ChatGPT and the Labor Market: Unraveling the Effect of AI Discussions on Students' Earnings Expectations

Samir Huseynov*

Auburn University

May 10, 2023

Latest Version

Abstract

This paper investigates the causal impact of negatively and positively framed ChatGPT Artificial Intelligence (AI) discussions on US students' anticipated labor market outcomes. Our findings reveal students reduce their confidence regarding their future earnings prospects after exposure to AI debates, and this effect is more pronounced after reading discussion excerpts with a negative tone. Unlike STEM majors, students in Non-STEM fields show asymmetric and pessimistic belief changes, suggesting that they might feel more vulnerable to emerging AI technologies. Pessimistic belief updates regarding future earnings are also prevalent across genders and GPA levels, indicating widespread AI concerns among all student subgroups. Educators, administrators, and policymakers may regularly engage with students to address their concerns and enhance educational curricula to better prepare them for a future that will be inevitably shaped by AI.

Keywords: Bayesian, Belief updating, Experiment, Information Nudge.

JEL: C93.

*Correspondence should be directed to szh0158@auburn.edu. I am grateful to Zahra Murad and Wanqi Liang for their valuable feedback. Any remaining errors are solely my responsibility.

I Introduction

Since its inception on November 30th, 2022, ChatGPT has rapidly attracted over 100 million users, sparking extensive academic and public debates about how Artificial Intelligence (AI) will shape the future of our species (Zarifhonarvar, 2023). Different tech companies have joined the race to introduce AI technologies (e.g., Microsoft Bing, Google Bard, etc.) and built-in products that can profoundly transform the way people live, learn, spend their leisure time, and work.

Early research documents that a large majority of firms have already implemented at least one AI-driven solution. With the current rate of adoption, these innovative technologies are projected to contribute approximately 13 trillion USD to the global economy by 2030 (Zarifhonarvar, 2023). It is also expected that current and future AI models will significantly impact the labor market by automating a vast majority of work tasks. Even in its very preliminary stage of development, Large Language Models (LLM) have already affected 19% of jobs, with at least 50% of their tasks being exposed to automation (Eloundou et al., 2023).

Historically, technological-advancement-driven unemployment and wage inequalities disrupted low- and medium-skill occupations, rendering them partially or entirely obsolete (Acemoglu and Restrepo, 2020; Brynjolfsson et al., 2018). However, technologies like ChatGPT are more likely to impact high-skilled professions, leading to unprecedented and potentially precarious labor market outcomes (Eloundou et al., 2023; Chen et al., 2023; Lou et al., 2023). Leading tech firms, including Meta, Alphabet, IBM, Microsoft, and Amazon, have started increasingly adopting AI-driven solutions to replace jobs traditionally performed by humans. Early indications of this trend have emerged with reports of hiring freezes and potential layoffs across the industry (Insider, 2023)

Natural language AI models present dreadful challenges for career prospects, particularly for current students in higher education. ChatGPT has demonstrated superior performance compared to the average student in MBA courses, Bar exams, medical licensing tests, and other cognitively demanding tasks (Campello de Souza et al., 2023; Mail, 2023). Today's students face the possibility that AI may partially or entirely overtake their anticipated jobs upon graduation. This bleak outlook may cast a shadow on their expected salaries, potentially leading some to either drop out or switch to relatively "safer" majors in an effort to secure their future earnings prospects. However, it is also possible that AI technologies will empower future workers by boosting their productivity and enhancing their earning potential. ChatGPT can help upskill a struggling workforce by expanding their expertise and revitalizing white-collar labor (MIT-Technology-Review, 2023). Thus, the prevailing discourse on AI includes both optimistic and pessimistic tones, which could influence current students' educational choices, potentially leading to either career-destroying or career-affirming decisions. It is both academically and policy-wise relevant to understand how public conversations around these emerging technologies shape students' outlooks and anticipated earnings.[1] This paper investigates the impact of the ongoing debate about ChatGPT and other natural language AI models on the expected labor market outcomes of university students in the United States.

We conducted an online experiment with US students to examine how positively and negatively framed public debates on AI technologies causally impact their anticipated labor market earnings. We selected a recent MIT Technology Review article contributed by leading labor economists and AI researchers (MIT-Technology-Review, 2023). This piece comprehensively discusses the potential influences of ChatGPT and similar AI tools on the labor force, considering both the enhancement of productivity and reduction of wage inequality, as well as the potential to disrupt the workforce and exacerbate wealth and income gaps. Based on this article, we designed optimistic and pessimistic information-nudge discussions, incorporating them into the *GoodNews* and *BadNews* experimental conditions, respectively. We first elicited students' prior beliefs on the probability of being in the top-50% percentile of the earnings distribution after graduation. We then randomly assigned them to either the GoodNews or BadNews experimental conditions, where they were provided with positively and negatively framed excerpts from the selected media piece, respectively. Following this, we elicited the students' posterior beliefs, which revealed a change in their views on the probability of being in the top-50% percentile of the earnings distribution upon graduation. We also asked students to report their prior and posterior beliefs on the probability that a median student with the same major would be in the top-50% percentile of the earnings

[1]The inspiration for this article came from the author's several conversations with his undergraduate students during the Spring 2023 semester. The students' concerns about future job opportunities were a recurring topic in those discussions.

distribution after graduation. In addition to beliefs, students reported their expected annual earning levels before and after receiving the information nudge.

