market signals 8
Contribution 8
Data 8
Variable definitions 9
Occupation AI exposure 10
Occupation measures for job mismatch 12
Sample sizes 13
Raw unemployment rates 13
Hypotheses 14
Chapter 3 - Methodology and Modeling 15
Job-mismatch specifications 16
Concerns for validity 17
Chapter 4 - Results 18
Part A: Enrollment by field of study 18
Longer-run trend in computing degrees (ACS) 18
Undergraduate enrollment in computing fields (NSC) 19
Part B: Labor market outcomes 22
Main results: computing graduates relative to all other majors 22
Within computing: computer science relative to computer engineering 24

Job mismatch among employed computing graduates 25
Sensitivity to the baseline period 28
Computing relative to the distribution across majors 29
Chapter 5 - Discussion 31
Chapter 6 - Conclusion 36
References 37

List of Figures

Figure 2.1. Average AI exposure of occupations held by employed graduates, by degree field (50 largest majors; computing highlighted) 11

Figure 4.1. Share of bachelor's degree holders aged 22-25 with a computer science degree, a computer engineering degree, or either (CS + CE), by survey year (2020 excluded; shaded = post-AI years) 19

Figure 4.2. Undergraduate enrollment in computer science, computer engineering, and both combined, fall 2020-2025 (fall 2022 = 100; shaded = post-AI years) 20

Figure 4.3. Share of employed CS and CE graduates working in the ten most common computing-graduate occupations, by age group, 2015-2024 (dashed line = start of the post-AI period) 28

Figure 4.4. Post-2022 change in the unemployment gap at ages 22-25 for 165 other majors, each treated as the comparison group of interest (red line = computing, 3.45 points) 30

List of Tables

Table 2.1. Variable definitions 10
Table 2.2. Ten most common occupations of employed computing (CS and CE) graduates, 2015-2019 and 2021-2022 12
Table 2.3. Number of observations by age group, field group, and period 13
Table 2.4. Unemployment rate (%) by field group and age group, 2010-2022 baseline versus 2023-2024 14
Table 4.1. Undergraduate enrollment in computer science and computer engineering at four-year institutions, fall 2020-2025 20
Table 4.2. State panel: undergraduate enrollment in computer and information sciences relative to fall 2022 (log points) 21
Table 4.3. Post-2022 change in the unemployment gap between computing graduates and graduates of other majors, by age group 23
Table 4.4. Post-2022 change in the CS-CE unemployment gap: age contrasts 25
Table 4.5. Post-2022 change in job mismatch for employed computing graduates relative to graduates of other majors, by age group 26
Table 4.6. Unemployment rates (%) and differences-in-differences under two baseline periods 29

Chapter 1 - Introduction

Since the release of generative artificial intelligence (AI) tools in late 2022 such as the release of ChatGPT, there has been growing public concern that entry-level jobs in tech-heavy fields could be replaced. Computer science (CS) and computer engineering (CE) graduates have described a difficult U.S. labor market (Dickler, 2025), but it is unclear whether this reflects simple tasks being replaced by AI, a tech-sector hiring freeze, or some form of ordinary variation across cohorts. Unemployment among recent college graduates was about 5.6 percent in 2026:Q2 (Federal Reserve Bank of New York, 2026).

This thesis asks whether unemployment rose more after 2022 for computing graduates (CS and CE combined) than for graduates of all other majors, and whether that difference depends on age. I look at an age gradient because recent, younger graduates should be more replaceable by AI than workers with more experience, who have accumulated more human capital, such as firm-specific skills, that AI cannot easily substitute for. The unemployment gap between computing graduates and other graduates rose after 2022, and the increase is significantly larger for ages 22-25 than for ages 26-30 and 31-40.

The thesis has two parts. Because AI might be able to replace workers, especially in computing jobs, students may have shifted away from computing majors, so Part A asks whether the number of CS and CE majors has fallen, both among recent graduates (ACS) and among students who are currently enrolled (National Student Clearinghouse, NSC). Part B asks whether labor market outcomes actually worsened: whether unemployment and non-employment rose more for computing graduates than for graduates of other fields, and how this differs by age. It also asks whether the job match of employed computing graduates changed in the post-AI

period, that is, whether they became more likely, relative to graduates of other majors, to work outside the occupations that are typical for computing graduates.

Young graduates looking for entry-level jobs are a natural place to look for the early effects of AI. Recent graduates are typically assigned well-defined tasks such as writing code, debugging, and documentation, which generative AI tools can perform readily (Eloundou et al., 2023), while older workers perform other tasks that insulate them from the risk of displacement. If AI is reducing the demand for this type of work, then the effect should appear first among the youngest workers, and it should fade with age as workers gain experience and transition through job promotions or specialization. Recent evidence is consistent with this pattern. Using payroll data, Brynjolfsson et al. (2025) find that employment for workers aged 22-25 in AI-exposed occupations is 19% below where it would be had it kept pace with less-exposed occupations, with no comparable gap for experienced workers.

I also examine what the difference does and does not indicate. The remainder of the thesis tests whether the pattern holds within computing, by comparing CS with CE, and whether employed computing graduates moved into different kinds of jobs. This contributes a focused, age-specific comparison of computing graduates against all other graduates that removes shocks common to all majors in each age group and year. I show where the post-2022 increase in unemployment appears and discuss how far it can be linked to AI.

To answer these questions, I use two data sources. For Part A, I use the ACS to measure the share of young bachelor's degree holders with a computing degree, and enrollment counts by major from the NSC to measure current undergraduate enrollment. For Part B, I use the ACS for 2010 to 2024 (excluding 2020), covering bachelor's degree holders aged 22-40, and estimate a triple-difference model that compares computing graduates with graduates of other majors, by

age group, before and after 2022. First, the main results indicate that computing degrees and enrollment grew through 2024, and then fell in 2025, almost entirely in computer science. Second, unemployment rose more for computing graduates than for graduates of other majors in every age group, by 3.45 percentage points at ages 22-25, compared with 1.38 points at ages 26-30 and 1.29 points at ages 31-40. Third, among employed computing graduates I find little evidence indicating that jobs moved away from computing, although young graduates were somewhat more likely to work in jobs that do not typically require a bachelor's degree. These results are consistent with AI reducing the labor market prospects of computing graduates, especially younger ones, and they may help rationalize the fall in computing enrollment in 2025.

Chapter 2 provides the background, reviews the literature, describes the data, and states the hypotheses. Chapter 3 presents the methodology and models. Chapter 4 presents the results. Chapter 5 discusses the implications of the results. Chapter 6 concludes and suggests directions for future work.

Chapter 2 - Background and Literature Review

Background

The initial question stems from the public concern that AI could displace entry-level work (Massenkoff & McCrory, 2026). Recent graduates are the group most likely to show effects because employers hire entry-level workers for well-defined tasks, which AI tools perform best. Older workers may be insulated from this effect because they have more experience and have moved into roles that require judgment.

The empirical evidence appears to be mixed. Massenkoff and McCrory (2026) find no systematic rise in unemployment among highly exposed workers since late 2022 and flat unemployment for workers aged 22-25 in exposed occupations, though they find suggestive evidence of slower hiring. The Yale Budget Lab (Gimbel et al., 2025) finds no discernible economy-wide disruption and notes that the faster widening of the occupational gap between recent and older graduates could reflect AI or a slowing labor market. By contrast, Brynjolfsson et al. (2025) find declines in junior employment concentrated in exposed occupations and at adopting firms. These studies classify workers by occupation or employer. This thesis classifies graduates by college major and measures unemployment directly, which includes graduates who are not working in any occupation.

To measure how exposed younger graduates are to AI, I use the exposure scores from "GPTs are GPTs" (Eloundou et al., 2023), which rate each occupation by the share of its tasks that a large language model (LLM) could perform. In that paper, about 80% of U.S. workers could have at least 10% of their tasks affected by LLMs, and about 19% could have at least half. In the major-level measure used in this thesis, CS and CE both rank approximately one standard deviation above the average major, with almost identical scores (1.090 for CS and 1.083 for CE).

Because these two fields are closely related, with similar coursework and overlapping occupations, I combine them into a single computing group. Chapter 4 later compares CS with CE to show whether the post-AI (2023-2024) rise in unemployment is concentrated in one of them.

Literature Review

This thesis draws on four strands of research: how technology changes the demand for tasks, early evidence on generative AI and employment, the cycle in technology hiring, and the consequences of entering the labor market under weak conditions. I also use work on how students choose a major to interpret the enrollment results in Part A.

Technology and the demand for tasks

A long literature treats jobs as bundles of tasks and asks which tasks a new technology can perform. Autor, Levy, and Murnane (2003) argue that computer capital substitutes for a limited and well-defined set of human activities, the routine cognitive and manual tasks that follow explicit rules, while it complements non-routine problem-solving and interactive tasks. Using task measures from 1960 to 1998, they find that computerization was associated with falling relative demand for routine tasks and rising demand for non-routine cognitive tasks within industries, occupations, and education groups. They estimate that task shifts within industries and occupations explain roughly 30 to 40 percent of the rise in the relative demand for college-educated workers between 1970 and 1998. Their findings imply that college graduates gained from earlier waves of automation because much of their work was non-routine.

