Adapting to AI: Experimental Evidence on Human Capital Responses to Technological Disruption*

Amanda Chuan Germán Reyes Joaquín Serrano Hanzhe Zhang

September 2026

Abstract

We conduct a randomized information experiment with undergraduates at a large public university in the U.S. to test whether students adapt their career plans to AI's expected effects on the labor market. Treated students explore an interactive visualization that ranks their desired occupations by AI exposure, while control students explore a similar visualization without information about AI exposure. Treated students are more likely to switch desired occupations toward less-exposed alternatives, consider changing majors, and cite avoiding technological displacement as a reason behind their major choices. They lower the employment rates and salaries they expect for the occupation (market beliefs) more than their own chances of being employed in it (own beliefs), and their own expected earnings do not move. The effects on choices and employment beliefs are stronger among students whose desired occupations are highly exposed at baseline. Treatment also changes which information students seek: treated students shift their exploration away from high-exposure occupations and ask a chatbot framed as a career advisor about AI's impact on jobs rather than about salaries. A mediation decomposition attributes between a quarter and a third of the major and occupation shifts to belief revisions and an increased weight on avoiding displacement. Information about AI's labor-market impact thus shifts market beliefs and career plans before students enter the labor market, but own beliefs move less, so information alone is unlikely to reallocate students at scale.

*Chuan: Michigan State University (achuan@msu.edu). Reyes: Middlebury College and IZA (greyes@middlebury.edu). Serrano: Universitat Autònoma de Barcelona (joaquin.serrano@uab.cat). Zhang: Michigan State University (hanzhe@msu.edu). We thank Christopher Boyette, Wenjia Cao, Tina Chiou, Grace Jin, and Jie Wu for excellent research assistance. This study was approved by the Michigan State University Institutional Review Board (STUDY202500204), and the experiment is registered at the AEA RCT Registry (AEARCTR-0018261), where the pre-analysis plan was filed before the data were analyzed. Financial support from the Anthropic Economic Futures Program and Schmidt Sciences (AI at Work) is gratefully acknowledged.

1 Introduction

For decades, technological advances substituted for less-educated workers in manual and routine tasks and complemented college-educated workers in nonroutine cognitive work (Autor et al., 2003; Goldin and Katz, 2008; Goos et al., 2014). Generative artificial intelligence (AI) may break this pattern. It can execute many of the nonroutine cognitive tasks that college graduates perform (Felten et al., 2023; Eloundou et al., 2024), and adoption concentrates in the occupations that are intensive in these tasks (Bick et al., 2026a,b). Consistent with this, employment of early-career workers in the most-exposed occupations has already begun to decline relative to less-exposed occupations (Brynjolfsson et al., 2025a; Hosseini Maasoum and Lichtinger, 2025).

How AI affects labor markets partly depends on how elastically people adjust their human capital investments (Acemoglu and Autor, 2011; Autor, 2019). The standard policy response is to retrain displaced workers, but retraining reaches them only after they have invested in the occupation-specific skills that technological change devalues (Jacobson et al., 1993), and its estimated gains are modest and gradual (Card et al., 2018). Even when retraining pays, the gains come mainly from transitions into less-exposed occupations (Hyman et al., 2025). Those transitions are less disruptive before people enter the labor market, when no occupation-specific investment has been sunk, which makes earlier interventions an attractive policy option. Yet evidence on whether such interventions are effective is scarce.

In this paper, we provide experimental evidence on how information about occupational AI exposure affects the beliefs and educational choices of college students. Students can still adapt in anticipation of labor-market shocks, by changing their field of study, their desired occupation, or their plans for graduate school, and so avoid the costs of retraining (Altonji et al., 2012; Stinebrickner and Stinebrickner, 2014).[1] A challenge in studying this response is that AI exposed every labor market at once, leaving no unexposed market for comparison. Thus, we randomize who receives AI-exposure information, which identifies the effect of the information on beliefs and choices. Because occupations differ sharply in AI exposure (Eloundou et al., 2024), the same facts are bad news for students whose

[1] Survey and enrollment evidence suggests that students may already be responding. In a national survey fielded in late 2025, 47 percent of U.S. undergraduates said AI had made them reconsider their major (70 percent among technology majors) and 16 percent had already changed it (Gallup and Lumina Foundation, 2026). Computer science enrollment at four-year institutions fell 8.4 percent from spring 2025 to spring 2026 (National Student Clearinghouse Research Center, 2026), a decline that news accounts attribute to fears that AI will automate entry-level coding work (The Washington Post, 2026).

desired occupation is highly exposed ("high-exposure" students) and neutral or good news for the rest, so we estimate effects for the full sample and separately for high-exposure students.

We conduct the experiment with 3,471 undergraduates at Michigan State University (MSU), a large public university. Treated students explore an interactive visualization that ranks occupations by their AI exposure, built from the Anthropic Economic Index (AEI; Handa et al., 2025). Control students explore an interactive visualization that ranks occupations alphabetically. Both arms end with a free-search screen and a chatbot framed as a career advisor. We elicit beliefs, intentions, and choices before and after the information treatment.

We document three sets of findings. First, stated choices respond in the expected direction on all three margins we track: occupational choice, educational plans, and the reasons students cite for their major choice. On the occupation margin, high-exposure treated students are 4.6 percentage points (pp) more likely to name a different top-choice occupation after the treatment, a 44 percent increase over the control rate, and the switches move toward less-exposed occupations. Treatment also raises the stated probability of changing majors by 1.35 pp and makes students 8.1 pp more likely to cite avoiding technological displacement as a reason for their major choice.

Second, the information moves both market and own beliefs, but market beliefs more. We distinguish what students believe about their top-choice occupation as a whole ("market beliefs": its employment rate and salaries) from what they believe about their own prospects in it ("own beliefs": the probability that they will be employed in it and their own expected earnings). Treatment lowers the occupation's expected employment rate by 2.40 pp at the 10-year horizon and 2.45 pp at the 20-year horizon, and it lowers the occupation's expected salaries at every horizon (2.3 percent at 10 years). The employment effects are larger for high-exposure students, while the salary effects are similar in the two groups. Own beliefs move less: the probability students assign to being employed in the occupation falls by about half as much as its employment rate at 10 years and four-fifths as much at 20 years, and own expected earnings do not move. As a result, treatment widens the gap between the own employment probability and the occupation's employment rate by 1.4 pp at 10 years. A mediation decomposition attributes less than a third of the choice response to the belief revisions and the added weight on avoiding displacement.

Third, treatment changes *which* information students seek more than *how much*. Treated students shift their free exploration in the interactive visualization away from high-exposure

occupations, and their questions to the chatbot move the same way. They are 9.3 pp more likely to ask about AI's impact on jobs, a roughly tenfold increase over the 1.0 percent control base, and less likely to ask about salaries. Search intensity also falls, by about a fifth: treated students enter fewer search terms and spend less time on occupation detail panels.

Taken together, our findings imply that career choices respond to information about occupational AI exposure, but modestly. Consistent with models of occupational choice (Wiswall and Zafar, 2015; Conlon, 2021), students' plans adjust in the direction the exposure ranking implies, move most for the students facing the worst news, and move before students enter the labor market. The modest response is consistent with our finding that own beliefs move less than market beliefs, and own earnings expectations not at all. Thus, information provision alone is unlikely to produce reallocation at scale.

Our results contribute to two literatures. First, we contribute to research on how people adapt to technological shocks. Existing evidence shows that human capital responds to technological change once the technology has diffused: college enrollment rose in U.S. labor markets that adopted industrial robots, mostly at community colleges (Di Giacomo and Lerch, 2026), and teenage boys dropped out of high school more where fracking raised the demand for less-educated male labor (Cascio and Narayan, 2022). Three experiments are closest to ours. Shen et al. (2026) give white-collar workers in China occupational AI-exposure rates and find that the information moves their displacement beliefs and policy preferences. Golin and Rauh (2023) give U.S. workers automation-risk probabilities and find that the information moves their policy preferences but not their retraining intentions. Huseynov (2025) gives U.S. college students qualitative assessments of ChatGPT that are not occupation-specific and finds that the information moves their beliefs about the earnings distribution but not their own expected earnings. We instead give college students quantitative exposure data for their own desired occupations, and trace which beliefs move, by how much, and which choices respond. To our knowledge, we provide the first experimental estimates of ex ante educational adjustment to occupational AI exposure.[2]

[2]A related literature measures the labor-market effects of AI and of earlier automation. Task-based studies examine which tasks machines can perform and how new tasks reinstate labor (Brynjolfsson and Mitchell, 2017; Acemoglu and Restrepo, 2019). Field experiments measure how AI changes the productivity of workers who use it (e.g., Noy and Zhang, 2023; Brynjolfsson et al., 2025b; Cui et al., 2026; Dell'Acqua et al., 2026), and early observational work tracks realized labor-market effects, which range from employment declines among exposed early-career workers (Brynjolfsson et al., 2025a; Hosseini Maasoum and Lichtinger, 2025) to precisely estimated near-zero effects on earnings and hours in Denmark (Humlum and Vestergaard, 2025). Studies of industrial robots measure displacement and wage effects after the technology diffuses (Graetz and Michaels, 2018; Acemoglu and Restrepo, 2020; Dauth et al., 2021;

Second, we contribute to research on the determinants of college major and occupational choice. Students choose fields according to their beliefs about earnings, employment, and the jobs a field leads to, and information experiments show that shifting these beliefs changes which field students choose (Wiswall and Zafar, 2015; Baker et al., 2018; Conlon, 2021; Conlon and Patel, 2026). Students also revise their choices as they learn about their own ability (Stinebrickner and Stinebrickner, 2014) and as labor-market conditions change (Blom et al., 2021). We add AI exposure to the attributes students weigh. Information about AI exposure moves the same choices, and avoiding technological displacement enters the reasons students cite for their major choice. Consistent with previous work showing that students who learn population earnings by major revise beliefs about their own earnings only partially (Wiswall and Zafar, 2015), we also find that students revise market beliefs more than own beliefs. The same gap between aggregate and personal risk appears in perceptions of technology (Barrera-Jimenez et al., 2025) and in households' macroeconomic expectations (Roth and Wohlfart, 2020).

The rest of the paper is organized as follows. Section 2 describes the experimental design. Section 3 covers the data and the estimation procedure. Sections 4, 5, and 6 report the effects on stated choices, beliefs, and demand for information. Section 7 concludes.

2 Experimental Design

This section describes the survey and the information treatment. We conducted the experiment in April 2026 and recruited MSU undergraduates through campus-wide emails. To avoid selection on AI attitudes and to keep AI content out of the control condition, we framed the survey as a study of how students make educational and career decisions. Appendix B provides additional implementation details, and the Supplementary Appendix contains the full survey instrument.

2.1 Survey Flow and Randomization

The experiment consists of three parts (Appendix Figure A1): pre-treatment baseline survey blocks, a randomized information block, and post-treatment blocks that repeat the pre-treatment questions and ask additional questions.

In the baseline blocks, students first report their current major, the reasons they cite for their major choice, their satisfaction with the major, and their probabilities of changing

Serrano, 2024).

major, applying to graduate school, and leaving college. They then give their occupation list (the occupations they are considering, ranked with their top choice first) and state their beliefs about their top-choice occupation: employment rates and expected salaries at four horizons, their own employment probability at three horizons, occupational tenure, and a subjective probability distribution over their own earnings at age 35. Finally, they report their own AI use and their beliefs about AI's effect on jobs. The post-treatment blocks repeat most these questions, so we observe beliefs and plans before and after the treatment.

We next randomize students to treatment or control with equal probabilities. Treated students explore an interactive visualization that ranks occupations by their exposure to AI, while control students complete a format-matched placebo with no AI content (both described below). After the visualization, students complete the post-treatment blocks, the demographic questions, and incentivized time- and risk-preference tasks. The survey closes by offering students help with scheduling a meeting with a career counselor. Students receive a $20 completion payment plus their earnings from the incentivized tasks.

2.2 The AI-Exposure Information Treatment

Treatment Visualization. The core of the treatment is an interactive visualization that ranks 775 six-digit Standard Occupational Classification (SOC) occupations into four tiers of AI exposure (Figure 1).[3] We measure AI exposure with the AEI, which maps a sample of conversations on Claude (Anthropic's AI model) to the tasks that define each occupation in the U.S. Department of Labor's O*NET database (Handa et al., 2025). The index scores each occupation by the share of its tasks that appear in those conversations as automation, that is, with AI performing the task rather than assisting a person with it. An occupation's exposure *rank* is its percentile in this ranking, from 1 (most exposed) to 100 (least exposed), and its *tier* is the group of ranks it falls in. We cut the ranking at the quintile boundaries at which Brynjolfsson et al. (2025a) document diverging employment trajectories for early-career workers: Very High (ranks 1 to 20), High (21 to 40), Moderate

[3]We emphasize to students that AI's impact on jobs is uncertain and is an active area of research. An introductory screen explains that exposure measures are based on whether AI helps or replaces workers in an occupation's tasks and that the results in the visualization come from researchers at Stanford University and Anthropic, with the two source papers cited on screen. The screen closes with an explicit disclaimer: "The impact of AI on employment remains an active area of research, and no conclusions are definitive. The data presented here primarily reflect past employment trends, which may not reliably predict future changes."

(41 to 60), and Lowest (61 to 100).[4] Throughout, "high-exposure students" are those whose baseline top-choice occupation falls in the Very High or High tier.

The visualization proceeds in three screens plus a detail panel. On the *selection* screen (Panel A of Figure 1), students type occupational titles into a search box. They can select up to six occupations. The *exposure display* (Panel B) then places each selected occupation in its exposure tier on a chart showing employment growth by tier from November 2022, the month of ChatGPT's release, to November 2025. Within this display, clicking any occupation opens a panel showing its exposure tier, the three most similar occupations and their tiers, and the occupation's education requirements and common fields of study (Panel C). On the final *free search* screen (Panel D), students can explore any occupation in the database, with filters by exposure tier, occupation group, education level, and field of study, and can open the same detail panel for each occupation they select.

Chatbot. After the visualization, students can interact with a chatbot framed as a career advisor. The treatment chatbot is configured to draw on the exposure data and the control chatbot to avoid AI topics (Appendix Figure A2). A front-end error caused the chatbot to return no replies in either arm, so we use the first question each student asks to capture demand for information.

Control: Placebo Information. The control condition holds the visualization's format fixed while stripping all AI-exposure information content (Appendix Figure A3). Control students see the same occupations ranked into four tiers, but the tiers group occupations alphabetically (A–F, G–M, N–S, T–Z) rather than by AI exposure. The three screens and the detail panel are otherwise identical. Students select up to six occupations and see them placed on the tiered chart. Clicking an occupation opens the same detail panel as in the treatment visualization, without the elements that convey AI exposure. On the *free search* screen, students explore any occupation in the database with the same filters, minus the exposure tier.

[4] We form one tier with the two least-exposed quintiles because nearly half of the occupations score zero (i.e., none of their tasks appear in the AEI conversations as automation).

3 Data and Research Design

3.1 Sample, Summary Statistics, and Balance

Our analysis sample consists of 3,471 students who reached the visualization screen. We exclude 383 students who dropped out before the visualization. Table 1 places this sample (column 1) alongside three benchmarks: the MSU undergraduate body, drawn separately from the Registrar (column 2) and from the Integrated Postsecondary Education Data System (IPEDS; column 3), and all U.S. four-year colleges (column 4). Relative to all three, our sample overrepresents women, white and Asian students (Panel A), and STEM fields (Panel B). Its share in health and education is close to both MSU benchmarks but well below the U.S. four-year average. Relative to MSU specifically, the sample underrepresents freshmen and overrepresents seniors (Panel C). Panel D shows that students engage with the information: on average, they spend 2.4 minutes on the visualization.

The sample is balanced on observable characteristics across treatment groups. Table 2 reports control-group means and treatment–control differences for student characteristics both for the full sample and for high-exposure students (71 percent of the sample).[5] We test balance jointly across the 37 pre-randomization covariates, setting missing values to the sample mean and adding an indicator for each covariate with missing values, so that every student in the sample enters. The joint F-test does not reject balance ($F = 1.13$, analytic $p = 0.24$). A randomization-inference version of the test, which permutes the treatment assignment 1,000 times and re-computes the joint statistic (Kerwin et al., 2024), yields a similar p-value ($p = 0.22$). Balance also extends to the variables most likely to moderate the treatment effect, such as students' baseline beliefs about AI.

3.2 Regression Model

We estimate linear models of the form:

$$Y_i = \alpha + \beta \, \text{InfoAI}_i + \mathbf{X}_i'\delta + \varepsilon_i, \tag{1}$$

where Y_i is an outcome, InfoAI_i indicates assignment to the AI-exposure information, $\mathbf{X}_i$ are controls included to improve precision, and ε_i is an error term. We select the controls by post-double-selection lasso (Belloni et al., 2014) from a candidate pool of demographics,

[5] This high share mirrors the pattern that motivates the paper: AI excels at the tasks performed in the occupations a college education leads to.

academic background, baseline career beliefs, AI knowledge, and the pre-treatment value of the outcome. We show robustness to alternative sets of controls (Appendix Table A1) and sample restrictions (Appendix Table A2), and we report treatment-effect heterogeneity along pre-specified and exploratory dimensions (Appendix Table A3). We compute heteroskedasticity-robust standard errors throughout.

The coefficient of interest, β, represents the average treatment effect of the AI-exposure information on a given outcome. Because control students engage with the same occupation-level information in a format-matched placebo, β isolates the effect of the AI-exposure content rather than of receiving occupational information per se. The information can move several intermediate outcomes at once, such as perceived exposure, beliefs about the occupation's prospects, and the weight placed on displacement risk, so we report the reduced-form effect of the information on each outcome and do not use it as an instrument for any one belief. Appendix Figure C1 illustrates the chain from the AI-exposure information through beliefs to choices. Section 5.4 decomposes, descriptively, how much of the choice response passes through the intermediate beliefs we elicit.

Whether the treatment provides good or bad news depends on the student's prior beliefs, but it is more likely to be bad news when the top-choice occupation ranks among the most exposed to AI and neutral or good news when it is less exposed. If students in exposed and insulated occupations revise their beliefs and plans in opposite directions, β averages the responses and understates the effect in each. We therefore also estimate equation (1) separately for high-exposure students, defined in Section 2. Every results table reports two estimates: $\hat{\beta}$ from equation (1) on the full sample, and $\hat{\beta}^{\text{high}}$ from the same equation among high-exposure students.[6]

3.3 Outcomes

We study three families of outcomes (see Appendix B.1 for variable definitions).

Stated choices. We measure the probabilities of changing major, applying to graduate school, and leaving college; interest in exploring alternative majors; and whether the student's preferred major or top-choice occupation differs after treatment.[7] We also measure

[6] We measure exposure from the top-choice occupation but verify robustness to using the other occupations in a student's ranking (Appendix Table A4).

[7] Students name the top-choice occupation in free text and the preferred major through program codes. Coders judged 30.4 percent of control-group occupation changes to be a different occupation rather than a rewording or specialization of the same one, and a coder judged 98.8 percent of control-group major changes to be a different program rather than the same major reached through a different college or department or listed under a different name or degree. Our two indicators equal one only for these judged changes.

how much students change their occupation consideration set. The first measure is the number of baseline occupations retained. The second is the summed rank movement of occupations kept on both lists.[8] The third is the Kendall τ rank correlation between the baseline and post-treatment rankings, computed over the union of the two lists, with an occupation missing from one list placed one position below that list's bottom. It falls below 1 when students reorder the list, drop an occupation from its middle, or leave more than one occupation off a list. Finally, we record the reasons students cite for their major choice.

Beliefs. We measure two kinds of post-treatment expectation. Market beliefs are the top-choice occupation's employment rate and expected salary at four horizons: today, the start of the full-time career, and 10 and 20 years ahead. Own beliefs are the student's own probability of being employed in the occupation at three horizons (on finishing education and 10 and 20 years ahead), expected occupational tenure, and the mean and variance implied by the earnings-at-35 distribution.

Demand for information. We measure how students explore the visualization's occupation database, whether they use the chatbot, and the topics they ask about. We classify each student's first message into one of eight topic categories, and each optional end-of-survey comment ("Do you have any general comments about the survey?") by its main point, using a language-model classifier whose labels human coders validated (Appendices B.4 and B.5). We also measure whether students request a meeting with a career counselor.

We recode probability scales to midpoints on a 0–100 scale and use bin midpoints for the earnings distribution. We winsorize salary expectations, expected years in the top-choice occupation, and survey and screen-time durations at the 95th percentile to limit the influence of implausible entries and idle screen time. Results are robust to alternative winsorization of salary expectations (Appendix Table A5) and to excluding long durations and the fastest completers outright (Appendix Table A2).

To limit the number of hypotheses we test, we pre-registered one summary z-score index per family: stated choices, beliefs, and demand for information. To construct each index, we standardize each component by the control-group mean and standard deviation, sign it so that positive values correspond to the direction we predicted in the pre-analysis plan, impute missing components at the treatment-arm mean, and average the components within each respondent (Appendix B.1 lists each index's components).

[8] A top-coded version also counts dropped occupations, each placed one position below the bottom of the post-treatment list.

4 Effects on Stated Choices

We start by examining the three margins of stated choice: occupational choice, educational plans (the probabilities of changing major, applying to graduate school, and leaving college), and the reasons students cite for their major choice. Figure 2 summarizes the treatment effects on each margin.

4.1 Occupation Switching

Treatment raises occupation switching (Table 3, Panel A). In the full sample, treatment raises the probability that the post-treatment top choice is a different occupation from the baseline one by $\hat{\beta} = 3.1$ pp on a 9.5 percent control baseline ($p = 0.006$). Among high-exposure students, the effect is $\hat{\beta}^{\text{high}} = 4.6$ pp, a 44 percent increase over the 10.4 percent control baseline ($p = 0.001$). For comparison, Conlon (2021) finds that correcting the salary beliefs of the undergraduates in the experiment, who on average underestimate salaries, makes them 9 pp (16 percent) more likely to major in a field they received information about. That response follows news that is favorable on average, whereas ours follows adverse news.

