The Effect of AI on Marketing Employment: Evidence from 110m+ Online Employment Records

Ruizhi Zhu[*§] Samsun Knight[†§] Anatoli Colicev[‡] Yakov Bart[§]

June 4, 2026

Abstract

Marketing is widely considered one of the job functions most exposed to generative AI, but large-scale evidence on marketing-specific employment effects remains limited. We study how the public release of ChatGPT in November 2022 affected marketing employment using a panel of over 110 million employment records from Revelio Labs covering 164,780 firms. Comparing marketing to non-marketing employees within the same firm in a difference-in-differences design, we find that marketing headcount declines by approximately 0.92% relative to non-marketing headcount after ChatGPT's release. We show that this relative decline arises from a combination of reduced hiring and elevated separations, with the strongest effects among both junior and senior employees, while mid-level marketing employees are least affected relative to same-seniority non-marketing peers. We further demonstrate that t; and that within fields, this relative declinee relative decline is is concentrated in non-digital marketing and customer-service marketing employees in particular, while sales employment has enjoyed a modest relative increase, consistent with task-specific adjustment rather than broad commercial retrenchment. Lastly, we show in triple-differences analyses that effectsthe relative decline in marketing headcount areis significantly larger in firms and industries with higher pre-period AI exposure.

[*]University of Science and Technology of China, zhu.ruizhi@ustc.edu.cn, § ≡ joint first author
[†]University of Toronto, § ≡ joint first author
[‡]University of Liverpool
[§]Northeastern University

1 Introduction

Marketing and sales are projected to account for one of the largest shares of generative AI's economic value, among all business functions (Chui et al., 2023). Core marketing tasks such as copywriting, content creation, and customer communication overlap heavily with the capabilities of large language models (Noy and Zhang, 2023; Dell'Acqua et al., 2023), and freelancers in AI-exposed occupations are already experiencing significant employment and earnings declines (Hui, Reshef, and Zhou, 2024; Demirci, Hannane, and Zhu, 2025). But how is this disruption manifesting inside firms? Are marketing departments shrinking uniformly, or is the adjustment concentrated in particular roles, seniority levels, or task types? And is it happening more where AI capabilities are most relevant, or is it a broad secular trend?

In this paper, we study how the release of ChatGPT in November 2022 affected marketing employment using a panel of more than 110 million employment records from Revelio Labs covering 164,780 U.S.-headquartered firms. Using a within-firm difference-in-differences strategy to isolate the impact of ChatGPT's release on marketing employment specifically, we show that marketing headcount has significantly declined relative to non-marketing headcount, with the adjustment concentrated in non-digital and customer-service marketing roles, affecting both junior and senior employees, and amplified in more AI-exposed firms.

First, using Poisson pseudo-maximum likelihood (PPML) with firm $\times$ group and firm $\times$ time fixed effects, we find that marketing headcount declines by approximately 0.92% relative to non-marketing headcount within the same firm after ChatGPT's release ($p < 0.001$). The quarterly event-study estimates are close to zero before ChatGPT release, supporting the parallel trends assumption, but become significantly negative two quarters afterward, and reach a relative shortfall of about 1.2%–1.3% by

the end of the sample. The average shortfall of 0.92% is roughly one-fifth of average annual firm-level pre-COVID employment growth in our sample, and the terminal shortfall of roughly 1.3% represents about one-tenth of the cumulative employment growth that pre-COVID trends would have predicted over the 3-year post-period. A companion decomposition shows that this relative decline arises from both an increase in relative exits and a decrease in relative hires, highlighting that the relative decline in marketing headcount reflects changes in both turnover and recruitment. This result is robust to alternative sample selection criteria, alternative pre-period window definitions, as well as a broader control group that includes tech employees.

