



Early Evidence on Employer Responses to College Student Use of Generative AI

Casey Boyd-Swan

Kent State University

C. Lockwood Reynolds

Kent State University

We conduct a survey experiment on professionals with hiring experience to investigate how employers respond to students' use of generative AI in higher education. Prior to respondents reviewing two resumes, they were randomized into a control group or one of two treatment groups. One treatment involved framing AI use as assisting students while the other involved framing AI use as substituting for student effort and ability. Results show that the framing increased respondents' concerns about evaluating job applicants, with a larger effect in the second treatment group. There is some evidence that the first treatment increased interview interest and willingness to negotiate starting salaries; the second treatment clearly decreased both outcomes. Lastly, respondents were asked what information from applicants would help allay their concerns. While a number of factors allayed the concerns for the first treatment group, only in-person tests were selected by both treatment groups.

VERSION: October 2026

Suggested citation: Boyd-Swan, Casey, and C. Lockwood Reynolds. (2026). Early Evidence on Employer Responses to College Student Use of Generative AI. (EdWorkingPaper: 26-1604). Retrieved from Annenberg Institute at Brown University: <https://doi.org/10.26300/13ds-ed72>

Early Evidence on Employer Responses to College Student Use of Generative AI*

Casey Boyd-Swan[†]

C. Lockwood Reynolds[‡]

October 2, 2026

Abstract

We conduct a survey experiment to investigate how employers respond to students' use of generative AI in higher education. Prior to respondents reviewing two resumes, they were randomized into a control group or one of two treatment groups. One treatment framed AI use as assisting students while the other framed AI use as substituting for student effort and ability. There is some evidence that the first treatment increased interview interest and willingness to negotiate starting salaries; the second treatment clearly decreased both outcomes. When asked what would ease their concerns, only in-person tests were selected by both treatment groups.

JEL classifications: I21, I23, I26, J20, J24, D83

Keywords: Higher education, generative artificial intelligence, human capital, signaling, college credentials

*Statements and Declarations: The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper. This study was approved by the Institutional Review Board of Kent State University (protocol number 2423).

[†]Department of Economics, Kent State University, Kent, OH. cboydswa@kent.edu

[‡]Corresponding Author. Department of Economics, Kent State University, Kent, OH. creynol9@kent.edu

1 Introduction

ChatGPT was released on November 30, 2022, and OpenAI claimed 100 million users by January 2023. While generative AI had existed previously, ChatGPT provided an accessible platform where individuals could ask questions conversationally. Survey research consistently shows that AI use among college students is now widespread (e.g., Gallup and Lumina Foundation (2026)). Within higher education there has been considerable handwringing about what that use means for academic integrity (e.g., McMurtrie (2025), McIvor (2026), Whitford (2026), Vance (2026)). Institutions are responding by considering a return to standardized admissions tests (University of California), requiring proctored exams (Princeton) (Belkin, 2026b), and banning technology from core courses (University of Chicago). Stories of student misuse of generative AI have begun to reach the popular press (e.g., Whitford (2026), Belkin (2026a), Belkin (2026b), Svrluga (2026), Gioino (2026)). What remains unclear is how employers, who ultimately price a college degree, view student AI use. Employers adopting AI in their own workplaces may value graduates who arrive with AI experience. Alternatively, if employers grow to believe that students used AI in place of their own skills and effort, the credentials colleges award could lose value. These questions are now pressing: the first cohort of students who had access to generative AI for most of their college careers entered the labor market in Spring 2026.

This paper combines a survey of U.S. professionals with hiring responsibilities with an information experiment about student AI use to provide early answers to four questions. First, how concerned are employers about whether AI use makes it more difficult to identify high-quality job applicants? Second, how are hiring outcomes (interview likelihood, salary negotiations) different if employers believe that student AI use is a complement versus a substitute for student ability and effort? Third, do these effects differ by the applicant's GPA, race, and sex? Finally, what information would employers want from applicants to ease their concerns about the academic ability of recent college graduates?

Existing research offers reasonable, and often competing, predictions for these ques-

tions. Much of the growing economics literature on AI focuses on productivity and labor demand. Workers report that AI has improved their productivity (Bick et al., 2025) while also raising concerns about AI-related job loss (Hartley et al., 2024). Evidence suggests that AI assistance increased productivity for customer support agents (Brynjolfsson et al., 2025b) and software developers (Cui et al., ahead of print), with the largest gains found in less-experienced workers. In a randomized trial of patent attorneys, Autor et al. (2026) find that AI use increased average productivity but these gains in productivity only carried over to higher offline (no use of AI) productivity for senior attorneys. Evidence generally supports the idea that AI is reorganizing tasks at work (Dillon et al., 2025; Hampole et al., 2025; Humlum and Vestergaard, 2026; Freund and Mann, 2026; Alekseeva et al., ahead of print); evidence on employment and wages is less clear (see Alekseeva et al. (2021), Webb (2025), Hartley et al. (2024), Humlum and Vestergaard (2026), Teutloff et al. (2025), Johnston and Makridis (2026), Chen and Wang (2026)).¹ Part of the disagreement in this literature reflects the fact that AI complements some occupations and tasks while replacing others (see Autor and Thompson (2025), Althoff and Reichardt (2026)). Taken together, this literature suggests one plausible prediction for our setting: if AI is a productivity tool that employers are actively adopting and there is rising demand for AI skills (Alekseeva et al., 2021), then employers might reward applicants when reminded that students use AI, treating that use as preparation for the modern workplace. Evidence on this point is thin; Danilov et al. (2026) find that listing advanced AI skills on a resume has no average effect on callbacks, with positive effects only for some positions.²

A related literature investigates the ways in which automation and generative AI influence the hiring process; this literature is arguably closer to the purpose of this study. The general focus has been on how applicants may use AI to help them prepare application materials and the subsequent effects that has on hiring. Cui et al. (2025) find that a newly-

¹A related set of papers measures occupation or industry AI-exposure based on job tasks. See, for example, Felten et al. (2021), Eloundou et al. (2024), Hampole et al. (2025), Massenkoff and McCrory (2026), and Verschuere and Cameron (2026).

²One additional note of caution: extant research is a mix of studies using older automation technologies that predate ChatGPT and similar platforms, as well as studies using newer generative AI tools.

deployed AI cover-letter tool on an online job-posting platform improved alignment with job descriptions and raised callbacks, but weakened the link between cover-letter alignment and callbacks. In other words, the *signal* value of cover letter “quality” appears to have decreased as a result of AI-written cover letters. Galdin and Silbert (2025) similarly find employers highly valued customized applications before generative AI but not after, and Cowgill et al. (ahead of print) show experimentally that AI-written cover letters and emails sent to potential employers compress the signal gap between high- and low-quality applicants. Wiles et al. (2025), in contrast, find that algorithmic writing assistance on resumes increases hiring and wages; this potential contradiction may not be a result of changes to signal value, but rather as a result of resumes becoming easier to read.

The focus of this study is an extension of this literature, moving away from application materials to investigate how AI use (or perceptions of AI use in college) affects hiring decisions. If AI use lowers the cost of producing good grades, it may undermine the logic that high GPA signals high ability. Widespread AI misuse in college would then reduce the signal value of a college degree. In a similar vein to previous work which found AI use in producing application materials weaken the value of cover letters, research can apply the use of AI to the entire college career - this would include not only academic performance but also major choice. AI use in this setting could effectively reduce the cost of effort (and thereby substitute for skills) such that there is no longer a separating equilibrium between high and low performers. The cost of reduced signal value from college credentials may fall hardest on new graduates, who have little work history to offset a weaker credential (e.g., Pallais (2014)); this would be consistent with evidence that AI adoption is particularly affecting employment of less-experienced workers (e.g., Brynjolfsson et al. (2025a), Maasoum and Lichtinger (2025), Teeselink (2025)).³

Signaling theory would further suggest that different groups may be differentially impacted by reduced college credential signal values. Models of statistical discrimination

³There is also evidence that AI adoption is decreasing labor demand especially for jobs held by those with college degrees (Liu et al., 2025), consistent with the task-based substitution found in other papers. Similarly, Orr et al. (2026) find that majors most exposed to AI saw decreases in employment and earnings - relative to less-exposed majors - after the introduction of ChatGPT.

imply that when individual productivity signals become noisier, employers place more weight on group-level priors. A well-developed literature using resume audit studies document that Black applicants receive fewer callbacks than otherwise identical white applicants (Bertrand and Mullainathan, 2004; Kline et al., 2022; Quillian et al., 2017). Given these two literatures, one would expect that if AI makes college credentials or GPA less informative, employers would lean more heavily on race and sex, and existing racial and gender gaps in hiring outcomes would widen. Moreover, because AI exposure varies sharply across industries (e.g., Felten et al. (2021), Massenkoff and McCrory (2026)), one might expect the strongest responses from employers in the most AI-exposed industries.

The final question this study focuses on concerns what employers would turn to instead of college credentials or GPA when making hiring decisions. When the value of one credential falls, employers might reasonably substitute toward others they already use: institutional selectivity and prestige, which audit studies show employers reward (Gaddis, 2015; Deming et al., 2016), or newer credentials such as online microcredentials, which have become a popular way for job seekers to signal up-to-date skills, including AI skills.

Whether student AI use is complementary to learning or substitutes for effort and skill formation is an active question in education research, but less so in other literatures. For example, a meta review by Alubthane (2026) suggests that undirected student AI use reduces higher-order cognition while directed use enhances it in some cases. A randomized control trial in economics has found that students may benefit from using AI to augment their learning while those who use it to automate their work do not see longer term learning gains (Contractor and Reyes, 2026). If institutions can figure out how to harness AI to improve student learning while simultaneously limiting misuse then the eventual equilibrium for employer responses may be more similar to the responses from our experimental framing that suggests student AI use increases student productivity. If institutions cannot do this, then employer responses may be more similar to the results of our experimental framing suggesting that student AI use is a substitute for effort or skill formation. Resolving that question is beyond the scope of this paper. Instead, our experiment is designed to show how

employers respond under either belief.