Most of the increase in switching is toward less-exposed occupations (Table 4). In the full sample, treatment raises switches into a lower-exposure tier by $\hat{\beta} = 2.3$ pp on a 3.1 percent control baseline ($p = 0.002$), raises within-tier switches by $\hat{\beta} = 1.6$ pp on a 3.3 percent baseline ($p = 0.029$), and does not detectably change switches into a higher-exposure tier, $\hat{\beta} = -0.9$ pp ($p = 0.12$). Among high-exposure students, the effect on switches into a lower-exposure tier is $\hat{\beta}^{\text{high}} = 2.9$ pp on a 4.1 percent control baseline ($p = 0.004$). Measured with the continuous exposure rank instead of tiers, so that a move to any less-exposed occupation counts, treated students are $\hat{\beta} = 2.6$ pp more likely to switch to a less-exposed occupation on a 4.4 percent control baseline ($p = 0.002$), and high-exposure treated students $\hat{\beta}^{\text{high}} = 3.1$ pp more likely on a control baseline of 6.0 percent ($p = 0.006$).[9]

Treatment also changes which occupations students list beyond their top choice. In the full sample, the share of listed occupations in the two most-exposed tiers falls, $\hat{\beta} = -2.2$ pp against a control share of 69.6 percent ($p < 0.001$). Treated students drop $\hat{\beta} = 0.13$ more

[9] Appendix Figure A4 splits these effects by the tier of the baseline top choice. High-exposure treated students are more likely to switch to an occupation in the Lowest tier. Treated students whose baseline top choice is already in the Lowest tier are more likely to keep it there and less likely to switch to the two most-exposed tiers. Students in the Moderate tier show no detectable response.

occupations from their lists on a control mean of 1.84 ($p = 0.011$), and the number of occupations listed falls by less, $\hat{\beta} = -0.08$ on a control mean of 4.3 ($p = 0.047$), so they replace part of what they drop. Because they do not drop high-exposure occupations at a greater rate than control students ($p = 0.28$), the shift toward less-exposed occupations comes from the occupations they add. Among the occupations kept on both lists, the order does not change: the rank distance between the two lists is $\hat{\beta} = 0.03$ positions on a control mean of 1.53 ($p = 0.72$). The list's mean exposure rank moves $\hat{\beta} = 1.3$ positions toward less-exposed occupations on a control mean of 35.5 ($p = 0.003$). Among high-exposure students, the share of listed occupations in the two most-exposed tiers falls, $\hat{\beta}^{\text{high}} = -2.3$ pp, and the share of baseline occupations retained falls, $\hat{\beta}^{\text{high}} = -3.8$ pp on a control mean of 62.1 percent ($p = 0.004$).

4.2 Educational Plans and Major Choice

Treated students report higher probabilities of changing their educational path (Table 5). The stated probability of changing major rises by $\hat{\beta} = 1.35$ pp on a control mean of 12.0 percent ($p = 0.005$), and by $\hat{\beta}^{\text{high}} = 1.75$ pp among high-exposure students ($p = 0.002$).[10] Whether the preferred major itself differs after treatment does not change detectably, $\hat{\beta} = 1.1$ pp on a control mean of 35.3 percent ($p = 0.49$) and $\hat{\beta}^{\text{high}} = 2.3$ pp among high-exposure students ($p = 0.22$; Table 3, Panel A). The stated probability of leaving college rises by $\hat{\beta} = 0.9$ pp on a 6.1 percent base ($p = 0.051$) and by $\hat{\beta}^{\text{high}} = 1.2$ pp among high-exposure students ($p = 0.023$).[11] Treatment does not detectably change the probability of applying to graduate school, in the full sample ($\hat{\beta} = 0.3$ pp, $p = 0.60$) or among high-exposure students ($\hat{\beta}^{\text{high}} = -0.2$ pp, $p = 0.79$). Interest in exploring alternative majors falls, $\hat{\beta} = -4.4$ pp on a control share of 20.1 percent ($p = 0.001$) and $\hat{\beta}^{\text{high}} = -4.5$ pp among high-exposure students ($p = 0.004$). Treated students thus say they are more likely to switch but less interested in exploring, a combination consistent with students who already know what they would switch into. The drop in exploration also fits the information avoidance we document in Section 6.

Treatment also changes the attributes students weigh in choosing a major (Table 3,

[10] For comparison, about one-third of bachelor's degree students change majors at least once within three years of enrollment (National Center for Education Statistics, 2017), and 47 percent of the fall 2017 entering cohort who completed a bachelor's degree by 2024 had changed their primary major at least once (American Academy of Arts and Sciences, 2025).

[11] For comparison, 23.5 percent of students who started college in the U.S. in fall 2022 had not returned to any institution a year later, and 22 percent of those who started at a public four-year institution had not returned to the same institution (National Student Clearinghouse Research Center, 2024).

Panel B). The share of students citing "avoiding technological displacement" as a reason for their major choice rises by $\hat{\beta} = 8.1$ pp on a control mean of 36.2 percent ($p < 0.001$). The average number of reasons cited is unchanged (2.31 in control and 2.32 in treated out of five total reasons, $p = 0.38$), so the information crowds out traditional reasons in favor of avoiding technological displacement. Three of the four other reasons fall (expected salary, job prospects, and developing diverse skills), while the fourth, innovative technology, rises. If these stated reasons carry into enrollment decisions, the substitution would shift demand toward majors students perceive as less exposed to AI, even among students who never switch desired occupations.

As a summary measure of the outcomes discussed in this section, we construct a z-score family index that pools the major- and occupation-switch indicators, the number of occupations listed and the rank distance, the probabilities of changing major and applying to graduate school, interest in exploring alternative majors, and the displacement reason (Appendix B.1), signed so that higher values run in the expected direction. Treatment increases this index by $\hat{\beta} = 0.03$ SD overall ($p = 0.019$) and by $\hat{\beta}^{\text{high}} = 0.04$ SD among high-exposure students ($p = 0.002$; Table 3, Panel C): choices move in the predicted direction, and they move more for students who chose highly exposed occupations at baseline.

5 Effects on Beliefs

Why do students switch occupations and majors in this way? We next study beliefs about employment and salaries. We distinguish what students believe about the occupation ("market beliefs") from what they believe about their own prospects ("own beliefs").

5.1 Beliefs about Employment

The Occupation's Employment Rate. Control students believe that 64.3 percent of people in their top-choice occupation are currently employed (Table 6, Panel A).[12] Looking forward, they expect the rate to rise relative to today: on average, they forecast that 67.2 percent of people in their top-choice occupation will be employed in 10 years and 66.1 percent in 20 years (Figure 3, Panels A–B). The 10- and 20-year expectations fall

[12] This figure is below the 96–98 percent that the 2025 Current Population Survey implies for these fields (U.S. Bureau of Labor Statistics, 2026). These low levels likely reflect students gauging how hard a field is to enter rather than how often its incumbents are unemployed. Consistent with this interpretation, students hold strongly stereotyped beliefs about the jobs education leads to: 65 percent of prospective art majors expect to become artists, whereas 17 percent of art graduates do (Conlon and Patel, 2026).

under treatment, $\hat{\beta} = -2.40$ pp at 10 years and $\hat{\beta} = -2.45$ pp at 20 years (both $p < 0.001$), with no detectable effect at the current and start-of-career horizons. The effects are larger for high-exposure students: $\hat{\beta}^{\text{high}} = -3.10$ pp at 10 years and $\hat{\beta}^{\text{high}} = -3.61$ pp at 20 years (both $p < 0.001$). These effects vary with the news each student receives (Figure 4): treated students whose top choice sits in the two most-exposed tiers revise down sharply, while those in the two safest tiers revise little.

The horizon profile of the effects—no revision at the current and start-of-career horizons and a larger one at 20 years than at 10—is consistent with students interpreting AI as a shock that builds over decades rather than one that arrives at once. Control students' baseline beliefs show the same profile. Their employment-rate forecasts decline only between the 10- and 20-year horizons.[13] The same profile appears in expert forecasts: economists and superforecasters surveyed by Karger et al. (2026) estimate that aggregate labor-market outcomes will stay near their historical trend in the short run, with larger effects building over decades as adoption diffuses.

Own Employment Probability. Treatment also lowers students' own probability of being employed in the top-choice occupation, $\hat{\beta} = -1.96$ pp at 20 years ($p < 0.001$; Table 6, Panel B, and Figure 3, Panels C–D). The worse prospects students admit for themselves do not change how long a career they expect in the occupation: control students expect to work 28.6 years in it, and treatment does not change this expectation, $\hat{\beta} = -0.01$ years ($p = 0.96$). The beliefs-about-employment family index, which pools the occupation-level and own-employment components, rises by $\hat{\beta} = 0.04$ SD ($p < 0.001$) overall and $\hat{\beta}^{\text{high}} = 0.05$ SD among high-exposure students ($p < 0.001$; Table 6, Panel D).[14]

Because students revise the occupation's employment rate down by more than their own probability of being employed in it, treatment widens the gap between beliefs about oneself and beliefs about the market (Appendix Table A6 and Appendix Figure A5). At baseline, students already expect to outperform their occupation's average: control students rate their own employment probability 6.5 pp above the occupation's employment rate at the

[13] As an additional measure of beliefs about AI's effect on jobs, we asked all students at baseline how generative AI will change the number of jobs in their top-choice occupation, currently and at three future horizons. Both measures point to job losses, but at different horizons: 38.0 percent of students believe AI is already destroying jobs in their top-choice occupation, and 64.0 percent expect fewer jobs in 20 years than today. In 20 years, 23.0 percent expect AI to eliminate more than half of the jobs in the occupation, and 13.6 percent expect it to add jobs.

[14] All belief indices are signed so that positive values denote the predicted direction (here, downward revisions): a positive index coefficient reflects a larger treatment-induced reduction.

start-of-career horizon, 15.4 pp at 10 years, and 18.9 pp at 20 years. The widening is $\hat{\beta} =$ 1.4 pp at 10 years ($p = 0.043$) in the full sample; the 20-year estimate is not distinguishable from zero.[15] Among high-exposure students the gap widens by $\hat{\beta}^{\text{high}} = 2.2$ pp at 10 years ($p < 0.05$) and by $\hat{\beta}^{\text{high}} = 1.8$ pp at 20 years, the latter significant at the 10 percent level (Appendix Table A6).

5.2 Beliefs about Salaries

The Occupation's Salaries. Salary expectations are lower among treated than control students at every horizon (Table 6, Panel C, and Figure 3, Panels E–F).[16] We measure salary expectations in logs, so coefficients read as approximate percent changes. Expected 10-year salary falls under treatment, $\hat{\beta} = -2.3$ percent against a control mean of \$128,605 ($p = 0.006$), with similar declines at the other horizons. Unlike the employment effect, which grows with the horizon, the salary effect is roughly constant across horizons. Among high-exposure students the point estimates are similar, $\hat{\beta}^{\text{high}} = -2.7$ percent at 10 years ($p = 0.005$).

Own Earnings at Age 35. Students' beliefs about their own earnings barely move. They report a subjective probability distribution over their earnings at age 35. Treatment changes neither the implied mean, $\hat{\beta} = -78$ dollars on a control mean of \$114,059 ($p = 0.86$) and $\hat{\beta}^{\text{high}} = -145$ dollars among high-exposure students ($p = 0.78$), nor the implied variance, $\hat{\beta} = -8.3$ on a control mean of 459.2 thousand dollars squared ($p = 0.37$) and $\hat{\beta}^{\text{high}} = -12.4$ among high-exposure students ($p = 0.28$), and the probability assigned to each bin moves little, in the full sample and among high-exposure students alike (Appendix Table A8 and Appendix Figure A7). Together with the salary declines of the

[15] The gap opens only under adverse news (Appendix Table A7 and Appendix Figure A6). Among high-exposure students, for whom the ranking is bad news, the occupation's 10-year employment expectation falls by more than the own probability of being employed in it, so the gap widens. Among students in the two safest tiers, for whom the same screen is neutral or good news, occupation-level revisions are not distinguishable from zero and the gap does not open. The two gap responses differ at the 10 percent level.

[16] Baseline salary beliefs are far better calibrated than employment-rate beliefs (footnote 12). We match each control student's top-choice occupation to a 6-digit SOC code and its Bureau of Labor Statistics median wage (U.S. Bureau of Labor Statistics, 2025). We exclude the 104 students whose occupation has a top-coded BLS median. The median student's belief about what a typical worker currently earns is 89 percent of the BLS median for their occupation, 56 percent of beliefs fall within 25 percent of the BLS median, and the mean belief (\$93,479) sits 10 percent below the mean BLS median wage (\$104,312). This calibration contrasts with the large earnings misperceptions documented in other settings: Chilean applicants to low-return degree programs overestimate past graduates' earnings by over 100 percent (Hastings et al., 2015), and undergraduates' errors about population earnings range from −17 to +28 percent (Wiswall and Zafar, 2015).

previous subsection, this null is the earnings counterpart of the own-versus-market gap of Section 5.1: students revise the occupation's prospects more than their own. The null is also consistent with Huseynov (2025), who gives college students positive or negative news about ChatGPT's labor-market impact and finds that expected own-earnings levels stay fixed in both arms, even as students revise their expected position in the earnings distribution.

The beliefs-about-earnings family index, which pools the occupation salary expectations with the earnings-at-35 distribution, moves $\hat{\beta} = 0.03$ SD ($p = 0.003$) overall and $\hat{\beta}^{\text{high}} = 0.02$ SD among high-exposure students ($p = 0.029$; Table 6, Panel D), an effect driven by the shift in market beliefs.

5.3 Summary of Results and Interpretation

Overall, students revise their beliefs in the way the information implies, and more about the market than about themselves. They revise the occupation's employment prospects down at long horizons and its salaries down at every horizon, revise their own employment probability down by less, and leave their expected tenure and own earnings expectations unchanged.

The asymmetry between beliefs about the market and beliefs about oneself is the common thread in our results, and it is a recurrent finding in population surveys of technology risk. Some 77 percent of Americans consider it realistic that robots and computers will eventually do much of the work now done by humans, yet only 30 percent of workers think their own job is at risk (Pew Research Center, 2017). Likewise, 64 percent expect AI to lead to fewer jobs over the next two decades (Pew Research Center, 2025), yet only 17 percent of workers expect AI to hurt them more than help them personally (Pew Research Center, 2023), and just 22 percent worry that technology will make their own job obsolete (Gallup, 2023). In the most recent data, 40 percent of Americans expect AI's impact on society to be negative while only 31 percent expect a negative impact on themselves, a gap that holds across age groups (Pew Research Center, 2026). Barrera-Jimenez et al. (2025) label this pattern "invulnerability bias": survey respondents rate AI's future impact on others' jobs as larger than on their own.

We extend this evidence from survey attitudes to an experiment. Here the asymmetry appears within students and in response to randomized information: the same students who revise down their occupation's employment prospects revise their own down by less,

so the gap between the two widens.[17]

5.4 Why Students Switch

We next quantify how much beliefs about earnings, beliefs about employment, and concerns regarding AI displacement contribute to the occupational and major switching documented in Section 4. We follow the mediation decomposition of Fagereng et al. (2021) and Finkelstein et al. (2024). The idea is to ask how much of the treatment effect remains once we hold a candidate mediator fixed. For each mediator, we first estimate the total treatment effect (equation 1) on two outcomes, the stated probability of changing major and the top-choice occupation switch. We then re-estimate the same regression with the mediator added as a flexible control, entered as bin fixed effects over its post-treatment revisions plus a separate bin for students who do not revise.[18] The share of the treatment effect that operates through the mediator is the percentage by which the treatment coefficient falls when we add the mediator.

We measure each channel with one mediator, a z-score index built as in Section 3.3 from post-minus-pre revisions rather than levels. The earnings mediator pools the revisions in the occupation's expected salary at the four horizons; the employment mediator pools the revisions in the student's own employment probability at its three horizons; and the displacement mediator is the change in whether the student cites avoiding technological displacement as a major-choice reason. Because students are not randomly assigned to particular revisions, the mediators are endogenous, so we interpret the decomposition as a descriptive accounting exercise. Measurement error in the mediators also biases the mediated shares toward zero, so the may likely understate the role of the measured channels.

Together the three channels account for between a quarter and a third of the two responses: 30.9 percent of the occupation response and 26.8 percent of the major response (Table 7). Treatment moves every mediator, but less than a third of either response passes through them. For the probability of changing major (Panel A), the displacement channel

[17] More generally, incomplete passthrough from aggregate news to personal expectations is a common finding in information experiments. Students told the truth about population earnings by major revise beliefs about their own earnings with an elasticity of only 0.08 (Wiswall and Zafar, 2015), and a professional recession forecast moves households' aggregate recession and unemployment expectations about eight times more than their own job-loss risk (Roth and Wohlfart, 2020). The passthrough rises when the information refers to people similar to the respondent: unemployment statistics for "people with similar characteristics" pass through to own job-loss expectations at about one-half (Roth et al., 2022).

[18] For the two belief mediators the bins are deciles of the nonzero post-minus-pre revision, plus the separate bin for students who do not revise; the displacement mediator, a change in an indicator, enters fully saturated. See notes to Table 7 for details.

is the largest single channel, accounting for 10.2 percent of the response. For occupation switching (Panel B), the ordering flips toward the belief channels: earnings, employment, and displacement account for 10.1, 11.5, and 7.9 percent.

6 Effects on Demand for Information

Information seeking may be an additional margin on which people adapt to AI. We therefore examine whether the treatment changes how much, and which, information students seek. We use three measures: how students search the visualization's free-search screen, what type of information they seek from the career-advisor chatbot, and the comments they leave at the end of the survey. Figure 5 summarizes the treatment effects.

6.1 Occupation Search

Treatment changes which occupations students search on the free-search screen (Table 8, Panel A). The share of freely viewed occupations in the High tier falls, $\hat{\beta} = -6.6$ pp against a control share of 30.8 percent ($p < 0.001$), and the shares of the other three tiers rise, the Lowest tier's the most. The effect is similar in the high-exposure subgroup, $\hat{\beta}^{\text{high}} = -6.8$ pp ($p = 0.002$; columns 3–4). Split by the student's own baseline tier (Figure 6 and Appendix Table A9, unadjusted treated-minus-control differences), treated students whose top choice sits in the Very High tier move views from the High tier, where their share falls by 6.9 pp ($p = 0.010$), to the Lowest tier, where it rises by 5.6 pp ($p = 0.010$). Those whose top choice sits in the Lowest tier concentrate their views in that tier: their share in the High tier falls by 14.5 pp and their share in the Lowest tier rises by 23.2 pp (both $p < 0.001$). Those in the Moderate tier move views out of the Lowest tier, where their share falls by 10.8 pp ($p = 0.003$). Measured relative to each student's baseline tier, in the full sample the share of freely viewed occupations in higher-exposure tiers falls, $\hat{\beta} = -2.9$ pp ($p = 0.072$), and the share in the student's own tier rises by an amount not distinguishable from zero, $\hat{\beta} = 3.1$ pp ($p = 0.12$). Neither effect is detectable in the high-exposure subgroup (Appendix Table A10).

Treatment also lowers how much students search on the free-search screen (Table 8, Panel A). Control students enter 1.4 search terms, spend 12.3 seconds on occupation detail panels, and explore 1.4 occupations on average. Treated students enter fewer search terms, $\hat{\beta} = -0.25$ terms ($p = 0.001$), spend less time on detail panels, $\hat{\beta} = -2.8$ seconds ($p < 0.001$), and explore fewer occupations, $\hat{\beta} = -0.15$ occupations ($p = 0.075$).

6.2 Career-Advisor Chatbot Questions

Treatment also changes what students ask the chatbot. About a quarter of students in each arm send at least one message, with no detectable difference between arms (Table 8, Panel B). Treated students are $\hat{\beta} = 9.3$ pp more likely to open with a question about AI's impact on jobs, a roughly tenfold increase over the 1.0 percent control base ($p < 0.001$); the point estimate is larger still among high-exposure students, $\hat{\beta}^{\text{high}} = 10.0$ pp (Table 8, Panel B). In contrast, questions about salary are less likely, $\hat{\beta} = -2.9$ pp against a 4.9 percent base ($p < 0.001$). Questions about careers are also less likely, $\hat{\beta} = -3.8$ pp against a control base of 8.4 percent ($p < 0.001$).[19]

The comments students leave at the end of the survey show a similar shift. Treated students are $\hat{\beta} = 1.0$ pp more likely to express anxiety about AI's impact on their own future ($p < 0.001$), whereas control students essentially never do (Appendix Table A11). Student interest shifts from pay and careers to AI risk. The move mirrors the change in reasons for major choice in Section 4.2, in which treated students replace generic concerns like job prospects with concern about AI displacement.

7 Conclusion

How much AI disrupts labor markets depends partly on whether workers, both present and future, adapt to the technology. This paper estimates how strongly career choices respond to information about occupational AI exposure. Using randomized information to shift perceived exposure, we find a response that is statistically significant but modest in size. Treated students shift stated choices toward less-exposed occupations, are more likely to consider changing majors, and more often cite avoiding AI displacement as a reason for their major choice. They revise employment and salary prospects in exposed occupations downward and seek information about AI risk rather than about salary or job prospects. The strongest responses come from students whose own plans are most exposed.

Because the visualization costs almost nothing per additional student and reaches students before they make occupation-specific investments, career-services offices could offer it at scale. Yet the effects are limited. A response of a few percentage points on stated

[19] Students could also request a meeting with a career counselor at the end of the survey. Treatment might raise requests if it increased demand for guidance and students faced frictions in booking on their own. Take-up is extremely low in both arms (below 2 percent), and treated students do not request the service more (Table 8, Panel C). After data collection closed, MSU's career-services office told us that students can and do easily make appointments on their own, which explains the low take-up in both arms.

choices, concentrated among the most-exposed students and larger for occupation choices than for majors, will not re-sort the supply of skills on its own. If AI's labor-market effects prove far-reaching, the supply-side adjustment that could blunt them will require more than providing information.