Generative AI challenges that assumption because it performs tasks such as writing, summarizing, and programming, which were long thought to be non-routine and cognitive. The occupation-level exposure scores of Eloundou et al. (2023) are the basis for the exposure

measures in this thesis. A score describes what a large language model could do and not whether employers use it, so an exposure effect on employment requires that firms adopt the technology and change how they hire. This is why the analysis looks for differences by age and by major and not only for a general effect.

Generative AI and early-career employment

Experimental work shows that generative tools can raise productivity in tasks that computing graduates perform. In a controlled experiment, Peng et al. (2023) find that developers with access to GitHub Copilot completed a programming task 55.8% faster than a control group. Faster work does not by itself reduce employment, because demand for output can rise, but it is one way in which fewer junior workers could be needed for a given workload.

Evidence on actual labor market effects is mixed and depends on the data and the unit of analysis. Brynjolfsson et al. (2025), using payroll data, find the 19% relative decline for young workers in exposed occupations described in Chapter 1, and no comparable decline for experienced workers. Hosseini and Lichtinger (2025) use résumé and job posting data on about 62 million workers and find that, from the first quarter of 2023, junior employment at firms that adopt generative AI fell relative to non-adopters while senior employment kept rising. Junior employment at adopters was about 7.7% lower six quarters later, the decline came from slower hiring and not from more separations, and it was concentrated in high-exposure occupations. Massenkoff and McCrory (2026) and Gimbel et al. (2025), discussed above, find little aggregate disruption, with suggestive evidence of slower hiring of young workers and an occupational shift that began before generative AI was widely used.

These findings fit together if AI affects the hiring of junior workers more than the total employment of existing workers. That is the margin this thesis studies, by looking at unemployment and comparing recent graduates with older graduates.

The technology hiring cycle

Any AI effect on computing graduates is hard to separate from the broader contraction in technology hiring. Indeed Hiring Lab (2025) reported in July 2025 that U.S. technology job postings remained about 36 percent below their pre-pandemic level. This contraction coincides in time with the release of generative AI tools in late 2022, so a rise in unemployment among computing graduates after 2022 is consistent with both explanations. The empirical strategy therefore compares computing graduates with graduates of other majors in the same age group and year. The finding of Gimbel et al. (2025) that the economy-wide occupational mix has changed little is a reminder that a contraction in the technology sector need not imply economy-wide effects of AI.

Entering the labor market under weak conditions

A weak market at graduation can have costs that outlast the first spell of unemployment. Kahn (2010) finds that graduating from college in a bad economy has large and persistent negative effects on wages. Oreopoulos, von Wachter, and Heisz (2012) follow Canadian graduates and find large initial earnings losses that fade over about eight to ten years, in part because affected graduates start at lower-paying firms and move to better ones over time. If young computing graduates faced a weak market in 2023 and 2024 for any reason, this literature suggests that the cost of the initial unemployment could extend beyond it. This thesis measures only initial outcomes and cannot say how long any effect will last.

College major, enrollment, and labor market signals

Earnings differ substantially across college majors, and the choice of major is an important determinant of careers (Altonji, Blom, & Meghir, 2012). If students respond to the signals they receive about job prospects, a weak entry-level market for computing graduates could reduce later enrollment in computing, which is the pattern Part A examines for fall 2025. The ACS records field of degree only for graduates, so it can track the longer-run trend in degrees but not enrollment by major, which requires administrative data such as that of the National Student Clearinghouse.

Contribution

This thesis adds to this work in three ways. First, it compares recent computing graduates with older computing graduates and with graduates of other majors in the same age group, which separates an age pattern from a general effect on computing. Second, it examines the occupations of the graduates who are employed, to see whether the adjustment happens through unemployment or through job mismatch. Third, it connects the labor market outcomes to the supply of computing graduates using ACS data on degrees and Clearinghouse data on enrollment. The main limit is that all of the evidence is observational and covers only two years after 2022.

Data

I use the one-year American Community Survey (ACS) from IPUMS USA (Ruggles et al., 2025) for 2010-2024, excluding 2020 because of data quality problems stemming from the disruption to survey collection during the COVID-19 pandemic. The ACS is a household survey that covers about one percent of the U.S. population each year. It records educational attainment, field of degree for people with a bachelor's degree, age, labor force status, enrollment, and

occupation. I keep persons aged 22-40 with a bachelor's degree or higher who report a field of degree, and I drop persons whose employment status is not applicable. The sample therefore includes graduates with advanced degrees as well as those whose highest degree is a bachelor's, and the field of degree refers to the bachelor's degree. I exclude ages 20-21 because many are still enrolled in undergraduate programs. The age groups are 22-25, 26-30, and 31-40. The post-AI period is 2023-2024, the first full survey years available after the release of generative AI tools in late 2022, and all earlier years (2010-2019 and 2021-2022) form the baseline. The main sample has 3,262,349 observations for unemployment and 3,644,350 for non-employment. Regressions are weighted by the ACS person weights.

Part A uses two sources. The first is the ACS, which records field of degree only for people who hold a bachelor's degree. I use bachelor's degree holders aged 22-25, whether or not they are still enrolled in school, to measure the share with a computer science or computer engineering degree in each survey year. Because the ACS cannot count current students by major, the second source is enrollment by major from the National Student Clearinghouse Research Center (2026b): undergraduate headcount at four-year institutions for fall 2020 through fall 2025, nationally and by state.

Variable definitions

Table 2.1 lists the variables used. Computing graduates are those whose detailed field of degree is computer science or computer engineering. I use the detailed field code, which separates these two fields from related ones such as electrical engineering and computer and information systems. The two main outcomes are indicators for an individual's labor force status. Unemployed equals one if a person is in the labor force and is looking for work without having a job, and zero if the person is employed; it is defined only for persons in the labor force. Not

employed equals one if a person does not have a job, whether unemployed or out of the labor force, and zero if the person is employed.

Table **2.1**. Variable definitions

Variable	Definition
Computing graduate	Detailed field of degree is computer science (DEGFIELDD 2102) or computer engineering (2407). The comparison group is every other field.
Age group	22-25, 26-30, and 31-40.
Post-AI period	Survey years 2023 and 2024. The baseline is 2010-2019 and 2021-2022.
Unemployed	Unemployed (EMPSTAT) among persons in the labor force; missing for persons not in the labor force.
Not employed	Unemployed or not in the labor force.
Enrolled in school	Currently attending school (SCHOOL).
Tech occupation	Computer and mathematical occupations (SOC 15) and computer and information systems managers (SOC 11-3021).
Top-10 occupations	The ten most common occupations of employed CS and CE graduates before 2023 (2015-2019 and 2021-2022).
Non-college occupation	An occupation in which fewer than half of workers aged 25-40 hold a bachelor's degree (2015-2019).
Exposure of the occupation held	Standardized AI-exposure score of the graduate's own occupation.

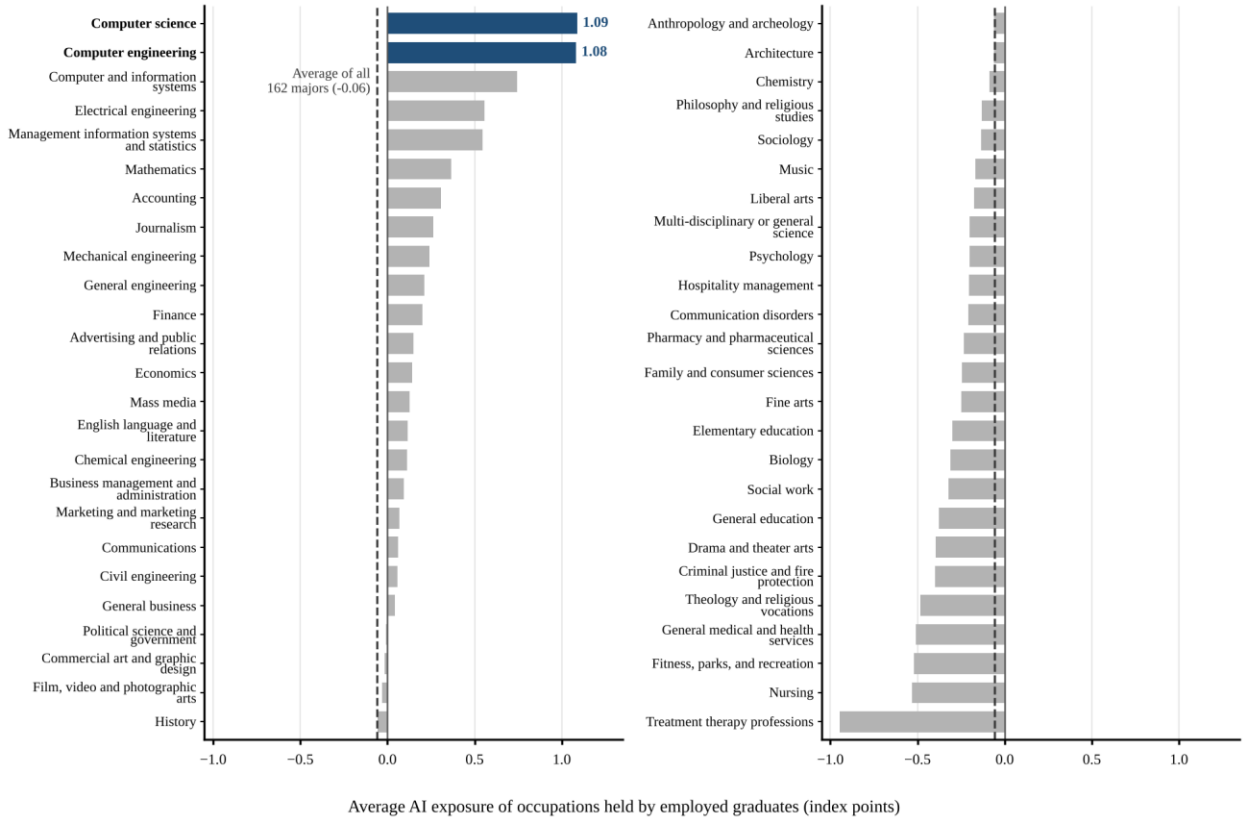
Notes: SOC = Standard Occupational Classification. IPUMS variable names are in capitals.