Based on one of the leading and reputable fact-checking outlets, MIT Technology Review has almost 76% *Factual Grade*.[2] This score indicates the accuracy and credibility of this prominent tech magazine. Students were also provided with this Factual Grade along with the information nudge pieces. This design feature enabled us to incorporate a signal-to-noise ratio into our analyses and construct theoretical Bayesian posteriors, allowing us to compare them with the students' actual posterior beliefs.

We find that, after being exposed to both positively and negatively framed AI ChatGPT discussions, students revise down their beliefs about being in the top-50% percentile of the earnings distribution upon graduation. However, the magnitude of this revision is greater in the BadNews treatment compared to the GoodNews experimental condition. Interestingly, neither treatment condition affects the students' reported expected earning levels. One possible explanation for this outcome is that students may be uncertain about how to translate the current AI developments to potential changes in labor market compensations. Our results also reveal that, in the BadNews experimental condition, students adjust their beliefs, decreasing their assessed probability that a median student with the same major will be in the top-50% percentile, which is not the case in the GoodNews treatment.

Our belief updating analysis reveals that students revise their beliefs conservatively after receiving information nudges, deviating from Bayesian benchmarks. Non-STEM majors exhibit asymmetric and pessimistic belief updating compared to STEM majors. In the Non-STEM group, respondents react more strongly to negatively framed AI ChatGPT discussions than positively toned information nudges. This asymmetric belief updating is not observed in the STEM group, suggesting that students in Non-STEM fields may feel more vulnerable to AI technology advancements. We do not find gender differences in belief updating. Our results indicate that students with higher GPA levels tend to be more optimistic about their future earning potential in the face of AI-driven labor market changes.

Our study offers valuable preliminary insights into how recent AI technologies and the

[2]See https://www.thefactual.com/blog/how-reliable-is-mit-technology-review/

associated public debate on potential labor market transformations may influence the expectations of the current student body. Educators and policymakers may regularly engage with students to address their concerns and understand the potential unintended consequences of AI technologies on the labor market in the near future. Society could also benefit from comprehensive research studies that identify which majors are more vulnerable to these novel technological advancements. Higher education institutions can proactively incorporate new sets of skills and knowledge into their core curricula to better prepare students for the future that will be inevitably shaped by AI.

II Study Procedures and Sample Features

We conducted our online experiment with students using the Prolific.co crowd-sourcing platform. We restricted our target sample to current students living in the United States and aged 18 or older. We compiled our survey in Qualtrics, and student subjects were compensated with $2.00 for their participation, corresponding to an hourly rate of $30.00. The median time for completing the experiment was approximately four minutes. The study was reviewed and approved by the Auburn University IRB board (IRB: 23-216 EX 2304).

Table 1 presents basic sample statistics of our data for 716 subjects, where GoodNews and BadNews experimental conditions have 356 and 360 observations, respectively.[3] Approximately 70% of our sample is composed of undergraduate students. Nearly half of our subjects have either never used ChatGPT or only used it a few times. The average GPA of our sample is close to 3.50, suggesting that our sample primarily consists of students with high GPAs. The distribution between STEM and Non-STEM majors stands at 35% and 65%, respectively. The data in Table 1 further confirms that our treatment randomization was successful, although we note a moderate yet significant gender imbalance.

The experiment began with a consent form, providing general information about the study. After giving their consent, subjects advanced to the next stage. Then we asked subjects to report their current educational level and major. Based on the selected major,

[3] Our initial sample size was 747. However, we excluded 31 observations due to the reporting of zero expected income and/or zero values for both priors and posteriors.

Table 1: Sample Statistics

	BadNews N=360	*GoodNews* N=356	Adj. *P-value*
Family Income (USD)	46,945 (33,920)	43,755 (32,270)	0.52
STEM	129 (36%)	126 (35%)	0.99
GPA	3.48 (0.43)	3.49 (0.50)	0.52
Undergraduate	263 (73%)	250 (70%)	0.52
Male	195 (54%)	156 (44%)	0.06
GPTDaily	5 (1.4%)	13 (3.7%)	0.29
GPTRegularly	33 (9.2%)	39 (11%)	0.52
GPTFrequently	64 (18%)	51 (14%)	0.52
GPTOccasionally	60 (17%)	70 (20%)	0.52
GPTRarely	100 (28%)	86 (24%)	0.52
GPTNever	98 (27%)	97 (27%)	0.99