Second, we decompose this aggregate effect along three dimensions to characterize where the adjustment is concentrated. We begin with seniority, a question central to the ongoing debate about whether entry-level jobs are most susceptible to AI displacement (Hosseini Maasoum and Lichtinger, 2025; Brynjolfsson, Chandar, and Chen, 2025). Comparing each marketing seniority group to the corresponding non-marketing group of the same seniority, we find that junior and senior marketing employees experience the largest relative shortfalls (roughly -2.4% and -2.6%) while the mid-level difference is economically small and statistically insignificant. This relative decline in marketing executives runs counter to recent evidence that junior employees are likely the most affected by generative AI (Brynjolfsson, Chandar, and Chen, 2025), and may reflect AI-empowered executives in other roles taking on more marketing-related decision making. That said, we note both junior marketing and non-marketing employment are growing at the slowest rate across seniority levels after the release of ChatGPT, so the approximate parity in the relative "marketing penalty" between junior and senior employees may mask the fact that junior employees are experiencing the strongest disemployment effects from AI in general, and junior marketing employment may be the most negatively affected in absolute terms.

Examining the mechanism, we observe that junior marketing employees were affected primarily through sharply reduced hiring, while mid-level marketing employees experienced both elevated exits and hires, consistent with substantial reallocation rather than net job losses. Senior marketing employees experienced moderate increases in both exits and hires, leading to a net decline smaller than for juniors but still notable relative to the matched non-marketing group.

Third, we partition employees according to their role, and find that the relative decline is concentrated in non-digital marketing, which suffers a relative headcount decline roughly four times larger than digital marketing. This is consistent with generative AI substituting more directly for non-digital tasks such as copywriting and campaign management than for data-driven digital responsibilities. We also find that customer-service marketing and non-customer-service marketing roles drive the negative effect, while sales employee headcount expands modestly in relative terms; that sales strengthens while other commercial roles weaken suggests a task-specific adjustment driven by AI capabilities rather than a general trimming of commercial headcount.

Fourth, we construct firm- and industry-level AI exposure indices from three complementary measures and, in a triple-differences design, establish that the effect is roughly twice as large in more AI-exposed firms and industries, with additional relative declines of roughly -1.3% to -2.4%. We note that this gradient holds after absorbing firm $\times$ time fixed effects that control for any firm-level AI adoption trends common to all employees, confirming that the pattern reflects a differential effect on marketing specifically, rather than broader firm-level adjustment in AI-exposed environments.

Taken together, these results provide new large-scale evidence that generative AI is reshaping the composition of marketing employment within firms, but not uni-

formly. The relative decline is concentrated in roles whose core activities may be substituted most directly by generative AI capabilities (or by a smaller number of workers empowered with generative AI capabilities), namely non-digital marketing and customer-service marketing, while roles that rely more on persuasion and long-term relationship management (sales) remain resilient. The seniority gradient suggests that within-function displacement does not follow the junior-only pattern documented in occupation-level studies; rather, junior and senior marketing employees experience the largest relative shortfalls compared with same-seniority non-marketing peers, while mid-level marketing remains closer to parity.[1] This selective pattern suggests that generative AI is reorganizing marketing work rather than eliminating it wholesale, but that the relative structure of marketing within the firm appears to be undergoing a meaningful shift that may only accelerate in the years to come.

The rest of this paper proceeds as follows. Section 2 reviews related literature. Section 3 describes the data. Section 4 presents the baseline empirical results and event study. Section 5 reports the heterogeneity analysis. Section 6 presents robustness checks. Section 7 concludes.

2 Related Literature

Our paper connects to three streams of research: first, the burgeoning new empirical literature on the labor market consequences of generative AI; second, the broader labor economics literature on the effect of technological change on workers of different types; and third, the marketing strategy literature on marketing's role within the firm

[1] That said, as noted above, all seniority groups exhibit positive headcount growth in levels after ChatGPT's release, so the negative DiD coefficients reflect marketing growing less than non-marketing, rather than absolute declines. We note that for most seniority groups, and especially for junior employees, post-period growth rates across marketing and non-marketing are both below what pre-COVID trends would have predicted over a comparable horizon.

and the role of AI in marketing.