Occupation AI exposure

One of the job-mismatch outcomes uses the AI exposure of the graduate's own occupation. Occupation exposure comes from Eloundou et al. (2023), who rate the tasks in each occupation by whether a large language model could reduce the time needed to complete them. I merge their scores to the ACS occupation codes with a crosswalk, using separate versions before and after 2018 because the ACS occupation classification changed that year. I standardize the score to mean zero and standard deviation one among employed graduates aged 22-40 in 2012-

2019, so “index points” in this thesis are standard deviations of that score. Computing graduates hold much more exposed occupations than graduates of other majors: before 2023, the average score of employed graduates aged 22-25 was 1.17 for computing graduates and -0.14 for graduates of other majors (Chapter 4). Figure 2.1 condenses the 162 majors in the sample to the 50 largest and shows the average score for each degree field. Computer science (1.09) and computer engineering (1.08) have the highest exposure among them, more than one index point above the average major, followed by computer and information systems (0.74) and electrical engineering (0.56).

Figure **2.1**. Average AI exposure of occupations held by employed graduates, by degree field (50 largest majors; computing highlighted)



Notes: Each bar is the average standardized AI exposure score (Eloundou et al., 2023) of the occupations held by employed graduates of that degree field, ages 22-40, 2012-2019; the score has mean zero and standard deviation one among employed graduates. Degree fields with fewer than 5,000 observations are omitted. The dashed line is the average across all 162 degree fields. Majors are ranked from highest to lowest exposure, first down the left panel and then down the right panel.

Occupation measures for job mismatch

The job-mismatch outcomes use the occupation of employed graduates. I classify a computing graduate as matched if they work in one of the ten most common occupations among employed computing graduates, and as mismatched if they work in any other occupation. Table 2.2 lists the ten most common occupations of employed computing graduates before 2023, which together account for 59.8% of them. Software developers alone account for 35.3%. Occupation codes differ across the earliest pooled ACS samples, and the share of computing graduates in these occupations shows an abrupt change between 2014 and 2015 for older workers that I attribute to coding and not to behavior. The mismatch outcomes therefore use survey years 2015 through 2024 and the reference list is built from 2015-2019 and 2021-2022.

Table **2.2**. Ten most common occupations of employed computing (CS and CE) graduates, 2015-2019 and 2021-2022

Rank	SOC code	Occupation	Share (%)
1	15-1252	Software developers	35.3
2	15-1251	Computer programmers	5.1
3	15-1299	Computer occupations, all other	3.5
4	11-91XX	Management occupations, grouped code	3.1
5	11-3021	Computer and information systems managers	2.6
6	15-1230	Computer support specialists	2.6
7	15-1211	Computer systems analysts	2.6
8	25-1000	Postsecondary teachers	1.7
9	17-21YY	Engineering occupations, grouped code	1.7
10	15-1253	Software quality assurance analysts and testers	1.6
		Total, top ten	**59.8**

Notes: Shares are weighted shares of employed CS and CE graduates aged 22-40. Codes ending in X or Y are grouped codes in the IPUMS occupation classification.

Sample sizes

Computing graduates are a small share of the sample: 124,959 of 3,644,350 observations, or 3.4%. Table 2.3 shows the number of observations by age group, field group, and period. The estimates for ages 22-25 rest on 15,323 computing graduates in the baseline and 5,128 in 2023-2024. The 2017-2019 column is the shorter baseline used in the sensitivity analysis in Chapter 4.

Table **2.3**. Number of observations by age group, field group, and period

Age group	Field group	2010-2022	2017-2019	2023-2024
22-25	Computing (CS/CE)	15,323	4,488	5,128
	All other majors	465,338	123,486	90,222
26-30	Computing (CS/CE)	25,909	6,749	7,285
	All other majors	794,707	215,751	155,333
31-40	Computing (CS/CE)	58,482	16,516	12,832
	All other majors	1,660,065	448,237	353,726

Notes: Counts are unweighted observations in the main sample (persons with a bachelor's degree or higher and a reported field of degree). The 2010-2022 column excludes 2020 and the 2017-2019 column is a subset of it.

Raw unemployment rates

Table 2.4 compares raw unemployment rates in 2010-2022 (excluding 2020) with 2023-2024. Unemployment rose for computing graduates in every age group, most sharply at ages 22-25 (2.5 points), while it fell for graduates of other majors. The difference between the two changes is 3.7 points at ages 22-25, 1.4 points at ages 26-30, and 1.3 points at ages 31-40, which closely match the regression estimates in Chapter 4. These numbers are suggestive of a change in unemployment among young computing graduates after the introduction of AI that differs from that of all other majors, which I test more formally in Chapter 4.

Table **2.4**. Unemployment rate (%) by field group and age group, 2010-2022 baseline versus 2023-2024

	Computing (CS/CE)			All other majors		
Age group	**2010-22**	**2023-24**	**Change**	**2010-22**	**2023-24**	**Change**
22-25	6.7	9.3	+2.5	6.4	5.2	-1.2
26-30	3.5	4.1	+0.7	3.5	2.7	-0.8
31-40	2.5	3.2	+0.7	2.9	2.3	-0.6

Source: ACS 2010-2024 (IPUMS), author's calculations; 2020 excluded; 2021-2022 included in the baseline. Unemployment is the share of the labor force that is unemployed. Changes are in percentage points and computed from unrounded rates.

Hypotheses

Based on the argument above, I test two hypotheses. H_1: unemployment rose more for computing graduates than for graduates of other majors after the introduction of AI, so the post-AI change in the unemployment gap is positive in all three age groups. This follows from the high exposure of computing occupations to AI (Figure 2.1). H_2: the increase is greater for younger, more recent graduates than for older workers, which follows from the argument that experience and firm-specific skills protect older workers. The same hypotheses are tested for non-employment as a secondary outcome.

Chapter 3 - Methodology and Modeling

This chapter presents the methods in the order of the results in Chapter 4: the enrollment model used in Part A, followed by the triple-difference model used in Part B.

The enrollment analysis in Part A uses a model, estimated on a state-by-year panel of undergraduate enrollment in computer and information sciences (CIP 11) from the NSC for Fall 2020 through Fall 2025:

$$ln\, E_{s,t} \;=\; \alpha_s \;+\; \delta_t \;+\; \varepsilon_{s,t} \qquad (1)$$

where $E_{s,t}$ is undergraduate enrollment in CIP 11 at four-year institutions in state s in fall of year t, α_s are state fixed effects, and δ_t are year effects relative to Fall 2022 ($\delta_{2022} \;=\; 0$). The regression is weighted by each state's 2020 enrollment and clusters standard errors by state. The coefficients δ_t are the log-point differences that are reported in Table 4.2 for enrollment relative to Fall 2022.

For Part B, I estimate the following linear probability model on the sample of bachelor's degree holders from all fields of study:

$$Y_i \;=\; \alpha_d \;+\; \gamma_{g,t} \;+\; \sum_g \beta_g(\mathbb{1}[AgeGroup_i = g] \times CP_i) \;+\; \sum_g \theta_g(\mathbb{1}[AgeGroup_i = g] \times CP_i \times Post_t) \;+\; \varepsilon_i \qquad (2)$$

where Y_i indicates that person i is unemployed (among the labor force) or not employed; α_d are degree-field fixed effects; $\gamma_{g,t}$ are age-group-by-year fixed effects; CP_i equals 1 if the person's field of study is computing (CS or CE), and 0 for every other field; and $Post_t$ equals 1 for 2023 and 2024 and 0 otherwise.