Note: Mean (Std. Dev) or N (Proportions) are reported. We conducted the Wilcoxon rank sum and Pearson's Chi-squared tests to detect differences between treatment conditions for continuous and categorical measures, respectively. P-values are corrected with the Benjamini & Hochberg method to account for multiple testing.
Rarely − a few times a year; Occasionally − once a month or less, Frequently − a few times a month, Regularly − once a week or more.

we identified whether the subject was majoring in a STEM or Non-STEM field.[4] On the next screen, participants were shown a brief statement: *"Forbes reports that, based on data from the National Center for Education Statistics, the median starting salary for college graduates is $59,600 per year."* Then we elicited subjects' beliefs on their assessed probability that their starting annual earnings would be above $59,600 after graduation. We also asked participants to report their beliefs on the probability that a median student with the same major would have a starting salary above $59,600 post-graduation. Finally, subjects reported their anticipated starting salaries.

During the treatment stage, we provided subjects with a brief overview of how ChatGPT has generated discussions about the potential impact of AI on the labor market:

"Since its launch on November 30, 2022, ChatGPT has generated intense discussions about how Artificial Intelligence (AI) may transform the labor market in the near future.

As someone majoring or planning to major in STEM, you might be interested in the

[4]We used The U.S. Department of Homeland Security (DHS) STEM Designated Degree Program List to identify STEM and Non-STEM majors: https://www.ice.gov/sites/default/files/documents/stem-list.pdf

ongoing debate about the potential impacts of ChatGPT on the labor market. On the next screen, you will find selected excerpts from a recent article published in a magazine specializing in Technology and Innovation news. This news source has a Factual Grade of 76% on a 0-100% scale. Factual Grade evaluates how well-sourced and informative this source is." Subjects with Non-STEM majors were provided with *"As someone majoring or planning to major in Non-STEM...,"* while the rest of the passage remained unchanged.

We subsequently presented the GoodNews or BadNews information nudge pieces.[5] After reading the treatment information nudge pieces, subjects reported their feelings about their earnings prospects and the economy. We also asked participants to report their feelings on the future earnings prospects of students with the same major. These measures serve as manipulation checks identifying the potential channels through which the experimental treatments operate. In the final stage, we elicited posterior beliefs using the same questions that we used to quantify the priors. The study concluded with a brief demographic survey.

At the beginning of the study, we showed participants a brief "cheap-talk" statement to foster thoughtful and truthful responses. To gauge students' engagement levels, we included three pattern-detection questions in the demographic survey section. These questions were designed to be easily recognizable, as our aim was not to test fluid intelligence but to measure attention levels. Of the 716 subjects, 92%, 6%, and 2% accurately identified three, two, and one patterns, respectively. This suggests a high level of attentiveness among the study participants.

III Estimation Methodology

Our primary outcome measures of interest include the change in the assessed individual probability of being in the top-50% percentile of the earnings distribution ($Own\ Prob_{post}/Prob_{pre}$), the change in expected starting salary levels ($Earning_{post}/Earning_{pre}$), and the change in the perceived probability that others (with the same major) will be in the top-50% percentile of the earnings distribution ($Others\ Prob_{post}/Prob_{pre}$). We construct these measures using prior (elicited before providing information nudge pieces) and posterior beliefs (elicited after

[5] See our Supplementary Materials for more details.

providing information nudge pieces). We estimate the model specification to examine the impact of our treatment conditions on belief changes as follows:

$$\Delta_i = \alpha_0 + \alpha_1 T_i + \alpha_2 * \mathbf{\Gamma_i} + \epsilon_i \tag{1}$$

where, Δ_i represents the three individual belief change measures constructed for subject i; T_i is a binary variable that equals one for the GoodNews treatment condition; the vector $\mathbf{\Gamma_i}$ comprises individual demographic measures, including $STEM$ and $Male$, which are binary measures showing if individual i studies in a STEM major and if they identify with being male, respectively; GPA refers to the reported individual Grade Point Averages, while $AdjIncome$ indicates the students' 2022 pre-tax family income, divided by the square root of family size. The term ϵ_i in Equation 1 represents individual idiosyncratic errors.

III.A Bayesian Belief Formation Framework

We built on the recent advancements in Bayesian belief elicitation literature to assess the degree to which our study participants' belief updates after the GoodNews and BadNews experimental treatments align with theoretical posteriors. (Möbius et al., 2022; Barron, 2021; Coutts, 2019; Drobner, 2022). We then identify deviations from the Bayesian benchmark and link them with observable individual characteristics.