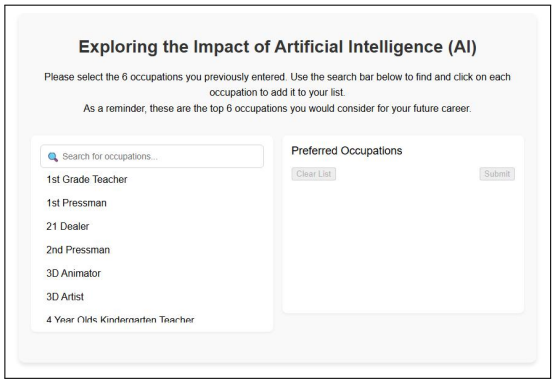
Why is the response so limited? Three patterns in our data suggest that students process the information selectively, thus dampening the effect. First, threatening news reduces search: treated students enter fewer search terms, spend less time on occupation detail panels, and explore fewer occupations, and they report more willingness to switch but less interest in exploring alternatives. Information avoidance of this kind recurs in high-stakes settings (e.g., Karlsson et al., 2009; Sicherman et al., 2016; Golman et al., 2017; Ganguly and Tasoff, 2017). Second, the information moves furthest for the students whose priors already accept it: those who expect AI to destroy jobs revise employment beliefs more, while those who do not discount the message, perhaps because people grant prior-congruent information more credibility (Rabin and Schrag, 1999; Gentzkow and Shapiro, 2006). Third, students separate what will happen to their occupation from what will happen to them: they acknowledge that their occupation will employ fewer people, yet revise their own employment prospects by less and their own earnings not at all. A survey of workers finds a similar invulnerability: respondents expect AI to affect other people's jobs more than their own (Barrera-Jimenez et al., 2025). Future work should try to assess the relative importance of these channels.

Our results point to a few ways in which AI may shape the future of higher education. Among enrollees, information about AI exposure lowers stated demand for college: treated students report a higher probability of leaving college and no change in their plans for graduate school. What the same information does to students who have not yet enrolled we cannot say, since every student in our sample is enrolled. Within college, it shifts stated demand across majors toward fields students perceive as less exposed to AI: treated students report a higher probability of changing major and justify their choice by avoiding technological displacement in place of the traditional reasons. If these stated responses carry into enrollment decisions, colleges would see fewer students overall, fewer in the majors that lead to exposed occupations, and more in those that lead to less-exposed ones.

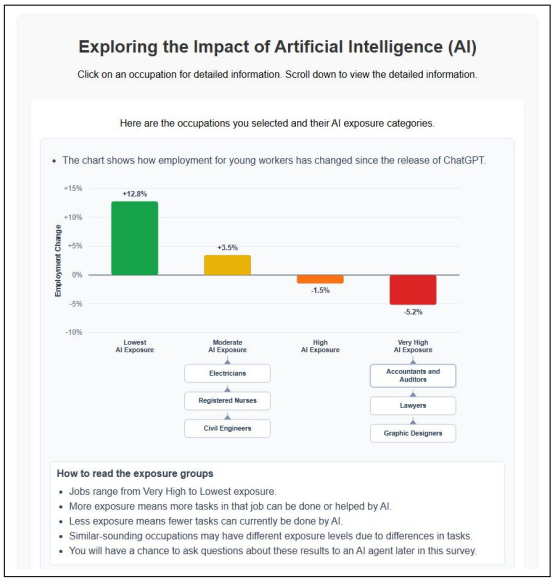
Figures and Tables

Figure 1: The AI-Exposure Information Treatment

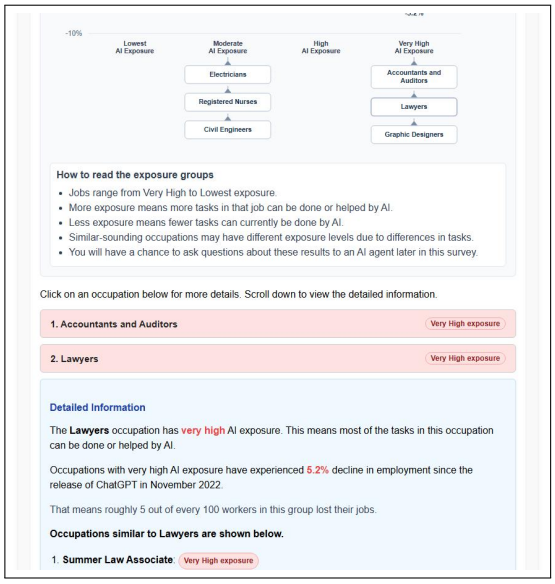
Panel A. Occupation selection



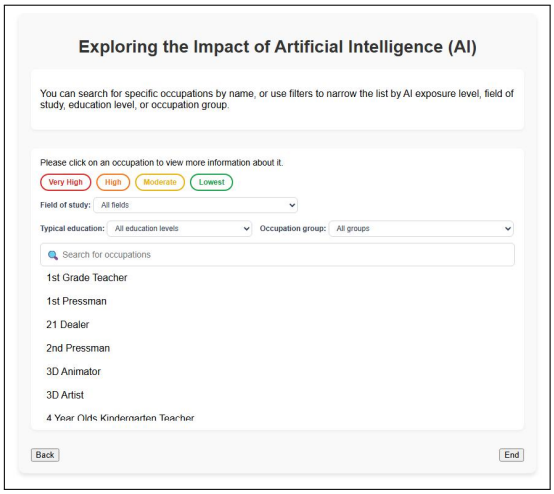
Panel B. Exposure display



Panel C. Occupation detail panel



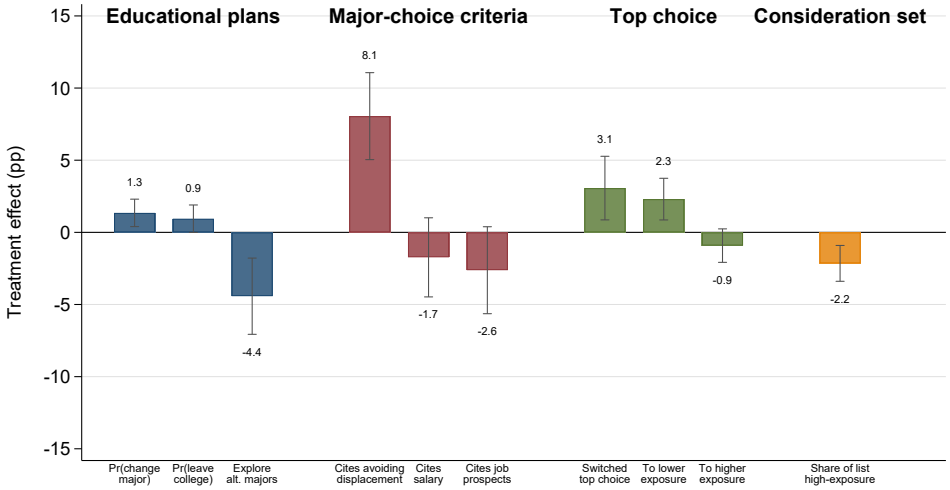
Panel D. Free search with filters



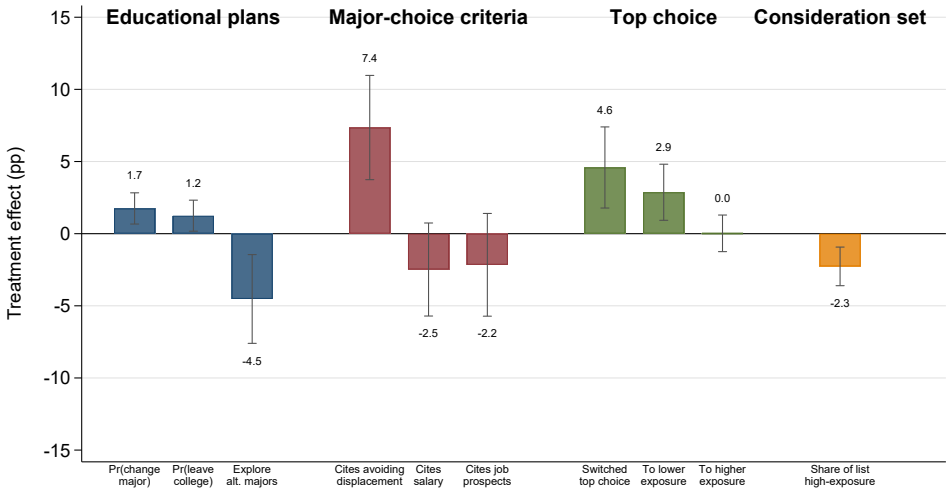
Notes: This figure shows the four screens of the treatment visualization, which ranks 775 six-digit SOC occupations into four AI-exposure tiers (Very High, High, Moderate, Lowest) derived from the Anthropic Economic Index. Panel A is the occupation-selection screen, Panel B the tiered exposure display, Panel C the detail panel for a selected occupation, and Panel D the free-search screen.

Figure 2: Stated Choices: Treatment Effects

Panel A. Full sample



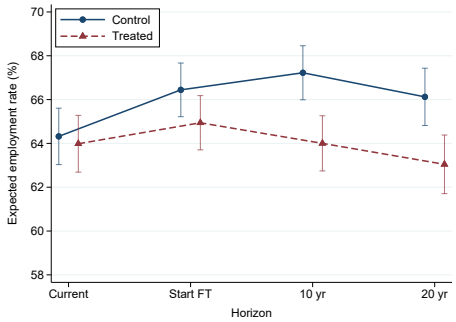
Panel B. High-exposure subgroup



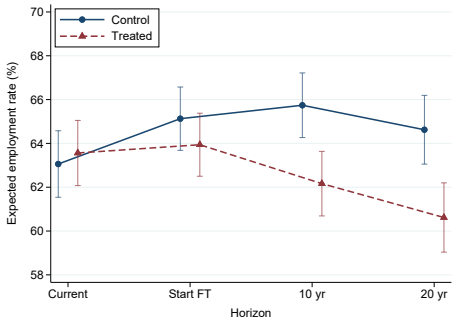
Notes: This figure plots the treatment effect, in percentage points, on each stated-choice outcome, estimated from equation (1) with double-lasso control selection and heteroskedasticity-robust standard errors. The stated probabilities of changing major and of leaving college are elicited on seven-point scales recoded to 0–100 midpoints; the remaining outcomes are indicators or shares scaled to percentage points. The switching-direction outcomes compare the exposure tier of the post-treatment top choice with the baseline tier, with non-switchers as zeros (Table 4); the consideration-set outcome is the share of listed occupations in the two most-exposed tiers. Panel A uses the full sample; Panel B restricts to students whose baseline top choice falls in the Very High or High exposure tiers. Capped vertical bars are 95 percent confidence intervals, and printed values are the point estimates.

Figure 3: Beliefs About Employment, Own Employment, and Salary by Horizon

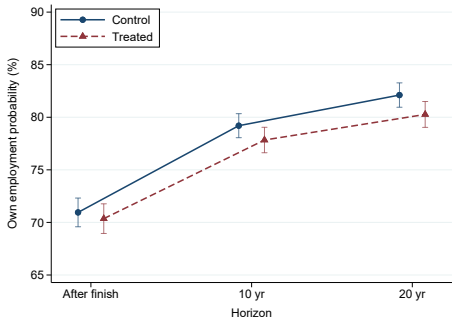
Panel A. Expected market employment rate, full sample



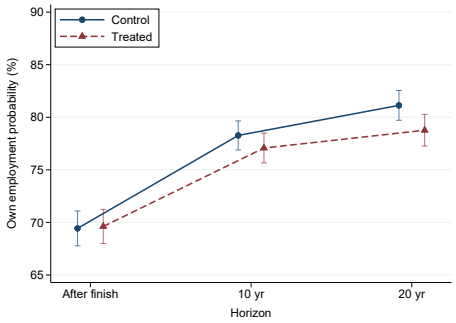
Panel B. Expected market employment rate, high-exposure sample



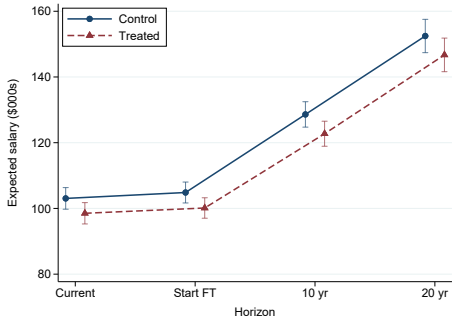
Panel C. Own employment probability, full sample



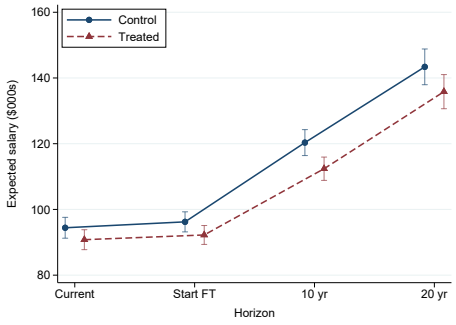
Panel D. Own employment probability, high-exposure sample



Panel E. Expected market salary ($000s), full sample

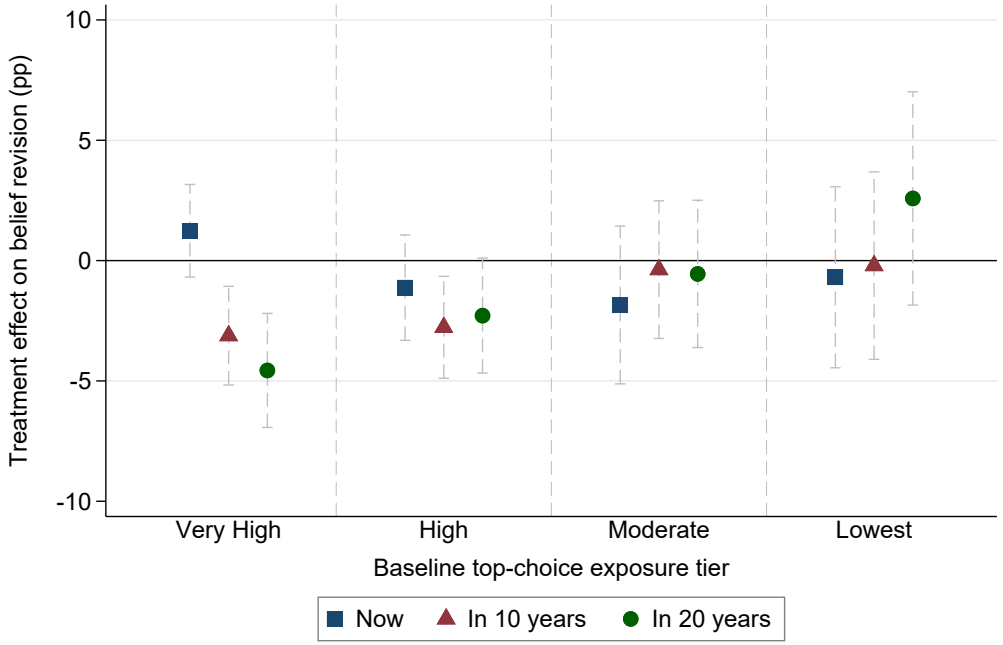


Panel F. Expected market salary ($000s), high-exposure sample



Notes: This figure plots post-treatment mean belief levels by forecast horizon for control students (navy circles) and treated students (maroon triangles). The left column uses the full sample; the right column restricts to students whose baseline top-choice occupation is in the Very High or High exposure tiers. Panels A–B show the expected employment rate (%) in the student's top-choice occupation; Panels C–D show the student's own probability of being employed in the top-choice occupation (%); Panels E–F show the expected salary (in $1,000s). Capped vertical bars are 95 percent confidence intervals for each group mean; the two series are offset horizontally at each horizon so the intervals remain legible. Means use the students who reached the visualization screen and answered the item.

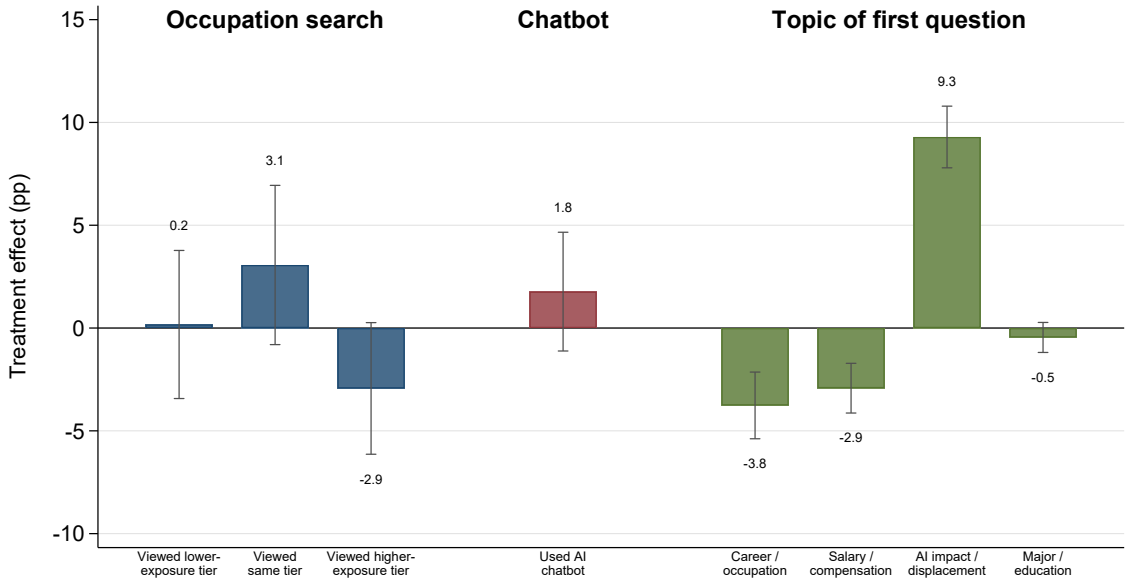
Figure 4: Belief Revision by Baseline Exposure Tier



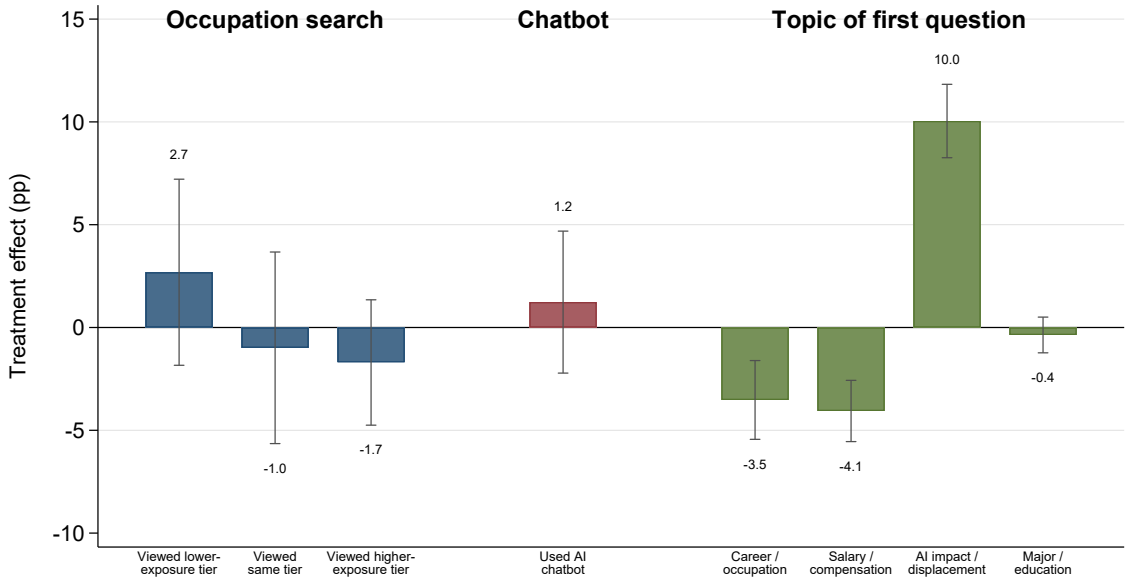
Notes: This figure plots the treatment effect, in percentage points, on the post-minus-pre change in the expected employment rate in the student's top-choice occupation, estimated separately for students whose baseline top choice falls in each AI-exposure tier and shown at three forecast horizons: now (squares), in 10 years (triangles), and in 20 years (circles). Estimates are unadjusted treated-minus-control differences with heteroskedasticity-robust standard errors; capped vertical bars are 95 percent confidence intervals. The three horizons are offset horizontally within each tier for legibility.

Figure 5: Demand for Information: Treatment Effects

Panel A. Full sample

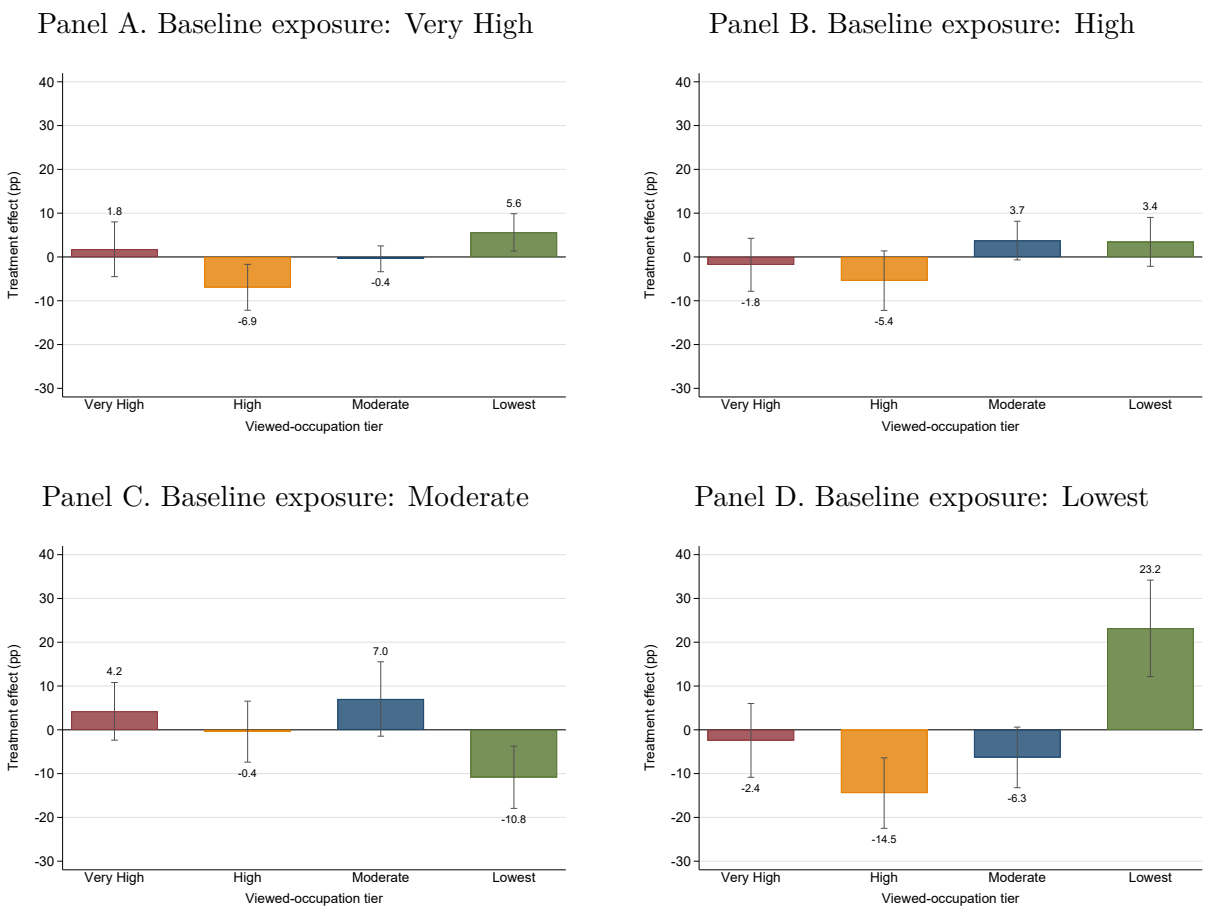


Panel B. High-exposure subgroup



Notes: This figure plots the treatment effect, in percentage points, on the main demand-for-information outcomes, estimated from equation (1) with double-lasso control selection and heteroskedasticity-robust standard errors. Panel A uses the full sample; Panel B restricts to students whose baseline top choice falls in the Very High or High exposure tiers. Capped vertical bars are 95 percent confidence intervals, and printed values are the point estimates.