In May and June 2026, we surveyed approximately 1,750 U.S. professionals with hiring responsibilities recruited through Prolific. After collecting baseline measures - including respondents' concerns about identifying high-quality applicants - we randomly assigned respondents to a control group or to one of the two treatment framings. All respondents then evaluated two otherwise similar resumes of recent college graduates, one with a high GPA and one with a low GPA, with applicant race and sex signaled through randomly assigned names. Prior to reviewing the resumes, the control group received only instructions. The mild framing group read a vignette acknowledging that some students misuse AI but that many use it responsibly and productively to improve performance. The strong framing group read a vignette emphasizing that AI can substitute for student effort and ability. Respondents then viewed each resume and rated how likely they were to interview each applicant and their willingness to negotiate a starting salary for each applicant. After evaluating the resumes, respondents rated what information would ease their concerns about applicants' academic ability.

Results suggest that most respondents at the time of the survey were not particularly worried that AI made it more difficult to identify high-quality applicants. While a minority of respondents were very concerned, this may be because they have seen some of the popular press articles - alternatively, it may be a reflection of their own experiences hiring recent college graduates. The results of our information experiment for this minority of responses are suggestive of employer responses if worry over student AI use becomes more widespread.

While literature may suggest that employers would reward AI familiarity, the mild framing produced no significant premium. Results from fixed effects models show no significant evidence that the mild AI framing decreased how high GPA resumes were evaluated. The strong AI framing sharply decreased respondent willingness to interview and negotiate salary, especially for high GPA applicants. This is consistent with employers discounting grades that may have been earned with AI, and complements evidence that generative AI weakens the signal value of application materials. Contrary to what signaling theory would

suggest, we find no systematic differences in these effects across applicant race and sex. There is a possible exception for white female applicants, who appear to avoid penalty. In additional specifications, we show that this result is driven by two specific industries and that removing respondents employed in these industries produces similar negative effects of the strong AI framing for all sub-groups of high GPA applicants. The lack of differential effect by race and sex may be due to respondent beliefs about differential AI use across these same groups. Chirikov et al. (2026) show variation in AI-based student cheating across academic disciplines, with lower use among female and minority students - though it is not clear that employers would be aware of these differences.

Finally, when asked what would ease their concerns about hiring newly-graduated job applicants, respondents in both framing groups were generally more positive about alternative in-person methods of validating applicant quality than the control group. Other signals of applicant quality, which previous research suggests improves hiring outcomes, were not viewed as particularly favorable by respondents in either AI framing group.

These findings carry serious implications for colleges and universities. If employers come to believe that AI substitutes for student effort and ability, our results suggest they will discount the credentials which institutions award. Interestingly, some of the same patterns around framing that we note among high GPA applicants can also be seen for the low GPA applicants, particularly for salary negotiation. The low GPA applicants are clearly less favored compared to the high GPA applicants, but our results suggest that both groups would be hurt if employers come to view student AI use as replacing student effort or ability. Institutions need to find ways to certify that students can and did do the work themselves. Their ability to do so may vary with sector, size and institutional resources.

This paper contributes to existing literatures in three ways. We extend research on AI and labor market signaling from application materials to the educational credential itself. We add to the resume audit literature by capturing a post-interview margin (salary negotiation) and by testing whether AI-induced signal noise widens known racial and gender gaps. Most pressing, we provide some of the first evidence on how employers are likely to

respond to the first cohort of AI-era graduates, at a moment when employer beliefs are still developing. The remainder of the paper proceeds as follows. Section 2 describes the survey, vignettes, and resumes. Section 3 presents results including a manipulation check, hiring outcomes, and potential solutions. Section 4 concludes.

2 Methodology

2.1 Respondent pool

We constructed a survey which was administered via Prolific, an online platform for survey research. Prolific inverts the typical survey recruitment process: instead of researchers creating surveys then using a sampling strategy to identify and recruit participants, the Prolific platform maintains a standing pool of people who are interested in completing surveys, from which researchers can target appropriate individuals for their survey. Participants are compensated for their time, at rates which vary across studies, but they are free to choose whether or not to participate in any particular survey available to them. Considerable research has found that the quality of responses and participation is higher than alternative platforms, such as Amazon’s MTurk (Peer et al., 2022; Douglas et al., 2023; Exley et al., 2026). The deliberately complex sign up process - including a video facial scan and government ID check - limits participation by AI bots and Prolific has implemented tools to prevent use of AI in responses.

One advantage of the Prolific platform is that participants answer a wide range of questions during enrollment that can then be used by researchers to target specific responder characteristics. We restricted the sample by using the following criteria: they reside in the United States, they are employed full-time, they work in the for-profit sector, they have worked for their current employer for at least 1 year, and they have “hiring” as part of their decision-making responsibilities.⁴ Every Prolific participant meeting all of these criteria

⁴As required by the platform, we also screened for willingness to participate in studies involving deception given that we did not initially disclose that our resumes were constructed. Note that participants are told that they matched the screening criteria but are not told what those criteria are so they were not aware of any deception until disclosure occurred at the end of the survey.

was then offered our survey in their dashboard. Dashboard information made available to participants included: survey name, description, compensation and estimated time to complete the survey. Depending on their account notification preferences, they may also have received an email announcement about the survey or a push notification from their app. The survey was named “Modern Hiring Practices” and the description stated that we were examining how hiring professionals evaluate entry-level job applicants. See Figure B-1 for an image of what potential participants were shown about the study. Note that participation in the survey was voluntary at every stage.

We estimated that the survey would take 11 minutes to complete and participants were paid \$3.50, equivalent to an hourly wage of \$19.09. Prolific also displays the actual median completion time and median hourly wage for each survey; ours were 9 minutes 25 seconds and \$22.30 per hour (respectively), a rate which the platform classified as competitive for our set of participants. Approximately 2,300 participants who were active on the platform in the prior 90 days matched our criteria. We collected complete responses from 1,750 between May 21 and June 8, 2026.⁵

2.2 Survey design and vignettes

The survey initially asked participants some background questions about demographics, work and participants’ use and knowledge of AI. Participants were then randomly assigned to see one of three primes.

All three primes shared a common set of instructions: participants were told that they would review application materials for two candidates applying for an entry-level position at a company similar to their own. Instructions further stated that, for confidentiality purposes, the researchers had reformatted the resumes. As discussed below, this reformatting included redacting the applicant’s last name along with any information that would identify the applicant’s university and major. Redacting last name ensures that the first names

⁵An additional 81 people opened the survey but then opted out on the initial consent screen. Another 31 people started but did not complete the survey. None of these individuals were compensated. Thus, among those we know saw the survey because they interacted with it, 94% completed it.

alone carry the race and gender signal. Redacting institution and major served a different purpose: due to sample size constraints and because we did not know *a priori* which industries respondents would be coming from, we could not tailor resumes to specific industries. Removing institution and major, and asking respondents to assume the applicants had attended universities and pursued majors typical of those the respondents normally hire from, allowed us to recruit across a broader range of industries. This departs from a conventional resume audit study, in which resumes are matched to real job openings; our design instead holds applicant credentials constant and isolates the effect of the AI framing across multiple industries. The tradeoff is that respondents may differ in the credentials they imagine, introducing unobserved variation - a limitation we acknowledge.

Participants in the control group received this instruction alone. However, participants randomly assigned to one of the two treatment primes received the same instructions, preceded by information about student use of AI in higher education. The first of these treatment primes, which we label the “mild AI” framing, reads as follows:

Educators and employers have long contended with challenges to academic integrity - for example, plagiarism, online essay sites, or for-pay cheating services. Recently, the introduction of AI writing tools has added a new dimension to these concerns. Tools such as ChatGPT and similar AI assistants have become widely accessible to students, making it easier to get help drafting papers, completing homework, or answering exam questions. While many students use such tools responsibly as study aids, instructors have documented cases in which students rely on AI to complete work they would otherwise do themselves. These concerns are ongoing and universities have not developed clear methods of addressing them.

This paragraph is designed to convey that AI can be helpful for students, but that its use may still violate academic integrity. Furthermore, it makes it clear that universities cannot effectively police this use. Participants assigned the second treatment prime, which we label “strong AI” framing, are instead shown this text:

In the last month, large news outlets such as the New York Times and the Wall Street Journal have reported on research showing grade inflation at US universities that corresponds to modern AI availability. Evidence suggests that AI technology has advanced to a point where automated systems - known as AI

agents - can now complete virtually any task routinely assigned in a university setting. This includes writing research papers, answering exam questions, solving problem sets, completing online coursework, and producing professional-quality work across nearly every undergraduate discipline. Unlike prior generations of AI tools that required significant human input to use effectively, newer AI agents can operate with minimal guidance. Universities have yet to successfully develop tools that can detect this use of AI.

This is designed to convey a stronger claim: it emphasizes how much AI could be used by students not just to aid them in their academics, but to actively do academic work for the student. In this case, the AI use is not a complement to the student’s own effort and ability, but potentially a substitute. While strong, the case presented in this vignette is not hyperbolic. A high profile report from MIT concluded that many assignments and tasks assigned to students can already be accomplished by AI (MIT’s Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training, 2026). Community colleges have been dealing with financial aid fraud associated with AI bots enrolling in institutions, signing up for classes and completing coursework (McMurtrie, 2025; McIvor, 2026). An AI program named Einstein was announced in February 2026 that students could grant access to their learning management systems and it would monitor and complete all activities for them (Heidt, 2026). While that program was taken down, more could easily follow.

2.3 Resumes and evaluation

After the prime, participants then reviewed the first resume and answered a series of questions, then they reviewed the second resume and answered the same series of questions. The primary question used to evaluate applicants was a 7-point Likert scale about the likelihood that the respondent would invite the applicant for an interview. This is our primary outcome of interest, and was designed to parallel the standard metric in resume audit studies where the outcome is whether or not a particular resume *was* invited to an interview. Since respondents were not actually filling a position, their responses may differ from behavior observed as a result of a live resume audit study. However, the results are still informative (including for future audit studies). First, randomization does the necessary work: an in-

dividual respondent may over- or understate their true interview likelihood based on how they view the question, but because both the prime and resumes were randomly assigned there should be no systematic differences across comparisons of interest. Second, we use a 7-point scale instead of a binary outcome, which permits finer differentiation in responses.