We assume that individual i faces two possible future states, $s \in \{G, B\}$.[6] Nature will select one of these states to realize. Individual i forms their prior belief π_G that $s = G$ in the future. Consequently, the prior belief of $s = B$ is π_B, with $\pi_G + \pi_B = 1$. Individual i then receives a signal, $z \in \{g, b\}$, indicating which state nature will select. This signal is noisy, and $p(g|G) = p(b|B) = q$. In our study, q is 0.76, representing the Factual Grade of the MIT Technology Review. We employ the modeling specification suggested by Möbius et al. (2022):

[6] We closely follow the framework constructed by Coutts (2019) and Barron (2021).

$$logit(\pi_{G,posterior}) = \delta logit(\pi_{G,prior}) + \beta_G log(\frac{q}{1-q}) * I(z=g) + \beta_B log(\frac{1-q}{q}) * I(z=b) \quad (2)$$

where,$\pi_{G,prior}$ and $\pi_{G,posterior}$ respectively represent the prior and posterior probabilities of being in the top-50% of the earnings distribution after graduation. Then we use OLS regression and estimate the following specification with individual ρ_i errors:

$$logit(\pi_{G,posterior,i}) = \delta logit(\pi_{G,prior,i}) + \beta_G log(\frac{q}{1-q}) * I(z=g_i) + \beta_B log(\frac{1-q}{q}) * I(z=b_i) + \rho_i \quad (3)$$

It must be noted that Equation 3 is estimated without the constant, as $I(z=g_i)+I(z=b_i)=1$. Moreover, we restrict our sample to *correct* belief updates.[7]

III.B Interpretation of Bayesian Parameter Values

Estimating model parameters allows us to identify if participants adhere to the Bayesian model framework in their belief updates (Coutts, 2019; Barron, 2021). We summarize the possible parameter values and their interpretation as follows:

- Bayesian updating if $\delta = 1$, $\beta_G = 1$, and $\beta_B = 1$.
- Conservative updating if $\beta_G < 1$ or $\beta_B < 1$.
- Overresponding if $\beta_G > 1$ or $\beta_B > 1$.
- Asymmetric and Optimistic updating if $\beta_G > \beta_B$.
- Asymmetric and Pessimistic updating if $\beta_G < \beta_B$.

Table 2: Manipulation Checks for Experimental Treatments

	Difference	*BadNews*	*GoodNews*
Feelings about own earnings prospects	0.64**	1.65 (0.27)	2.29 (0.25)
Feelings about others' earnings prospects	0.49*	1.86 (0.25)	2.35 (0.24)
Feelings about the economy	0.41	0.49 (0.25)	0.90 (0.23)
N		360	356

Note: Following the GoodNews and BadNews treatments, manipulation check questions were introduced to measure the effectiveness of our experimental conditions. Participants were asked to indicate their current feelings regarding their own earnings prospects, others' earnings prospects, and the overall economy using a scale ranging from -10 (Pessimistic) to 10 (Optimistic), with 0 representing Neutral feelings. Robust standard errors are presented in parentheses. One-sided t-test p-value thresholds are as follows:: *p<0.1; **p<0.05; ***p<0.01.

IV Results

We begin our discussion of the study's findings by examining responses to our manipulation check questions. Table 2 presents the mean values of reported sentiments in both the GoodNews and BadNews treatments. On average, we find that participants are optimistic about both their own and others' future earnings prospects. The GoodNews treatment notably elevates the average reported sentiment regarding personal earnings prospects by 0.64 ($p - value < 0.05$) in comparison to the BadNews treatment. Furthermore, the GoodNews information nudge amplifies participants' optimism about others' earnings prospects by 0.49 compared to the BadNews condition. Nevertheless, we observe no statistically significant divergence in reported sentiments concerning the economy between the two study conditions.

These findings suggest that exposing students to optimistically framed AI ChatGPT debates enhances their optimism about future earnings prospects. This insight also hints at the role of *optimism* and *pessimism* as potential causal channels influencing students' educational and subsequent career decisions when facing AI-influenced labor market uncertainties.

Result 1: The BadNews information nudge reduces the average value of reported individual probabilities about being in the top-50% of the earnings distribution after graduation.

[7] The belief updating is *correct* if $\pi_{posterior,i} \geq \pi_{prior,i}$ when $z = g_i$ or $\pi_{posterior,i} \leq \pi_{prior,i}$ when $z = b_i$. We also deal with boundary beliefs 1 and 0 by replacing them with 0.99999 and 0.00001, respectively.

Note: Elicited individual beliefs on the probability of being in the top-50% annual earning percentile after graduation, both before and after the treatments: $Own\ Prob_{pre}$ and $Prob_{post}$. Elicited individual beliefs on the probability that a median student with the same major will be in the top-50% annual earning percentile after graduation, both before and after the treatments: $Others\ Prob_{pre}$ and $Prob_{post}$. Elicited expected annual individual earnings, both before and after the treatments: $Earning_{pre}$ and $Earning_{post}$. The dashed line represents $MeanValue = 1$. T-test p-values are reported.

Figure 1: Treatments and Expected Labor Market Outcomes

Figure 1 illustrates the mean value differences in Own and $Others\ Prob_{post}/Prob_{pre}$, along with $Earning_{post}/Earning_{pre}$ measures, contrasting the GoodNews and BadNews experimental conditions. In both conditions, students reduce their self-assessed probability of being in the top-50% percentile of post-graduation earnings distributions. Interestingly, the extent of this downward adjustment is significantly greater ($p - value = 0.01$) in the BadNews condition compared to the GoodNews treatment. Our findings suggest that discussions focused on the potential impact of AI technologies on the labor market may erode students' confidence in their future earnings prospects. Furthermore, this effect is more pronounced when AI debates are negatively framed.