First, we contribute to the literature on the employment effects of AI and generative AI. Theoretical work models automation as displacing labor in existing tasks while the creation of new tasks reinstates it (Acemoglu and Restrepo, 2018), and a parallel body of work develops task-based measures of occupational AI exposure (Felten, Raj, and Seamans, 2019; Felten, Raj, and Seamans, 2023; Eloundou et al., 2024). On the empirical side, a growing number of studies exploit ChatGPT's November 2022 release as a natural experiment and document disemployment patterns in AI-exposed occupations, including reductions in freelance employment and job postings (Hui, Reshef, and Zhou, 2024; Demirci, Hannane, and Zhu, 2025), productivity gains that reduce task completion time by 25–40% (Noy and Zhang, 2023; Dell'Acqua et al., 2023), and employment declines in AI-exposed occupations in payroll microdata (Acemoglu, Autor, et al., 2022; Brynjolfsson, Chandar, and Chen, 2025). That said, Humlum and Vestergaard (2025) find precise null effects on aggregate earnings and hours in Danish administrative data, suggesting that early labor-market adjustment may in some cases operate through task reallocation rather than aggregate employment contraction. We contribute a focused analysis of a single high-exposure function (marketing) using a within-firm design that identifies effects by comparing marketing and non-marketing employees in the same firm, allowing us to isolate function-specific effects while controlling for firm-level shocks.

Second, our paper relates to the literature on the incidence of technological change across worker seniority and skill levels. The routine-biased technical change hypothesis predicts that automation disproportionately displaces middle-skill workers performing routine tasks (Autor, Levy, and Murnane, 2003; Acemoglu and Autor, 2011), while recent evidence on generative AI suggests that early-career and junior workers bear the largest employment declines (Hosseini Maasoum and Lichtinger, 2025; Bryn-

jolfsson, Chandar, and Chen, 2025). Our findings suggest a more nuanced picture within marketing, since junior and senior marketing employees experience the largest relative decline, while mid-level employees remain much closer to parity, consistent with a function-specific seniority gradient that does not map neatly onto either pattern. That said, the finding that endpoints of the seniority distribution are most affected aligns with the theoretical prediction of Cheng, Dogan, and Yildirim (2025) that AI optimally replaces workers at the top and bottom of a team hierarchy while retaining those in the middle, since middle workers serve to sustain peer monitoring and internal information flows.

Third, our paper contributes to the marketing literature on marketing organization and the emerging role of AI in marketing. A body of work has studied how marketing departments are organized within firms and how their influence has evolved (Verhoef and Leeflang, 2009; Moorman and Day, 2016), while a parallel stream develops conceptual frameworks for AI task replacement in marketing, predicting that AI substitutes for tasks in a progression from mechanical to analytical to intuitive and may serve more as an augmenter than a replacement (Huang and Rust, 2018; Huang and Rust, 2021; Davenport et al., 2020). These frameworks generate clear predictions about how AI should reshape the marketing workforce, yet the empirical evidence has remained largely survey- and case-based. We contribute large-scale causal evidence on how a specific AI shock reshapes the composition of marketing employment within firms, complementing these conceptual frameworks with granular data on which roles, seniority levels, and task types are most affected.

3 Data

3.1 Job Position Data: Revelio Labs

Our measures of marketing employment are constructed from Revelio Labs' individual-level position histories.[2] Revelio Labs aggregates workforce information from publicly accessible sources, including online professional profiles and resumes. The data cover both public and private companies and have been used increasingly in academic work to measure job transitions and turnover (Li et al., 2022; Cai et al., 2024). The position histories include employer identifiers, start and end dates, job role and category, seniority, and O*NET occupation codes. We restrict attention to firms headquartered in the United States and focus on 2021Q4–2025Q3.[3]