A second question asked how much more, above the offered starting salary, the respondent would be willing to pay if a candidate wanted to negotiate a higher salary. Conceptually, a respondent should be willing to negotiate a higher starting salary for stronger candidates. This provides a second dimension to the analysis: rather than a hypothetical interview decision, it monetizes applicant evaluation. To the extent the interview and salary responses diverge, this may provide information on post-interview discrimination, which normally cannot be observed in a conventional resume audit study.⁶

While respondents were led to believe that the resumes were real, they were constructed by the researchers. Because each respondent saw only two resumes, we designed the pair to be similar along several dimensions. They had similar work histories, similar listed skills, and each had a student group and a volunteer experience listed (see Figures B-2 and B-3 for examples of the resumes). The remaining resume attributes were then randomized, yielding 8 comparisons. Each respondent saw a high GPA (3.8/4.0) and a low GPA (2.7/4.0) resume but the order in which they saw them was randomized (some see high GPA first and some see low GPA first). First names were randomized to signal race and sex. Jake and Tyler signaled white male, Emily and Lauren signaled white female, Darnell and Jamal signaled Black male, and Keisha and Aaliyah signaled Black female. To select names that signal race and sex, we rely on the New York City Health Department’s Bureau of Vital Statistics birth records from 1994 to 2014.

⁶A traditional resume audit study can only observe whether a resume received a request for an interview so what might happen conditional on an interview is not observed. Some studies use follow-up surveys to learn more about the decision-making process (e.g., Boyd-Swan and Herbst (2019)).

2.4 Manipulation check and potential solutions

After reviewing the resumes, respondents were asked a question about whether they had been given information about student use of AI (checking whether they remember the prime they were shown). They were then asked a set of questions about concerns they had with regard to academic credentials given student use of AI. Lastly, they were asked to evaluate what information - if any - provided by the applicant would ease their concerns about the student's academic ability. They could rate a variety of pre-identified characteristics, such as small class sizes, in-person classes, in-person exams, independent studies, senior theses, paid or unpaid internships, selectivity of university and microcredentials earned online or in-person. They were also provided space to write other suggestions. The goal of these questions is to see what information individuals can communicate to employers - or actions that institutions of higher education can take - to limit concerns about academic integrity due to AI.

3 Results

3.1 Summary Statistics

Basic summary statistics about the sample are presented in Table 1. Approximately 73% of the sample is white, non-Hispanic and about 52% of respondents are male. Consistent with their role in hiring, the sample is relatively highly educated with 67% having a bachelors degree or more. The respondents work at organizations across a wide range of sizes and are spread across a number of industries. 92.2% of respondents indicated some use of AI at work, with a majority using AI at work from 1 to 10 hours a week.⁷ These variables are mostly balanced across treatment, with the exception of race where white/non-Hispanic are somewhat higher in the control group than in the two treatment groups (see also Table

⁷Respondents also indicate that AI is used regularly at their organization (Figure A-1); existence of organization AI policies is mixed (Figure A-2); and the respondents are generally supportive of AI in the workforce (Figure A-3).

A-1).⁸ Our primary models utilize individual fixed effects which account for any sample differences and we also explicitly explore the role of respondent race.

Table 2 presents the means of the outcomes for the four sets of race/sex resumes in our survey. The average score for interview likelihood is 4.69 on a 7-point scale with white male applicants the lowest average (4.61) followed by Black male applicants (4.70), Black female applicants (4.71) and white female applicants with the highest average (4.77). The average willingness to negotiate is 1.69 percent above the starting salary. Again, the lowest percent is for white male applicants (1.61) followed by Black male applicants (1.64). However, white and Black female applicants switch with white female applicants averaging 1.75 and Black female applicants averaging 1.77.

Lastly, we asked two questions about the level of concern respondents have about both (1) identifying high-quality employees/applicants and (2) employees/applicants using AI as a substitute for own skills. These questions are posed prior to the information experiment so that we have a baseline of concern. Figure 1 shows histograms of the responses. In both cases, the mode is at a 5 on the 7-point scale, so just above the middle of the scale. However, there is a long tail to the left indicating that large percentages of respondents are not particularly concerned - at this point in time and prior to the information experiment - about AI use in either case. This provides room for the information experiment to affect perceptions, but also means that as of June 2026 many individuals with hiring experience are not particularly concerned about AI use. That could change as more stories are published in the popular press and as these individuals experience hiring new applicants who have been able to access AI throughout their college career.

The level of concern reported for both questions appears to increase with the AI exposure of the respondent's industry (Figure A-4), using the AI exposure measure of Felten et al. (2021). This suggests that high AI exposed industries are more worried about the potential workforce problems either because of their knowledge about AI use at work or

⁸There is also balance across other variables including how much AI is used at their place of work, how they view AI (positive/negative), how frequently they hire entry-level candidates, whether a bachelors degree is typically required for entry-level positions and how important undergraduate GPA is in applications. See Table A-1.

because of their recent experience hiring students versed in AI.

3.2 Manipulation check

Table 3 shows the results of our manipulation check across the three primes. Each column shows the results of a regression with an indicator variable for the mild AI framing and an indicator for the strong AI framing. The coefficients represent the difference for the individuals with each framing relative to the control (the excluded category) which means that the constant represents the average level of response in the control group.

The first column shows the results of the 7-point scale asking if any text that they had read before evaluating resumes had mentioned how AI tools may affect academic credentials. The coefficient on the mild AI framing is 1.26 and the coefficient on the strong AI framing is 1.67, both statistically different increases from the control (the constant in the regression). This indicates that individuals do remember seeing the framing.⁹

The second column of the table shows the results of a second question gauging their level of concern that academic credentials of recent college graduates may have been produced with the help of AI tools. Again, both the mild and strong AI framing coefficients are positive and statistically significant and, furthermore, the strong AI framing is larger in magnitude and statistically different than the mild AI framing.

This evidence suggests two things. First, the information experiment was successful in manipulating respondents and that the strong AI framing led to more concern than the mild AI framing. Second, the fact that individuals were manipulated means that employers are - as of this writing - not fully aware of the potential problems that student AI use may have on hiring. So while stories about student AI use are beginning to appear in newspapers and online media, and concerns about student AI use are highly discussed in higher education, information about student AI use may not yet be widespread.

⁹Prior to the prime, all respondents were asked two questions about whether use of AI was making it more difficult to evaluate employees or job applicants to provide a baseline for what these respondents believed about AI. Thus, even the control group was shown something about AI use among students which is why the constant is not 1 (the lowest possible value).

3.3 Interview and salary effects

To investigate differences in the primary hiring outcomes - willingness to hire and willingness to negotiate higher salary - we estimate a model where we triple interact a set of categorical variables: treatment (3 levels), binary GPA variable (high/low) and a variable for the race/sex combination of the applicant (4 levels: white male, white female, Black male, Black female). Specifically,

$$\begin{aligned}
Y_{ir} = & \alpha + \beta_1 Treat_{ir} + \beta_2 GPA_{ir} + \beta_3 RaceGen_{ir} + \\
& \beta_4 Treat_{ir} \times GPA_{ir} + \beta_5 Treat_{ir} \times RaceGen_{ir} + \beta_6 GPA_{ir} \times RaceGen_{ir} + \quad (1) \\
& \beta_7 Treat_{ir} \times GPA_{ir} \times RaceGen_{ir} + \phi_i + \epsilon_{ir}
\end{aligned}$$

where Y_{ir} is the outcome for respondent i for resume r , $Treat_{ir}$ is the set of binary variables for treatment (control is excluded category), GPA_{ir} is binary variable for high GPA (low GPA is excluded category) and $RaceGen_{ir}$ is the set of dummy variables for the race/sex combinations of the resumes (white male is the excluded category). Because each respondent views both a high and a low GPA resume with the race and sex randomized there are two observations per respondent. We therefore include respondent fixed effects ϕ_i and cluster standard errors at the respondent level. We include one extra covariate which is an indicator capturing which resume the respondent saw first. Due to the individual fixed effects $Treat_{ir}$ is dropped from the estimation.

It is worth noting that we compare many treatment, GPA, and applicant race-sex categories across our hiring outcomes, respondent subgroups, industry groupings, and potential solutions, so some of the individually significant differences we report could arise by chance. For the main results, we apply the Romano-Wolf correction (Romano and Wolf, 2005; Clarke et al., 2020) for multiple hypothesis tests and find that the statistical significance holds, while insignificant results remain insignificant. Nevertheless, we still acknowledge that due to our small sample size we have power issues. We therefore place the most weight on patterns that

are consistent across outcomes and specifications rather than on any single given estimate.

Interpreting the coefficients with this many interactions can be difficult so we instead estimate the predicted margin following the regression and then graph the differences in these margins. Figure 2 shows the results for the predicted interview likelihoods (measured on a 7-point scale where 1 is very unlikely and 7 is very likely). Exact point estimates and standard errors are presented in Table A-2. The predicted margins are all relative to the low-GPA white male results, as they are the excluded category, whose predicted interview likelihood is 4.18. The subgraph on the left shows the results for the low GPA resumes (Figure 2a). There is little evidence that either AI framing affected respondents' likelihood of interviewing candidates for any set of resumes, particularly for white and Black males. The strong AI framing is associated with a small decrease in interview likelihood for white and Black females (-0.29 relative to control for white females and -0.19 for Black females relative to control) which brings their estimate down to a similar level as white males but the results are not statistically significant.

As expected, the likelihoods of interview are higher in the right subgraph (Figure 2b) consistent with having a higher GPA on their resume. The interview likelihood for high GPA white males in the control group is 5.22, a full point higher than the low GPA (approximately a 25 percent increase relative to the baseline). Similarly, the other high GPA resumes in the control group also have likelihoods that are 0.8 to 0.9 points higher (Black male 5.18, white female 5.22, Black female 5.25). In this case, however, there is some evidence that AI framing matters. The mild AI framing is similar to the control group for all sets of resumes. For white male, Black male and Black female resumes, however, the strong AI framing produces a lower likelihood of interviewing, although the results are not quite statistically significant at the 95% level. Additionally, the gap between the high GPA and low GPA likelihoods in the strong AI prime has decreased relative to the control prime. For white male resumes, the high to low GPA gap decreases from 1.04 in the control prime to 0.62 in the strong AI prime (approximately a 40 percent decrease). For Black males this gap decreases from 0.89 to 0.63 and for Black females the gap decreases from 0.95 to 0.75.