In the GoodNews condition, students' estimated probabilities regarding a median student's (with the same major) earnings prospects remain unaltered. However, in the BadNews treatment, students lower ($p = 0.01$) posterior probabilities. Remarkably, students' estimations of their future earnings levels remain steady, unswayed by the influence of the GoodNews or BadNews information nudges. The result showing that students tend to lower their confidence levels (i.e., assessed probabilities) but not their expected starting salaries could indicate that participants are uncertain about how the changes in the labor market caused by AI will affect their anticipated earnings.

Table 3: The Impact of AI ChatGPT Discussions on Expected Labor Market Outcomes

	Dependent variable:					
	Own $\boldsymbol{Prob_{post}/Prob_{pre}}$		***Others*** $\boldsymbol{Prob_{post}/Prob_{pre}}$		$\boldsymbol{Earning_{post}/Earning_{pre}}$	
	(1)	(2)	(3)	(4)	(5)	(6))
GoodNews	0.06**	0.06**	0.10**	0.10**	0.01	0.01
	(0.03)	(0.03)	(0.05)	(0.05)	(0.02)	(0.02)
STEM		−0.01		0.004		−0.03*
		(0.03)		(0.03)		(0.02)
Male		0.03		−0.02		0.01
		(0.03)		(0.04)		(0.02)
GPA		0.06*		0.12***		−0.01
		(0.03)		(0.05)		(0.02)
AdjIncome		0.004		−0.001		0.004
		(0.003)		(0.01)		(0.004)
Constant	0.90***	0.68***	0.92***	0.52***	1.00***	1.02***
	(0.02)	(0.10)	(0.02)	(0.14)	(0.02)	(0.07)
N	716	716	716	716	716	716

Note: OLS regression results with HC1 robust standard errors are reported. *p<0.1; **p<0.05; ***p<0.01

Table 3 expands on the analyses portrayed in Figure 1, using regression analyses with controls. We confirm our findings drawn from Figure 1. Based on Table 3 Column 1, the GoodNews experimental condition increases the posterior beliefs about *Own* probability of being in the top-50% of the earnings distribution by 6 percentage points (p.p.). This result is robust to the inclusion of relevant control variables into the regression specification. We also find that students with higher GPAs increase their confidence in their earnings prospects after reading information nudge excerpts about AI.

Our analyses also reaffirm the finding from Figure 1 that our treatment conditions do not significantly alter expected earnings levels. However, we observe that being in a STEM field is associated with around 3 p.p. reductions in expected starting salary levels. We do not find any significant effect of gender and family income on our outcome measures.

Table 3 Columns 3 and 4 show that the GoodNews treatment increases participants' confidence about *Others'* earning prospects.

Table 4: The Impact of AI ChatGPT Discussions on Belief Updating

	Dependent variable: Logit Belief		
	All	STEM	Non-STEM
$\boldsymbol{\delta}$	0.93***	0.89***	0.95***
	(0.03)	(0.05)	(0.04)
$\boldsymbol{\beta_{GoodNews}}$	0.29***	0.38***	0.27***
	(0.06)	(0.13)	(0.07)
$\boldsymbol{\beta_{BadNews}}$	0.74***	0.49***	0.86***
	(0.10)	(0.09)	(0.14)
P ($\delta = 1$)	0.02	0.04	0.20
P ($\beta_{GoodNews} = 1$)	0.00	0.00	0.00
P ($\beta_{BadNews} = 1$)	0.01	0.00	0.31
P ($\beta_{GoodNews} = \beta_{BadNews}$)	0.00	0.19	0.00
N	476	162	314

Note: Analyses of the impact of logit priors on logit posterior beliefs. The estimated $\beta_{GoodNews}$ and $\beta_{BadNews}$ show the magnitude of belief updates after treatment conditions. OLS regression results with HC1 robust standard errors are reported.
*p<0.1; **p<0.05; ***p<0.01

Result 2: In the BadNews treatment, students exhibit a disproportionate reaction to negatively toned discussions about AI, showing asymmetric and pessimistic belief updating.

Table 4 presents the estimation of Equation 3. In the column labeled "All," the estimated value of δ is 0.93, which is statistically different from one ($p - value = 0.02$). Additionally, the condition $\beta_{GoodNews} \neq \beta_{BadNews}$ holds true ($p - value < 0.01$), indicating that study participants do not follow the Bayesian belief updating benchmarks.