The coefficients β_g capture the baseline gap between computing graduates and graduates of other majors in age group g. The coefficients of interest are θ_g, each of which is the change after 2022 in the unemployment gap between computing graduates and graduates of other

majors, for age group g. A value of zero means that unemployment among computing graduates moved in line with other majors, and a positive value means it rose by more. I estimate θ_1 (ages 22-25), θ_2 (ages 26-30), and θ_3 (ages 31-40) separately so that the youngest group can be compared directly with each older group.

In terms of the model, the hypotheses are H_1: $\theta_1 > 0$, $\theta_2 > 0$ and $\theta_3 > 0$; and H_2: $\theta_1 > \theta_2$ and $\theta_1 > \theta_3$, where the first is tested with each coefficient separately and the second with linear combinations of the coefficients ($\theta_1 - \theta_2$ and $\theta_1 - \theta_3$). Standard errors are clustered by survey year. The age-group-by-year fixed effects absorb shocks common to all majors in each age group and year, so θ_g is identified from the differential change for computing graduates, not from overall labor market conditions. This assumes that, without the post-2022 shock, computing and other majors would have followed parallel paths within each age group.

Two supporting specifications follow. First, I restrict the sample to CS and CE graduates and repeat the model (2) with CS in place of the computing group, to see how the increase is split within the computing group itself. Second, I estimate the same model for job-mismatch outcomes among employed graduates: working outside the ten most common occupations of computing graduates, working in a tech occupation, working in a non-college occupation, and the AI exposure of the occupation held (survey years 2015-2024).

Job-mismatch specifications

The job-mismatch models use the same triple-difference specification, estimated on employed graduates in survey years 2015-2024. The outcomes are an indicator for working outside the ten most common occupations of computing graduates, an indicator for working in a tech occupation, an indicator for working in a non-college occupation, and the standardized AI-exposure score of the occupation held. Because the sample is restricted to the employed, these

estimates describe how jobs changed among those who had jobs, and not how the probability of any job changed.

Concerns for validity

The triple-difference design identifies the post-2022 change for computing graduates relative to other majors in the same age group and year, so it requires that computing and other majors would have moved in parallel within each age group without the post-2022 shock. Four concerns follow. First, the composition of computing graduates changed: their share among young graduates rose from 2.2% to 5.6% between 2010 and 2024, so the marginal graduate may differ from earlier graduates. Second, the age groups in the post-period are different cohorts, so an age pattern could reflect cohort differences and not experience. Third, the mismatch outcomes describe only those who are employed, so selection into employment can bias them. Fourth, the post-period has only two survey years. I address the first two concerns with the not-enrolled sample, and with comparisons across baselines, and I discuss the others where they affect interpretation.

Chapter 4 - Results

Part A: Enrollment by field of study

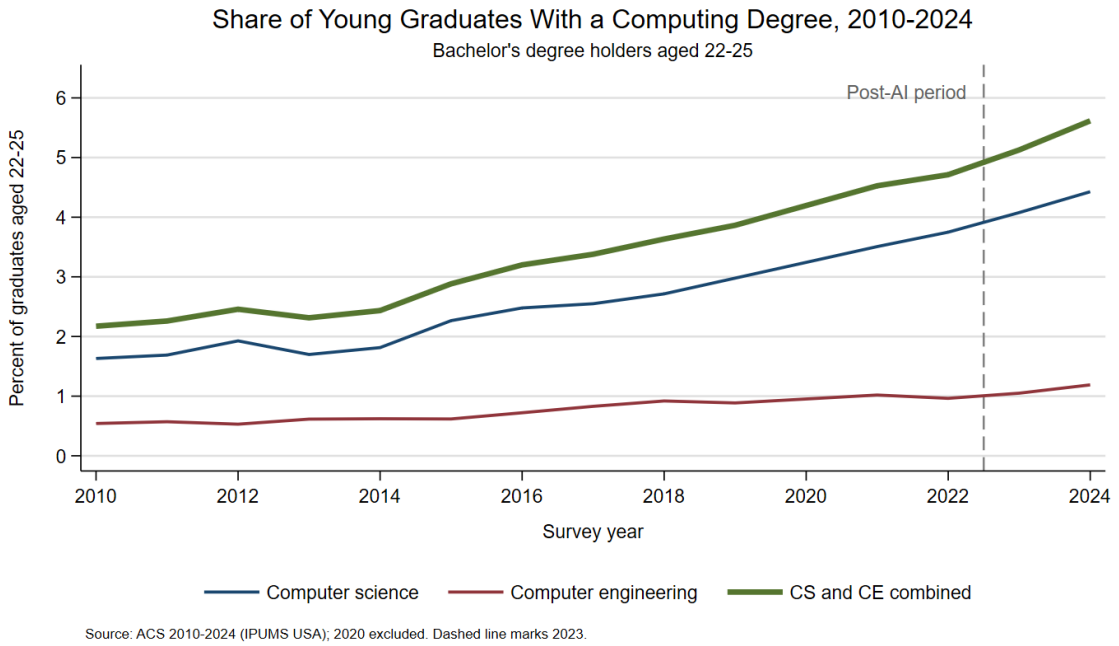
This section describes how enrollment in computing fields has changed over time. If AI is expected to replace computing workers, students may respond by choosing other majors, so I ask whether the number of computing majors has fallen. I first use the ACS to show the longer-run trend in the share of young graduates with a computing degree, and then use National Student Clearinghouse enrollment data to show the trend in current undergraduate enrollment, including a recent decline.

Longer-run trend in computing degrees (ACS)

The ACS records field of degree only for people who already hold a bachelor's degree, so it cannot count students who are still enrolled by major. It can show how many young graduates have a computing degree. Among bachelor's degree holders aged 22-25, whether or not they are still enrolled in school (for example in graduate programs), the share with a computer science degree rose from 1.6% in 2010 to 3.0% in 2019 and 4.4% in 2024. The computer engineering share rose from 0.5% to 0.9% to 1.2% (Figure 4.1). Combined, the share of young graduates with a computer science or computer engineering degree rose from 2.2% in 2010 to 3.9% in 2019 and 5.6% in 2024, and the pattern is similar among graduates who are not enrolled in school (2.2% in 2010 to 6.0% in 2024). The share of young graduates with a computing degree therefore grew steadily over the full period and kept rising through 2024.

This series is unlikely to capture much of an AI effect. Graduates in 2023 and 2024 chose their majors before generative AI tools were released, and because the ACS observes the major only after graduation, it does not capture students who are still enrolled. I use it to document the longer-run growth in computing, which motivates the labor market analysis in Part B.

Figure **4.1**. Share of bachelor's degree holders aged 22-25 with a computer science degree, a computer engineering degree, or either (CS + CE), by survey year (2020 excluded; shaded = post-AI years)



Source: ACS 2010-2024 (IPUMS USA); 2020 excluded. Dashed line marks 2023.

Undergraduate enrollment in computing fields (NSC)

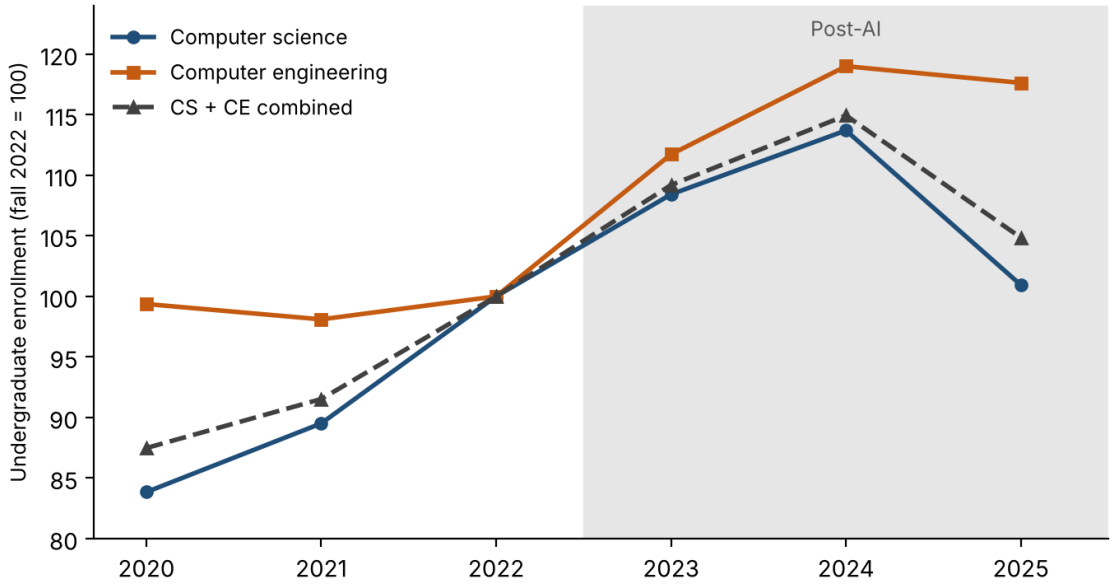
To measure current enrollment by major, I use National Student Clearinghouse Research Center (NSC) counts of undergraduates at four-year institutions by field of study, for fall 2020 through fall 2025 (Table 4.1 and Figure 4.2). Enrollment in computer science and computer engineering combined rose every year from 2020 to 2024, from 233,583 to 306,924 students, which raised the combined share of all undergraduates from 2.43% to 3.15%. This upward trend matches the ACS series in Figure 4.1. Enrollment then fell in fall 2025: the combined count dropped 8.8% to 279,968 (2.86% of undergraduates), almost entirely in computer science (-11.2%), while computer engineering fell only 1.2%. Computer science enrollment in fall 2025 was back at its 2022 level.