White female resumes however show no effect of the strong AI framing (but we will show that this is sensitive to respondent industry). These results are consistent with respondents being more hesitant about student GPA when they received the strong AI framing about AI substituting for student ability or effort.

The effects of framing are more pronounced when looking at the willingness of respondents to negotiate a starting salary higher than the initial offer. Again, respondents are more willing to negotiate higher starting salaries for high GPA applicants compared to low GPA applicants. On average, respondents with the control prime are willing to negotiate a 1.97 percent increase in salary for high GPA white male resumes compared to 1.23 for the low GPA white male resumes. The equivalent numbers are 2.16 instead of 1.47 for Black male resumes, 1.99 instead of 1.33 for white female resumes, and 2.18 instead of 1.16 for Black female resumes.

For the low GPA resumes, there is some evidence of differential effects of framing for the Black male, white female and Black female resumes. There is a fairly consistent pattern throughout that mild AI framing leads to higher - although not statistically different - negotiated starting salaries compared to the control group. The strong AI framing reverses the pattern and lowers the negotiated starting salary. Relative to the mild AI prime, the salary negotiation percent decreases by 0.51 for Black male resumes, decreases by 0.31 for white female resumes, and decreases by 0.41 for Black female resumes. These are decreases of approximately 20 to 30 percent relative to the mild AI baseline.

The priming seems to generally have even larger effects for the high GPA resumes in the subgraph on the right (Figure 3b). There is again some evidence of an increase for the mild AI prime versus the control prime. For all groups the strong AI framing is lower, and with the exception of white female resumes, statistically different from the mild AI framing. For high GPA white male resumes the point estimate drops from 2.28 to 1.76 between the mild and strong AI primes. The equivalent numbers are 2.27 to 1.72 for Black male resumes, 2.15 to 1.87 for white female resumes and 2.51 to 1.78 for Black female resumes.

The effects may be more pronounced for salary negotiation because the outcome

(percentage above initial salary offer) was more salient than the interview scale. If so, it would suggest that our interview results may underestimate the effects compared to what would be found in a field resume audit study where actual interview decisions are being made. Alternatively, it may suggest that the interview likelihood is actually less affected than post-interview outcomes which are normally missed in resume audit studies.

The fact that high GPA white female applicants are less affected by the strong AI framing is driven entirely by two industries in our data: retail trade and information technology & software. The fixed effects specification we use will not fully account for these industry effects unless the respondent sees both a high and low GPA application from white female applicants. In Figure A-5 we use our fixed effects estimation separating respondents into two groups. The first group are respondents in retail trade and information technology & software. The other group is all other respondents. The top panel shows the first group and the positive effect for high GPA white female applicants with the strong AI framing is apparent. In the bottom panel using all other industries, the high GPA white female applicants show the same pattern as other high GPA applicants of the strong AI framing, lowering interview likelihoods and negotiated salary.

We also estimated correlated random effects models which allow us to include respondent controls and allows us to estimate different margins for the low GPA white male applicants under different framings. As covariates we include which resume was viewed first (as in the fixed effects models), binary demographic variables of the respondent (male, Black, Asian, Hispanic, multiethnic) and categorical variables for respondent educational attainment, organization size, stated importance of GPA in a job application, and the respondent's industry. The results from these models are shown in Figure A-6 and mostly follow similar patterns to Figures 2 and 3.

3.3.1 Respondent race and sex

We next turn to exploring how outcomes vary by different subsets of respondents. The graphs do not include confidence intervals to make them easier to read but we note that not

all of the results are statistically significant. That is mostly due to the relatively small size of our sample so that subgroups of respondents can be too small to generate enough power to detect statistically meaningful differences. The results, however, do show patterns that are informative if appropriate caution is used with interpretation.

Figure 4 presents the predicted margins for willingness to interview for male respondents. There are four subgraphs. The graphs in the top row (4a and 4b) show the predicted margins for the mild AI framing compared to the control while the graphs in the bottom row (4c and 4d) show the same for the strong AI framing relative to the control. The graphs in the left column show the results for low GPA resumes and the graphs in the right column show the results for high GPA resumes. Within each subgraph is a set of arrow charts where a closed circle indicates the predicted margin for the control group, an open symbol represents the mild AI framing, and an open square represents the strong AI framing. Colors are used to denote the four race/sex combinations of the resumes. Finally, each subgraph shows the results for white male respondents and then nonwhite male respondents.

Male respondents show an increased likelihood of interviewing high GPA white male and white female applicants with the mild AI framing relative to the control. They are somewhat less likely to interview Black male applicants and results for Black female applicants depend on respondent race/ethnicity. Relative to the control, the strong AI framing decreases the likelihood of interviewing high GPA Black male and Black female applicants (regardless of respondent race) while increasing the likelihood of interviewing white female applicants. White male respondents do not change their likelihood of interviewing white male applicants but nonwhite respondents lower their likelihood.

The results for female respondents in Figure 5 show a slightly different pattern. The mild AI framing decreases the likelihood of interview for white male and white female applicants among white female respondents while nonwhite female respondents decrease their likelihood of interviewing white male applicants but increase their likelihood of interviewing Black male applicants. The strong AI framing leads to lower likelihoods of interviews for all groups although magnitudes vary.

Switching to the salary negotiation results for male respondents in Figure 6 the mild AI framing leads to white male respondents to negotiate salaries increases of 0.5 to 1 percentage points for high GPA white male and female applicants but increases by less for Black male and female applicants. Nonwhite male respondents show no effect of the mild AI framing except for a decrease in negotiated percentage for Black male applicants. The strong framing decreases the negotiated salary for Black male and Black female applicants but the white male respondents show increases for white male and female while nonwhite male respondents show decreases for white male applicants and no effect for white female applicants.

Finally, for female respondents in Figure 7 there are fewer positive effects of the mild AI framing than male respondents, and these are concentrated on Black male and Black female applicants. This is a bit different than the interview likelihoods which either did not change or decreased with the mild AI framing. However, like the interview results, the strong AI framing reduces the negotiated salary levels for all groups on average particularly for white female respondents.

3.3.2 Industry

We next consider the industry of the respondent, in particular how much exposure their industry has to AI. On the one hand, industries with high AI exposure may have more interest in hiring students who can use AI technologies compared to industries where AI use is low. On the other hand, high AI exposure industries may be more aware of the potential problems of AI use.

In our primary estimates we map the industries of respondents to the AI exposure classification of Massenkoff and McCrory (2026). We categorize industries into “low exposure”, “mixed exposure” and “high exposure” (see Table A-3 for the exact mapping). We acknowledge that the categorization is somewhat ad hoc but they are meant to capture broad sets of industries; we only have 19 broad industries in our survey so we are forced to use broad classifications. We find similar results when we use terciles of the continuous AI exposure measure of Felten et al. (2021) but prefer our categorization because it is based

on more recent data about industry AI exposure. Future work may want to explore smaller industry classifications; alternatively, detailed *occupations* is likely more important as AI use is more associated with the occupation than the industry. In our survey, we do not specify the occupation so we cannot utilize that dimension but this would be something that a future resume audit study should consider.

One final limitation is that having a measure at the industry level means that it could be capturing other industry effects that are correlated with AI exposure. One example is the frequency of remote work: people using computers are more likely to be able to work remotely and more likely to have a use for AI. Suppose that it is more difficult to supervise workers in a remote environment making employers more concerned about new hires in these industries. This could lead employers in these industries to become more worried about the possibility that student AI use has substituted for effort or ability.

Using data on 18-64 year olds in the 2022-2026 monthly CPS files provided by IPUMS (Flood et al., 2025), we construct the percent of workers in each of our industries who work remotely at least some of the time. For our analysis, we categorize industries into “low/middle/high”, where low is defined as less than 15% remote work, middle is 15 to 49.9% remote work, and high is 50% or more remote work. These categories are admittedly ad hoc but are similar to terciles. We prefer this categorization over terciles because the cut points are cleaner but we will show that we get similar results using terciles.

As expected, remote work is correlated with AI exposure (Figure A-7). On average, 13.8% of workers in our low AI exposure measure work remotely. That rises to 23.0% for the mixed AI exposure and then 55.4% for the high AI exposure. Similarly, the correlation between percent remote work in the industry and the Felten et al. (2021) measure of AI exposure is 0.83. The connection is particularly strong for the highest AI exposure industries: the 5 industries in our “high AI exposure category” are also the top 5 for remote work and four of the five have over 50% remote work.¹⁰

¹⁰ As shown in Table A-3, the 5 “high AI exposure” industries are “legal services”, “finance and insurance”, “information technology and software”, “professional, scientific and technical services” and “media, publishing and communications”. Only “legal services” has remote work less than 50% but it is close at 47%.

Overall, our industry-based estimates cannot completely disaggregate between AI exposure and remote work. Therefore, we first present overall results - not separating by applicant race and sex - for low/middle/high categories of AI exposure and percent remote work in Figure 8 (results are similar when using tercile categories of both measures in Figure A-8). These results are estimated using a version of our individual fixed effects model where instead of interacting by candidate type (race/sex) we just control for candidate type. This means that we are only estimating the average effect of the treatments for the high GPA relative to the low GPA (low GPA is the left out category).

We find similar results when using industry AI exposure measures or industry remote work measures, particularly for our “high” category which is essentially the same industries whether measured by AI exposure or remote work. While higher GPAs are associated with higher interview likelihoods for all measures and categories, there is not much variation across treatment within GPA group for the lowest and highest categories. Instead, the strong AI treatment reduces the likelihood of interview for the “middle” industries under either measure. In terms of salary negotiation, there is less overall variation within measure and industry but some suggestion of lower values for high GPA applicants in the strong AI framing condition.

Of course, these estimates may mask important variation across resume types. For simplicity, we present estimates by applicant race/sex across industries using the low/middle/high AI exposure based on Massenkoff and McCrory (2026) in Figure 9, but note that results are similar using other measures (Figures A-9 and A-10). The most striking results are again in the middle category of AI exposure. There is a clear decrease in the interview likelihood for all resumes following the strong AI treatment. In fact, there is a decrease for the low GPA as well but it is not statistically significant. The strong AI framing also seems to reduce willingness to negotiate for all high GPA resumes in the middle AI-exposure industries. The strong framing also has negative effects for low GPA female applicants in the high AI-exposed industries.