Figure 6: Reallocation of Exploration by Baseline Exposure Tier



Notes: This figure plots treatment effects, in percentage points, on the share of freely viewed occupations falling in each exposure tier, for students whose baseline top-choice occupation is in the tier named in the panel caption. Estimates are unadjusted treated-minus-control differences with heteroskedasticity-robust standard errors; capped vertical bars are 95 percent confidence intervals and printed values are the point estimates. The sample is the 2,015 students who reached the visualization screen, have a ranked baseline top choice, and have defined shares. Shares are defined only for students who opened at least one occupation on the free-search screen. Treatment lowers that probability (55 against 61 percent), so the estimates compare students who chose to search. Appendix Table A9 reports the underlying group means.

Table 1: Summary Statistics

	Survey sample (1)	MSU Registrar (2)	MSU (IPEDS) (3)	U.S. 4-year (IPEDS) (4)
Panel A. Demographics				
Female	0.642	0.515	0.514	0.568
White	0.763	0.668	0.669	0.478
Black	0.077	0.068	0.065	0.117
Hispanic	0.069	0.072	0.062	0.197
Asian	0.158	0.078	0.080	0.074
Panel B. Field of study				
STEM	0.437	0.382	0.371	0.243
Business & social sciences	0.325	0.332	0.332	0.362
Arts & humanities	0.136	0.209	0.156	0.148
Health & education	0.102	0.077	0.093	0.199
Other	0.000	0.000	0.048	0.049
Panel C. Class year				
Freshman	0.205	0.247	–	–
Sophomore	0.241	0.236	–	–
Junior	0.253	0.250	–	–
Senior or beyond	0.302	0.267	–	–
Panel D. Engagement with the visualization				
Time on visualization (sec)	144.2	–	–	–
Occupations explored	1.4	–	–	–
Search terms entered	1.3	–	–	–
N (students)	3,471	41,415	40,483	11,164,684

Notes: This table presents summary statistics for our survey sample alongside three benchmarks. Column (1) reports means for the 3,471 students who reached the visualization screen. Column (2) reports figures from the MSU Office of the Registrar: demographics (Panel A) and class year (Panel C) are both Fall 2025 enrollment; field of study (Panel B) is Spring 2026 bachelor's degrees conferred, classified into four broad areas using department and major names. Column (3) reports MSU figures from IPEDS (Fall 2023 enrollment; field of study from 2023–24 bachelor's degrees conferred). Column (4) covers all four-year, degree-granting, Title IV institutions: sex and race/ethnicity from IPEDS fall enrollment, and field of study from bachelor's degrees conferred by two-digit CIP code (IPEDS Completions). Demographics (Panel A), field of study (Panel B), and class year (Panel C) are proportions; engagement (Panel D) is measured in seconds or counts. Class year (Panel C) has no IPEDS benchmark, and engagement (Panel D) has no institutional benchmark in any column; both are marked "–". The final row reports the survey sample size in column (1) and total undergraduate enrollment in columns (2)–(4). In column (2), the Panel B figures come from Spring 2026 bachelor's graduates ($N = 6{,}940$) and the Panel C figures from the companion class-year report's own total ($N = 40{,}869$), both drawn from different Registrar reports than the enrollment total shown.

Table 2: Balance by Treatment Assignment and Top-Occupation Exposure

	Full sample			High-exposure subgroup		
	Mean control	$\hat{\beta}$	(SE)	Mean control	$\hat{\beta}$	(SE)
Panel A. Demographics						
Female	0.645	−0.006	(0.018)	0.611	−0.013	(0.022)
White	0.760	0.005	(0.016)	0.743	0.021	(0.019)
Black	0.077	−0.001	(0.010)	0.073	0.000	(0.011)
Hispanic	0.068	0.003	(0.009)	0.074	0.000	(0.012)
Asian	0.163	−0.009	(0.013)	0.182	−0.023	(0.017)
Panel B. Field of study and major						
STEM	0.432	0.009	(0.017)	0.387	−0.012	(0.020)
Business & social sciences	0.335	−0.019	(0.016)	0.391	−0.012	(0.020)
Arts & humanities	0.135	0.002	(0.012)	0.154	0.006	(0.015)
Health & education	0.098	0.008	(0.010)	0.067	0.018*	(0.011)
Has declared major	0.958	0.000	(0.007)	0.952	0.004	(0.008)
Major satisfaction (1–5)	4.127	0.023	(0.028)	4.120	0.007	(0.033)
Job relevance to major (1–5)	4.140	−0.035	(0.029)	4.137	−0.025	(0.035)
Prob. change major (%)	12.367	−0.127	(0.702)	12.543	0.199	(0.846)
Panel C. Class year and academic background						
Freshman	0.208	−0.006	(0.014)	0.211	−0.012	(0.016)
Sophomore	0.241	0.001	(0.015)	0.230	−0.002	(0.017)
Junior	0.254	−0.002	(0.015)	0.252	0.003	(0.018)
Senior or beyond	0.298	0.008	(0.016)	0.307	0.010	(0.019)
Self-ranked ability (10=top, 90=bottom)	49.002	−0.520	(0.816)	49.456	−0.687	(0.977)
Years until full-time work	2.900	0.005	(0.078)	2.634	0.030	(0.088)
Prob. grad school (%)	61.950	−0.876	(1.207)	59.004	0.068	(1.423)
Prob. leave college (%)	4.994	−0.531	(0.457)	4.658	0.042	(0.541)
Num. occupations listed (1–6)	4.678	−0.076	(0.053)	4.683	−0.058	(0.062)
Panel D. Baseline career beliefs						
Empl. rate, current (%)	62.169	0.025	(0.873)	61.398	0.576	(0.998)
Empl. rate, start FT (%)	64.933	−1.106	(0.848)	63.610	−0.618	(0.985)
Empl. rate, 10yr (%)	66.686	−0.824	(0.838)	65.222	−0.326	(0.986)
Empl. rate, 20yr (%)	65.500	−0.678	(0.907)	63.791	−0.239	(1.073)
Expected salary, current (thousands)	102.981	−2.947	(2.261)	94.731	−2.431	(2.172)
Expected salary, start FT (thousands)	101.660	−3.182	(1.981)	93.546	−1.918	(1.912)
Expected salary, 10yr (thousands)	131.504	−3.552	(2.720)	123.466	−4.825*	(2.722)
Expected salary, 20yr (thousands)	158.255	−1.429	(3.713)	150.518	−4.116	(3.924)
Own empl. prob., finish ed. (%)	71.448	0.030	(0.926)	70.560	0.297	(1.101)
Own empl. prob., 10yr (%)	82.102	0.066	(0.737)	81.870	0.092	(0.868)
Own empl. prob., 20yr (%)	84.418	0.826	(0.751)	84.185	0.635	(0.899)
Expected years in top occupation	28.414	−0.692	(0.444)	27.804	−0.950*	(0.530)
Expected earnings at 35 (thousands)	112.281	−1.601	(1.244)	110.685	−1.385	(1.405)

Notes: The table continues in the next page.

Table 2: Balance by Treatment Assignment and Top-Occupation Exposure (continued)

	Full sample			High-exposure subgroup		
	Mean control	$\hat{\beta}$	(SE)	Mean control	$\hat{\beta}$	(SE)
Panel E. AI knowledge and use						
Uses generative AI weekly or more	0.419	−0.019	(0.017)	0.449	0.004	(0.020)
Believes AI destroys jobs, current	0.375	0.009	(0.016)	0.397	0.033*	(0.020)
Believes AI destroys jobs, start FT	0.466	0.001	(0.017)	0.493	0.007	(0.020)
Believes AI destroys jobs, 10yr	0.617	0.004	(0.016)	0.649	0.007	(0.019)
Believes AI destroys jobs, 20yr	0.639	0.001	(0.016)	0.662	0.007	(0.019)
Panel F. Baseline reasons for the major choice						
Selected 'salary'	0.482	−0.002	(0.017)	0.507	−0.017	(0.020)
Selected 'job prospects'	0.540	−0.009	(0.017)	0.549	−0.013	(0.020)
Selected 'innovative tech.'	0.156	−0.028**	(0.012)	0.172	−0.031**	(0.015)
Selected 'diverse skills'	0.503	−0.005	(0.017)	0.496	−0.003	(0.020)
Selected 'avoid tech displ.'	0.233	−0.007	(0.014)	0.248	−0.020	(0.017)
N	1729	3471		1216	2468	
Joint F-test (p-value)		1.13	[0.24]		1.05	[0.37]

Notes: This table presents summary statistics and balance for the 3,471 students who reached the visualization screen. Demographics (Panel A) are self-reported at the end of the survey; the other characteristics are measured before treatment or fixed at baseline. Field of study groups each student's baseline major into broad areas. Panel E rows are indicators for using AI at least weekly and for believing AI will on net reduce the number of jobs in the top-choice occupation at each horizon. Each $\hat{\beta}$ is from a regression of the characteristic on the treatment indicator, with heteroskedasticity-robust standard errors in parentheses. The high-exposure subgroup comprises students whose top-choice occupation falls in the Very High or High AEI exposure tiers. Probabilities are recoded to 0–100 midpoints; salaries are in thousands of dollars, winsorized at the 95th percentile. The joint F-test covers 37 covariates: the academic, belief, AI, and reason items of Panels B to F (class year as one cohort variable, the AI items on their original scales, plus four baseline earnings bins not shown), not the demographic or field-of-study rows. Missing covariate values are set to the sample mean and flagged with an indicator, so every student enters the test. Its variance is the classical one, which is valid under the null of random assignment. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table 3: The Effect of AI-Exposure Information on Choices and Major-Choice Criteria

Outcome	Full sample		High-exposure subgroup	
	Mean control	$\hat{\beta}$	Mean control	$\hat{\beta}^{\text{high}}$
	(1)	(2)	(3)	(4)
Panel A. Major and occupation choices and list stability				
Switched preferred major	0.353	0.011	0.351	0.023
		(0.016)		(0.019)
Switched top choice occupation	0.095	0.031***	0.104	0.046***
		(0.011)		(0.014)
Consideration set size	4.302	−0.084**	4.304	−0.089*
		(0.042)		(0.049)
Share of occupations retained	0.618	−0.020*	0.621	−0.038***
		(0.011)		(0.013)
Num. occupations retained	2.805	−0.124**	2.811	−0.204***
		(0.051)		(0.061)
Rank distance (pre vs post)	1.528	0.026	1.500	0.099
		(0.073)		(0.086)
Rank distance, top coded	5.874	0.325**	5.812	0.651***
		(0.157)		(0.190)
Kendall τ (pre vs post)	0.226	−0.024	0.235	−0.068***
		(0.020)		(0.023)
Panel B. Reasons cited for the major choice				
Reason: expected salary	0.645	−0.017	0.667	−0.025
		(0.014)		(0.016)
Reason: job prospects	0.634	−0.026*	0.646	−0.022
		(0.015)		(0.018)
Reason: innovative technology	0.189	0.027**	0.205	0.022
		(0.011)		(0.014)
Reason: develop diverse skills	0.481	−0.024	0.467	−0.023
		(0.015)		(0.018)
Reason: avoid tech displacement	0.362	0.081***	0.368	0.074***
		(0.015)		(0.018)
Panel C. Family *z*-score index				
Family z-score: Stated choices	0.000	0.027**	−0.009	0.042***
		(0.011)		(0.014)
N	1729	3471	1216	2468

Notes: This table presents treatment effects on stated choices and reasons behind major choice for the 3,471 students who reached the visualization screen; columns and layout as in Table 8. The N row reports the design sample; estimation samples differ by row. Each outcome is a post-treatment value. For "Consideration set size" and the reason items, the pre-treatment value enters the double-lasso candidate pool. The switching and list-stability outcomes and the index have no pre-treatment counterpart. "Switched preferred major" and "Switched top choice occupation" are changes of program and of occupation as judged by coders, with rewordings and menu-path artifacts as zeros and the few pairs the coders could not judge as missing (Appendix B.5); the rank-distance measures sum absolute rank changes over the occupation list, and the top-coded version assigns dropped occupations one rank below the bottom of the post list. Panel C reports the stated-choices family z-score index. It averages the components listed in Appendix B.1, from this table and Table 5, each standardized and signed in the predicted direction. All specifications use double-lasso control selection (Belloni et al., 2014). *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table 4: Direction of Top-Choice Switching and List Reshuffling

Outcome	Full sample		High-exposure subgroup	
	Mean control	$\hat{\beta}$	Mean control	$\hat{\beta}^{\text{High}}$
	(1)	(2)	(3)	(4)
Switched top choice	0.095	0.031***	0.104	0.046***
		(0.011)		(0.014)
Switched to less-exposed rank	0.044	0.026***	0.060	0.031***
		(0.008)		(0.011)
Switched to lower exposure tier	0.031	0.023***	0.041	0.029***
		(0.007)		(0.010)
Switched within same tier	0.033	0.016**	0.042	0.020**
		(0.007)		(0.010)
Switched to higher exposure tier	0.031	−0.009	0.022	0.000
		(0.006)		(0.006)
Rank distance, top coded	5.874	0.325**	5.812	0.651***
		(0.157)		(0.190)
Rank distance, top coded per occ.	1.142	0.047	1.136	0.104***
		(0.031)		(0.037)
N	1729	3471	1216	2468

Notes: This table presents treatment effects on the direction of top-choice switching and list reshuffling. The N row reports the design sample; estimation samples differ by row. Rows 3 to 5 are indicators for switching the top choice to a lower, the same, or a higher AI-exposure tier; students who keep their top choice are zeros in all three, so the three rows sum to the switching rate in row 1 up to switches that involve an occupation without an exposure rank. Row 2 uses the continuous ranking instead of tiers and equals one when the new top choice is strictly less exposed than the old one (a higher rank number). The top-coded rank distance counts the rank positions each baseline occupation moves, charging a dropped occupation the distance to one rank below the bottom of the post-treatment list; the per-occupation version divides by the number of baseline occupations. Each outcome occupies two rows: coefficient (top) and heteroskedasticity-robust standard error (bottom); columns as in Table 8. All specifications use double-lasso control selection (Belloni et al., 2014). *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table 5: The Effect of AI-Exposure Information on Educational Plans

Outcome	Full sample		High-exposure subgroup	
	Mean control	$\hat{\beta}$	Mean control	$\hat{\beta}^{\text{high}}$
	(1)	(2)	(3)	(4)
Prob. change major (pp)	12.023	1.350***	11.806	1.750***
		(0.485)		(0.552)
Prob. grad school (pp)	61.963	0.264	59.429	−0.157
		(0.509)		(0.581)
Prob. leave college (pp)	6.149	0.947*	5.541	1.247**
		(0.485)		(0.549)
Interested in exploring majors	0.201	−0.044***	0.192	−0.045***
		(0.013)		(0.016)
N	1729	3471	1216	2468

Notes: This table presents treatment effects on educational plans for the 3,471 students who reached the visualization screen; columns and layout as in Table 8. The N row reports the design sample; estimation samples differ by row. Each outcome is a post-treatment value. For the three probabilities, the pre-treatment value enters the double-lasso candidate pool. "Interested in exploring majors" was asked only after treatment. Probabilities are stated on seven-point scales recoded to 0–100 midpoints; except for the probability of leaving college, which has no predicted direction, these outcomes enter the stated-choices family z-score index reported in Table 3. All specifications use double-lasso control selection (Belloni et al., 2014). *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table 6: The Effect of AI-Exposure Information on Beliefs

Outcome	Full sample		High-exposure subgroup	
	Mean control	$\hat{\beta}$	Mean control	$\hat{\beta}^{\text{high}}$
	(1)	(2)	(3)	(4)
Panel A. Employment-rate beliefs				
Empl. rate, current (pp)	64.320	0.043	63.059	0.445
		(0.580)		(0.681)
Empl. rate, start FT (pp)	66.444	−0.754	65.129	−0.832
		(0.536)		(0.636)
Empl. rate, 10yr (pp)	67.224	−2.404***	65.741	−3.105***
		(0.563)		(0.671)
Empl. rate, 20yr (pp)	66.125	−2.453***	64.624	−3.615***
		(0.629)		(0.756)
Panel B. Own employment probability and occupational tenure				
Own empl. prob., finish ed. (pp)	70.951	−0.243	69.430	0.069
		(0.560)		(0.693)
Own empl. prob., 10yr (pp)	79.199	−1.138**	78.273	−1.115*
		(0.533)		(0.660)
Own empl. prob., 20yr (pp)	82.111	−1.956***	81.130	−2.238***
		(0.573)		(0.716)
Expected years in top occ.	28.619	−0.008	28.100	−0.099
		(0.181)		(0.216)
Panel C. Salary expectations				
Log expected salary, current	4.457	−0.037**	4.402	−0.023
		(0.015)		(0.018)
Log expected salary, start FT	4.512	−0.016*	4.459	−0.013
		(0.009)		(0.010)
Log expected salary, 10yr	4.715	−0.023***	4.674	−0.027***
		(0.008)		(0.010)
Log expected salary, 20yr	4.853	−0.021**	4.809	−0.014
		(0.010)		(0.012)
Panel D. Family *z*-score indices				
Family z-score: Beliefs about employment	0.000	0.043***	0.048	0.053***
		(0.013)		(0.016)
Family z-score: Beliefs about earnings	0.000	0.029***	0.066	0.025**
		(0.010)		(0.011)
N	1729	3471	1216	2468

Notes: This table presents treatment effects on post-treatment beliefs for the 3,471 students who reached the visualization screen; columns and layout as in Table 8. The *N* row reports the design sample; estimation samples differ by row. Each outcome's pre-treatment value enters the double-lasso candidate pool (for the log salaries, the salary level). The two indices have no pre-treatment counterpart. Employment rates and own employment probabilities are in percentage points; salary outcomes are logs of expected salary in thousands of dollars (winsorized at the 95th percentile), so coefficients read as approximate percent changes. Panel D reports the two belief family *z*-score indices (components in Appendix B.1), signed so positive values denote the pre-specified direction (downward revisions, and a higher variance of earnings at age 35). *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table 7: Mediation: Channels Behind the Major- and Occupation-Switching Responses

Mediator	Treatment → mediator ($\hat{\alpha}_1$)	Indirect effect ($\hat{\beta} - \tilde{\beta}$, pp)	Share of total effect (%)
Panel A. Probability of changing major			
Downward salary revision (z)	0.053**	0.06	6.7
	(0.025)		
Downward own-employment revision (z)	0.072**	0.07	6.5
	(0.029)		
Weight on avoiding displacement	0.081***	0.13	10.2
	(0.015)		
All three channels jointly		0.24	26.8
Panel B. Top-choice occupation switching			
Downward salary revision (z)	0.053**	0.27	10.1
	(0.025)		
Downward own-employment revision (z)	0.072**	0.36	11.5
	(0.029)		
Weight on avoiding displacement	0.081***	0.24	7.9
	(0.015)		
All three channels jointly		0.83	30.9

Notes: This table presents the mediation decomposition of the treatment effects on the stated probability of changing major (Panel A, measured on a 0–100 scale) and on top-choice occupation switching (Panel B), following the approach of Fagereng et al. (2021) and Finkelstein et al. (2024): the total effect $\hat{\beta}$ on the panel's outcome is re-estimated on each mediator's estimation sample, $\tilde{\beta}$ is the treatment coefficient after adding the mediator to the same specification flexibly—bin fixed effects with a separate bin for no revision plus deciles of the nonzero revisions for the belief mediators, and full saturation for the displacement mediator, which takes three values—and the indirect effect $\hat{\beta} - \tilde{\beta}$ (in percentage points of the outcome) is expressed in the final column as a share of the total. The first column reports the treatment effect on the mediator ($\hat{\alpha}_1$), which does not involve the outcome and is therefore identical across panels; the last row of each panel enters all three mediators jointly. The belief mediators are z-score indices over the post-minus-pre revisions, signed so that revision in the direction of the information is positive; the displacement mediator is the change in citing technological displacement as a major-choice factor. Both panels use the full sample. All equations use double-lasso controls and heteroskedasticity-robust standard errors (in parentheses). Because mediators are not randomly assigned, the decomposition is descriptive. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table 8: Demand for Information: Search Behavior, Chatbot Use, and Question Topics

	Full sample		High-exposure subgroup	
	Mean control	$\hat{\beta}$	Mean control	$\hat{\beta}^{\text{high}}$
	(1)	(2)	(3)	(4)
Panel A. Search intensity				
Occupations explored (viz)	1.430	−0.151*	1.394	−0.141
		(0.085)		(0.109)
Detail panel time (sec)	12.289	−2.750***	11.974	−2.888***
		(0.562)		(0.666)
Search terms entered	1.378	−0.245***	1.354	−0.282***
		(0.076)		(0.090)
Share of viewed occupations in each exposure tier				
Viewed Very High exposure tier	0.339	0.016	0.416	0.012
		(0.019)		(0.024)
Viewed High exposure tier	0.308	−0.066***	0.343	−0.068***
		(0.017)		(0.022)
Viewed Moderate exposure tier	0.135	0.020	0.079	0.012
		(0.014)		(0.013)
Viewed Lowest exposure tier	0.219	0.031*	0.162	0.044**
		(0.016)		(0.018)
Panel B. Chatbot engagement				
Used AI chatbot	0.252	0.018	0.260	0.012
		(0.015)		(0.018)
Topic of student's first question				
Career / occupation	0.084	−0.038***	0.081	−0.035***
		(0.008)		(0.010)
Salary / compensation	0.049	−0.029***	0.058	−0.041***
		(0.006)		(0.008)
AI impact / displacement	0.010	0.093***	0.007	0.100***
		(0.008)		(0.009)
Major / education planning	0.014	−0.005	0.014	−0.004
		(0.004)		(0.004)
Other	0.094	−0.004	0.100	−0.011
		(0.010)		(0.012)
Panel C. Requests for further help				
Wants career counselor appointment	0.018	−0.001	0.020	−0.003
		(0.005)		(0.006)
N	1729	3471	1216	2468

Notes: This table presents treatment effects on demand-for-information outcomes for the 3,471 students who reached the visualization screen. The N row reports the design sample; estimation samples differ by row. The four tier-share rows of Panel A use the 2,023 students who have defined shares. Shares are defined only for students who opened at least one occupation on the free-search screen. Treatment lowers that probability (55 against 61 percent), so the estimates compare students who chose to search. Panel B classifies the primary topic of the student's first chatbot message with a language-model classifier (Appendix B.4); non-messagers are coded zero on every topic. Each outcome occupies two rows: coefficient and heteroskedasticity-robust standard error. Columns 1–2 give the full-sample control mean and treatment effect $\hat{\beta}$ from equation (1) with double-lasso controls (Belloni et al., 2014); columns 3–4 repeat for the high-exposure subgroup (AEI ranking ≤ 40). *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

References

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics*, volume 4, pages 1043–1171. Elsevier.