Table **4.1**. Undergraduate enrollment in computer science and computer engineering at four-year institutions, fall 2020-2025

Fall	Computer science	Computer engineering	CS + CE	CS + CE as % of undergraduates	Change in CS + CE
2020	171,472	62,111	233,583	2.43	n/a
2021	183,068	61,319	244,387	2.59	+4.6%
2022	204,520	62,505	267,025	2.85	+9.3%
2023	221,745	69,837	291,582	3.09	+9.2%
2024	232,539	74,385	306,924	3.15	+5.3%
2025	**206,442**	**73,525**	**279,968**	**2.86**	**-8.8%**

Notes: Source is National Student Clearinghouse Research Center (2026b), fall 2025 enrollment trends, CIP appendix. Enrollment is undergraduate headcount at four-year institutions by 4-digit CIP code (Computer Science, 11.07; Computer Engineering, 14.09). The share is combined computer science and computer engineering enrollment divided by all undergraduate enrollment at four-year institutions.

Figure **4.2**. Undergraduate enrollment in computer science, computer engineering, and both combined, fall 2020-2025 (fall 2022 = 100; shaded = post-AI years)



To check that the 2025 decline is not driven by a few states, I estimate a state-by-year panel of undergraduate enrollment in computer and information sciences (CIP 11, since the state file does not separate computer engineering). The model equation (1) from Chapter 3 includes state fixed effects and year effects relative to fall 2022, is weighted by 2020 enrollment, and clusters standard errors by state (49 states; Table 4.2). Enrollment was 7.2 log points higher than

in 2022 in fall 2023 and 9.7 log points higher in fall 2024 (both $p < 0.001$), and 0.8 log points higher in fall 2025 ($p = 0.65$), so enrollment in 2025 returned to its 2022 level. The decline from 2024 to 2025 is 8.8 log points, about 8.5% ($p < 0.001$), and enrollment fell in 47 of 50 states. The share of each state's undergraduates enrolled in computer and information sciences fell by 0.59 percentage points between 2024 and 2025 ($p < 0.001$).

Table **4.2**. State panel: undergraduate enrollment in computer and information sciences relative to fall 2022 (log points)

Fall (relative to 2022)	**Estimate (log points)**	**Std. error**	**p-value**
2020	-11.2	1.4	<0.001
2021	-7.9	0.7	<0.001
2023	7.2	1.0	<0.001
2024	9.7	1.6	<0.001
2025	0.8	1.8	0.650
2025 minus 2024	**-8.8**	**1.0**	**<0.001**

Notes: The outcome is the log of undergraduate enrollment in CIP 11 at four-year institutions, with state fixed effects, weighted by 2020 enrollment, with standard errors clustered by state. Source is National Student Clearinghouse Research Center (2026b).

This is suggestive evidence of a recent decline, not evidence about its cause. The decline appears in a single post-2022 year, and the data start in 2020, so there is little pre-period. A weaker tech hiring market could produce the same pattern as AI. More recent NSC reporting points the same way: computer science enrollment fell at every award level in fall 2025, and in spring 2026 it fell 8.4% at four-year institutions, 9.3% at primarily associate-baccalaureate institutions, and 11.2% at two-year colleges, while total undergraduate enrollment grew 1.3% (National Student Clearinghouse Research Center, 2026a). Computer science also had the lowest second-fall persistence among the ten largest majors (85.0% for the fall 2024 entering cohort, down 1.3 points from the fall 2022 cohort), while engineering had the highest (93.1%). These

figures describe computer science and engineering broadly and do not separate computer engineering.

Part B: Labor market outcomes

Computing degrees and enrollment grew rapidly through 2024 and have recently begun to fall. The rest of the chapter turns to labor market outcomes for computing graduates after 2022, using the ACS sample of bachelor's degree holders aged 22-40. Having asked whether students are moving away from majoring in computing, I now ask whether the labor market outcomes for computing graduates have actually worsened. The first hypothesis (H_1) is answered by the three unemployment estimates in Table 4.3, and the second (H_2) by the contrasts between age groups in the same table; the non-employment, within-computing, job-mismatch and robustness results then check how far these answers hold.

Main results: computing graduates relative to all other majors

Table 4.3 reports the estimates of equation (2) for unemployment: the post-2022 change in the unemployment gap for each age group (θ_1, θ_2, and θ_3) and the contrasts between the youngest group and the two older groups. Unemployment among computing graduates aged 22-25 rose by 3.45 points more than among graduates of other majors ($p < 0.001$). The corresponding increases were 1.38 points at ages 26-30 ($p = 0.007$) and 1.29 points at ages 31-40 ($p < 0.001$). All three estimates are positive and statistically significant, consistent with my first hypothesis that unemployment rose more for computing graduates than for graduates of other majors following the introduction of AI. These estimates are close to the raw differences in Table 2.4.

The increase for the youngest group is 2.07 points larger than for ages 26-30 and 2.16 points larger than for ages 31-40 (both $p < 0.001$), which is consistent with my second

hypothesis that the increase is greater for younger graduates than for older ones. The two older groups were not distinguishable from each other with a difference of 0.09 points and p-value of 0.75, so the break in the pattern is between the youngest group and both older age groups.

Thus, my results are consistent with younger computing graduates experiencing a greater level of unemployment after the introduction of AI. These results might help to rationalize the apparently large reduction in CS and CE enrollment in 2025 compared with previous years, since a weaker labor market would dissuade students from entering these computing majors.

Table **4.3**. Post-2022 change in the unemployment gap between computing graduates and graduates of other majors, by age group

	Estimate	**Std. error**	**p-value**
Unemployed (share of labor force); N = 3,262,349			
θ_1: ages 22-25	3.45	0.55	<0.001
θ_2: ages 26-30	1.38	0.43	0.007
θ_3: ages 31-40	1.29	0.30	<0.001
$\theta_1 - \theta_2$	2.07	0.39	<0.001
$\theta_1 - \theta_3$	2.16	0.37	<0.001
Not employed (share of population); N = 3,644,350			
θ_1: ages 22-25	1.72	0.78	0.045
θ_2: ages 26-30	0.07	0.64	0.919
θ_3: ages 31-40	2.01	0.36	<0.001
$\theta_1 - \theta_2$	1.66	0.83	0.066
$\theta_1 - \theta_3$	-0.28	0.88	0.752

Notes: Estimates are in percentage points. Computing = computer science or computer engineering; comparison group = all other fields. Standard errors clustered by survey year. Baseline = 2010-2019 and 2021-2022 (2020 excluded); post-AI = 2023-2024. Weighted by ACS person weights.

The non-employment results are weaker. The contrast between ages 22-25 and 26-30 is 1.66 points, significant at the 10% level. The contrast between ages 22-25 and 31-40 is -0.28 points ($p = 0.75$), because non-employment rose for older computing graduates as well, at ages 31-40 ($\theta_3 = 2.01$, $p < 0.001$). The age pattern is clearer for unemployment than for non-

employment. As a check on the definition for the computing group, I defined CS alone as the treated group, dropped CE, and re-estimated the model. The age contrasts for unemployment remained significant: $\theta_1 - \theta_2 = 1.22$ (p = 0.041) and $\theta_1 - \theta_3 = 1.47$ (p = 0.008). The non-employment contrasts are not significant.

As a further check, I re-estimate the unemployment model on graduates who are not enrolled in school (N = 2,855,263). The age pattern is slightly stronger: $\theta_1 = 3.61$ points ($p < 0.001$), $\theta_2 = 1.10$ points (p = 0.029), and $\theta_3 = 1.33$ points (p = 0.001), with $\theta_1 - \theta_2 = 2.51$ points and $\theta_1 - \theta_3 = 2.28$ points (both $p < 0.001$). The unemployment result is therefore not driven by graduates who are still in school.

Within computing: computer science relative to computer engineering

I combine CS and CE into a single computing group for the main analysis for three reasons: the two degrees prepare students for largely the same occupations, their graduates have almost identical AI exposure (Chapter 2), and CE is much smaller than CS (1.2% of young graduates in 2024, against 4.4% for CS), so separate estimates for CE by age group would be imprecise.