There is some evidence that the mild AI framing may have increased application out-

comes for some resumes in some industries. In the middle AI-exposed industries, interview likelihoods appear to increase for high GPA white female applicants and Black female applicants although neither is statistically different from the control. Willingness to negotiate salary in these industries may also have increased for low GPA Black female applicants and high GPA white male, white female and Black female (again, none are statistically different from the control value). In the high AI-exposure industries the mild framing may have increased the willingness to interview and negotiate for all low GPA resumes relative to white males. Among high GPA applicants in these industries, Black male applicants see a large increase in likelihood of being interviewed and Black male and Black female applicants see large increases in willingness to negotiate salaries.

3.4 Employer views of potential solutions

Lastly, respondents were asked to rate how much different pieces of information may help ease their concerns about the academic ability of applicants. The scale ranged from “Not well at all” (1) to “Extremely well” (5). Point estimates were produced from a regression estimated for each category with indicators for which priming the respondent received. Figure 10 presents the point estimates across treatment groups with 95% confidence intervals, ordered by decreasing effectiveness as rated by the control group.¹¹

While differences are not always statistically different across treatment several patterns and results stand out. First, the control group consistently shows lower beliefs in the effectiveness of each option, consistent with them not having been prompted with information about student use of AI. Second, the strong AI framing group indicates less effectiveness for most options than the mild AI framing group. So across most categories, the mild AI framing is associated with more positive views that extra information can be useful for evaluating job candidates but the strong AI framing is more pessimistic. The mild framing group - who were prompted about how AI can help students - looks more positively on things like projects (independent studies/senior thesis) and internships (paid/unpaid) than the strong

¹¹Full distributions for each response are presented in Figures A-11 to A-20.

AI framing who were prompted about how AI could be a substitute for student ability and effort. Similarly, the mild AI framing group seem more positive about in-person classes and selectivity of schools than the strong AI group.

Third, online vs. in-person differences appear important. Both the mild AI and the strong AI framing are more reassured by in-person exams than the control group and the difference between the two groups is small. Relatedly, no group was very reassured by online microcredentials compared to in-person microcredentials. And while not statistically different, the means for online microcredentials decreases with AI framing. The difference between the microcredentials could be beliefs about the inferior quality of online programs, but when compared with other responses is suggestive that proctoring of exams is an important component of proving academic quality. While the educational quality of small classes or in-person classes could be viewed as superior to online courses they would similarly be relatively easier for faculty to validate “in-person” work. However, those courses could still have online components and assessments. It is only in-person exams - which are more likely in small, in-person classes - that both framing groups find most reassuring. This implies that AI is likely to further worsen the outcomes for graduates of online programs who were significantly less likely to receive callbacks in 2017 (Lennon, 2021) long before generative AI was available.

Respondents were given an option to provide a free response for “other” along with a rating. Not all respondents provided additional responses and the responses included options that would ease concerns as well as options that respondents do not think would ease their concerns much. There were ultimately 267 responses that could be organized by theme into categories. We then repeated our regression approach for each of those categories and the results for all categories with at least 20 responses are shown in Figure 11. Note that the confidence intervals here are much larger due to the smaller sample sizes (relatively few respondents providing each category) and noise from the categorization of responses. The most common response was volunteer activities which was rated higher by the strong AI framing than the other two groups whose mean is approximately the same. This is consistent

with previous evidence: the group shown the most concerning information about student AI use is willing to look more at information that is non-academic but which cannot be done by AI. There is also a suggestion that the interview may be more important among the two treatment groups, also consistent with respondents wanting to evaluate the individual in person and rely less on their academic record.

4 Conclusion

The first cohort of college students with access to generative AI for their college careers entered the labor market in 2026. This paper explores how the people who make hiring decisions in this labor market are likely to respond. Using a survey experiment deployed to approximately 1,750 U.S. professionals with hiring responsibilities, we randomly exposed respondents to information framing student AI use either as an aid that can nonetheless be misused (mild framing) or as a technology capable of doing students' academic work for them (strong framing), and then asked them to evaluate a pair of otherwise similar resumes that differed in GPA and the applicant's race and sex.

Four findings stand out from this experiment. First, as of mid-2026, most hiring professionals were only moderately concerned that student AI use makes it harder to identify high-quality applicants, although a meaningful minority were very concerned. Second, those beliefs are malleable: both AI framings raised concern to a certain degree about the credentials of recent graduates, but the strong framing raised it significantly more than the mild framing. The fact that this shift in framing moved respondents suggests two things: one, that complete information about student AI use is not widespread among employers, and two, that employer concerns are likely to increase as press coverage accumulates and hiring managers gain direct experience with AI-era graduates.

The third finding which stands out is the connection between employer hiring beliefs about AI and downstream hiring outcomes. The evidence largely suggests that survey respondents shown mild framing did not penalize applicants relative to the control group; in

some cases, respondents shown mild framing were somewhat more willing to negotiate higher starting salaries, though these differences are not statistically significant. In contrast, the strong framing was associated with decreased interview likelihood scores and, more starkly, willingness to negotiate salary especially for high-GPA applicants. We may expect employers would discount the signal value of grades if they believe they may have been earned with the use of AI (e.g., a higher GPA is reflective of AI use instead of higher student ability); evidence from the experiment suggests this is occurring, with high GPA applicants, regardless of race/sex, receiving significantly larger penalties on interview likelihood and salary negotiations than low GPA applicants. This finding complements evidence that generative AI weakens the signal value of cover letters and other job application materials (Cui et al., 2025; Galdin and Silbert, 2025; Cowgill et al., ahead of print). It is also consistent with concerns that AI adoption may fall hardest on less-experienced workers or those newly entering the labor market, since they have little besides their educational credentials to signal ability (Brynjolfsson et al., 2025a; Maasoum and Lichtinger, 2025; Teeselink, 2025). That the strong framing moved salary negotiation more than interview likelihood does suggest that some of the cost of the prevalence of AI on job-seekers may be borne after the interview stage, a margin that conventional resume audit studies cannot observe by their very design. It is worth noting that, unlike findings from most conventional resume audit studies, we do not find differential interview likelihood or salary negotiations effects across applicant race and sex characteristics signaled on the resumes; this is not necessarily surprising given the lower-stakes nature of the survey experiment (hypothetical instead of actual hiring decisions). In addition, work from Chirikov et al. (2026) suggests that there is uneven access to AI by race and sex, so employers may respond differently to applicants depending on who they believe is most likely to use AI.

Finally, when asked what would ease their concerns, respondents in both treatment groups rated in-person tests as more reassuring than the control group did; indeed, in-person tests were the option on which the two treatment groups agreed most closely. Those in the mild framing group were generally also more receptive to in-person and alternative ways

of capturing ability signals such as internships, independent studies, senior theses, and in-person classes. In contrast, the strong framing group was more skeptical of most academic signals and, in open-ended responses, pointed toward capturing information which cannot be easily (re)produced by AI: volunteer and extracurricular activities, work experience, and in-person interviews. Importantly, online microcredentials - which have been popular post-degree supplements to signal ability - and small class sizes - which have historically been signals of higher-quality learning environments - ranked lowest across all groups; attending a more selective school ranked near the bottom as well. This suggests that signals which job-seekers have previously relied on to improve their prospects in the labor market may be increasingly less useful in the age of generative AI.

Taken together, these findings have several implications for colleges and universities. The first is that institutions of higher education should not treat AI purely as a threat to academic integrity which needs to be policed. More than 90 percent of our respondents report using AI at work, and the mild framing, which acknowledged that many students use AI responsibly or as a productivity tool, did not reduce employers' evaluations of applicants. Employers who expect new hires to work alongside AI are unlikely to reward graduates who have never learned to use it well, so integrating AI into curricula and teaching students how best to use AI appears consistent with what the labor market demands.

The second implication is that integration alone is not sufficient. If employers come to believe that AI can substitute for student effort and ability, our results suggest that they will discount all credentials that institutions award, including (and perhaps predominately so) those of their strongest students. Institutions, therefore, need credible ways to certify that students can do the work without relying on AI. Our respondents point to in-person, proctored assessment as the most persuasive of such signals. The challenge then becomes that in-person assessment at scale is costly: it is most attainable for low-enrollment programs, small institutions, and well-resourced institutions, and least attainable for large public universities, access institutions, and online programs that serve many working and nontraditional students. There is a real risk that the burden of verification widens already

existing gaps across institutions. On the other hand, employers in our survey did not simply fall back on selectivity; this suggests there is room for institutions of all types to compete on the issue of assessment credibility assuming scalable approaches, such as proctored testing centers, shared assessment infrastructure, or verified capstone or internship work, can be developed. For students, the results suggest that experiences outside the classroom which are difficult to outsource to AI, including internships, work experience, and demonstrated projects, become increasingly valuable complements to the transcript.

Several limitations in this study point to directions for future work. Our outcomes are hypothetical evaluations rather than real hiring decisions observed in conventional resume audit studies; although random assignment of both the vignettes and the resumes mean that any general over- or understatement should not bias our comparisons, magnitudes likely differ if compared to expected audit study estimates. To recruit across many industries we redacted institution and major reported on the resumes, so respondents would have imagined different credentials. We observe industry that the respondent is employed in but not occupation, and our industry measures of AI exposure and remote work are too highly correlated to separate. Field audit studies that vary AI-related signals on resumes (for example, AI coursework, AI certifications, or verified in-person assessment) and that match resumes to real job openings by occupation would more cleanly establish whether the responses we document here appear in actual hiring decisions. Future research should also take advantage of the fact that employer beliefs about student (and worker) AI use are still evolving. The next few graduating cohorts offer a narrow window in which to measure these beliefs as they change, for institutions to deploy countermeasures to guard against credential discounting, and for research to gauge the degree to which these countermeasures are successful.