Using the individual-by-position records, we construct a firm-by-quarter panel of marketing and non-marketing employment headcounts. Revelio classifies each position into one of ten job roles; we define marketing employees as those in the "Sales and Marketing" role, excluding sales occupations identified using SOC/O*NET information.[4] We define non-marketing employees as those outside "Sales and Marketing," excluding "Software Engineer" and "Project and IT Specialist" roles.[5] Sales occupations, which fall within the "Sales and Marketing" role but are excluded from the marketing group, do not enter either group in the baseline and are examined sepa-

[2]For details, see https://www.data-dictionary.reveliolabs.com/data.html#individual-level-data and https://wrds-www.wharton.upenn.edu/pages/about/data-vendors/revelio-labs/.

[3]We do not additionally restrict employees by work location, so our estimates should be interpreted as the effect of ChatGPT's release on the workforce composition of U.S.-headquartered firms, rather than on U.S.-located workers alone; we report the baseline analysis for only U.S.-located workers in Section B.6.

[4]See Online Appendix A.1 for detailed descriptions of all subgroup classifications.

[5]Similar to Acemoglu, Autor, et al. (2022), employees in these roles may be directly involved in developing or implementing generative AI tools. We report robustness checks including these employees in Online Appendix B.5.

rately in Section 5.2. To ensure meaningful within-firm variation, we require at least 25 total employees and at least 5 marketing and 5 non-marketing employees at some point during the pre-period.[6]

After applying these restrictions, the final dataset is a panel of 164,780 firms observed over 16 quarters, covering 80,977,034 unique individuals and 112,678,568 unique employment records. Summary statistics are reported in Online Appendix Tables A1 and A2.

3.2 Construction of Subgroups

For the heterogeneity analyses, we construct subgroups along three dimensions. First, we aggregate Revelio's seven seniority levels into three tiers: junior (entry through associate), mid-level (manager and director), and senior (executive and senior executive).[7] Second, we use the O*NET database to classify marketing positions as "digital" or "non-digital" using a composite score constructed from O*NET technology skills, task statements, and structured descriptors of skills, knowledge, and work activities; occupations above the median score are classified as digital.[8] Third, we use O*NET occupation codes to partition the broad "Sales and Marketing" role into sales, customer-service marketing, and non-customer-service marketing positions.

3.3 AI Exposure Measures

We construct three complementary pre-period firm- and industry-level AI exposure measures: 1) GPT-β (Eloundou et al., 2024; Labaschin et al., 2025), which measures the extent to which occupational tasks can be assisted or substituted by large lan-

[6] We also use different sample selection criteria for robustness checks in Section 6; the results are quantitatively similar.

[7] We also present results for each of the seven levels separately in Online Appendix Table B13.

[8] Full construction details are provided in Online Appendix A.1.

guage models (see also Brynjolfsson, Chandar, and Chen, 2025); 2) FRS-AIOE/AIIE (Felten, Raj, and Seamans, 2019), a capability-based index linking AI-relevant capabilities to O*NET work activities and occupational requirements (see also Acemoglu, Autor, et al., 2022); and 3) AEI usage, which captures observed intensity of AI tool usage by occupation, complementing the capability-based measures above with a realized-adoption dimension (see als (Handa et al., 2025; Brynjolfsson, Chandar, and Chen, 2025)).

All three indices are measured at the occupation level and aggregated to firm- and industry-level baselines using only pre-period workforce composition, so they are predetermined relative to firms' post-ChatGPT decisions. We use median splits on these baselines for the heterogeneity analyses. Full construction and aggregation details are provided in Online Appendix A.