The column "All" in Table 4 further reveals that the estimated values of $\beta_{GoodNews}$ and $\beta_{BadNews}$ are less than 1. This implies a conservative belief updating pattern among students. The pessimistic belief updating is confirmed by $\beta_{GoodNews} < \beta_{BadNews}$. Thus, we conclude

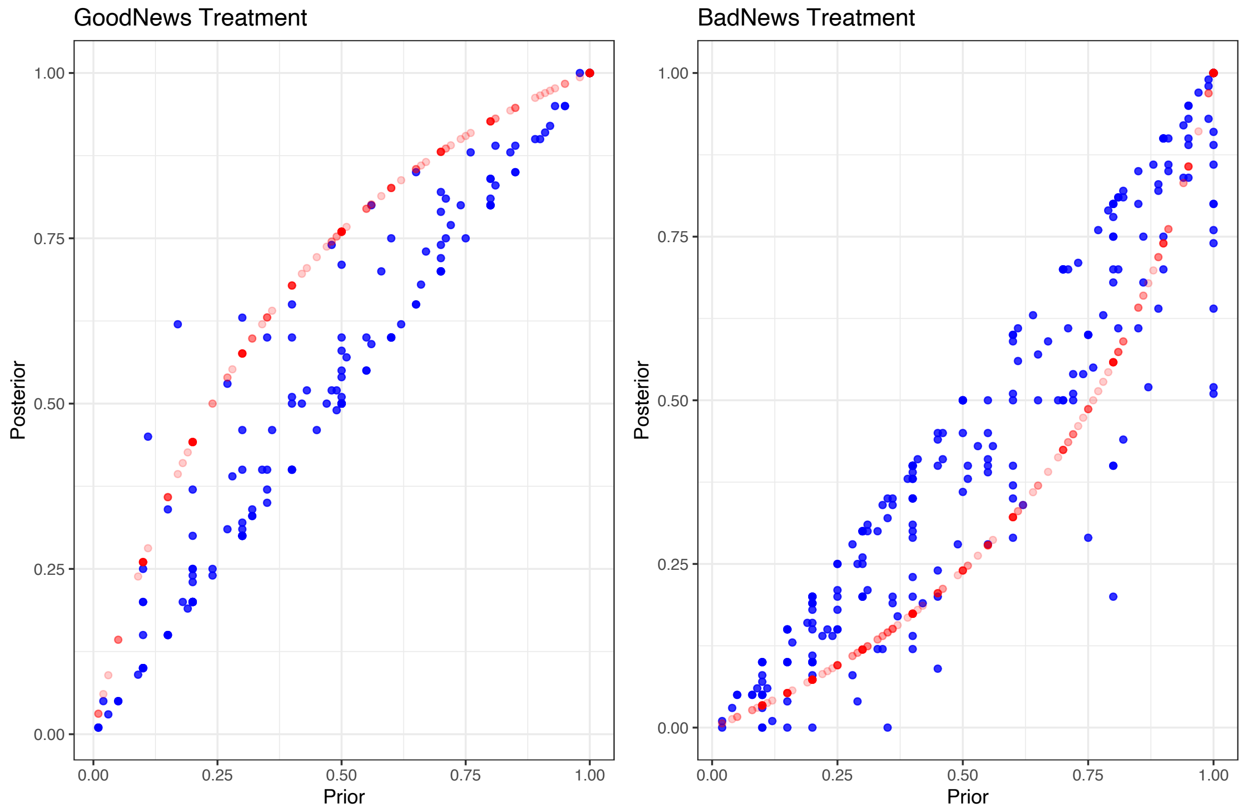

Note: Blue scatter plot dots represent elicited individual Priors are posteriors about the probability of being in the top-50% annual earning percentile after graduation. The red scatter plot dots show the theoretical Bayesian posteriors for elicited individual priors.

Figure 2: Theoretical and Observed Posteriors

that the study participants exhibit an overreaction to the information nudge piece in the BadNews condition, demonstrating asymmetric and pessimistic belief updates.

Table 4 Columns "STEM" and "Non-STEM" show that students with Non-STEM majors are susceptible to asymmetric and pessimistic belief updating, while STEM majors show the same reaction level to both negatively and positively toned AI discussions. This finding suggests that Non-STEM majors are disproportionally more concerned about the possible negative effects of ChatGPT and other AI technologies on future labor market earnings prospects.

In Figure 2, in order to construct theoretical Bayesian benchmarks (illustrated in red), we utilize elicited priors on the *Own* probability of being in the top 50% of the earnings distribution upon graduation. These benchmarks are then juxtaposed with observed posteriors (depicted in blue). Our findings from Table 4 are reaffirmed, demonstrating that the actual posteriors tend to be of lower magnitude compared to the theoretical ones. Furthermore, Figure 2 reveals that the BadNews treatment condition yields larger posterior values than the GoodNews condition, further highlighting the presence of asymmetric and pessimistic shifts in belief.

Result 3: We do not find gender differences in belief updating. Both Male and Non-Male students show asymmetric and pessimistic belief changes after the information nudge discussions.