References

- Alekseeva, Liudmila, José Azar, Mireia Giné, and Sampsa Samila**, “Artificial intelligence adoption and the demand for managerial expertise,” *Strategic Management Journal*, ahead of print.
- , —, —, —, —, and **Bledi Taska**, “The demand for AI skills in the labor market,” *Labour Economics*, 2021, *71*, 102002.
- Althoff, Lukas and Hugo Reichardt**, “Task-Specific Technical Change and Comparative Advantage,” Working Paper 35353, National Bureau of Economic Research June 2026.
- Alubthane, Fawzia Omer**, “Amplifier or substitute? A systematic review of generative AI’s impact on higher-order cognitive skills among university students,” *Frontiers in Psychology*, 2026, *Volume 17 - 2026*.
- Autor, David and Neil Thompson**, “Expertise,” *Journal of the European Economic Association*, 08 2025, *23* (4), 1203–1271.
- , **Tanya Rodchenko, Josh Martin, Zanna Iscenko, Scott Strand, David Pearl, and Melissa Ferere**, “Does AI Assistance Enhance or Erode Expertise? Evidence from a Three-Month Field Experiment in Patent Drafting,” Working Paper 35720, National Bureau of Economic Research September 2026.
- Belkin, Douglas**, “Princeton Changes Its 133-Year-Old Honor Code Over AI Cheating Fears,” *Wall Street Journal*, 2026. Accessed on 8/26/2026, <https://www.wsj.com/us-news/education/princeton-cheating-ai-proctors-2a1cf62e>.
- , “University of California Professors are Begging Schools to Reinstate the SAT,” *Wall Street Journal*, 2026. Accessed on 8/26/2026, <https://www.wsj.com/us-news/education/university-california-sat-requirement-reinstate-c3e32712>.
- Bertrand, Marianne and Sendhil Mullainathan**, “Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination,” *American Economic Review*, 2004, *94* (4), 991–1013.
- Bick, Alexander, Adam Blandin, and David Deming**, “The Impact of Generative AI on Work Productivity,” Technical Report, Federal Reserve Bank of St. Louis 2025. Accessed at <https://www.stlouisfed.org/on-the-economy/2025/feb/impact-generative-ai-work-productivity> on 8/19/2026.
- Boyd-Swan, Casey and Chris M. Herbst**, “Racial and Ethnic Discrimination in the Labor Market for Child Care Teachers,” *Educational Researcher*, 2019, *48* (7), 394–406.
- Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen**, “Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence,” Technical Report November 2025.
- , **Danielle Li, and Lindsey Raymond**, “Generative AI at Work,” *The Quarterly Journal of Economics*, 05 2025, *140* (2), 889–942.

- Chen, Mark A. and Joanna (Xiaoyu) Wang**, “Displacement Versus Augmentation: The Effects of AI on Employment Dynamics and Firm Value,” SSRN Working Paper 5246388 June 2026. Accessed at dx.doi.org/10.2139/ssrn.5246388 on 8/27/2026.
- Chirikov, Igor, Ivan Smirnov, and René F. Kizilcec**, “Generative AI use and misuse call for assessment reform in higher education,” *Science*, 2026, *392* (6800), 818–820.
- Clarke, Damian, Joseph P. Romano, and Michael Wolf**, “The Romano–Wolf multiple-hypothesis correction in Stata,” *The Stata Journal*, 2020, *20* (4), 812–843.
- Contractor, Zara and Germán Reyes**, “Experimental Evidence on the Learning Impact of Generative AI,” IZA Discussion Paper 18792, Bonn Jul 2026.
- Cowgill, Bo, Pablo Hernández-Lagos, and Nataliya Langburd Wright**, “Does AI Cheapen Talk? Theory and Evidence from Global Entrepreneurship and Hiring,” *Management Science*, ahead of print.
- Cui, Jingyi, Gabriel Dias, and Justin Ye**, “Signaling in the Age of AI: Evidence from Cover Letters,” 2025. Accessed at <https://arxiv.org/abs/2509.25054> on 8/19/2026.
- Cui, Kevin Zheyuan, Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz**, “The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers,” *Management Science*, ahead of print.
- Danilov, Anastasia, Teo Firpo, and Lukas Niemann**, “The Elusive Returns to AI Skills: Evidence from a Field Experiment,” *Academy of Management Proceedings*, 2026, *2026* (1), 17112.
- Deming, David J., Noam Yuchtman, Amira Abulafi, Claudia Goldin, and Lawrence F. Katz**, “The Value of Postsecondary Credentials in the Labor Market: An Experimental Study,” *American Economic Review*, March 2016, *106* (3), 778–806.
- Dillon, Eleanor W, Sonia Jaffe, Nicole Immorlica, and Christopher T Stanton**, “Shifting work patterns with generative AI,” Working Paper 33795, National Bureau of Economic Research May 2025.
- Douglas, Benjamin D., Patrick J. Ewell, and Markus Brauer**, “Data quality in online human-subjects research: Comparisons between MTurk, Prolific, CloudResearch, Qualtrics, and SONA,” *PLOS ONE*, 03 2023, *18* (3), 1–17.
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock**, “GPTs are GPTs: Labor market impact potential of LLMs,” *Science*, 2024, *384* (6702), 1306–1308.
- Exley, Christine, Can Celebi, Soren Harrs, Hannu Kivimaki, Marta Serra-Garcia, and Jeffrey Yusof**, “Mission Possible: The Collection of High-Quality Data,” CESifo Working Paper 12535, CESifo 2026.
- Felten, Edward, Manav Raj, and Robert Seamans**, “Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses,” *Strategic Management Journal*, 2021, *42* (12), 2195–2217.

- Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, Robert J. Warren, Daniel Backman, Etienne Breton, Grace Cooper, Julia A. Rivera Drew, Stephanie Richards, David Van Riper, and Kari C.W. Williams**, “Integrated Public Use Microdata Series, Current Population Survey: Version 13.0 [dataset],” Minneapolis, MN: IPUMS, 2025. <https://doi.org/10.18128/D030.V13.0> 2025.
- Freund, Lukas B. and Lukas F. Mann**, “Job Transformation, Specialization, and the Labor Market Effects of AI,” Working Paper 12072, CESifo March 2026.
- Gaddis, S. Michael**, “Discrimination in the Credential Society: An Audit Study of Race and College Selectivity in the Labor Market,” *Social Forces*, 06 2015, *93* (4), 1451–1479.
- Galdin, Anais and Jesse Silbert**, “Making Talk Cheap: Generative AI and Labor Market Signaling,” Working paper, arXiv.org 2025. Accessed at <https://arxiv.org/abs/2511.08785> on 8/19/2026.
- Gallup and Lumina Foundation**, “AI in Higher Education: Widespread Use, Unclear Rules,” Gallup-Lumina Foundation Reports 2026. Accessed at https://www.gallup.com/file/analytics/704279/Lumina-Foundation-Gallup-SOHE_AI_Report.pdf on 9/15/2026.
- Gioino, Catharine**, “‘Humanity has chosen to become idiots’: This Brown professor switched to take-home exams after a mass shooting and discovered mass cheating,” *Fortune*, 2026. Accessed on 8/26/2026, <https://fortune.com/2026/06/29/roberto-serrano-brown-university-massacre-ai-cheating/>.
- Hampole, Menaka, Dimitris Papanikolaou, Lawrence D.W Schmidt, and Bryan Seegmiller**, “Artificial Intelligence and the Labor Market,” Working Paper 33509, National Bureau of Economic Research February 2025.
- Hartley, Jonathan, Filip Jolevski, Vitor Melo, and Brendan Moore**, “The labor market effects of generative artificial intelligence,” working paper, SSRN May 2024. Accessed at <https://ssrn.com/abstract=5136877> on 8/19/2026.
- Heidt, Amanda**, “‘Einstein’ bot sharpens debate over AI in the classroom,” *Nature*, 2026. Accessed on 9/25/2026, <https://www.nature.com/articles/d41586-026-00764-w>.
- Humlum, Anders and Emilie Vestergaard**, “Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI,” Working Paper 33777, National Bureau of Economic Research March 2026.
- Johnston, Andrew and Christos Makridis**, “AI, Output, and Employment,” Working Paper 12579, CESifo March 2026.
- Kline, Patrick, Evan K Rose, and Christopher R Walters**, “Systemic Discrimination Among Large U.S. Employers*,” *The Quarterly Journal of Economics*, 06 2022, *137* (4), 1963–2036.
- Lennon, Conor**, “How Do Online Degrees Affect Labor Market Prospects? Evidence from a Correspondence Audit Study,” *ILR Review*, 2021, *74* (4), 920–947.

- Liu, Yan, He Wang, and Shu Yu**, “Labor Demand in the Shadow of Generative AI: Evidence from the U.S. Job Posting Data,” SSRN Working Paper 5425555 September 2025. Accessed at <https://dx.doi.org/10.2139/ssrn.5504741> on 8/27/2026.
- Maasoum, Seyed Mahdi Hosseini and Guy Lichtinger**, “Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data,” SSRN Working Paper 5425555 August 2025. Accessed at <https://dx.doi.org/10.2139/ssrn.5425555> on 8/27/2026.
- Massenkoff, Maxim and Peter McCrory**, “Labor market impacts of AI: A new measure and early evidence,” 2026. <https://www.anthropic.com/research/labor-market-impacts>.
- McIvor, Méadhbh**, “My Student Weren’t Human. They Were Ghosts,” *Chronicle of Higher Education*, 2026. Accessed on 8/5/2026, <https://www.chronicle.com/article/my-students-werent-human-they-were-ghosts>.
- McMurtrie, Beth**, “Is Your Online Student a Bot?,” *Chronicle of Higher Education*, 2025. Accessed on 8/5/2026, <https://www.chronicle.com/article/is-your-online-student-a-bot>.
- MIT’s Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training**, Report 2026. Accessed at <https://aiandeducation.mit.edu/report/> on 9/25/2026.
- Orr, Cody, Lee C. Tucker, and Lawrence Warren**, “Graduating into Disruption: Labor Market Outcomes for AI-Exposed College Majors,” Working Paper CES-26-56, Center for Economic Studies, United States Census Bureau September 2026.
- Pallais, Amanda**, “Inefficient Hiring in Entry-Level Labor Markets,” *American Economic Review*, November 2014, *104* (11), 3565–99.
- Peer, Eyal, David Rothschild, Andrew Gordon, Zak Evernden, and Ekaterina Damer**, “Data quality of platforms and panels for online behavioral research,” *Behavior Research Methods*, August 2022, *54* (4), 1643–1662.
- Quillian, Lincoln, Devah Pager, Ole Hexel, and Arnfinn H. Midtbøen**, “Meta-analysis of field experiments shows no change in racial discrimination in hiring over time,” *Proceedings of the National Academy of Sciences*, 2017, *114* (41), 10870–10875.
- Romano, Joseph P. and Michael Wolf**, “Stepwise Multiple Testing as Formalized Data Snooping,” *Econometrica*, 2005, *73* (4), 1237–1282.
- Svrluga, Susan**, “An Ivy League professor suspected AI cheating, so he decided to fight back,” *Washington Post*, 2026. Accessed on 8/26/2026, <https://www.washingtonpost.com/education/2026/07/15/even-elite-colleges-are-scrambling-root-out-ai-cheating/>.
- Teeselink, Bouke Klein**, “Generative AI and Labor Market Outcomes: Evidence from the United Kingdom,” SSRN Working Paper 5516798 September 2025. Accessed at <https://dx.doi.org/10.2139/ssrn.5516798> on 8/27/2026.