Combining CS and CE into one computing group could hide differences between the two fields, and the main result would be less convincing if it came from only one degree field. As a check on how I define the computing group, Table 4.4 restricts the sample to CS and CE graduates, with CS as the treated group. For ages 22-25, unemployment rose from 7.2% to 9.0% for CS and from 5.2% to 10.2% for CE. The negative contrasts in Table 4.4 therefore reflect a larger increase for CE, not an improvement for CS. The rise in computing unemployment among young graduates is not limited to CS, and it was larger for CE.

Table **4.4**. Post-2022 change in the CS-CE unemployment gap: age contrasts

	Estimate (pp)	Std. error	p-value
Unemployed: $\theta_1 - \theta_2$	-3.72	1.19	0.008
Unemployed: $\theta_1 - \theta_3$	-3.03	0.82	0.003
Not employed: $\theta_1 - \theta_2$	-3.52	1.74	0.064
Not employed: $\theta_1 - \theta_3$	-4.30	2.30	0.084

Notes: Sample restricted to CS and CE graduates (N = 113,897 for unemployment; 124,959 for non-employment). Standard errors clustered by survey year.

Job mismatch among employed computing graduates

Unemployment only captures graduates who are looking for work and do not have a job. To see whether computing graduates who did find work were placed in jobs that fit their degree, I use the occupation of employed workers. The reference list is the one described in Chapter 2 (Table 2.2). Besides software developers, it includes other computer occupations, computer and information systems managers, computer systems analysts, computer support specialists, software testers, and postsecondary teachers, as well as one miscellaneous management group and one engineering group that are not computing occupations but are common destinations. I construct three outcomes for employed workers: working outside these ten occupations, working in a tech occupation (any computer occupation or computer and information systems manager), and working in a non-college occupation, defined as one in which fewer than half of workers aged 25-40 hold a bachelor's degree. Occupation codes are not consistent across the earliest pooled ACS samples, so these results use survey years 2015-2024 and the baseline is shorter than in the main results. I estimate the same triple-difference model as before.

Table 4.5 and Figure 4.3 show little sign that young computing graduates who found jobs moved away from computing work. The share working in a tech occupation did not change for ages 22-25 ($\theta_1 = -0.05$ points, p = 0.96), while it fell by about 2.1 points at ages 26-30 and 1.6 points at ages 31-40. Relative to the older groups, the young were therefore better able to keep

tech jobs (contrasts of 2.06 points, p = 0.005, and 1.50 points, p = 0.033). The share working outside the ten reference occupations rose by 1.4 to 2.2 points in every age group, with no difference by age (p = 0.47 and p = 0.37), so this is not a change specific to young graduates. Among ages 22-25 the share in a top-ten occupation was between 58% and 63% before 2023 and 57.0% in 2024, which is the lowest value for that age group, but the series is noisy and 2023 (58.9%) is within the earlier range.

Table **4.5**. Post-2022 change in job mismatch for employed computing graduates relative to graduates of other majors, by age group

Outcome	**Ages 22-25 (θ_1)**	**Ages 26-30 (θ_2)**	**Ages 31-40 (θ_3)**	$\theta_1 - \theta_2$	$\theta_1 - \theta_3$
Job outside top-10 computing occupations	1.37	2.22**	1.97***	-0.85	-0.60
	(0.82)	(0.86)	(0.52)	(1.11)	(0.63)
Job in a tech occupation	-0.05	-2.11	-1.55	2.06***	1.50**
	(1.02)	(1.20)	(0.98)	(0.53)	(0.58)
Job in a non-college occupation	2.70***	0.33	1.09*	2.37*	1.61*
	(0.65)	(0.95)	(0.53)	(1.08)	(0.71)
AI exposure of occupation held (index points)	0.15	0.13	0.18	0.02	-0.03
	(0.13)	(0.13)	(0.13)	(0.04)	(0.03)

Notes: Estimates are in percentage points; standard errors clustered by survey year are in parentheses. * p < 0.10, ** p < 0.05, *** p < 0.01. Employed bachelor's degree holders aged 22-40, survey years 2015-2024 (2020 excluded), N = 2,206,786. Top-10 occupations are the ten most common occupations of employed CS and CE graduates in 2015-2019 and 2021-2022. Non-college occupation = fewer than half of workers aged 25-40 in the occupation hold a bachelor's degree (2015-2019). The first three rows are in percentage points; the last row is in index points of occupation-level AI exposure (Eloundou et al., 2023) and uses N = 1,429,473 employed workers in occupations with an exposure score.

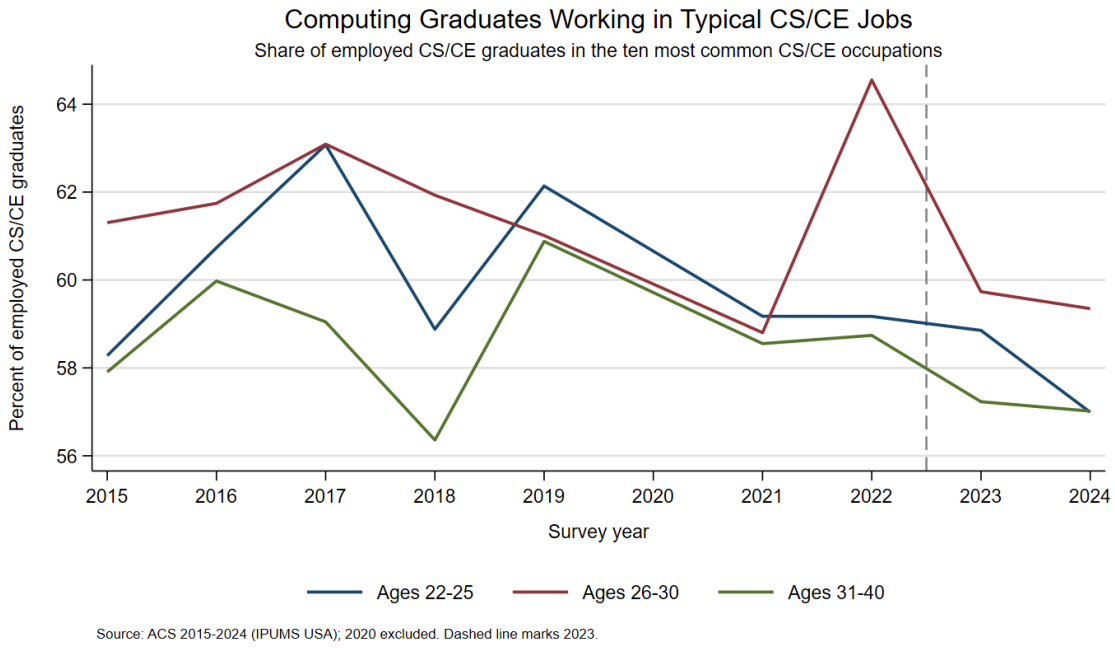
The one measure with a larger change for the youngest group is working in a non-college occupation. It rose by 2.70 points for ages 22-25 (p = 0.003), compared with 0.33 points at ages 26-30 and 1.09 points at ages 31-40. The contrasts are 2.37 and 1.61 points and are significant only at the 10% level (p = 0.060 and p = 0.054). Part of this estimate comes from the other

majors: the share of young computing graduates in non-college jobs rose from 20.4% to 22.3%, while for graduates of other majors it fell from 42.4% to 40.9%.

I also use the AI exposure of the occupation each employed graduate holds as an outcome, to see whether graduates moved into jobs with more tasks that generative AI can perform. I use the occupation-level scores of Eloundou et al. (2023), and the estimates are in index points. As described in Chapter 2, the average score before 2023 was 1.17 for computing graduates aged 22-25, against -0.14 for other majors. After 2022 the score rose for computing graduates in every age group (by 0.13 at ages 22-25, 0.14 at ages 26-30, and 0.21 at ages 31-40) and changed little for other majors (-0.01, 0.04, and 0.04).

The triple-difference estimates in the last row of Table 4.5 are 0.15 at ages 22-25 (standard error 0.13), 0.13 at ages 26-30, and 0.18 at ages 31-40. None is statistically different from zero ($p = 0.27$, 0.35, and 0.21), and the age contrasts are small (0.02 and -0.03; $p = 0.59$ and 0.36). The 95% confidence interval for the youngest group runs from -0.10 to 0.40, so I cannot rule out a moderate increase, relative to a gap of about 1.3 index points between computing and other occupations. The point estimates are not larger for the young, however, and the raw increase was largest at ages 31-40. I therefore find no evidence that young computing graduates moved into more AI-exposed jobs than older ones. Because the exposure of the job held is close to the exposure of the degree, this measure mostly confirms that computing graduates remained in the same kind of work.