Acemoglu, D. and Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30.

Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6):2188–2244.

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

American Academy of Arts and Sciences (2025). From matriculation to completion: How college students move between majors. Bulletin of the American Academy of Arts and Sciences, Winter 2025; Humanities Indicators analysis of National Student Clearinghouse data.

Autor, D. H. (2019). Work of the past, work of the future. In *AEA Papers and Proceedings*, volume 109, pages 1–32.

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

Baker, R., Bettinger, E., Jacob, B., and Marinescu, I. (2018). The effect of labor market information on community college students' major choice. *Economics of Education Review*, 65:18–30.

Barrera-Jimenez, F., Arroyo-Barrigüete, J. L., Garrido-Merchán, E. C., and Grinda-Luna, G. (2025). Invulnerability bias in perceptions of artificial intelligence's future impact on employment. *Scientific Reports*, 15:28698.

Becker, G. S. (1962). Investment in human capital: A theoretical analysis. *Journal of Political Economy*, 70(5, Part 2):9–49.

Belloni, A., Chernozhukov, V., and Hansen, C. (2014). High-Dimensional Methods and Inference on Structural and Treatment Effects. *Journal of Economic Perspectives*, 28(2):29–50.

Bick, A., Blandin, A., and Deming, D. J. (2026a). The rapid adoption of generative AI. *Management Science*. Articles in Advance.

Bick, A., Blandin, A., Deming, D. J., Fuchs-Schündeln, N., and Jessen, J. (2026b). Mind the gap: AI adoption in Europe and the US. Brookings Papers on Economic Activity, conference draft.

Blom, E., Cadena, B. C., and Keys, B. J. (2021). Investment over the business cycle: Insights from college major choice. *Journal of Labor Economics*, 39(4):1043–1082.

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

Brynjolfsson, E., Li, D., and Raymond, L. (2025b). Generative AI at work. *The Quarterly Journal of Economics*, 140(2):889–942.

Brynjolfsson, E. and Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science*, 358(6370):1530–1534.

Card, D., Kluve, J., and Weber, A. (2018). What works? A meta-analysis of recent active labor market program evaluations. *Journal of the European Economic Association*, 16(3):894–931.

Cascio, E. U. and Narayan, A. (2022). Who needs a fracking education? The educational response to low-skill-biased technological change. *ILR Review*, 75(1):56–89.

Conlon, J. J. (2021). Major malfunction: A field experiment correcting undergraduates' beliefs about salaries. *Journal of Human Resources*, 56(3):922–939.

Conlon, J. J. and Patel, D. (2026). What jobs come to mind? Stereotypes about fields of study. *Quarterly Journal of Economics*. forthcoming.

Cui, Z. K., Demirer, M., Jaffe, S., Musolff, L., Peng, S., and Salz, T. (2026). The effects of generative AI on high-skilled work: Evidence from three field experiments with software developers. *Management Science*. Articles in Advance.

Dauth, W., Findeisen, S., Suedekum, J., and Woessner, N. (2021). The adjustment of labor markets to robots. *Journal of the European Economic Association*, 19(6):3104–3153.

Dell'Acqua, F., McFowland III, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., and Lakhani, K. R. (2026). Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality. *Organization Science*. Articles in Advance.

Di Giacomo, G. and Lerch, B. (2026). Automation and human capital adjustment: The effect of robots on college enrollment. *Journal of Human Resources*, 61(3):733–760.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702):1306–1308.

Fagereng, A., Mogstad, M., and Rønning, M. (2021). Why do wealthy parents have wealthy children? *Journal of Political Economy*, 129(3):703–756.

Felten, E., Raj, M., and Seamans, R. (2023). How will language modelers like ChatGPT affect occupations and industries? arXiv preprint arXiv:2303.01157.

Finkelstein, A., Notowidigdo, M. J., Schilbach, F., and Zhang, J. (2024). Lives vs. livelihoods: The impact of the great recession on mortality and welfare. Working Paper 32110, National Bureau of Economic Research.

Gallup (2023). More U.S. workers fear technology making their jobs obsolete. September 11, 2023.

Gallup and Lumina Foundation (2026). AI in higher education: Widespread use, unclear rules. State of Higher Education Study, fielded October 2025. Accessed September 2026.

Ganguly, A. and Tasoff, J. (2017). Fantasy and dread: The demand for information and the consumption utility of the future. *Management Science*, 63(12):4037–4060.

Gentzkow, M. and Shapiro, J. M. (2006). Media bias and reputation. *Journal of Political Economy*, 114(2):280–316.

Goldin, C. and Katz, L. F. (2008). *The Race between Education and Technology*. Harvard University Press, Cambridge, MA.

Golin, M. and Rauh, C. (2023). The impact of fear of automation. CEPR Discussion Paper 17816.

Golman, R., Hagmann, D., and Loewenstein, G. (2017). Information avoidance. *Journal of Economic Literature*, 55(1):96–135.

Goos, M., Manning, A., and Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review*, 104(8):2509–2526.

Graetz, G. and Michaels, G. (2018). Robots at work. *Review of Economics and Statistics*, 100(5):753–768.

Handa, K., Tamkin, A., McCain, M., Huang, S., Durmus, E., Heck, S., Mueller, J., Hong, J., Ritchie, S., Belonax, T., Troy, K. K., Amodei, D., Kaplan, J., Clark, J., and Ganguli, D. (2025). Which economic tasks are performed with AI? Evidence from millions of Claude conversations. Working paper, Anthropic.

Hastings, J., Neilson, C. A., and Zimmerman, S. D. (2015). The effects of earnings disclosure on college enrollment decisions. Technical report, National Bureau of Economic Research.

Hosseini Maasoum, S. M. and Lichtinger, G. (2025). Generative AI as seniority-biased technological change: Evidence from U.S. résumé and job posting data. SSRN Working Paper No. 5425555.

Humlum, A. and Vestergaard, E. (2025). Large language models, small labor market effects. NBER Working Paper 33777.

Huseynov, S. (2025). ChatGPT and the labor market: Unraveling the effect of AI discussions on students' earnings expectations. *Journal of Economic Psychology*, 108:102803.

Hyman, B. G., Lahey, B., Ni, K., and Pilossoph, L. (2025). How retrainable are AI-exposed workers? NBER Working Paper 34174.

Jacobson, L. S., LaLonde, R. J., and Sullivan, D. G. (1993). Earnings losses of displaced workers. *American Economic Review*, 83(4):685–709.

Jensen, R. (2010). The (perceived) returns to education and the demand for schooling. *The Quarterly Journal of Economics*, 125(2):515–548.

Karger, E., Kuusela, O., Abaluck, J., Bryan, K., Halperin, B., Jones, T., Murphy, C., Trammell, P., Reynolds, M., Mayland, D., Viswanathan, R., Mittal, A., Ceppas de Castro, R., Rosenberg, J., and Tetlock, P. E. (2026). Forecasting the economic effects of AI. Working paper, Forecasting Research Institute.

Karlsson, N., Loewenstein, G., and Seppi, D. (2009). The ostrich effect: Selective attention to information. *Journal of Risk and Uncertainty*, 38(2):95–115.

Kerwin, J., Rostom, N., and Sterck, O. (2024). Striking the right balance: Why standard balance tests over-reject the null, and how to fix it. IZA Discussion Paper No. 17217.

Kirkebøen, L. J., Leuven, E., and Mogstad, M. (2016). Field of study, earnings, and self-selection. *Quarterly Journal of Economics*, 131(3):1057–1111.

Kling, J. R., Liebman, J. B., and Katz, L. F. (2007). Experimental analysis of neighborhood effects. *Econometrica*, 75(1):83–119.

Kogan, L., Papanikolaou, D., Schmidt, L. D., and Seegmiller, B. (2023). Technology and labor displacement: Evidence from linking patents with worker-level data. Technical report, National Bureau of Economic Research.

Moore, D. A. and Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review*, 115(2):502–517.

National Center for Education Statistics (2017). Beginning college students who change their majors within 3 years of enrollment. Data Point NCES 2018-434, U.S. Department of Education.

National Student Clearinghouse Research Center (2024). Persistence and retention: Fall 2022 beginning postsecondary student cohort. Report, June 27, 2024.

National Student Clearinghouse Research Center (2026). Final spring enrollment trends 2026. Clearinghouse Enrollment Insights, released June 4, 2026.

Noy, S. and Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654):187–192.

Pew Research Center (2017). Automation in everyday life. October 4, 2017.

Pew Research Center (2023). Which U.S. workers are more exposed to AI on their jobs? July 26, 2023; chapter "Workers' views on the risk of AI to their jobs".

Pew Research Center (2025). How the U.S. public and AI experts view artificial intelligence. April 3, 2025.

Pew Research Center (2026). Americans and AI 2026: Chatbots, smart devices and views on impact. Survey report, June 17, 2026; fielded February 2026.

Pizzinelli, C., Panton, M., Tavares, M., Cazzaniga, M., and Li, L. (2023). Labor market exposure to AI: Cross-country differences and distributional implications. IMF Working Paper 23/216, International Monetary Fund.

Rabin, M. and Schrag, J. L. (1999). First impressions matter: A model of confirmatory bias. *Quarterly Journal of Economics*, 114(1):37–82.

Roth, C., Settele, S., and Wohlfart, J. (2022). Risk exposure and acquisition of macroeconomic information. *American Economic Review: Insights*, 4(1):34–53.

Roth, C. and Wohlfart, J. (2020). How do expectations about the macroeconomy affect personal expectations and behavior? *Review of Economics and Statistics*, 102(4):731–748.

Roy, A. D. (1951). Some thoughts on the distribution of earnings. *Oxford Economic Papers*, 3(2):135–146.

Schultz, T. W. (1961). Investment in human capital. *American Economic Review*, 51(1):1–17.

Serrano, J. (2024). Do robots necessarily displace workers? Evidence from Mexican local labor markets. Working paper.

Shen, M., Xiao, Q., Chen, X., Fang, S., and Yue, J. (2026). Generative AI, perceived job displacement, and policy preferences: Experimental evidence from China. *Journal of Economic Behavior & Organization*. Forthcoming.

Sicherman, N., Loewenstein, G., Seppi, D. J., and Utkus, S. P. (2016). Financial attention. *Review of Financial Studies*, 29(4):863–897.

Stinebrickner, R. and Stinebrickner, T. R. (2014). A major in science? initial beliefs and final outcomes for college major and dropout. *Review of Economic Studies*, 81(1):426–472.

The Washington Post (2026). Computer science majors are disappearing. Is AI to blame? April 13, 2026.

U.S. Bureau of Labor Statistics (2025). Occupational projections and worker characteristics, 2024–34. Employment Projections program, Table 1.2.

U.S. Bureau of Labor Statistics (2026). Unemployed people by occupation and sex, 2025 annual averages. Current Population Survey, Household Data Annual Averages, Table 25.

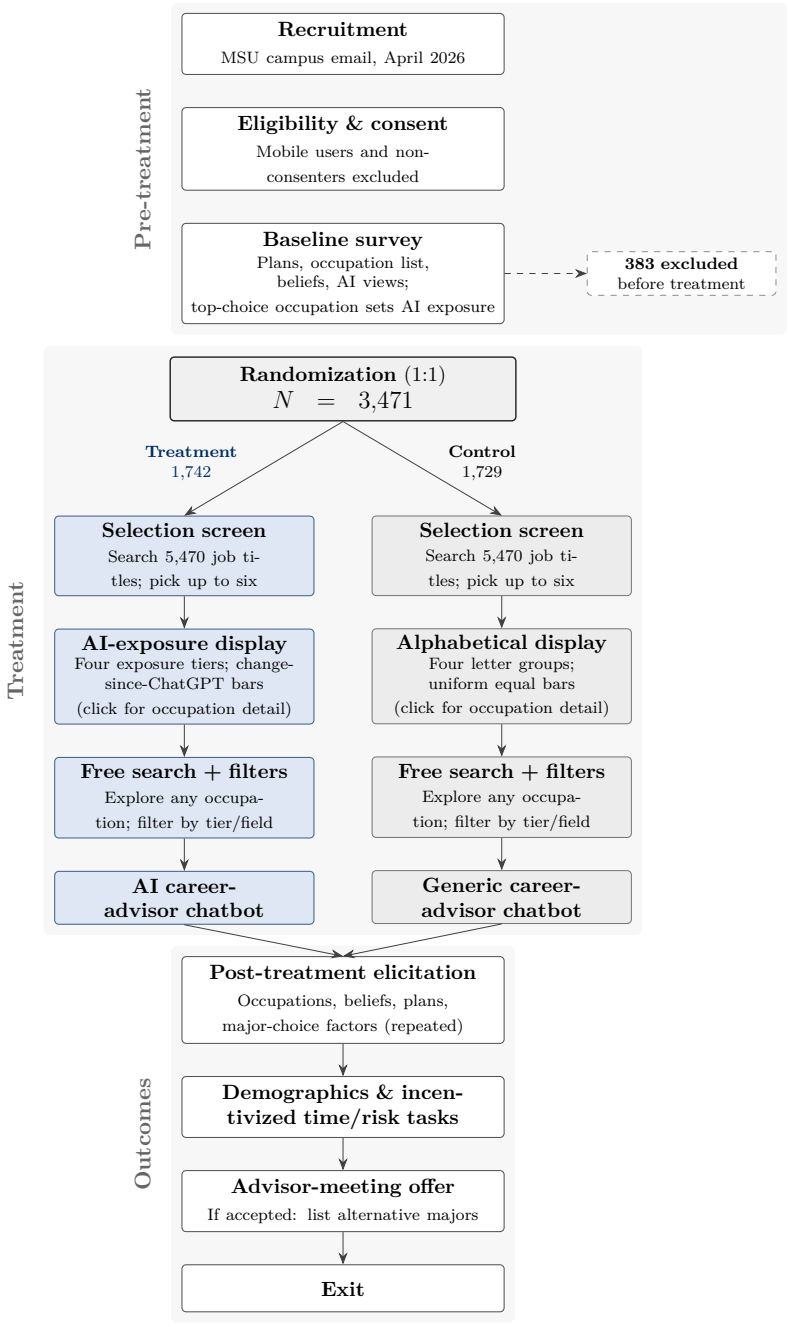
Weinstein, N. D. (1980). Unrealistic optimism about future life events. *Journal of Personality and Social Psychology*, 39(5):806–820.

Willis, R. J. and Rosen, S. (1979). Education and self-selection. *Journal of Political Economy*, 87(5, Part 2):S7–S36.

Wiswall, M. and Zafar, B. (2015). Determinants of college major choice: Identification using an information experiment. *The Review of Economic Studies*, 82(2):791–824.

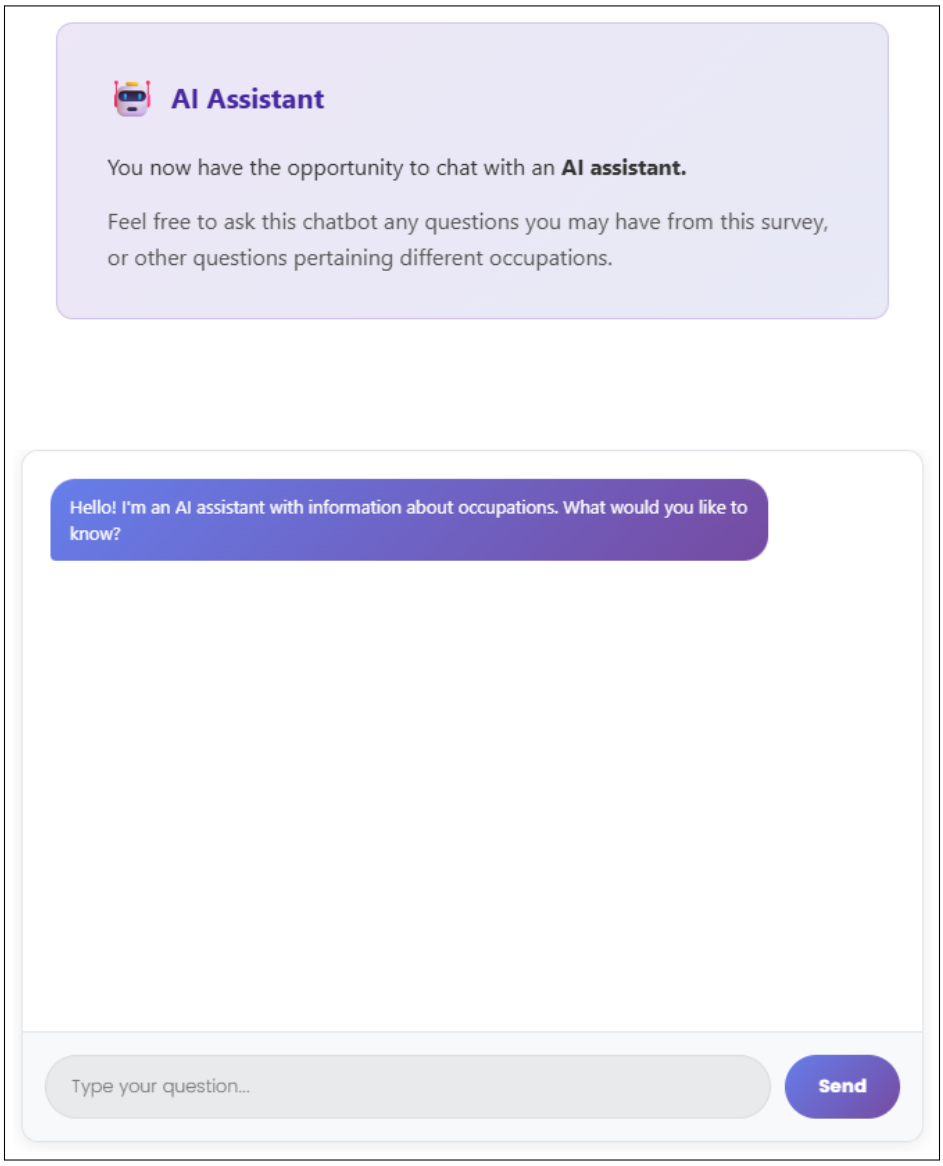
A Appendix Figures and Tables

Figure A1: Survey Flow and Randomization



Notes: This figure shows the flow of the experiment. Appendix Figures A2 and A3 show the chatbot and control interfaces.

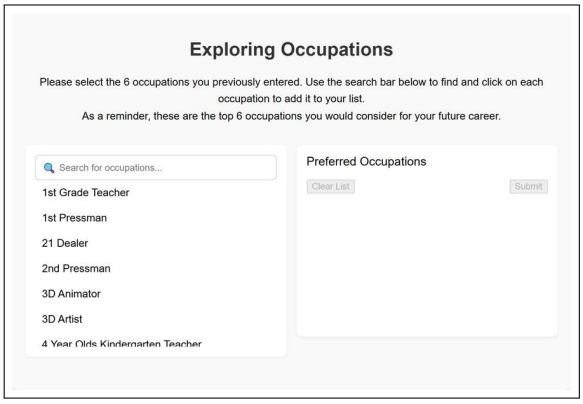
Figure A2: The Chatbot Interface



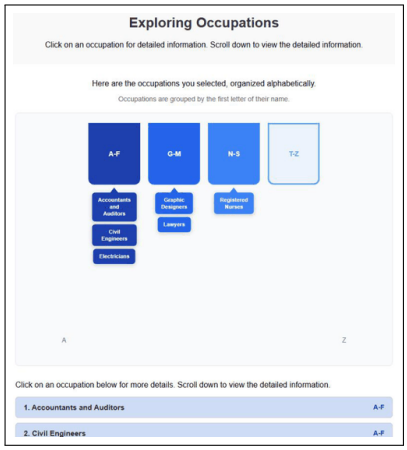
Notes: This figure shows the chatbot interface, which is identical in both arms. The treatment chatbot was configured as a career advisor who specializes in AI's labor-market impact, with backend access to occupation-level exposure data; the control chatbot was a generic career advisor who was prohibited from discussing AI.