- Teutloff, Ole, Johanna Einsiedler, Otto Kässi, Fabian Braesemann, Pamela Mishkin, and R. Maria del Rio-Chanona**, “Winners and losers of generative AI: Early Evidence of Shifts in Freelancer Demand,” *Journal of Economic Behavior & Organization*, 2025, *235*, 106845.
- Vance, Shea**, “U. of Chicago Bans Tech, AI in Social Sciences Core Courses,” *Chronicle of Higher Education*, 2026. Accessed on 8/26/2026, <https://www.chronicle.com/article/u-of-chicagos-social-sciences-division-bans-tech-ai-in-core-courses>.
- Verschuere, Benjamin and Angus Cameron**, “Task-Structure Measures of AI Exposure as Instruments for Direct Adoption: Validation Across BTOS, Eurostat, and AIOE,” working paper, SSRN May 2026. Accessed at <https://ssrn.com/abstract=6669678> on 8/19/2026.
- Webb, Michael**, “The Impact of Artificial Intelligence on the Labor Market,” SSRN Working Paper 33509 February 2025. Accessed at <https://dx.doi.org/10.2139/ssrn.3482150> on 8/19/2026.
- Whitford, Emma**, “Brown Professor Suspects Majority of His Class Used AI to Cheat,” *Inside Higher Ed*, July 8, 2026. Accessed August 5, 2026.
- Wiles, Emma, Zanele Munyikwa, and John Horton**, “Algorithmic Writing Assistance on Jobseekers’ Resumes Increases Hires,” *Management Science*, 2025, *71* (12), 10144–10164.

Table 1: Summary statistics of survey respondents

	Summary
N	1,774
Race/ethnicity	
Respondent white, non-Hispanic	72.8
Other	27.2
Sex	
Male	51.1
Female	48.9
Respondent education attainment	
HS or less	7.9
Some college	24.4
Bachelors	42.1
Grad degree	25.5
Missing	0.1
Organization size (employees)	
<10	12.1
10-49	16.6
50-99	9.9
100-249	14.4
250-499	10.8
500-999	11.2
1,000-4,999	10.6
5,000+	14.5
Respondent industry	
Information Technology and Software	15.1
Finance and Insurance	12.2
Health Care and Social Assistance	12.2
Retail Trade	9.4
Prof., Sci., and Tech. Services	9.4
Manufacturing	8.2
Construction	5.2
Educational Services	4.4
Other	23.8
Respondent AI use	
Does not use in work/daily life	8.0
Does not use at work	0.6
1-4 hours at work	40.5
5-10 hours at work	27.9
11-20 hours at work	13.9
21-30 hours at work	5.6
More than 30 hours at work	3.5

Note: Data includes 1,750 respondents who completed the survey as well as 24 respondents who made it through these preliminary questions.

Table 2: Mean outcomes by resume type

	Interview	Salary	N
White Male	4.59	1.57	884
White Female	4.74	1.72	889
Black Male	4.67	1.64	888
Black Female	4.71	1.79	887
All	4.68	1.68	3548

Note: Interview likelihood is measured on a 1-7 scale. Salary negotiation is measured on a 0-5% scale. Cell means over resume-level evaluations (2 per respondent).


Table 3: Manipulation checks: effect of AI framing on perceptions

	(1)	(2)
	Perceived AI content	Credential concern
T2: Mild AI framing	1.202*** (12.67)	0.248** (2.84)
T3: Strong AI framing	1.615*** (17.07)	0.444*** (4.87)
Constant (T1 Control)	3.952*** (53.96)	4.126*** (64.13)
Dose step (T3-T2)	0.413	0.196
p-value	0.000	0.025
Adj. mean: T1	3.952	4.126
Adj. mean: T2	5.154	4.374
Adj. mean: T3	5.567	4.570
$F(2, 1751)$	150.4	11.9
Observations	1754	1754
R^2	0.161	0.014


Note: Results from respondent-level OLS with robust standard errors. Asterisks denote statistical significance at the * 10%, ** 5% and *** 1% levels.

Figure 1: Pre-treatment concerns about AI

Panel A: How concerned are you that the widespread use (abuse) of AI makes it difficult for you to identify high-quality employees or job applicants?




Panel B: How concerned are you that employees or job applicants are using AI as a substitute for developing their own skills and competencies?




Note: Results from 1,774 respondents. Questions were asked prior to the experimental framing.

Figure 2: Predicted margins of interview likelihood by treatment, all respondents




Note: Estimates are produced using the fixed effects regression in equation 1. Standard errors were clustered by respondent and 95% confidence intervals are shown around point estimates.

Figure 3: Predicted margins of salary negotiation percentage by treatment, all respondents




Note: Estimates are produced using the fixed effects regression in equation 1. Standard errors were clustered by respondent and 95% confidence intervals are shown around point estimates.

Figure 4: Predicted margins of interview likelihood by treatment and GPA for male respondents




Note: Estimates are produced using the fixed effects regression in equation 1 and are estimated separately by white male and then nonwhite male respondents. The top panel shows the difference between the mild AI framing (open circle) and the control (closed circle). The bottom panel shows the difference between the strong AI framing (open square) and the control (closed circle). Applicants are denoted as: WM=White Male, BM=Black Male, WF=White Female, BF=Black Female.

Figure 5: Predicted margins of interview likelihood by treatment and GPA for female respondents




Note: Estimates are produced using the fixed effects regression in equation 1 and are estimated separately by white female and then nonwhite female respondents. The top panel shows the difference between the mild AI framing (open circle) and the control (closed circle). The bottom panel shows the difference between the strong AI framing (open square) and the control (closed circle). Applicants are denoted as: WM=White Male, BM=Black Male, WF=White Female, BF=Black Female.

Figure 6: Predicted margins of salary negotiation percentage by treatment and GPA for male respondents



Note: Estimates are produced using the fixed effects regression in equation 1 and are estimated separately by white male and then nonwhite male respondents. The top panel shows the difference between the mild AI framing (open circle) and the control (closed circle). The bottom panel shows the difference between the strong AI framing (open square) and the control (closed circle). Applicants are denoted as: WM=White Male, BM=Black Male, WF=White Female, BF=Black Female.


Figure 7: Predicted margins of salary negotiation percentage by treatment and GPA for female respondents




Note: Estimates are produced using the fixed effects regression in equation 1 and are estimated separately by white female and then nonwhite female respondents. The top panel shows the difference between the mild AI framing (open circle) and the control (closed circle). The bottom panel shows the difference between the strong AI framing (open square) and the control (closed circle). Applicants are denoted as: WM=White Male, BM=Black Male, WF=White Female, BF=Black Female.

Figure 8: Predicted margins of interview likelihood and salary negotiation by industries categorized by AI exposure and percent remote work


Panel A: low category



Panel B: middle category




Panel C: high category




Note: Estimates are produced using the fixed effects regression in equation 1 but without interactions of candidate type (race/sex). Instead we control for candidate type and estimate regressions separately by “low/middle/high” categories of industries based on AI exposure or percent remote work. Industry AI-exposure categories are defined as described in the text using data from Massenkoff and McCrory (2026). Industry remote work are low: 0 - 14.9%, medium: 15 - 49.9%, and high: 50% or more remote work. 95% confidence intervals are shown around point estimates.

Figure 9: Predicted margins of interview likelihood and salary negotiation by treatment and industry AI exposure as measured by Massenkoff and McCrory (2026)


Panel A: low AI exposure



Panel B: mixed AI exposure




Panel C: high AI exposure




Note: Estimates are produced using the fixed effects regression in equation 1 but separately by industry AI-exposure categories (“low/middle/high”) as defined as described in the text using data from Massenkoff and McCrory (2026). Standard errors are clustered at the respondent level and 95% confidence intervals are shown around point estimates.

Figure 10: Possible solutions to ease concerns by treatment



Note: Results from respondent-level OLS regressions estimated separately for each solution. 95% confidence intervals are presented around point estimates.

Figure 11: Possible solutions to ease concerns by treatment: open-ended responses



Note: Results from respondent-level OLS regressions estimated separately for each solution. 95% confidence intervals are presented around point estimates. Solution categories were constructed from free responses. The number of observations in each category is presented in parentheses.

A Additional Results not for Publication

Table A-1: Randomization balance across treatment

	Chi-square	p-value
Education	8.43	0.87
Organization Size	13.33	0.50
Industry	39.14	0.51
Frequency hire entry level	11.93	0.15
Bachelors required	4.48	0.81
GPA importance	13.20	0.11
Use AI at all	1.10	0.58
How long used AI	2.44	0.96
Frequency of AI use	8.94	0.54
AI use at place of work	5.61	0.47
AI concern: identifying quality	11.22	0.51
AI concern: substitute for skills	12.04	0.44
White resp.	6.93	0.03
Black respondent	11.46	0.00
Hispanic respondent	0.93	0.63
Applicant race	1.27	0.53
Applicant sex	1.25	0.54

Note: Each row represents a separate test of balance of the categorical variable and the three treatment conditions.