4 Empirical Results

4.1 Descriptive Evidence

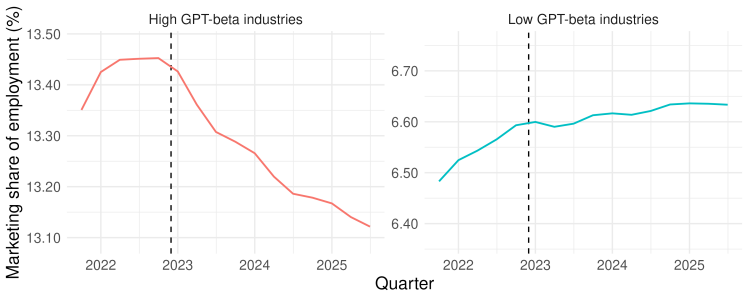
Before turning to the formal analysis, we present descriptive evidence on the marketing share of firm headcount split by AI exposure.[9] Figure 1 shows that both high- and low-exposure firms had increasing or flat marketing-share trajectories prior to ChatGPT's release. After the release, the trajectories diverge: in high-exposure firms, the marketing share peaks and then declines steadily, while in low-exposure firms it continues to grow slowly. The industry-level split tells a similar story, suggesting that the post-ChatGPT decline is concentrated where generative AI is most relevant.

[9]For this descriptive exercise, we use "GPT-β" (Eloundou et al., 2024; Labaschin et al., 2025) to measure AI exposure, with high/low exposure defined by the median split.

Figure 1: Aggregate marketing share across high- and low- AI exposed firms and industries



(a) AI exposure at firm level



(b) AI exposure at industry level

Note: Figure 1a plots aggregate marketing share of employment over time, separately for firms with high versus low AI exposure. Figure 1b reports analogous series splitting firms by high versus low industry-level AI exposure.

4.2 Econometric Specification

We now turn to formalizing the descriptive patterns documented above, using the public release of ChatGPT on November 30, 2022 as the focal event. We adopt a within-firm difference-in-differences design that compares marketing and non-marketing employees observed in the same firm and quarter, using non-marketing employees as a within-firm control group. Relative to cross-sectional comparisons of more- versus less-exposed occupations or industries used in some of the recent literature (Hui, Reshef, and Zhou, 2024; Demirci, Hannane, and Zhu, 2025), this approach has the advantage that firm $\times$ time fixed effects absorb any firm-specific, time-varying shock common to all employee groups, including firm-level AI adoption decisions and demand fluctuations; the identifying variation therefore comes only from differential changes in marketing versus non-marketing headcount within the same firm-quarter.[10] The key identifying assumption is that, absent ChatGPT's release, marketing and non-marketing headcount within the same firm would have evolved along parallel trajectories. We argue that this assumption is plausible because the fixed-effect structure absorbs the most natural threats to it, namely any firm-specific shock that affects all departments equally; what would be required to violate it is a shock that (1) coincides with ChatGPT's release, (2) differentially affects marketing relative to non-marketing employees within the same firm, and (3) is itself unrelated to generative AI. We assess the parallel trends assumption directly in the event study below.

Two additional concerns merit discussion. First, firms may have been restructuring marketing departments for reasons unrelated to generative AI, such as shifting spending toward programmatic advertising or outsourced services, or simply trimming non-essential commercial headcount. Three pieces of evidence weigh against this

[10] Our approach is similar in spirit to Brynjolfsson, Chandar, and Chen (2025), who also compare within-firm employee groups, although using different data and a different set of exposed occupations.