Figure **4.3**. Share of employed CS and CE graduates working in the ten most common computing-graduate occupations, by age group, 2015-2024 (dashed line = start of the post-AI period)



These results are conditional on being employed. If the graduates who lost out left employment, which the unemployment results suggest, a comparison among those who remained employed could understate mismatch. With that limit in mind, the adjustment for young computing graduates appears to have occurred mostly through unemployment and less through moving into jobs outside computing.

Sensitivity to the baseline period

This section asks whether the results depend on which years form the baseline, since the early 2010s followed the 2007-2009 recession and may not be typical of the years just before 2023. The regressions use every survey year before 2023 as the baseline, except 2020. Table 4.6 compares this with a shorter baseline of 2017-2019 using raw unemployment rates. With the longer baseline the unemployment difference-in-differences is 3.7 points at ages 22-25, 1.4 points at ages 26-30, and 1.3 points at ages 31-40, close to the regression estimates. With the

shorter baseline it is 2.7, 0.3, and 1.0 points. The reason is that unemployment of other majors at ages 22-25 was higher in the earlier baseline years (6.4% on average, against 5.2% in 2017-2019), so the longer baseline gives a larger decline for other majors. The age ordering is the same under both baselines, with the largest increase at ages 22-25, but the size of the gap depends on the baseline. The contrast with ages 31-40 is 2.4 points under the longer baseline and 1.7 points under the shorter one. The non-employment differences are 2.2, 0.3, and 2.0 points under the longer baseline and 3.2, -1.3, and 1.0 points under the shorter one, which confirms that the non-employment results are less stable than those for unemployment.

Table **4.6**. Unemployment rates (%) and differences-in-differences under two baseline periods

Age group	Baseline	Computing: baseline	Computing: 2023-24	Other: baseline	Other: 2023-24	Unemployment DiD	Non-employment DiD
22-25	2010-2022	6.7	9.3	6.4	5.2	3.7	2.2
	2017-2019	6.7	9.3	5.2	5.2	2.7	3.2
26-30	2010-2022	3.5	4.1	3.5	2.7	1.4	0.3
	2017-2019	3.7	4.1	2.7	2.7	0.3	-1.3
31-40	2010-2022	2.5	3.2	2.9	2.3	1.3	2.0
	2017-2019	2.1	3.2	2.1	2.3	1.0	1.0

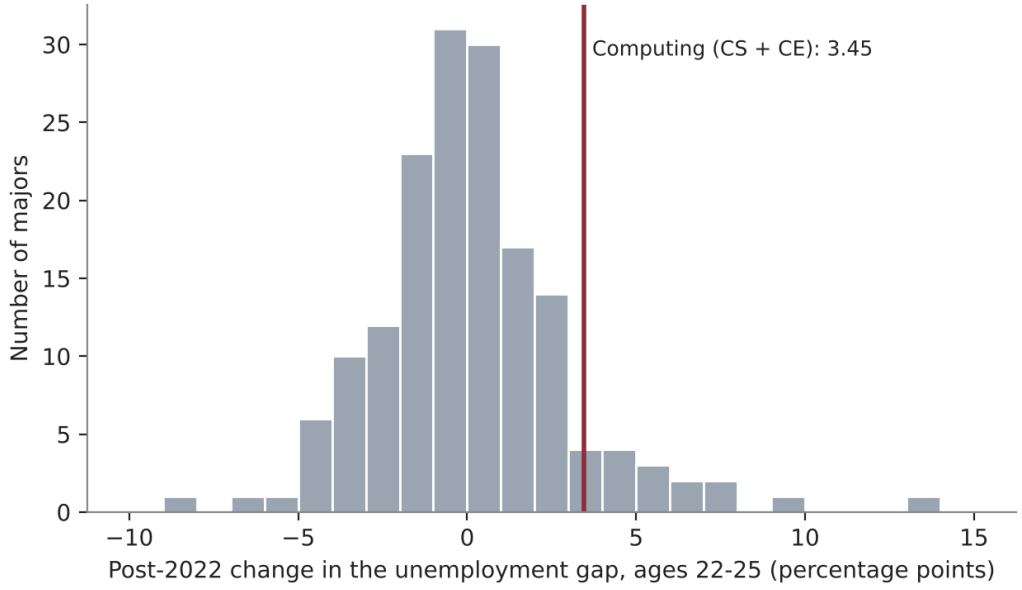
Notes: Rates are weighted unemployment rates (share of the labor force). The difference-in-differences (DiD) is the change for computing graduates minus the change for other majors, in percentage points; the last column is the same calculation for non-employment. 2023-2024 is the post period in both cases.

Computing relative to the distribution across majors

A concern with any comparison of one group against all others is that, with many degree fields, some field may be expected to show a large change by chance alone. Figure 4.4 addresses this by repeating the same age 22-25 estimate of model equation (2) for each of the 165 other degree fields in turn. For example, in model equation (2), I replace CP_i with $ECON_i$, where

$ECON_i$ is a dummy variable that equals one for economics graduates and zero otherwise. I repeat this exercise for each degree field, which gives a unique θ_1 for each degree field, and Figure 4.4 plots the distribution of those values. The computing estimate of 3.45 points is larger than the estimate for 148 of the other majors. The 17 majors with an estimate at least as large are mostly small majors with imprecise estimates, and none is among the 48 majors with 5,000 or more observations. The figure is clipped to the range -10 to 15 points and omits two small majors with estimates of 34.3 and -10.1 points. The increase for computing is therefore unusual compared with the other degree fields, and especially among large majors.

Figure **4.4**. Post-2022 change in the unemployment gap at ages 22-25 for 165 other majors, each treated as the comparison group of interest (red line = computing, 3.45 points)



Chapter 5 - Discussion

The results show that unemployment among young computing graduates increased more than for other majors after 2022, and that this increase is larger for younger computing graduates than for older computing graduates. These results are consistent with AI reducing the demand for entry-level work it has been shown to perform, which contributed to weaker hiring of younger graduates. A broader downturn in tech-sector hiring may also have affected computing majors, and the comparison used here does not separate the two, but the pattern across age groups and degree fields is what an AI effect on entry-level work would predict. The exposure results from Chapter 2 point in the same direction. Computing is the most AI-exposed field among the 50 largest majors (Figure 2.1), and the post-2022 rise in unemployment for young computing graduates was larger than for any of the 48 largest other majors (Figure 4.4), which is the pattern expected if the increase reflects AI exposure instead of a general labor market shock.

This reading is consistent with evidence that AI operates mainly through hiring: Brynjolfsson et al. (2025) describe declines in junior employment that come mainly from slower hiring, and Massenkoff and McCrory (2026) find suggestive evidence of a slower job-finding rate for young workers entering exposed occupations. Additionally, the finding of the Yale Budget Lab (Gimbel et al., 2025) that the economy-wide shift in the occupational mix is small fits effects that are concentrated among particular groups, such as young workers in exposed fields.

Part A adds a further supply-side analysis. The share of bachelor's degree holders aged 22-25 with a computing degree rose from 2.2% in 2010 to 3.9% in 2019 and 5.6% in 2024, and undergraduate computing enrollment grew about 9% in each of fall 2022 and fall 2023 before falling 8.8% in fall 2025. Larger graduating classes mean more young computing graduates

searching for entry-level jobs at the same time, which could raise their unemployment even if employer demand had not changed. The triple-difference design cannot remove this, because the increase in supply is specific to computing. The 2025 decline in enrollment is in line with students responding to weaker job prospects for computing graduates.

Changes in the enrollment of foreign students may also help explain the enrollment results in Part A, but the enrollment data cannot test this directly. The Clearinghouse data used here do not record country of birth or immigration status, so the enrollment of foreign-born and U.S.-born students cannot be separated in any state, field, or year. The evidence below is therefore indirect, and it focuses on undergraduates, the group studied in Part A.

The evidence on international undergraduates does not point to a large foreign-student effect. The Institute of International Education's fall 2025 snapshot, based on 828 institutions, reports that undergraduate international enrollment rose 2%, even though new international enrollment, which covers all levels, fell 17% (Baer & Ekin, 2025). The Clearinghouse likewise reports that international undergraduate enrollment rose 3.2% (about 5,000 students) in fall 2025, which is less than half the growth of the previous fall (8.4%) (National Student Clearinghouse Research Center, 2026c). International students are also a small share of undergraduates: in the IPEDS fall 2024 files, nonresident aliens were 2.9% of undergraduate enrollment (National Center for Education Statistics, 2025). Engineering is more international, at 6.0% of undergraduates, although the IPEDS major-field file reports engineering only as a whole and does not report computer science. A slowdown in the growth of international undergraduates could therefore explain only a small part of the 8.8% decline in undergraduate computing enrollment unless they were heavily concentrated in computing and fell sharply within it, which

the overall pattern does not suggest. I cannot rule this out with the data used here, and changes in international graduate enrollment, which this thesis does not study, are outside its scope.