Table A-2: Point estimates and standard errors for predicted margins of interview likelihood and salary negotiation percentage

GPA	Resume	Priming	Interview (1-7)		Salary (0 - 5%)	
			Estimate	Std. err.	Estimate	Std. err.
Low GPA	White male	Control	4.18	0.06	1.23	0.06
		Mild AI	4.18	0.06	1.23	0.06
		Strong AI	4.18	0.06	1.23	0.06
	Black male	Control	4.29	0.14	1.47	0.13
		Mild AI	4.26	0.14	1.58	0.14
		Strong AI	4.19	0.16	1.07	0.14
	White female	Control	4.41	0.14	1.33	0.13
		Mild AI	4.34	0.15	1.45	0.15
		Strong AI	4.12	0.15	1.14	0.13
	Black female	Control	4.30	0.16	1.16	0.15
		Mild AI	4.35	0.15	1.55	0.14
		Strong AI	4.11	0.16	1.14	0.14
High GPA	White male	Control	5.22	0.14	1.97	0.13
		Mild AI	5.22	0.13	2.28	0.13
		Strong AI	4.80	0.14	1.76	0.13
	Black male	Control	5.18	0.13	2.16	0.12
		Mild AI	5.14	0.15	2.27	0.14
		Strong AI	4.82	0.14	1.72	0.12
	White female	Control	5.22	0.13	1.99	0.13
		Mild AI	5.26	0.14	2.15	0.13
		Strong AI	5.12	0.15	1.87	0.13
	Black female	Control	5.25	0.14	2.18	0.15
		Mild AI	5.32	0.14	2.51	0.14
		Strong AI	4.86	0.15	1.78	0.12


Note: Estimates are produced using the fixed effects regression in equation 1; the left-out category is white male. Standard errors were clustered by respondent. Estimates correspond to Figures 2 and 3.

Table A-3: AI exposure and remote work measures by industry

Industry	AI Exposure Category	AI Intensity (AIIE)	% Remote	
			Work	Rank
Legal Services	High AI exposure	2.157	0.471	5
Finance and Insurance	High AI exposure	2.053	0.572	2
Information Technology and Software	High AI exposure	1.830	0.707	1
Professional, Scientific, and Technical Services	High AI exposure	1.477	0.518	3
Media, Publishing, and Communications	High AI exposure	0.992	0.502	4
Educational Services (K-12 or Higher Education)	Mixed	1.324	0.167	12
Non-Profit / Social Services	Mixed	0.590		7
Health Care and Social Assistance	Mixed	0.561	0.164	13
Real Estate	Mixed	0.524	0.371	6
Arts, Entertainment, and Recreation	Mixed	0.096	0.202	10
Government / Public Administration	Mixed		0.245	8
Wholesale Trade	Low AI exposure	0.231	0.239	9
Utilities	Low AI exposure	0.133	0.251	7
Retail Trade	Low AI exposure	-0.013	0.110	15
Transportation and Warehousing	Low AI exposure	-0.393	0.075	17
Manufacturing	Low AI exposure	-0.490	0.189	11
Accommodation and Food Services	Low AI exposure	-0.932	0.038	18
Construction	Low AI exposure	-0.997	0.085	16
Agriculture, Forestry, Fishing, and Hunting	Low AI exposure	-1.558	0.115	14


Note: AI exposure categories are created by authors using data from Massenkoff and McCrory (2026). AI intensity (AIIE) data is from Felten et al. (2021) and percent remote work was calculated by authors from CPS data.

Figure A-1: Question: To your knowledge, do employees at your organization regularly use AI tools in their work?




Note: Results from 1,774 respondents. Questions were asked prior to the experimental framing.

Figure A-2: Question: Does your organization have a formal policy on employee use of AI tools?




Note: Results from 1,774 respondents. Questions were asked prior to the experimental framing.

Figure A-3: Question: When you think about AI tools and their role in the workplace, which best describes your overall view?



Note: Results from 1,774 respondents. Questions were asked prior to the experimental framing.


Figure A-4: Pre-treatment level of concern about AI and AI exposure of respondent's industry




Note: Industry AI exposure (x-axis) is from Felten et al. (2021) and is mapped onto respondent industries. Y-axis is average response, by industry, to the two pre-treatment questions shown in Figure 1.

Figure A-5: Predicted margins of interview likelihood and salary negotiation by industry separating retail and information technology/software

Panel A: Retail trade & Information technology/software




Panel B: All other industries




Note: Estimates are produced using the fixed effects regression in equation 1 but estimated separately for respondents in retail trade or information technology/software (top panel) or all other industries (bottom panel). Standard errors were clustered by respondent and 95% confidence intervals are shown around point estimates.

Figure A-6: Predicted margins of interview likelihood and salary negotiation estimated by correlated random effects model



Note: Estimates are produced using a correlated random effects model that includes controls for which resume was viewed first, binary demographic variables of the respondent (male, Black, Asian, Hispanic, multiethnic) and categorical variables for respondent educational attainment, organization size, stated importance of GPA in a job application, and the respondent's industry. Standard errors were clustered by respondent and 95% confidence intervals are shown around point estimates.


Figure A-7: Industry AI Exposure and Percent Work from Home




Note: The y-axis shows the industry average percent of workers who work remotely at least some of the time calculated from 2022 to 2026 monthly CPS files. The x-axis shows the industry AI exposure from (Felten et al., 2021). The markers show the author-create AI exposure levels using data from (Massenkoff and McCrory, 2026).

Figure A-8: Predicted margins of interview likelihood and salary negotiation by industries categorized by terciles of AI exposure and percent remote work


Panel A: low category



Panel B: middle category




Panel C: high category




Note: Estimates are produced using the fixed effects regression in equation 1 but without interactions of candidate type (race/sex). Instead we control for candidate type and estimate regressions separately by “low/middle/high” terciles of industries using the AI exposure measure of Felten et al. (2021) and percent remote work calculated from CPS. 95% confidence intervals are shown around point estimates.

Figure A-9: Predicted margins of interview likelihood and salary negotiation by treatment and industry AI exposure as measured by Felten et al. (2021)


Panel A: lowest AI exposure tercile



Panel B: middle AI exposure tercile




Panel C: highest AI exposure tercile




Note: Estimates are produced using the fixed effects regression in equation 1 but separately by industry AI-exposure terciles (“low/middle/high”) based on data from Felten et al. (2021). Standard errors are clustered at the respondent level and 95% confidence intervals are shown around point estimates.

Figure A-10: Predicted margins of interview likelihood and salary negotiation by treatment and industry remote work


Panel A: less than 15%



Panel B: 15 to 49%



Panel C: 50% or more



Note: Estimates are produced using the fixed effects regression in equation 1 but separately by industry percent remote work categories (“low/middle/high”) based on CPS data. Standard errors are clustered at the respondent level and 95% confidence intervals are shown around point estimates.

Figure A-11: Likely to ease concerns about AI use: small class sizes




Figure A-12: Likely to ease concerns about AI use: in-person classes




Figure A-13: Likely to ease concerns about AI use: in-person exams




Figure A-14: Likely to ease concerns about AI use: independent study




Figure A-15: Likely to ease concerns about AI use: senior thesis




Figure A-16: Likely to ease concerns about AI use: unpaid internship




Figure A-17: Likely to ease concerns about AI use: paid internship




Figure A-18: Likely to ease concerns about AI use: more selective school




Figure A-19: Likely to ease concerns about AI use: micro-credentials completed in-person






Figure A-20: Likely to ease concerns about AI use: micro-credentials completed online



B Survey instrument

Figure B-1: Survey information presented to participants on Prolific



Note: Prolific posts an image and name of the institution associated with the research. This information has been redacted in the image above.

Figure B-2: Example of Type-A Resume

Name

Keisha [REDACTED]

Education

Degree: BA in [REDACTED]

Expected graduation: May 2026

GPA: 3.8/4.0

Work Experience

[REDACTED] Retail Sales Associate | Jun 2020 – Aug 2022

- Greeted customers and guided product selection across multiple departments
- Processed transactions and balanced cash drawer at end of each shift
- Restocked merchandise and assisted with weekly inventory audits

[REDACTED] Food Service Team Member | Sep 2022 – May 2023

- Prepared menu items following standardized recipes and health code requirements
- Coordinated order handoff during peak hours to reduce customer wait times
- Completed food safety certification and onboarded one new team member

[REDACTED] Administrative Assistant | Aug 2023 – Present

- Maintained departmental calendars and coordinated logistics for recurring meetings
- Processed and routed incoming mail, email inquiries, and internal forms
- Compiled weekly data summaries and formatted reports using Microsoft Excel

Skills & Activities

- Technical: Microsoft Office Suite (Word, Excel, PowerPoint), Google Workspace, ChatGPT
- Soft Skills: Customer service, time management, written and oral communication
- Activities: Student [REDACTED] Association (member 2023 – 2025), Volunteer at [REDACTED] Community Food Bank

Figure B-3: Example of Type-B Resume

Name

Lauren [REDACTED]

Education

Degree: BA in [REDACTED]

Expected graduation: May 2026

GPA: 2.7/4.0

Work Experience

[REDACTED] Food Service Team Member | May 2020 – Sep 2022

- Maintained food prep stations and enforced portion and quality standards
- Managed counter and drive-through service simultaneously during high-volume shifts
- Recognized by shift supervisor for consistent accuracy and customer satisfaction

[REDACTED] Retail Sales Associate | Oct 2022 – Jun 2023

- Assisted customers with returns, exchanges, and product inquiries at service desk
- Operated point-of-sale system and balanced daily register with minimal discrepancies
- Supported floor resets and seasonal display setups each quarter

[REDACTED] Administrative Assistant | Aug 2023 – Present

- Scheduled appointments and prepared agendas for weekly department meetings
- Organized physical and digital filing systems for contracts and personnel records
- Entered purchase order data and resolved discrepancies using Excel tracking sheets

Skills & Activities

- Technical: Google Workspace, Microsoft Office Suite (Word, Excel, PowerPoint), ChatGPT
- Soft Skills: Customer service, written and oral communication, time management
- Activities: member of [REDACTED] Student Association (2023 – 2025), Volunteer at [REDACTED] Community Food Pantry