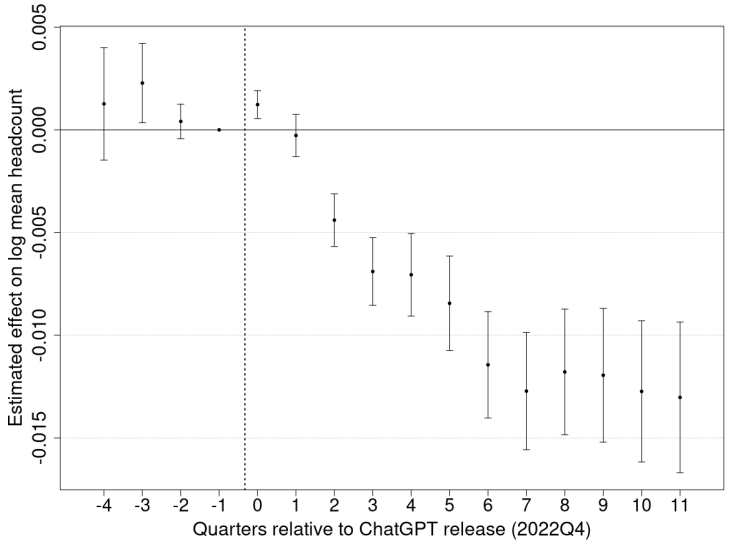
interpretation, namely (1) the event-study estimates show no differential pre-period trajectory between marketing and non-marketing employees; (2) the sales comparison in Section 5.2 shows that sales headcount actually expands relative to non-marketing after ChatGPT's release, which is difficult to reconcile with a general commercial-headcount reduction; and (3) the triple-differences analysis in Section 5 shows that the relative decline is stronger in firms and industries with higher pre-period AI exposure. Second, Revelio Labs data are based on online profiles rather than payroll records, so the data may contain coverage differences and profile-update lags. Our fixed-effects design absorbs time-invariant coverage differences and firm-level shocks to reporting behavior; the remaining threat would require a sharp post-ChatGPT change in profile coverage that differentially affects marketing relative to non-marketing employees. This concern is further mitigated by the finding that non-digital marketing roles exhibit a relative decline roughly four times larger than digital roles (Section 5); if the results were driven by differential profile-updating behavior, the effect should, if anything, be larger for digital workers, who are presumably more active on online platforms.

Stacking the data at the firm-quarter-group level, we implement the following Poisson pseudo-maximum likelihood (PPML) specification:

$$\log\left(E\left[Y_{fgt}\right]\right) = \beta \cdot \text{Post}_t \times \mathbb{1}\{g = M\} + \alpha_{fg} + \delta_{ft} + \epsilon_{fgt}, \tag{1}$$

where $g \in \{M, N\}$ denotes the employee group (marketing versus non-marketing), Y_{fgt} is headcount for firm f, group g, and quarter t, α_{fg} are firm $\times$ group fixed effects, and δ_{ft} are firm $\times$ time fixed effects. The coefficient $\exp(\beta) - 1$ captures the differential post-ChatGPT proportional change in marketing mean headcount relative to non-marketing mean headcount within the same firm. Standard errors are

Figure 2: Quarterly event study for log headcount changes



Note: Each point plots estimated coefficient $\hat{\beta}_s$ from specification (2). Vertical bars show 95% confidence intervals. Omitted reference quarter: 2022Q3 ($s = -1$). Standard errors clustered at firm level.

clustered at the firm level. We use PPML rather than OLS in log headcount because headcounts can be zero, especially in subgroup analyses (Santos Silva and Tenreyro, 2006; Chen and Roth, 2024; Wooldridge, 2023).

For the event study, we replace the single $\text{Post}_t \times \mathbb{1}\{g = M\}$ interaction in (1) with quarter-specific event-time indicators interacted with the marketing group indicator, omitting $s = -1$ (2022Q3) as the reference category:

$$\log\left(E\left[Y_{fgt}\right]\right) = \sum_{s \neq -1} \beta_s \, \mathbb{1}\{s(t) = s\} \times \mathbb{1}\{g = M\} + \alpha_{fg} + \delta_{ft} + \epsilon_{fgt}. \tag{2}$$

Under parallel trends, the pre-period coefficients should be close to zero for $s \leq -2$.

Table 1: Marketing versus non-marketing employee headcount after GPT release

	Headcount, PPML log link
Post × Marketing	-0.00923***
	(0.00145)
Firm-quarter observations	2,636,480
Firm × group FE	✓
Firm × time FE	✓
Pseudo R^2	0.993

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Note: PPML difference-in-differences estimates for employee headcount. Marketing defined as non-sales marketing employees. Sample restricted to firms with at least 25 total employees and at least 5 marketing and 5 non-marketing employees in pre-period. Standard errors clustered at firm level.