Figure A3: The Control Condition: Alphabetical Placebo

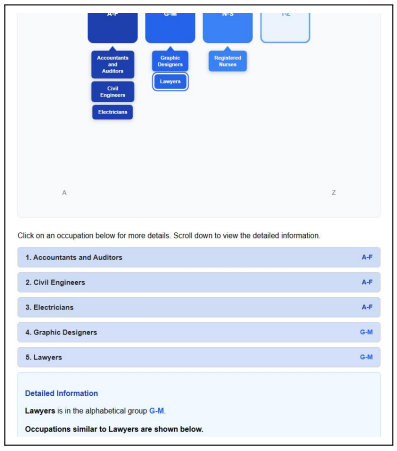
Panel A. Occupation selection



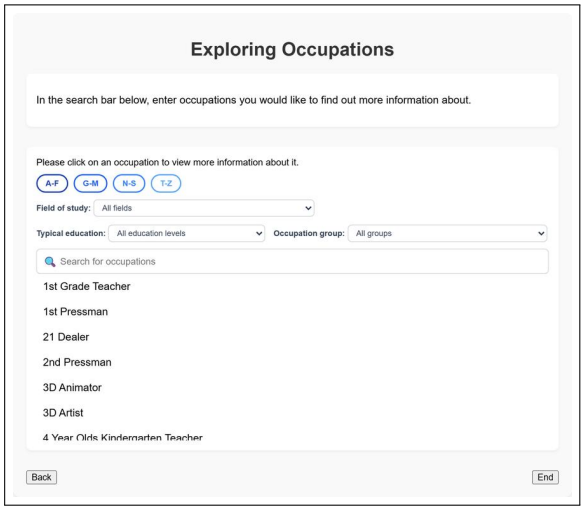
Panel B. Alphabetical display



Panel C. Occupation detail panel

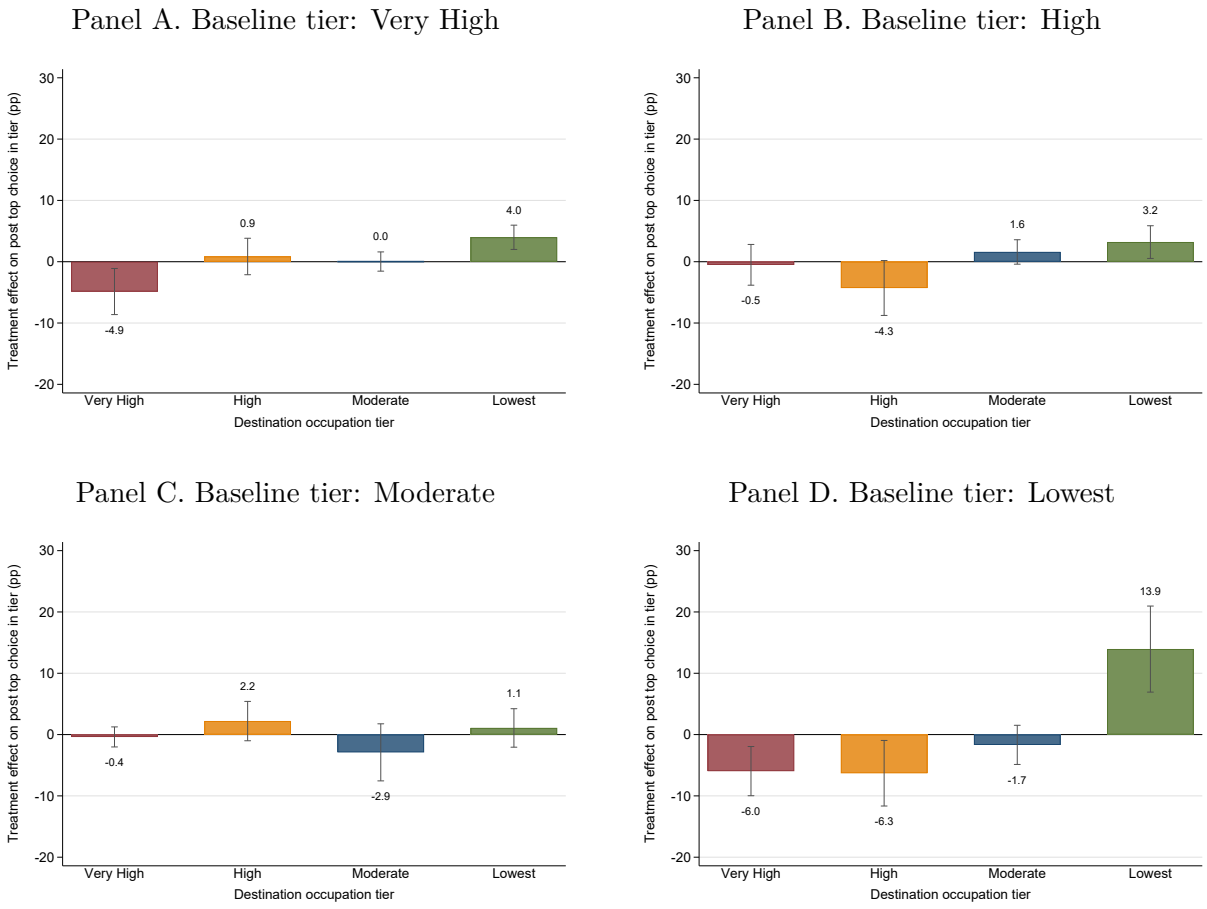


Panel D. Free search with filters



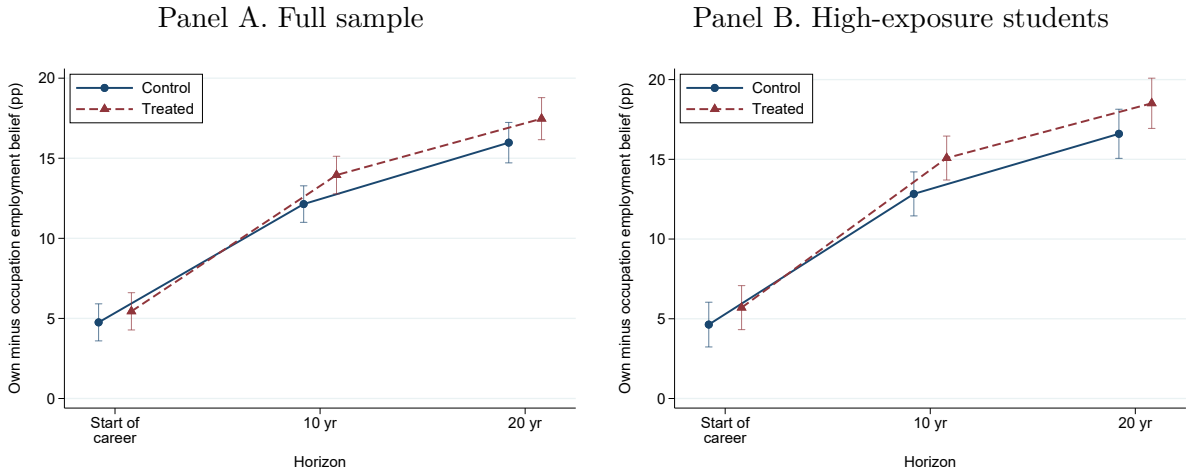
Notes: This figure shows the four screens of the control visualization, a format-matched placebo that strips all AI-exposure content. It displays the same occupations as the treatment but groups them alphabetically into four letter-range tiers (A–F, G–M, N–S, T–Z) rather than by AI exposure, with the tiers shaded in a single blue hue that carries no exposure information. Panel A is the occupation-selection screen, Panel B the alphabetical display, Panel C the detail panel for a selected occupation (with no exposure tier or employment-trend chart), and Panel D the free-search screen (with filters for alphabetical tier, occupation group, education, and field of study, but no AI-exposure filter).

Figure A4: Destination Exposure Tier of the Top-Choice Occupation, by Baseline Tier



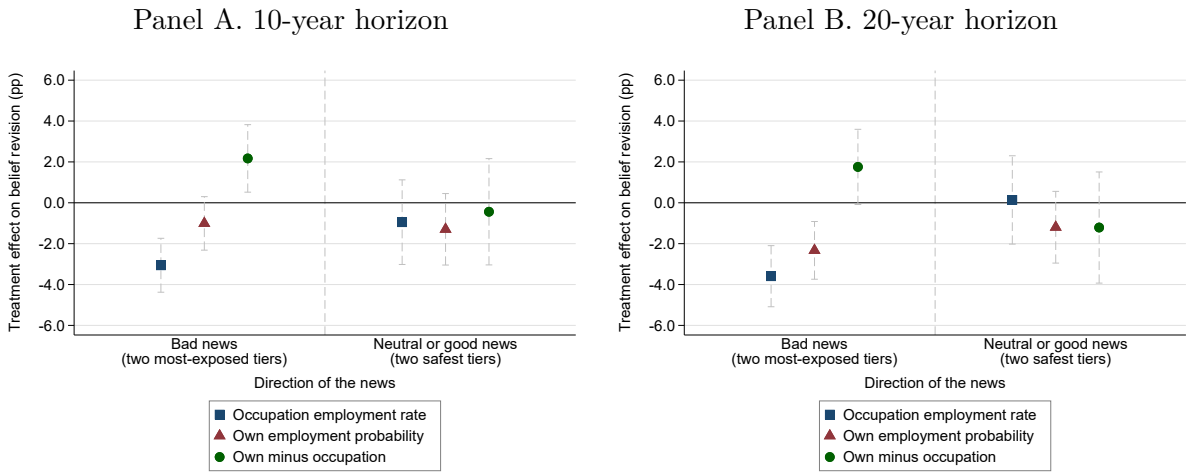
Notes: This figure plots treatment effects, in percentage points, on indicators for the exposure tier of the student's post-treatment top-choice occupation, for students whose baseline top choice is in the tier named in each panel caption. Estimates are unadjusted treated-minus-control differences with heteroskedasticity-robust standard errors; capped vertical bars are 95 percent confidence intervals and printed values are the point estimates. Control/treated counts are 545/590 (Very High), 463/482 (High), 266/269 (Moderate), and 166/150 (Lowest). The sample is students with a matched baseline and post-treatment top-choice exposure tier.

Figure A5: The Own-versus-Market Employment Gap by Horizon



Notes: This figure plots the mean difference between a student's own employment probability and the occupation's employment rate (own minus market, in percentage points) at each elicitation horizon, separately for control students (solid line, circles) and treated students (dashed line, triangles); the own probability on finishing education is paired with the occupation's start-of-career rate. Whiskers are 95-percent confidence intervals for the group means. Panel A uses the full sample and Panel B the high-exposure subgroup. Appendix Table A6 reports the corresponding regression estimates.

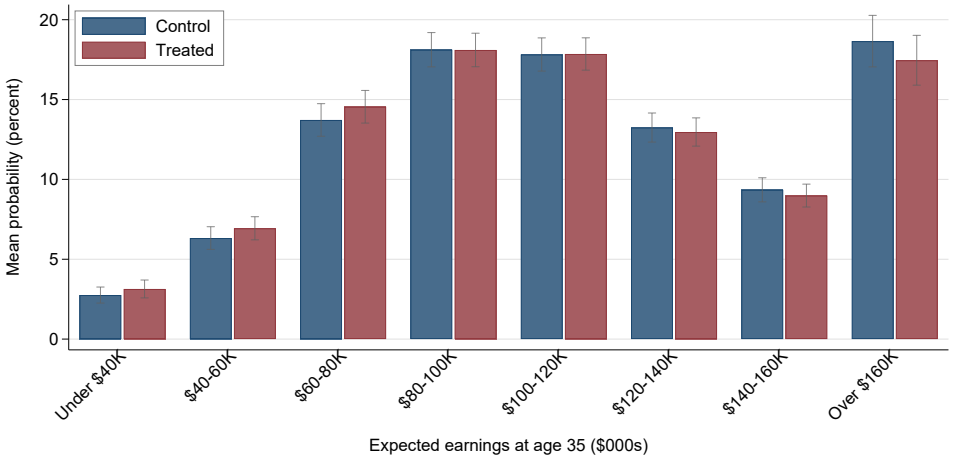
Figure A6: Belief Revisions by the Direction of the News



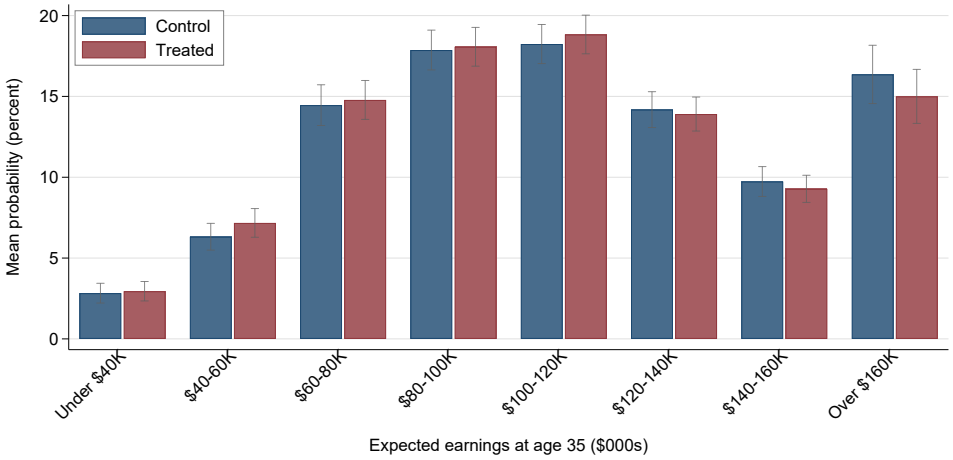
Notes: This figure plots the treatment effects underlying Panel A of Appendix Table A7: post-minus-pre belief revisions, in percentage points, of the occupation's expected employment rate, the student's own employment probability, and their difference (own minus occupation), estimated from equation (1) with double-lasso controls separately for students whose baseline top choice makes the ranking bad news (two most-exposed tiers) and neutral or good news (two safest tiers). Panel A uses the 10-year revisions and Panel B the 20-year revisions. Whiskers are 95-percent confidence intervals from heteroskedasticity-robust standard errors. The own-employment-prior split appears in Panel B of Appendix Table A7.

Figure A7: Implied Distributions of Expected Earnings at Age 35

Panel A. Full sample



Panel B. High-exposure subgroup



Notes: This figure plots the implied distribution of the student's own expected earnings at age 35 for control and treated students: bars are the mean post-treatment probability assigned to each of the eight earnings bins, with 95 percent confidence intervals. The estimation sample is students who respond in all eight bins, as in Appendix Table A8. Panel A uses the full sample; Panel B restricts to students whose baseline top-choice occupation is in the Very High or High exposure tiers.

Table A1: Robustness to Alternative Control Specifications

Index	Full sample		High-exposure subgroup	
	Mean control	$\hat{\beta}$	Mean control	$\hat{\beta}^{\text{High}}$
	(1)	(2)	(3)	(4)
Panel A. Double-lasso controls (benchmark)				
Family z-score: Stated choices	0.000	0.027**	−0.009	0.042***
		(0.011)		(0.014)
Family z-score: Beliefs about employment	0.000	0.043***	0.048	0.053***
		(0.013)		(0.016)
Family z-score: Beliefs about earnings	0.000	0.029***	0.066	0.025**
		(0.010)		(0.011)
Panel B. No controls				
Family z-score: Stated choices	0.000	0.016	−0.009	0.035**
		(0.014)		(0.017)
Family z-score: Beliefs about employment	0.000	0.065**	0.048	0.067**
		(0.026)		(0.031)
Family z-score: Beliefs about earnings	0.000	0.073***	0.066	0.073**
		(0.027)		(0.028)
Panel C. Full control set				
Family z-score: Stated choices	0.000	0.027**	−0.009	0.039***
		(0.011)		(0.014)
Family z-score: Beliefs about employment	0.000	0.043***	0.048	0.052***
		(0.013)		(0.016)
Family z-score: Beliefs about earnings	0.000	0.027***	0.066	0.023**
		(0.010)		(0.011)
N	1729	3471	1216	2468

Notes: This table presents the family z-score indices estimated under three control specifications, in the column layout of Table 8 (full-sample and high-exposure ITTs). Panel A uses post-double-selection lasso, the benchmark used throughout the paper, so its estimates reproduce the family-index rows of Tables 3 and 6; Panel B uses no controls; and Panel C uses the full pre-treatment control set (the lasso candidate pool, mean-imputed with missingness indicators, entered without selection). Each index occupies two rows: coefficient (top) and heteroskedasticity-robust standard error (bottom). *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table A2: Robustness to Sample and Design Restrictions

Index	Full sample		High-exposure subgroup	
	Mean control	$\hat{\beta}$	Mean control	$\hat{\beta}^{\text{High}}$
	(1)	(2)	(3)	(4)
Panel A. Recruitment-wave fixed effects				
Family z-score: Stated choices	0.000	0.027**	−0.009	0.041***
		(0.011)		(0.014)
Family z-score: Beliefs about employment	0.000	0.044***	0.048	0.053***
		(0.013)		(0.016)
Family z-score: Beliefs about earnings	0.000	0.029***	0.066	0.025**
		(0.010)		(0.011)
N	1729	3471	1216	2468
Panel B. Excluding durations over two hours				
Family z-score: Stated choices	−0.009	0.029**	−0.019	0.044***
		(0.013)		(0.015)
Family z-score: Beliefs about employment	−0.008	0.049***	0.046	0.055***
		(0.014)		(0.017)
Family z-score: Beliefs about earnings	0.001	0.027**	0.066	0.021*
		(0.011)		(0.013)
N	1397	2819	996	2024
Panel C. Excluding attention-check failures				
Family z-score: Stated choices	−0.031	0.025	−0.029	0.029
		(0.018)		(0.022)
Family z-score: Beliefs about employment	0.003	0.030	0.037	0.072***
		(0.021)		(0.025)
Family z-score: Beliefs about earnings	0.018	0.029**	0.099	0.035**
		(0.014)		(0.015)
N	613	1215	439	864
Panel D. Excluding the fastest completers				
Family z-score: Stated choices	0.011	0.030**	0.001	0.047***
		(0.012)		(0.014)
Family z-score: Beliefs about employment	0.010	0.044***	0.061	0.051***
		(0.014)		(0.017)
Family z-score: Beliefs about earnings	0.002	0.027***	0.069	0.026**
		(0.010)		(0.012)
N	1584	3137	1112	2230

Notes: This table presents the family z-score indices estimated under four sample or design restrictions, in the column layout of Table 8 (full-sample and high-exposure ITTs), on top of the double-lasso control selection used throughout the paper. Panel A adds recruitment-wave fixed effects (April 6 versus April 7-onward). Panel B excludes students whose uncapped survey duration exceeds two hours (19.2 percent of the control arm). Panel C keeps only students who passed the attention check. Panel D excludes the fastest 9.6 percent of students, defined as those below the fifth percentile of the control distribution of visualization time or of total survey duration. The panel samples therefore differ, and the N row gives each panel's sample; each index occupies two rows, coefficient (top) and heteroskedasticity-robust standard error (bottom). *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table A3: Heterogeneity of Treatment Effects by Student Characteristics

	Outcome: z-score index for...		
Moderator	Beliefs: employment (1)	Beliefs: earnings (2)	Stated choices (3)
Treated ×			
First or second year	−0.007	0.012	0.007
	(0.027)	(0.020)	(0.023)
Female	0.041	0.009	−0.021
	(0.030)	(0.022)	(0.027)
Declared major	0.098	−0.059	−0.084
	(0.076)	(0.050)	(0.061)
Believes AI destroys jobs	0.082***	−0.012	0.046*
	(0.027)	(0.022)	(0.023)
Own employment belief above median	−0.045*	−0.014	−0.033
	(0.026)	(0.020)	(0.023)
Occupation employment belief above median	−0.021	−0.003	−0.047**
	(0.026)	(0.020)	(0.024)
STEM	−0.006	−0.021	−0.007
	(0.026)	(0.020)	(0.023)
Ability above median	−0.020	0.023	−0.031
	(0.026)	(0.020)	(0.023)
First generation	0.003	−0.017	−0.008
	(0.034)	(0.025)	(0.032)
Risk averse	0.012	0.030	−0.021
	(0.028)	(0.020)	(0.025)
Patience above median	−0.057	−0.026	0.019
	(0.038)	(0.031)	(0.035)
N	3286	3130	3418

Notes: This table presents treatment-effect heterogeneity for the three family z-score indices in the full sample, along six pre-specified dimensions (first or second year, female, declared major, believes AI destroys jobs, risk averse, and above-median patience) and five exploratory moderators: baseline own and occupation employment beliefs (split at the median), STEM, above-median ability, and first generation. Each cell reports the coefficient on the interaction between treatment and the row characteristic from a separate regression that includes treatment, the characteristic, their interaction, and the index's double-lasso controls; heteroskedasticity-robust standard errors in parentheses. First or second year indicates freshmen and sophomores; believes AI destroys jobs indicates expecting AI to destroy at least some jobs at the 20-year horizon; the employment-belief moderators indicate above-median baseline beliefs about the student's own 20-year employment probability and about the occupation's 20-year employment rate; above-median splits are strict and defined only where the input is observed; risk averse indicates a below-median incentivized risk-seeking score and above-median patience an above-median incentivized patience score. The N row reports each index regression's sample; estimation samples differ by row because some moderators are missing for some students. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table A4: Robustness to the Exposure Definition

Index	Full sample		High-exposure subgroup	
	Mean control	$\hat{\beta}$	Mean control	$\hat{\beta}^{\text{High}}$
	(1)	(2)	(3)	(4)
Panel A. High-exposure from the top-choice occupation (baseline)				
Family z-score: Stated choices	0.000	0.027**	−0.009	0.042***
		(0.011)		(0.014)
Family z-score: Beliefs about employment	0.000	0.043***	0.048	0.053***
		(0.013)		(0.016)
Family z-score: Beliefs about earnings	0.000	0.029***	0.066	0.025**
		(0.010)		(0.011)
N	1729	3471	1216	2468
Panel B. High-exposure from the top-three-average rank				
Family z-score: Stated choices	0.000	0.027**	−0.017	0.041***
		(0.011)		(0.014)
Family z-score: Beliefs about employment	0.000	0.043***	0.045	0.058***
		(0.013)		(0.016)
Family z-score: Beliefs about earnings	0.000	0.029***	0.022	0.020
		(0.010)		(0.012)
N	1729	3471	1125	2260
Panel C. High-exposure from the top-five-average rank				
Family z-score: Stated choices	0.000	0.027**	−0.011	0.035**
		(0.011)		(0.014)
Family z-score: Beliefs about employment	0.000	0.043***	0.060	0.051***
		(0.013)		(0.016)
Family z-score: Beliefs about earnings	0.000	0.029***	0.014	0.024*
		(0.010)		(0.013)
N	1729	3471	1144	2268