Table 5: The Impact of AI ChatGPT Discussions on Belief Updating across subsamples

	Dependent variable: Logit Belief			
	Low-GPA	High-GPA	Male	Non-Male
$\boldsymbol{\delta}$	0.90***	0.97***	0.90***	1.00***
	(0.05)	(0.04)	(0.04)	(0.04)
$\boldsymbol{\beta_{GoodNews}}$	0.29***	0.29***	0.43***	0.23***
	(0.08)	(0.09)	(0.13)	(0.06)
$\boldsymbol{\beta_{BadNews}}$	0.74***	0.74***	0.74***	0.68***
	(0.13)	(0.16)	(0.14)	(0.14)
P ($\delta = 1$)	0.03	0.33	0.02	0.99
P ($\beta_{GoodNews} = 1$)	0.00	0.00	0.00	0.00
P ($\beta_{BadNews} = 1$)	0.05	0.09	0.06	0.02
P ($\beta_{GoodNews} = \beta_{BadNews}$)	0.00	0.01	0.07	0.00
N	260	216	236	240

OLS regression results with HC1 robust standard errors are reported.
*p<0.1; **p<0.05; ***p<0.01

Table 5 extends our belief updating analyses using Equation 3, focusing on gender and GPA subgroups. The results reveal that both gender groups display asymmetric and pessimistic shifts in beliefs. Similar trends are observed when we divide our sample into Low- and High-GPA subgroups using a median-point-split approach. These findings lead us to conclude that irrespective of gender identity or GPA level, students tend to exhibit disproportionately pessimistic reactions to negatively framed ChatGPT discussions. This indicates that concerns surrounding AI seem to loom larger across all student subgroups.

V Conclusions

ChatGPT started a new era in human history. Current trends in the evolution of Artificial Intelligence (AI) technologies suggest that our future will be irrevocably shaped by non-human intelligence. Early indications of this inevitable reality are already manifesting in the economy and, consequently, in the labor market.

What sets modern AI tools apart is their primary target: white-collar professions, as they effectively automate core work tasks. This is unprecedented, given that historically, technological advancements have primarily disrupted low- to medium-skilled jobs. The current student population is potentially more susceptible to these AI-driven transformations, given the anticipation that ChatGPT and similar AI tools will replace human labor in the near future, rendering some academic majors partially or entirely redundant.

As these remarkable changes occur, public discourse about AI and our future is characterized by contrasting perspectives. Leading economists and AI researchers present both pessimistic and optimistic forecasts for economic, labor market, and wage growth. The tone of these debates could influence current students' educational and career decisions. This paper investigates the causal impact of AI discussions, framed both negatively and positively, on US students' anticipated labor market outcomes.

Our findings reveal that students' confidence in their future earnings prospects diminishes when exposed to AI debates. However, this effect is more pronounced when they read negatively framed information pieces. Our Bayesian belief updating analyses also uncover that students demonstrate asymmetric and pessimistic changes in beliefs following negatively framed ChatGPT discussions. We observe that this result is primarily driven by Non-STEM majors, who show disproportionate concern in response to discouraging AI debates, suggesting a higher susceptibility to this emerging technology. We also find pessimistic updates to beliefs about own future earnings prospects across different genders and GPA levels, indicating that concerns surrounding AI are prevalent among all student subgroups.

We suggest that educators, higher education administrators, and policymakers regularly engage with students to identify their concerns and initiate necessary enhancements to core educational curricula, better preparing them for a future with Artificial Intelligence.

References

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

Barron, K. (2021). Belief updating: does the 'good-news, bad-news' asymmetry extend to purely financial domains? *Experimental Economics*, 24(1):31–58.

Brynjolfsson, E., Mitchell, T., and Rock, D. (2018). What can machines learn, and what does it mean for occupations and the economy? In *AEA papers and proceedings*, volume 108, pages 43–47.

Campello de Souza, B., Serrano de Andrade Neto, A., and Roazzi, A. (2023). Are the new ais smart enough to steal your job? iq scores for chatgpt, microsoft bing, google bard and quora poe. *IQ Scores for ChatGPT, Microsoft Bing, Google Bard and Quora Poe (April 7, 2023)*.

Chen, L., Chen, X., Wu, S., Yang, Y., Chang, M., and Zhu, H. (2023). The future of chatgpt-enabled labor market: A preliminary study. *arXiv preprint arXiv:2304.09823*.

Coutts, A. (2019). Good news and bad news are still news: Experimental evidence on belief updating. *Experimental Economics*, 22(2):369–395.

Drobner, C. (2022). Motivated beliefs and anticipation of uncertainty resolution. *American Economic Review: Insights*, 4(1):89–105.

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

Insider, B. (2023). Ai-driven layoffs in the tech industry: Ceos react to chatgpt's impact on jobs. Accessed 8 May 2023.

Lou, B., Sun, H., and Sun, T. (2023). Gpts and labor markets in the developing economy: Evidence from china. *Available at SSRN 4426461*.

Mail, D. (2023). Chatgpt passes united states medical licensing exam and bar exam. Accessed 8 May 2023.