The ACS offers a second view because it records citizenship and place of birth, which the enrollment data do not. Foreign-born persons, defined here as naturalized citizens and non-citizens (so a broader group than international students), were 44.8% of computing graduates aged 22-40 in 2010-2022 and 39.8% in 2023-2024, compared with 16.5% of graduates in other fields in both periods. Computer engineering (51.6%) and computer science (42.7%) rank third and sixth among the 172 degree fields with at least 200 observations in 2010-2022. Most foreign-born computing graduates arrived in the United States at age 18 or older (82.2% in 2010-2022, compared with 66.9% for foreign-born graduates in other fields), which is consistent with people who came as students or workers. Other data show that foreign-born workers are concentrated in computing, especially at higher degree levels: in the 2021 National Survey of College Graduates, foreign-born persons were 22.3% of bachelor's degree holders working as computer and mathematical scientists, compared with 18.7% across all science and engineering occupations, and 46.8% and 57.6% of master's and doctoral degree holders in computer and mathematical occupations (National Science Board, 2024). A larger foreign-born share means that immigration policy and international student flows can affect the supply of computing graduates and their job search more than in other fields, so a change in those flows could move enrollment and the labor market outcomes in Part B at the same time. The composition of the computing group also changed after 2022 in a way that other fields did not. The foreign-born share of computing graduates fell by 3.6 points at ages 22-25 (29.9% to 26.3%), 7.8 points at ages 26-30 (43.0% to 35.2%), and 1.7 points at ages 31-40 (49.8% to 48.0%), while the share in other fields was unchanged. Part of the post-AI change in the unemployment gap could therefore reflect who is in

the group and not AI exposure, and the direction of any bias depends on how unemployment differs by nativity. The largest fall in the foreign-born share is at ages 26-30, whereas the largest rise in unemployment is at ages 22-25, so the two do not line up one-for-one. I do not test this here, and estimating equation (2) separately for U.S.-born and foreign-born graduates is a natural extension.

The job-mismatch results suggest where the adjustment took place. Employed young computing graduates were as likely as before to hold tech jobs, and the share in jobs that do not typically require a bachelor's degree rose by 2.7 points. The point estimates for older computing graduates show a fall of 1.6 to 2.1 points in tech employment, although these are not individually significant. Together with the unemployment results, this is consistent with young computing graduates ending up either in a tech job or without a job, and less often in a lower-level job elsewhere. The caveat is that the mismatch measures describe only the graduates who were employed. If the graduates who would have been poorly matched are now unemployed, the measures understate mismatch, and the results also depend on the list of reference occupations.

For the hypotheses tested, H_1 is supported because unemployment rose more for computing graduates than for graduates of other majors in all three age groups ($\theta_1, \theta_2, \theta_3 > 0$). H_2 is supported for unemployment: the increase for ages 22-25 exceeds that at ages 26-30 and 31-40 by about two points ($p < 0.001$). Both are only partly supported for non-employment, where the contrast with ages 26-30 is significant at the 10% level and the contrast with ages 31-40 is not. The result does not depend on the definition of computing, since CS alone gives contrasts of 1.2 and 1.5 points, and does not appear to be driven by graduates who are still enrolled, since excluding them gives contrasts of 2.5 and 2.3 points.

The results for computing are unusual compared with the other degree fields: only 17 of 165 other majors, and none of the 48 largest showed as large an increase in unemployment after 2022 as computing did. Although I cannot rule out chance, this is suggestive that something unique may be occurring in this field, and the pattern is consistent with an AI effect on young computing graduates.

These results also bear on students and programs. If hiring of junior workers stays weak, the entry-level market for computing graduates, and not only the number of degrees awarded, is what has changed, and enrollment may continue to fall. The cohorts that graduated in 2023 and 2024 may also bear costs beyond the first spell of unemployment (Kahn, 2010; Oreopoulos et al., 2012), although this thesis cannot measure them.

Several limits apply to this study. The post-AI window only covers two years, since 2025 data are not available yet, and the linear probability model has only 14 survey-year clusters. The age groups in the post-period are also different cohorts, so the age pattern could partly reflect differences between cohorts, such as the size of graduating classes. Non-employment also rose for computing graduates aged 31-40, so the age pattern is clearest when looking at unemployment. Finally, the job-mismatch outcomes describe only employed graduates and depend on the list of reference occupations. The enrollment data also do not record country of birth or immigration status, so changes in the enrollment of foreign-born and U.S.-born students cannot be separated.

Extending the sample to later years of the ACS as they are released, adding measures of actual AI adoption by occupation or employer in place of task-based exposure, and following individual graduates over time would help to distinguish the explanations discussed here.

Chapter 6 - Conclusion

Part A shows that computing degrees and enrollment grew rapidly through 2024, and that combined computer science and computer engineering enrollment then fell 8.8% in the Fall of 2025, almost entirely in computer science. This is suggestive evidence of a recent decline, which cannot yet be seen among graduates. Unemployment rose more after 2022 for young computing graduates than for graduates of other majors, and the increase was substantially larger for ages 22-25 than for ages 26-30 and 31-40. This age pattern is clear for unemployment and weaker for non-employment. Among employed graduates, young computing graduates did not move out of tech occupations, but they were somewhat more likely to hold jobs that do not typically require a bachelor's degree, so the adjustment appears mostly in unemployment rather than in job mismatch. The results are consistent with AI reducing the labor market prospects of computing graduates, especially younger ones, although a broader tech downturn may also play a role. The analysis is also limited to a short window (2023 and 2024), because 2025 data are not yet available and a larger effect may appear later. Extending the analysis as the new data arrives, and using occupation-level measures of AI adoption, would help determine how far these differences reflect the influence of AI in the U.S. labor market.

References

Altonji, J. G., Blom, E., & Meghir, C. (2012). Heterogeneity in human capital investments: High school curriculum, college major, and careers. *Annual Review of Economics, 4*(1), 185-223.

Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics, 118*(4), 1279-1333.

Brynjolfsson, E., Chandar, B., & Chen, R. (2025). *Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence*. Stanford Digital Economy Lab.

Dickler, J. (2025, May 16). College majors with the best and worst job prospects — art history beats finance. CNBC.

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). *GPTs are GPTs: An early look at the labor market impact potential of large language models*. arXiv:2303.10130.

Federal Reserve Bank of New York. (2026). *The labor market for recent college graduates (2026:Q2)*. https://www.newyorkfed.org/research/college-labor-market

Gimbel, M., Kinder, M., Kendall, J., & Lee, M. (2025). *Evaluating the impact of AI on the labor market: Current state of affairs*. Yale Budget Lab.

Hosseini, S. M., & Lichtinger, G. (2025). *Generative AI as seniority-biased technological change: Evidence from U.S. résumé and job posting data*. SSRN. https://ssrn.com/abstract=5425555

Indeed Hiring Lab. (2025, July 30). *The US tech hiring freeze continues*. https://www.hiringlab.org/2025/07/30/the-us-tech-hiring-freeze-continues/

Institute of International Education. (2025). *Fall 2025 snapshot on international student enrollment: Key findings*. Open Doors. https://opendoorsdata.org/wp-content/uploads/2025/11/IIE_Fall-2025-Snapshot_Key-Findings.pdf

Kahn, L. B. (2010). The long-term labor market consequences of graduating from college in a bad economy. *Labour Economics, 17*(2), 303-316.

Massenkoff, M., & McCrory, P. (2026, March 5). *Labor market impacts of AI: A new measure and early evidence*. Anthropic.

National Center for Education Statistics. (2025). *Integrated Postsecondary Education Data System (IPEDS), fall enrollment 2024: EF2024A and EF2024CP* [Data set]. U.S. Department of Education. https://nces.ed.gov/ipeds

National Science Board, National Science Foundation. (2024). *The STEM labor force: Scientists, engineers, and skilled technical workers* (NSB-2024-5). https://ncses.nsf.gov/pubs/nsb20245/foreign-born-stem-workers

National Student Clearinghouse Research Center. (2026a, August 11). *Computer science enrollment is cooling*. https://www.studentclearinghouse.org/nscblog/computer-science-enrollment-is-cooling/

National Student Clearinghouse Research Center. (2026b). *Fall 2025 enrollment trends: CIP appendix*[Data set]. National Student Clearinghouse.

National Student Clearinghouse Research Center. (2026c, January 15). *Clearinghouse final fall enrollment trends* [Enrollment Insights Series]. https://nscresearchcenter.org/final-fall-enrollment-trends/

Oreopoulos, P., von Wachter, T., & Heisz, A. (2012). The short- and long-term career effects of graduating in a recession. *American Economic Journal: Applied Economics, 4*(1), 1-29.

Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). *The impact of AI on developer productivity: Evidence from GitHub Copilot*. arXiv:2302.06590.

Ruggles, S., Flood, S., Sobek, M., Brockman, D., Cooper, G., Richards, S., Rodgers, R., Schroeder, J., & Williams, K. *IPUMS USA: Version 16.0* [usa00013.dta]. Minneapolis, MN: IPUMS, 2025. https://doi.org/10.18128/D010.V16.0