Notes: This table presents the family z-score indices estimated under three definitions of the high-exposure subgroup, in the column layout of Table 8. The full-sample columns (1–2) are the standard ITT and do not depend on the exposure definition, so they are identical across panels. The high-exposure columns (3–4) reclassify students by exposure: Panel A uses the baseline top-choice occupation (rank ≤ 40, as in the main tables); Panel B the average AEI rank of the top three listed occupations; and Panel C the average of the top five (all-slots exposure match, slot-level match rates between 93 and 100 percent). The correlation between the top choice's rank and the average rank of the other listed occupations is 0.42, and 18.7 percent of students reclassify between the top-choice and top-three definitions. Each index occupies two rows: coefficient (top) and heteroskedasticity-robust standard error (bottom). Double-lasso controls. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table A5: Salary Results Under Alternative Winsorization Schemes

Outcome	Full sample		High-exposure subgroup	
	Mean control	$\hat{\beta}$	Mean control	$\hat{\beta}^{\text{High}}$
	(1)	(2)	(3)	(4)
Panel A. Winsorize at the 95th percentile (baseline)				
Salary, current (thousands)	103.048	−1.446*	94.420	−1.374
		(0.765)		(0.836)
Salary, start FT (thousands)	104.843	−0.642	96.217	−1.041
		(0.787)		(0.814)
Salary, 10yr (thousands)	128.605	−1.687*	120.336	−3.378***
		(0.867)		(0.989)
Salary, 20yr (thousands)	152.471	−2.125*	143.385	−2.157
		(1.276)		(1.561)
Beliefs about earnings	0.000	0.029***	0.066	0.025**
		(0.010)		(0.011)
N	1729	3471	1216	2468
Panel B. Winsorize at the 99th percentile				
Salary, current (thousands)	107.493	−1.819*	97.007	−1.877*
		(0.996)		(1.130)
Salary, start FT (thousands)	109.869	−1.012	98.875	−1.500
		(1.028)		(1.069)
Salary, 10yr (thousands)	136.996	−1.571	126.258	−3.213*
		(1.432)		(1.677)
Salary, 20yr (thousands)	162.567	−1.072	151.491	−0.650
		(2.413)		(3.022)
Beliefs about earnings	0.000	0.031***	0.068	0.028**
		(0.010)		(0.012)
N	1729	3471	1216	2468
Panel C. Drop entries above $500,000, then winsorize at the 99th percentile				
Salary, current (thousands)	104.766	−1.325	94.200	−0.919
		(0.880)		(0.909)
Salary, start FT (thousands)	105.469	−0.898	95.938	−1.876**
		(0.807)		(0.801)
Salary, 10yr (thousands)	125.673	−1.799**	117.557	−3.238***
		(0.842)		(0.885)
Salary, 20yr (thousands)	143.259	−1.859	136.804	−2.404
		(1.272)		(1.461)
Beliefs about earnings	0.000	0.033***	0.064	0.025**
		(0.010)		(0.011)
N	1729	3471	1216	2468

Notes: This table presents the salary outcomes (in thousands of dollars), re-parsed from the raw responses and re-estimated under three cleaning schemes with the sample and specification held fixed, in the column layout of Table 8: Panel A winsorizes at the 95th percentile (the pipeline's scheme), Panel B at the 99th percentile, and Panel C drops entries above $500,000 before winsorizing at the 99th percentile. The full-sample columns (1–2) are the standard ITT; the high-exposure columns (3–4) restrict to students whose baseline top-choice occupation is in the Very High or High exposure tiers. The last outcome row of each panel—above the observation count—is the beliefs-about-earnings family z-score index, rebuilt under each scheme with the salary components in logs, as in the main analysis (the earnings-at-35 components do not vary with the scheme); the rows above it are the salary levels. Each outcome occupies two rows: coefficient (top) and heteroskedasticity-robust standard error (bottom). Double-lasso controls. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table A6: Effects on the Own-versus-Market Employment and Earnings Gaps

Outcome	Full sample		High-exposure subgroup	
	Mean control	$\hat{\beta}$	Mean control	$\hat{\beta}^{\text{High}}$
	(1)	(2)	(3)	(4)
Panel A. Employment gap (percentage points)				
Gap change, finish ed. (pp)	−1.864	0.355	−2.327	0.694
		(0.719)		(0.862)
Gap change, 10yr (pp)	−3.585	1.439**	−4.271	2.172**
		(0.712)		(0.843)
Gap change, 20yr (pp)	−3.349	0.888	−4.351	1.755*
		(0.779)		(0.938)
Panel B. Log earnings gap				
Log earnings gap change, 10yr	0.030	0.014	0.031	0.022**
		(0.009)		(0.010)
Log earnings gap change, age 35	0.033	0.012	0.039	0.014
		(0.008)		(0.010)
Log earnings gap change, 20yr	0.042	0.013	0.053	0.006
		(0.010)		(0.012)
N	1729	3471	1216	2468

Notes: This table presents treatment effects on the own-versus-market employment gap—the student's own employment probability minus the expected employment rate in their top-choice occupation, at matched horizons (own "finish education" paired with the occupation's "start full-time"). Outcomes are post-minus-pre changes in the gap, in percentage points, so positive coefficients mean the treatment widens the gap. The earnings rows repeat the exercise for pay in logs: the log of the implied mean of the student's earnings-at-35 distribution (Appendix Table A8) minus the log of the expected typical-worker salary in the top-choice occupation, at the 10- and 20-year horizons and at the student's age-35 horizon (the occupation salary interpolated between the 10- and 20-year elicitations); coefficients multiplied by 100 read as approximate percent changes in the own-to-market earnings ratio. Sample sizes are smaller for the earnings rows because the implied mean requires a complete probability distribution and the salary beliefs must be answered at both waves. The N row reports the design sample; estimation samples differ by row. Each outcome occupies two rows: coefficient (top) and heteroskedasticity-robust standard error (bottom); columns as in Table 8. All specifications use double-lasso control selection (Belloni et al., 2014). *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table A7: Belief Passthrough by the Direction of the News

Panel A. News direction from the baseline top-choice tier

Outcome (revision, pp)	Bad news (most-exposed tiers) (1)	Reassurance (safest tiers) (2)	Interaction (3)
Occupation employment revision, 10yr	−3.057***	−0.949	2.146*
	(0.673)	(1.056)	(1.247)
Own employment revision, 10yr	−1.008	−1.297	−0.328
	(0.668)	(0.893)	(1.117)
Gap: Own minus occupation revision, 10yr	2.172**	−0.439	−2.649*
	(0.843)	(1.327)	(1.571)
Occupation employment revision, 20yr	−3.591***	0.139	3.734***
	(0.763)	(1.103)	(1.344)
Own employment revision, 20yr	−2.328***	−1.196	1.118
	(0.721)	(0.895)	(1.149)
Gap: Own minus occupation revision, 20yr	1.755*	−1.211	−3.020*
	(0.938)	(1.388)	(1.666)
N	2189	890	3079

Panel B. Baseline own-employment prior

Outcome (revision, pp)	Below-median own prior (1)	Above-median own prior (2)	Interaction (3)
Occupation employment revision, 10yr	−3.312***	−1.367	1.904*
	(0.777)	(0.835)	(1.140)
Own employment revision, 10yr	−1.680**	−0.357	1.348
	(0.784)	(0.722)	(1.071)
Gap: Own minus occupation revision, 10yr	1.623	1.079	−0.402
	(0.997)	(1.000)	(1.417)
Occupation employment revision, 20yr	−3.098***	−1.805*	1.228
	(0.852)	(0.955)	(1.276)
Own employment revision, 20yr	−2.720***	−1.291*	1.580
	(0.884)	(0.679)	(1.121)
Gap: Own minus occupation revision, 20yr	0.884	0.811	0.069
	(1.083)	(1.106)	(1.552)
N	1695	1402	3097

Notes: This table tests whether the passthrough from occupation-level news to own-outcome beliefs depends on the direction of the news. Outcomes are post-minus-pre belief revisions in percentage points: the occupation's expected employment rate, the student's own employment probability, and their difference (own minus occupation), at the 10- and 20-year horizons. Panel A splits students by the baseline top-choice exposure tier: the ranking is bad news when the top choice sits in the two most-exposed tiers and neutral or good news when it sits in the two safest. Students whose top choice has no exposure rank are in neither group. Panel B splits by the baseline own 20-year employment probability (strict median split, defined only where the input is observed). Columns 1–2 report the treatment effect within each group from equation (1) with double-lasso controls; column 3 reports the coefficient on the treatment-by-group interaction from the pooled regression. Each outcome occupies two rows: coefficient and heteroskedasticity-robust standard error. The *N* row reports estimation samples for the 20-year own-minus-occupation regressions. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table A8: Treatment Effects on the Subjective Earnings Distribution at Age 35

Outcome	Full sample		High-exposure subgroup	
	Mean control	$\hat{\beta}$	Mean control	$\hat{\beta}^{\text{High}}$
	(1)	(2)	(3)	(4)
Panel A. Implied moments				
Mean (thousands)	114.059	−0.078	112.489	−0.145
		(0.430)		(0.529)
Variance	459.156	−8.260	472.154	−12.401
		(9.300)		(11.472)
Panel B. Probability bins (Δ post − pre, pp)				
Pr(earnings under 40K)	0.081	−0.061	0.124	−0.079
		(0.214)		(0.258)
Pr(40K–60K)	−0.424	0.131	−0.534	0.254
		(0.278)		(0.317)
Pr(60K–80K)	−0.750	0.403	−0.735	0.549
		(0.427)		(0.502)
Pr(80K–100K)	−0.482	0.088	−0.318	−0.077
		(0.492)		(0.588)
Pr(100K–120K)	0.698	−0.594	0.553	−0.480
		(0.520)		(0.631)
Pr(120K–140K)	−0.068	−0.373	−0.093	−0.724
		(0.494)		(0.611)
Pr(140K–160K)	−0.069	0.183	−0.099	0.041
		(0.384)		(0.455)
Pr(over 160K)	1.015	0.273	1.103	0.580
		(0.436)		(0.558)
N	1729	3471	1216	2468

Notes: This table presents treatment effects on the subjective distribution of the student's own earnings at age 35. Panel A reports the post-treatment mean and variance implied by bin midpoints (in thousands of dollars); Panel B reports post-minus-pre changes in the probability (percentage points) assigned to each bin, with bin edges following the fielded instrument. The N row reports the design sample; every row uses the students who answered the post-treatment distribution, a smaller sample. Each outcome occupies two rows: coefficient (top) and heteroskedasticity-robust standard error (bottom); columns as in Table 8. Double-lasso controls. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table A9: Viewed-Tier Composition by Baseline Exposure Tier

	Control mean	Treated mean	Difference	(SE)
Panel A. Baseline tier: Very High exposure				
Viewed: Very High exposure	0.546	0.564	0.018	(0.032)
Viewed: High exposure	0.259	0.190	−0.069***	(0.027)
Viewed: Moderate exposure	0.068	0.064	−0.004	(0.015)
Viewed: Lowest exposure	0.126	0.182	0.056**	(0.022)
Panel B. Baseline tier: High exposure				
Viewed: Very High exposure	0.269	0.251	−0.018	(0.031)
Viewed: High exposure	0.437	0.383	−0.054	(0.035)
Viewed: Moderate exposure	0.091	0.128	0.037*	(0.023)
Viewed: Lowest exposure	0.203	0.237	0.034	(0.028)
Panel C. Baseline tier: Moderate exposure				
Viewed: Very High exposure	0.166	0.208	0.042	(0.034)
Viewed: High exposure	0.218	0.213	−0.004	(0.036)
Viewed: Moderate exposure	0.333	0.403	0.070	(0.043)
Viewed: Lowest exposure	0.284	0.176	−0.108***	(0.036)
Panel D. Baseline tier: Lowest exposure				
Viewed: Very High exposure	0.181	0.157	−0.024	(0.043)
Viewed: High exposure	0.242	0.098	−0.145***	(0.041)
Viewed: Moderate exposure	0.142	0.079	−0.063*	(0.035)
Viewed: Lowest exposure	0.434	0.666	0.232***	(0.056)

Notes: This table presents the share of a student's freely viewed occupations (on the free-search screen, among crosswalk-matched views) falling in each exposure tier, with panels splitting students by the AEI tier of their baseline top-choice occupation (tier 1 = Very High = rank ≤ 20; tier 4 = Lowest = rank > 60). Differences come from arm regressions conditional on opening at least one occupation on the free-search screen, with heteroskedasticity-robust standard errors within each baseline-tier panel. The sample is the 2,015 students who reached the visualization screen, have a ranked baseline top choice, and have defined shares. Shares are defined only for students who opened at least one occupation on the free-search screen. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table A10: Exploration Relative to the Baseline Exposure Tier

	Full sample		High-exposure subgroup	
	Mean control	$\hat{\beta}$	Mean control	$\hat{\beta}^{\text{High}}$
	(1)	(2)	(3)	(4)
Viewed higher exposure tier	0.230	−0.029*	0.126	−0.017
		(0.016)		(0.016)
Viewed same tier as baseline	0.457	0.031	0.495	−0.010
		(0.020)		(0.024)
Viewed lower exposure tier	0.313	0.002	0.379	0.027
		(0.018)		(0.023)
N	1058	2015	725	1383

Notes: This table presents treatment effects on the composition of a student's freely viewed occupations measured relative to the baseline top choice: the share of viewed occupations (on the free-search screen, among crosswalk-matched views) falling in a higher-exposure tier than, the same tier as, or a lower-exposure tier than the student's baseline top-choice occupation. The sample is the students who reached the visualization screen, have a ranked baseline top choice, and have defined shares (the N row gives the counts). Shares are defined only for students who opened at least one occupation on the free-search screen. Each outcome occupies two rows: coefficient (top) and heteroskedasticity-robust standard error (bottom). Columns 1–2 give the full-sample control mean and treatment effect $\hat{\beta}$ from equation (1) with double-lasso controls (Belloni et al., 2014); columns 3–4 repeat for the high-exposure subgroup (AEI ranking ≤ 40). *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table A11: End-of-Survey Comments as Outcomes

Outcome	Full sample		High-exposure subgroup	
	Mean control	$\hat{\beta}$	Mean control	$\hat{\beta}^{\text{High}}$
	(1)	(2)	(3)	(4)
Left any comment	0.382	−0.005	0.368	0.010
		(0.015)		(0.018)
Comment: found it informative	0.078	0.015	0.085	0.012
		(0.009)		(0.011)
Comment: states intended action	0.002	0.001	0.003	−0.001
		(0.002)		(0.002)
Comment: AI-related worry	0.001	0.010***	0.002	0.011***
		(0.003)		(0.003)
Comment: skepticism / criticism	0.009	0.010**	0.007	0.015***
		(0.004)		(0.005)
Comment: clarifies an answer	0.032	−0.002	0.025	0.005
		(0.006)		(0.007)
Comment: confusion / technical issue	0.078	−0.007	0.076	−0.006
		(0.009)		(0.010)
Comment: none / filler (no, n/a, thanks)	0.183	−0.022*	0.169	−0.018
		(0.013)		(0.015)
Comment: other	0.019	0.002	0.021	0.000
		(0.005)		(0.006)
N	1729	3471	1216	2468

Notes: This table presents treatment effects on end-of-survey comment categories. The final survey question invites general comments, and 1,351 students in the deidentified file left one; comments are classified by the same language-model pipeline as the chat messages into eight multi-label categories (Appendix B.4). Indicators are unconditional—students without a comment are zeros. Columns as in Table 8. Double-lasso controls, heteroskedasticity-robust standard errors. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

B Empirical Appendix

This appendix defines the outcome variables and documents the construction of the AI-exposure measure, the mapping from free-text occupations to exposure rankings, and the language-model classification of chatbot messages and end-of-survey comments.

B.1 Outcome Definitions

This section defines the outcomes, organized by family. We elicit every post-treatment outcome with the same wording and response format as its baseline counterpart, and apply the transformations described in Section 3.3: we winsorize salary, time, and tenure variables at the 95th percentile, recode probability questions to range midpoints on a 0–100 scale, and summarize each family with a standardized z-score index.

Stated Choices

Occupation switch. An indicator equal to one if coders judged the student's post-treatment free-text top-choice occupation to be a different occupation from the one named at baseline (Appendix B.5); a rewording or a broader or narrower specialization of the same occupation is zero, as is a post-treatment answer of "same as before," and the indicator is missing for the pairs the coders could not judge. The pre-analysis plan registered the *text* version, equal to one whenever the lower-cased, whitespace-trimmed text differs.

List exposure composition. Among the occupations on the student's list (Appendix B), the share whose AEI ranking falls in the two most-exposed tiers (ranking ≤ 40) and the mean AEI ranking of the list. Both are computed for the baseline and post-treatment lists and are missing when no listed occupation is ranked.

Consideration-set size and stability. Students rank up to six occupations they are considering, in order of preference. Write the baseline list as the ordered set of normalized occupation strings $P = (p_1, \ldots, p_{n^{\text{pre}}})$ and the post-treatment list as $Q = (q_1, \ldots, q_{n^{\text{post}}})$, where an occupation's position in a list is its rank (p_1 and q_1 are the top choices) and occupations are matched as lower-cased, whitespace-trimmed text. An occupation entered in more than one slot counts once, at its first position. *Size* is n^{post}, the number of occupations in the post list (missing, not zero, when the list is empty). We summarize *stability* against

the baseline list with three kinds of measures, retention, rank distance, and Kendall's τ, all defined only when both lists are non-empty.

Retention. Let $C = \{o : o \in P \text{ and } o \in Q\}$ be the occupations appearing in both lists. The count retained is $|C|$ and the share retained is $|C|/n^{\text{pre}}$.

Rank distance. Writing $r^{\text{pre}}(o)$ and $r^{\text{post}}(o)$ for the position of occupation o in P and Q, the baseline rank distance D (reported as "Rank distance (pre vs post)") sums absolute position changes over occupations listed in *both* periods,

$$D \;=\; \sum_{o \in C} \Big| r^{\text{pre}}(o) - r^{\text{post}}(o) \Big|,$$

so occupations that are dropped or newly added contribute nothing. The top-coded variant D^{tc} (reported as "Rank distance, top coded") instead assigns every dropped baseline occupation the post-rank $n^{\text{post}} + 1$—one position below the bottom of the post list—and sums over the *entire* baseline list,

$$D^{\text{tc}} \;=\; \sum_{o \in P} \Big| r^{\text{pre}}(o) - \tilde{r}^{\text{post}}(o) \Big|, \qquad \tilde{r}^{\text{post}}(o) = \begin{cases} r^{\text{post}}(o) & o \in Q, \\ n^{\text{post}} + 1 & o \notin Q, \end{cases}$$

so that dropping an occupation registers as a large change; the per-occupation version, $D^{\text{tc}}/n^{\text{pre}}$, divides this by the number of baseline occupations. Occupation entries are used as written after trimming and lower-casing; rare free-text filler ("none," "n/a"—0.2 percent of listed slots) is retained, and excluding it leaves every reported estimate unchanged at reported precision.

Kendall's τ. A rank correlation computed over the *union* $U = P \cup Q$ of the two lists ($n^U = |U|$ distinct occupations), with each occupation absent from one list ranked just off that list's bottom: $\rho^{\text{pre}}(o) = r^{\text{pre}}(o)$ if $o \in P$ and $n^{\text{pre}} + 1$ otherwise, and $\rho^{\text{post}}(o)$ defined analogously from Q. Over the $\binom{n^U}{2}$ unordered pairs, a pair $\{o, o'\}$ is concordant if $\big(\rho^{\text{pre}}(o) - \rho^{\text{pre}}(o')\big)\big(\rho^{\text{post}}(o) - \rho^{\text{post}}(o')\big) > 0$ and discordant if that product is negative; with n_c concordant and n_d discordant pairs,

$$\tau \;=\; \frac{n_c - n_d}{n^U\,(n^U - 1)/2}.$$

This is Kendall's τ_a: pairs tied in either ordering—including any pair of occupations that both sit off one list—count toward neither n_c nor n_d but remain in the denominator. It is defined when $n^U \geq 2$. Higher retention and τ and lower rank distance all indicate a more stable list. Of these, only the rank distance enters the stated-choices index, signed so that more reordering counts as change; retention, τ, and the top-coded variant appear only as rows of Tables 3 and 4.

Educational plans. The student's reported probability of changing major, applying to graduate school, and leaving college. Each is elicited on a seven-option scale that we map to the midpoint of its range on a 0–100 scale.

Interest in exploring alternatives. An indicator equal to one if offered the chance to explore other majors, the student either accepts help exploring options or reports interest but says they will arrange it themselves, and zero if the student declines.

Major switch. An indicator for whether the student's preferred program changes after treatment. Students name the program along three nested fields—college, department, and major. The pre-analysis plan registered an indicator equal to one if any post-treatment field differs from its baseline value, comparing each field only when it is answered in both periods. The paper's indicator applies the coder verdicts of Appendix B.5 to these pairs: it equals one only for pairs a coder judged a different program, zero for menu-path or naming artifacts and for non-switchers, and is missing when unsure.

Reasons for major choice After treatment, students select the reasons behind their major choice from a fixed multi-select list. Each reason—expected salary, job prospects, working with innovative technology, developing diverse skills, and avoiding technological displacement—is an indicator for whether the student's selection includes it, defined for students who answer the question.

Beliefs

Expected employment rate. The percentage of people in the student's top-choice occupation whom the student expects to be employed, asked at four horizons: currently, when the student first works full time, in ten years, and in twenty years. Students answer on a

seven-category scale: 0, 1–20, 21–40, 41–60, 61–80, 81–99, and 100 percent. We use the midpoint of each category: $\{0, 10, 30, 50, 70, 90, 100\}$.

Expected salary. The salary the student expects a typical worker in the top-choice occupation to earn, at the same four horizons. Students enter a dollar amount as free text; we parse it to a number, recode any range they type ("X–Y") to its midpoint $(X+Y)/2$, apply the audited parsing rules of Appendix B.5 (per-hour and per-day rates converted to annual amounts, "million" expanded, bare values between 35 and 999 read in thousands, values below 35 and non-numeric entries set to missing, malformed comma grouping and double-encoded dashes repaired), and winsorize each horizon at the 95th percentile. The salary family enters the analysis in logs: the outcome is the log of the winsorized post-treatment salary in thousands of dollars, defined for strictly positive responses.