4.3 Baseline Results

Figure 2 plots the quarterly event-study coefficients from (2). The pre-period coefficients are close to zero, providing little evidence of a differential pre-trend. Beginning about two quarters after release, the estimates turn negative and become increasingly pronounced, reaching roughly 1.2–1.3% by the end of the sample. To benchmark this magnitude, weighted firm-average annual employment growth from 2018Q1 to 2020Q1 was roughly 4–5%, implying cumulative expected growth of approximately 14% over the 3-year post-period; the terminal relative shortfall therefore represents about one-tenth of the employment growth that pre-COVID trends would have predicted. The event-study evidence thus points to a growing and persistent relative decline in marketing employment following ChatGPT's introduction.

Pooling the post-period into a single coefficient, Table 1 reports that the average relative shortfall across the full post-period is $\hat{\beta} = -0.00923$, implying that expected marketing headcount is approximately $\exp(-0.00923) - 1 = -0.92\%$ lower relative to non-marketing headcount after ChatGPT's release. As shown below in Section 6, this result is robust to alternative sample selection criteria, dropping the transition

quarter (2022Q4), and broadening the control group to include tech employees.

To better understand the baseline effect, we decompose it into exits and hires using a PPML specification with lagged headcount as an offset (Online Appendix Table B1). Marketing employees experienced an approximately 4.2% increase in exits and a 1.7% reduction in hires relative to non-marketing employees, indicating that the baseline headcount effect reflects both elevated turnover and reduced recruitment.

5 Heterogeneity Analysis

We next examine which subgroups within marketing are most affected, decomposing the baseline effect by seniority, job type, and firm-level AI exposure.

5.1 Heterogeneity by Marketing Subgroups

To decompose the baseline effect, we stack the K marketing subgroups together with their corresponding non-marketing comparison subgroups and estimate subgroup-specific PPML difference-in-differences specifications with firm-by-subgroup and firm-by-time fixed effects; the coefficients measure subgroup-specific post-ChatGPT changes in marketing headcount relative to the corresponding non-marketing subgroup.

We begin with seniority, a dimension of particular interest given the emerging view that generative AI primarily reduces demand for junior workers (Hosseini Maasoum and Lichtinger, 2025; Brynjolfsson, Chandar, and Chen, 2025). Results are presented in Table 2.

Panel A presents within-seniority difference-in-differences estimates that compare each marketing subgroup (junior, mid-level, and senior) with the corresponding non-marketing subgroup of the same seniority. The estimates show negative relative post-ChatGPT changes for junior and senior marketing employees, but not

Table 2: Seniority heterogeneity in marketing versus non-marketing employee headcount

	Relative post-ChatGPT change (marketing vs. corresponding non-marketing subgroup)	Within-group post-ChatGPT change (no control group)
Panel A. Difference-in-differences estimates		
Junior marketing	-0.0241***	
	(0.00187)	
Mid-level marketing	-0.00116	
	(0.00151)	
Senior marketing	-0.0262***	
	(0.00216)	
Panel B. Within-group post-ChatGPT changes		
Junior Marketing		0.0121*
		(0.00569)
Mid-level Marketing		0.0387***
		(0.00261)
Senior Marketing		0.0765***
		(0.00589)
Junior non-marketing		0.0218***
		(0.00236)
Mid-level non-marketing		0.0463***
		(0.00175)
Senior non-marketing		0.11***
		(0.0027)
Firm-quarter observations	2,636,480	2,636,480
Firm × group FE	✓	✓
Firm × time FE	✓	
Pseudo R^2	0.988	0.990

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Note: Panel A: PPML difference-in-differences estimates for each marketing seniority subgroup relative to corresponding non-marketing subgroup of same seniority. Panel B: subgroup-specific post-ChatGPT changes from separate specification with firm × group fixed effects and no control group. Standard errors clustered at firm level.