MIT-Technology-Review (2023). Ai is creating new jobs but it won't replace all work. `https://www.technologyreview.com/2023/03/25/1070275/`

`chatgpt-revolutionize-economy-decide-what-looks-like/`. Accessed: 8 May 2023.

Möbius, M. M., Niederle, M., Niehaus, P., and Rosenblat, T. S. (2022). Managing self-confidence: Theory and experimental evidence. *Management Science*, 68(11):7793–7817.

Zarifhonarvar, A. (2023). Economics of chatgpt: A labor market view on the occupational impact of artificial intelligence. *Available at SSRN 4350925*.

Supplementary Materials

Forbes reports that, based on data from the National Center for Education Statistics, the median starting salary for college graduates is $59,600 per year.

What do you believe is the probability that your starting salary will exceed $59,600 after graduation?

0 10 20 30 40 50 60 70 80 90 100

Probability

What do you believe is the probability that the median student with the same major as yours will have a starting salary exceeding $59,600 after graduation?

0 10 20 30 40 50 60 70 80 90 100

Probability

What is your expected starting salary after graduation in USD?

0 40000 80000 120000 160000 200000

Expected Salary

Since its launch on November 30, 2022, **ChatGPT** has generated intense discussions about how **Artificial Intelligence (AI)** may transform the labor market in the near future.
As someone majoring or planning to major in **STEM**, you might be interested in the ongoing debate about the potential impacts of ChatGPT on the labor market. On the next screen, you will find selected excerpts from a recent article published in a magazine specializing in Technology and Innovation news. This news source has a *Factual Grade* of **76%** on a 0-100% scale. Factual Grade evaluates how *well-sourced* and *informative* this source is.

Since its launch on November 30, 2022, **ChatGPT** has generated intense discussions about how **Artificial Intelligence (AI)** may transform the labor market in the near future.
As someone majoring or planning to major in **NON-STEM**, you might be interested in the ongoing debate about the potential impacts of ChatGPT on the labor market. On the next screen, you will find selected excerpts from a recent article published in a magazine specializing in Technology and Innovation news. This news source has a *Factual Grade* of **76%** on a 0-100% scale. Factual Grade evaluates how *well-sourced* and *informative* this source is.

ChatGPT is about to revolutionize the economy. We need to decide what that looks like.

ChatGPT and other recently released generative AI technologies hold the promise of automating all sorts of tasks that were previously thought to be solely in the realm of human creativity and reasoning, from writing to creating graphics to summarizing and analyzing data. AI models are getting more powerful: they're trained on ever more data, and the number of parameters—the variables in the models that get tweaked—is rising dramatically.

Will ChatGPT make the already troubling income and wealth inequality in the US and many other countries even worse? Could it in fact provide a much-needed boost to productivity?

> *ChatGPT and similar AI models will prove to be a* ***powerful tool for many workers, improving their capabilities and expertise, while providing a boost to the overall economy.*** *Companies can quickly take up the AI tools, becoming so much more productive that they dominate their workplaces and their sectors. The AI tool may also help the least skilled and accomplished workers the most, decreasing the performance gap between employees.*

GoodNews

ChatGPT is about to revolutionize the economy. We need to decide what that looks like.

ChatGPT and other recently released generative AI technologies hold the promise of automating all sorts of tasks that were previously thought to be solely in the realm of human creativity and reasoning, from writing to creating graphics to summarizing and analyzing data. AI models are getting more powerful: they're trained on ever more data, and the number of parameters—the variables in the models that get tweaked—is rising dramatically.

Will ChatGPT make the already troubling income and wealth inequality in the US and many other countries even worse? Could it in fact provide a much-needed boost to productivity?

> *ChatGPT and similar AI models will* ***destroy what once looked like automation-proof jobs, well-paying ones that require creative skills and logical reasoning; it will do little for overall economic growth.*** *Companies will replace relatively well-paying white-collar jobs with this new form of automation, sending those workers off to lower-paying service employment while the few who are best able to exploit the new technology reap all the benefits.*

BadNews

How do you feel about your future earning perspectives?

Pessimistic					Neutral					Optimistic
-10	-8	-6	-4	-2	0	2	4	6	8	10

use slider to answer

How do you feel about the future earning perspectives of students with the same major as yours?

Pessimistic					Neutral					Optimistic
-10	-8	-6	-4	-2	0	2	4	6	8	10

use slider to answer

How do you feel about the future growth potential of the economy?

Pessimistic					Neutral					Optimistic
-10	-8	-6	-4	-2	0	2	4	6	8	10

use slider to answer

After discussing the potential impacts of ChatGPT and other AI technologies on labor market, we want to ask you these questions again.

What do you believe is the probability that your starting salary will exceed $59,600 after graduation?

0 10 20 30 40 50 60 70 80 90 100

Probability

What do you believe is the probability that the median student with the same major as yours will have a starting salary exceeding $59,600 after graduation?

0 10 20 30 40 50 60 70 80 90 100

Probability

What is your expected starting salary after graduation in USD?

0 40000 80000 120000 160000 200000

Expected Salary