Own employment probability. The student's own probability of being employed in the top-choice occupation, asked at three horizons: on finishing their education, in ten years, and in twenty years. Students answer on the same seven-category probability scale as the employment-rate question, recoded to the midpoints $\{0, 10, 30, 50, 70, 90, 100\}$. Set against the expected employment rate, it isolates how far students separate their own prospects from the occupation's.

Expected occupational tenure. The number of years the student expects to work in the top-choice occupation over their career, entered as a free-text number of years and winsorized at the 95th percentile.

Earnings at 35. Students describe what they expect to earn at age 35 by spreading probability across eight earnings bins: below \$40K, \$40–60K, \$60–80K, \$80–100K, \$100–120K, \$120–140K, \$140–160K, and \$160K or more. Writing p_k for the probability (in percent) the student places on bin k and $m = (20, 50, 70, 90, 110, 130, 150, 180)$ for the bin midpoints in thousands of dollars—the open bottom and top bins take the conventional midpoints \$20K and \$180K—we summarize each student's distribution by its implied mean and variance,

$$\mu = \sum_{k=1}^{8} \frac{p_k}{100}\, m_k, \qquad \sigma^2 = \sum_{k=1}^{8} \frac{p_k}{100}\, (m_k - \mu)^2.$$

Both moments are defined only for students who place a probability on all eight bins and whose probabilities sum to within 5 percentage points of 100.

Own-versus-market employment gap. The student's own employment probability minus the expected employment rate of the top-choice occupation, in percentage points, at three matched horizons: finishing education (paired with the occupation's rate when the student first works full time), ten years, and twenty years. The outcome in Panel A of Appendix Table A6 is the post-minus-pre change in this gap, so positive coefficients mean the treatment widens the gap between beliefs about oneself and beliefs about the market.

Own-versus-market earnings gap. The log of the implied mean of the earnings-at-35 distribution minus the log of the expected typical-worker salary in the top-choice occupation, at the 10-year and 20-year horizons and at the student's age-35 horizon, where the occupation salary is interpolated linearly between the 10- and 20-year elicitations to the student's years-to-35. The outcome in Panel B of Appendix Table A6 is the post-minus-pre change in this log gap—coefficients multiplied by 100 read as approximate percent changes in the ratio of own to occupation earnings—so positive coefficients mean the treatment widens the gap between beliefs about oneself and beliefs about the market. Observations with nonpositive salary entries are excluded from the logs.

Demand for information

Search behavior. Four outcomes come from the visualization's logs. *Occupations explored* is the number of occupations whose detail panel the student opened, and *detail-panel time* is the number of seconds those panels stayed open, winsorized at the 95th percentile. *Search terms entered* counts the distinct terms the student typed in the search box, after trimming and lower-casing. The *tier shares* are the shares of the explored occupations in each of the four exposure tiers. The first three are missing when the visualization recorded no log for the student, and the tier shares are missing when the student explored no occupation.

Chatbot use. An indicator equal to one if the student sends at least one message to the chatbot.

Chatbot question topic. After the visualization, students may type questions to a chatbot framed as a career advisor. We classify each student's first question into one of eight topics

with a language-model classifier (Appendix B.4) and use the topic indicators as outcomes.

Career counselor meeting request. An indicator equal to one if the student asks for help scheduling a meeting with a career counselor at the end of the survey, and zero if the student declines or says they will arrange it themselves.

Family indices. We summarize each outcome family with a standardized index following Kling et al. (2007). For respondent i and a family whose component outcomes are indexed by $\mathcal{K}$, we standardize each component on the control group, sign it so that a positive value denotes the direction predicted in the pre-analysis plan, impute missing components at the treatment-arm mean, and average.

The paper reports three family indices; the same indices appear in the main tables, the prose, and every robustness and heterogeneity table. The *stated-choices* index takes one component per outcome in its family, each signed in the predicted direction: the stated probabilities of changing major and of applying to graduate school, interest in exploring alternatives, the major- and occupation-switch indicators, the number of occupations listed, the rank distance between the baseline and post-treatment lists, and citing technological displacement as a reason for major choice; we exclude the probability of leaving college because we assigned it no direction in the pre-analysis plan. The *beliefs-about-employment* index pools the expected employment rate, the student's own employment probability, and expected occupational tenure, signed so that downward revisions count as positive. The *beliefs-about-earnings* index pools expected salary (in logs) and the mean of the earnings-at-35 distribution, both signed so that downward revisions are positive, together with the variance of that distribution, signed so that a higher variance is positive. There are three departures from the pre-analysis plan: we report the registered belief family as its employment and earnings margins rather than as a single pooled index. We also registered a *demand-for-AI-information* index, which we did not use since the chatbot malfunction made some measures of demand unusable. And the two switch indicators in the stated-choices index are the coder-adjudicated ones defined above rather than the registered text and code comparisons.

B.2 The AI-Exposure Measure

The visualization's content and the paper's exposure rankings come from the Anthropic Economic Index (Handa et al., 2025), which maps millions of real-world conversations with

the Claude language model to the work tasks of the O*NET taxonomy and scores each occupation by the share of its tasks that appear in those conversations as automation, with AI performing the task rather than assisting a person with it. We aggregate the task-level data to six-digit SOC occupations, convert scores to percentile rankings from 1 (most exposed) to 100 (least exposed), and show 775 of these occupations in the visualization. The four display tiers (Very High, High, Moderate, Lowest) cut the ranking at 20, 40, and 60—the exposure-quintile boundaries at which Brynjolfsson et al. (2025a) document diverging employment trajectories for early-career workers since the diffusion of generative AI—with the two least-exposed quintiles pooled into the Lowest tier. The high-exposure indicator used throughout the paper equals one when a student's baseline top-choice occupation ranks in the top two tiers (ranking ≤ 40). Appendix Table A4 shows robustness to defining exposure from the top three listed occupations instead.

B.3 Matching Free-Text Occupations to Exposure Rankings

Students enter their considered occupations as free text. We map each student's top-choice occupation, both pre-treatment and post-treatment, to the exposure ranking through a four-stage matching procedure: an exact match against the crosswalk's SOC titles, their curated synonyms, and a small set of hand-verified free-text variants; a layer of 279 coder-audited corrections keyed on the exact text (Appendix B.5); a hand-built dictionary of several hundred common variants, each mapped to the closest crosswalk occupation; and a keyword-substring fuzzy fallback, scanned in order with the first match winning. The exact stage skips the generic title "Doctor," which the crosswalk lists only among the reported titles of podiatrists, so the dictionary maps it to physicians. Post-treatment entries such as "same as before" carry the pre-treatment value forward. The same procedure, including the audited layer, ranks the other occupations on each student's list for the list-composition outcomes and for the alternative exposure definitions of Appendix Table A4. One safeguard applies throughout: a ranking of 100 (the least-exposed bucket) is kept only when it comes from the crosswalk itself; dictionary or fuzzy hits that would impute 100 are set to missing rather than guessed, so informally typed bottom-tier titles stay unranked unless the crosswalk resolves them.

Among the 3,471 students in the analysis sample, 99.5 percent (3,453) have a successfully matched pre-treatment top choice: 57.6 percent via exact match, 8.2 percent via the audited corrections, 32.4 percent via the dictionary, and 1.2 percent via keyword matching; the remaining 0.5 percent could not be assigned a ranking. For the post-treatment occu-

pation, 84.6 percent are matched (53.4 percent exact, 5.8 percent audited, 24.7 percent dictionary, 0.6 percent keyword, and 0.1 percent carried forward from the pre-treatment match); 15.0 percent left the field blank and 0.4 percent could not be ranked. The main limitations are that keyword matching can misfire on generic titles ("analyst," "specialist"), that the rank-100 safeguard selectively censors least-exposed occupations typed informally, and that genuinely ambiguous titles ("technician") are set to missing by design.

B.4 Text Classification of Chat Messages and Comments

Two parts of the survey collect free text that we code into categories: the messages students type to the career-advisor chatbot, and the open-ended comment at the end of the survey. We classify both with the same language-model pipeline, described here once, and then give the category scheme and the aggregation to analysis variables for each task in turn.

Classification pipeline. We use Anthropic's Claude Haiku 4.5, queried through the Messages API, as the coder for both tasks. The model receives a system prompt that serves as the full codebook: a description of the study and of the kind of text being coded, the fixed list of categories, a one- or two-sentence definition of each category with explicit inclusion and exclusion rules, and a set of routing conventions for the very short messages that dominate both samples—how to code a bare occupation name, a lone salary figure, or a one-word non-answer, for example. The codebook asks the model to return, for each item, both a single primary category—the best label for the item as a whole—and a multi-label set containing every category the item touches; the primary label is chosen by a chat message's explicit request or a comment's main point, then by the category that occupies the most of the text, with ties broken by the codebook order. One category in each scheme is exclusive: the off-topic code for chat messages and the no-substantive-content code for comments may be used only when no other category applies.

We collect the model's output through a required tool call whose schema pins each item to one primary code and a list of codes drawn from the fixed category set, so the model cannot return free-form labels. Items are sent in batches of about forty, each numbered so the returned rows can be matched back to their inputs, and the codebook is marked for prompt caching so that it is processed once rather than re-sent with every batch. We validate every returned record against the codebook: an unrecognized code is reset to the scheme's fallback category, the primary code is forced into the multi-label set, the exclusive off-topic or no-content code is dropped whenever a substantive code is also present, and

duplicate codes are removed. Each item is classified once and its result is written to an on-disk cache keyed to the item, so the released classifications are fixed and a re-run reuses the stored labels rather than re-querying the model. If a call fails or returns malformed output, the affected items are assigned the fallback code and flagged, and flagged items are dropped before we count classified items; no item failed in either task.

Chatbot messages. The classifier is applied to the 1,486 user messages students sent to the AI career advisor. These come from 916 distinct students in the deidentified file, each of whom may send more than one message. The eight categories are career or occupation questions, covering what a job involves, its demand, and how to enter it; salary or compensation questions; questions about AI's impact, automation, or displacement; questions about majors, degrees, or training; personal background or self-disclosure; questions about the survey, payment, or instructions; tests, greetings, jokes, or other off-topic content; and other career-relevant questions that fit none of the first four. For the main analysis (Table 8) we take each student's first message and its primary label, so each message counts in exactly one topic, and use the four substantive question types—career, salary, AI impact, and major—as separate outcomes, pooling personal background, survey or payment, off-topic, and other career-relevant messages into a single "Other" category; students who never messaged are coded zero on every topic.

End-of-survey comments. The same pipeline codes the 1,351 non-empty responses to the survey's closing question, which invites students to clarify or expand on their answers or to flag anything they found unclear; blank responses are excluded before classification. The eight categories are: found the survey or information informative or interesting; states an intended action, such as contacting an advisor or researching careers; expresses AI-related worry or anxiety about one's own future; voices skepticism or criticism of the information or the survey's design; clarifies or expands an earlier answer; reports confusion or a technical problem; contains no substantive content; and other substantive comments. The paper uses the multi-label indicators, defined over all students so that those who left no comment are zeros; the treatment effects in Appendix Table A11 therefore combine the decision to comment with the content of the comment.

B.5 Measurement and Validation

Several outcomes and covariates rest on free-text answers or on model-generated labels. Research assistants audited each of these constructions, and this section reports the design of the audits, the agreement rates, and the definitions of the audited indicators. Two coders classified each sample independently, without seeing the model's labels or each other's answers, and their disagreements were adjudicated without reference to the model's output; a third coder handled the three tasks that involve a single judgment per item. Every verdict became one of three things: a rule applied to every entry with the same text, a validation statistic reported here, or an adjudicated indicator. The two adjudicated switching indicators are the paper's outcomes, because a rewording or a typo in a free-text answer should not count as a switch. The other audits changed parsing rules and matches, not definitions.

Text classification. For each classification task (Appendix B.4), the two coders labeled a random sample of 200 chatbot messages and 200 end-of-survey comments with the codebook the model received. Against the adjudicated consensus, the model's primary label agrees on 87.0 percent of messages (Cohen's $\kappa = 0.84$) and 88.5 percent of comments ($\kappa = 0.85$), comparable to the coders' own agreement with the consensus (86.4 and 92.3 percent for messages, 83.5 and 87.5 percent for comments). The validation covers primary labels only. The model's labels remain the outcomes: sampled errors are reported, not corrected.

Occupation switching. The pre-registered indicator counts any change in the normalized free-text top choice. To separate changes of occupation from rewordings, the two coders classified every pre/post pair among the 962 students whose text changed (890 distinct pairs) as a different occupation, a rewording of the same occupation, a broader or narrower specialization, or unsure. They give the same verdict for 62.7 percent of pairs ($\kappa = 0.46$) and agree on different-versus-same for 78.1 percent of the pairs with two definite verdicts ($\kappa = 0.56$). The 337 pairs on which the verdicts differ or both coders are unsure were resolved under a written rubric (name variants and career-ladder moves are rewordings, moves within a discipline are specializations, and crossings of function, discipline, or credential are different occupations). The adjudicated indicator equals one for a different occupation, zero for a rewording or specialization, and is missing for the 7 unsure students; students who answered the post-treatment question with "same as before," whom

the matcher carries forward, are zeros although the pre-registered indicator counts them as switches. Only 34.7 percent of text changes are changes of occupation. Among the 327 adjudicated switches with both occupations ranked, 92.7 percent change exposure rank, whereas 59.3 percent of the 619 rewordings and specializations keep their rank. The adjudicated indicator is the paper's measure. In the control group, 30.4 percent of pre-registered switches are adjudicated switches; the control means of the pre-registered and adjudicated indicators are 31.1 and 9.5 percent.

Salary expectations. The two coders audited the salary parser on three sheets drawn from the 26,777 salary entries in the raw file: a random sample of 150 entries, the 158 entries the parser could not read, and the 391 entries that parsed below $1,000. On the random sample, both coders confirm the parser's reading of 96.7 percent of entries, and their verdicts agree on 97.3 percent. Of the unparsed entries, 156 are uncodable to both coders. The audit found one bug, ranges typed with a double-encoded dash that parsed as concatenated digits, and a family of unit errors. Rather than correcting sampled entries, we encode the verdicts as rules applied to every entry with the same text: per-hour rates are multiplied by 2,080 and per-day rates by 260, "million" is expanded, bare values between 35 and 999 are read in thousands, bare values below 35 and non-numeric entries are set to missing, malformed comma grouping ("60,00") is read as a three-digit group (60,000), and double-encoded dashes are normalized before parsing. The rules change 922 of the 26,777 raw entries (3.4 percent) and 866 of the 25,923 entries in the analysis sample (3.3 percent); the largest classes are values set to missing (288 entries) and bare thousands (286). Winsorization at the 95th percentile follows the rules; Appendix Table A5 reports the salary results under alternative cleaning schemes.

Major switching. A coder reviewed all 837 distinct pre/post program-code pairs behind the 1,200 students whose codes differ and judged whether each pair is a different program or an artifact of the menu, a same-major path through a different college or department or a naming or degree variant: 98.9 percent of switching students changed program, 0.3 percent are artifacts, and 0.8 percent are unsure. The adjudicated indicator, the paper's measure, is one for a different program, zero for an artifact and for non-switchers, and missing when unsure; in the control group, 98.8 percent of pre-registered switches are adjudicated switches (control means 35.7 and 35.3 percent).

Years to full-time work. Students state how many years remain until they expect to work full time; the answer is a candidate control in the double-lasso step (Section 3.2). Numeric answers parse directly. A coder assigned a value to each of the 1,078 distinct free-text answers ("2–3 years," "after grad school") where one is implied, leaving 145 uncodable. This recovers 2,125 of the 2,283 free-text answers in the raw file (93.1 percent) and 1,936 of the 2,078 in the analysis sample (93.2 percent), which were previously missing and mean-imputed.

Exposure matching. A coder checked every keyword match and a random sample of the dictionary matches against the crosswalk title each was assigned, and every unmatched text against the crosswalk. Exact matches, 57.6 percent of the pre-treatment top choices, were not reviewed. The audit traced most errors to the hand-built dictionary, whose ranks had been typed by hand rather than looked up by title, and produced 279 corrections: 213 texts re-matched to a different crosswalk occupation and 66 previously unmatched texts placed (9 suggested re-matches were set aside as defensible either way). The corrections enter the matcher as an audited stage between the exact and dictionary stages (Appendix B.3) and now resolve 8.2 percent of pre-treatment and 5.8 percent of post-treatment top choices. Because they reassign exposure ranks, the corrections move a small number of students across the high-exposure boundary; every result in the paper uses the corrected matches.

C Conceptual Framework

This appendix sets out the framework that organizes the experiment's outcomes and collects the predictions the results sections examine.

Economists have long viewed education as an investment in human capital: students incur schooling costs today in exchange for skills that command returns in the labor market of tomorrow (Schultz, 1961; Becker, 1962), and they sort into the fields and occupations where those skills earn the most (Roy, 1951; Willis and Rosen, 1979; Kirkebøen et al., 2016). In Goldin and Katz (2008)'s account of the race between education and technology, technological change raises the demand for cognitive work as a class, so its incidence falls between education groups and widens the college premium. GenAI is predicted to break this pattern: exposure measures agree that generative models substitute for tasks in some cognitive occupations while augmenting others (Felten et al., 2023; Kogan et al., 2023; Pizzinelli et al., 2023; Eloundou et al., 2024), so exposure varies at least as much across fields of study and target occupations as between education groups. Whatever happens to the average return to college, the occupation a student prepares for becomes a margin of adjustment in its own right, and it is the margin this experiment studies.

Acting on that margin requires knowing which occupations the technology touches. Students hold sizable misperceptions about labor-market returns even in settings where information is abundant, and correcting them changes schooling decisions (Jensen, 2010; Hastings et al., 2015; Wiswall and Zafar, 2015). Such frictions are plausibly larger for AI, whose occupational incidence is uncertain even to experts (Karger et al., 2026). The experiment supplies this input directly: an occupation-level ranking of AI exposure.

We read the experiment through a standard model of occupational choice (Wiswall and Zafar, 2015; Conlon, 2021). Student i values occupation j according to

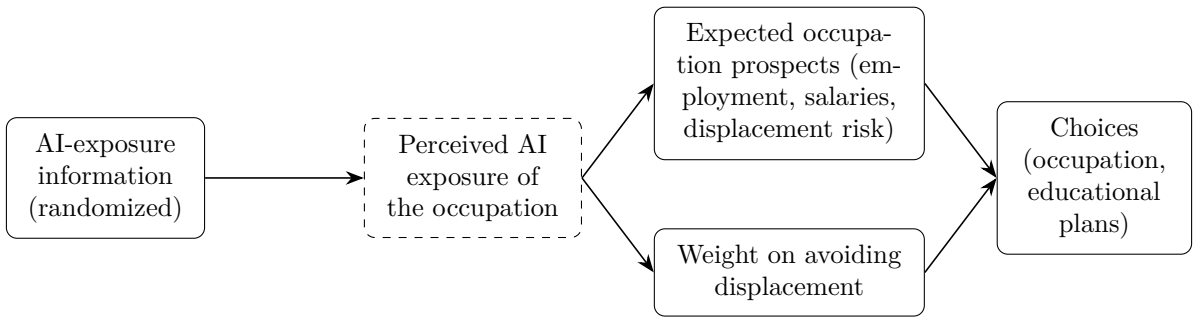
$$V_{ij} \;=\; \pi_{ij}\, u(w_{ij}) \;+\; a_{ij}, \tag{C1}$$

where π_{ij} is the student's subjective probability of being employed in the occupation, w_{ij} is expected earnings conditional on employment, and a_{ij} collects non-pecuniary attributes, among them insulation from technological displacement. The student ranks occupations by V_{ij}; the top-ranked occupation is the top choice we elicit.

The treatment provides information about occupations, not about students, so it can move the occupation-level inputs of equation (C1): beliefs about the occupation's employment rate, beliefs about its salaries, and the weight on displacement as an attribute worth

considering. What matters for choice, however, is partly personal: π_{ij} is the student's *own* employment probability, which responds to occupation-level news only insofar as the student ties her own draw to the occupation's average. Writing the revision as $\Delta\pi_{ij} = \lambda_i \Delta P_j$, the passthrough $\lambda_i \in [0, 1]$ measures how much of the market-level news the student internalizes. Unrealistic optimism about one's own outcomes and overplacement relative to others (Weinstein, 1980; Moore and Healy, 2008) push λ_i below one: bad news about the occupation is accepted for the market but discounted for oneself. Figure C1 summarizes the chain from the information to choices.

Figure C1: The Causal Chain from Information to Choices



Notes: This figure illustrates the chain that organizes the interpretation of the treatment effects (Section 3.2). Perceived exposure is not directly elicited (dashed node); the survey elicits the intermediate beliefs and the final choices (solid nodes). Section 5.4 decomposes, descriptively, how much of the choice response passes through the elicited intermediate nodes.

How far the occupation-level beliefs themselves move depends on the priors of the students. The signal concerns exposure—which occupations' tasks AI can perform—while the payoff-relevant object is employment, so the employment implication of a given exposure signal scales with the student's belief that AI will prove potent enough to matter; students who arrive expecting AI to destroy jobs should revise more from the same information, a complementarity that requires no departure from Bayesian updating. Receptivity works in the same direction: students weigh prior-congruent information as more credible (Rabin and Schrag, 1999; Gentzkow and Shapiro, 2006), so pessimists take the message at face value while skeptics discount it.

The framework yields four predictions, each examined in the body of the paper:

1. Occupation-level beliefs move in the direction of the information, and most where the news is worst—for students whose top choice is highly exposed (Section 5.1).

2. Own-outcome beliefs move by only λ_i times as much, so adverse occupation-level news widens the gap between what students believe about the market and what they believe about themselves (Section 5.1).

3. Choices respond in the direction of lower exposure: the stated top choice and the composition of the consideration set shift toward less-exposed occupations; and because the response runs through occupation-level objects rather than revisions to own-outcome beliefs, less than a third of the switching is mediated by the belief revisions we elicit (Sections 4.1–5.4).

4. Treatment effects are stronger for students whose priors accept AI's potency, on beliefs and on choices alike (Section 5.3).