for mid-level marketing employees. The junior coefficient implies an approximately $\exp(-0.0241) - 1 = -2.38\%$ relative decline in marketing headcount compared with junior non-marketing mean headcount, while the senior coefficient implies an approximately 2.59% relative decline. By contrast, the mid-level coefficient is economically close to zero and statistically insignificant. Thus, relative to matched non-marketing employees, the post-ChatGPT evolution of marketing headcount is weaker for junior and senior workers, while mid-level marketing employment evolves close to parity with its non-marketing counterpart. This pattern is consistent with the theoretical prediction in Cheng, Dogan, and Yildirim (2025) that a principal deploying AI in a sequential team optimally replaces workers at the beginning and end of the workflow while retaining middle workers, who serve to sustain internal information flows and coordination.

Panel B reports within-group post-ChatGPT changes without firm $\times$ time fixed effects. All six coefficients are positive, indicating headcount growth within every seniority group after ChatGPT's release, with growth monotone in seniority. At every seniority level, however, marketing growth is lower than the corresponding non-marketing growth, and at most seniority levels, especially junior, both marketing and non-marketing growth rates fall well short of what pre-COVID trends would have predicted. The negative DiD estimates in Panel A therefore reflect marketing growing less than non-marketing rather than absolute declines.[11]

A finer decomposition across all seven Revelio seniority levels (Online Appendix Table B13) reveals that within the junior tier, the relative decline is largest among associate-level employees (-4.29%), the most senior of the three junior ranks, and

[11] Using weighted firm-average annual growth rates from 2018Q1 to 2020Q1 as a benchmark, the junior result represents a reversal of a pre-period differential that had favored marketing (+1.2 pp), the senior result reflects an acceleration of a pre-existing shortfall (-0.85 pp), and mid-level marketing and non-marketing were growing at nearly identical rates pre-period, consistent with the near-zero post-period DiD.

within the senior tier, the effect is driven by executives (−2.92%) while the most senior rank (senior executives) shows no significant relative change.

5.1.1 Exits versus Hires Decomposition

A decomposition into exits and hires by seniority (Online Appendix Table B2) clarifies the underlying dynamics: juniors experience a sharp reduction in hires (−6.45%) with a modest increase in separations, indicating that the junior shortfall is largely driven by reduced recruitment; mid-level marketing employees show elevated flows in both directions, consistent with reallocation rather than net displacement; and seniors exhibit moderate increases in both hires and exits, reflecting ongoing restructuring. The exit/hire decomposition for digital versus non-digital roles (Online Appendix Table B3) confirms that the non-digital shortfall is also driven primarily by reduced recruitment.

5.1.2 Digital versus Non-Digital Marketing

Online Appendix Table B4 decomposes the effect by digital intensity. The relative post-ChatGPT decline is present in both digital and non-digital marketing roles, but the non-digital shortfall (−3.13%) is roughly four times as large as the digital shortfall (−0.79%), consistent with generative AI capabilities overlapping more with non-digital tasks such as copywriting and campaign management than with data-driven digital roles.

5.2 Sales, Customer-Service Marketing, and Non-Customer-Service Marketing

We next decompose the broader "Sales and Marketing" role into sales, customer-service marketing, and non-customer-service marketing, with non-marketing employees as the comparison group throughout (Online Appendix Table B5). Customer-service marketing and non-customer-service marketing together constitute the baseline marketing group used above; sales, which is excluded from the baseline entirely, is reintroduced here to permit a direct comparison across all three subgroups.

The estimates reveal sharply divergent patterns. Customer-service marketing exhibits the largest relative decline (-2.6%), followed by non-customer-service marketing roles (-1.1%). Sales moves in the opposite direction, with a statistically significant $+0.59\%$ relative increase. We note that the baseline effect is not driven solely by customer-service marketing; non-customer-service marketing roles exhibit a meaningful -1.1% relative decline on their own, indicating that the adjustment extends across non-sales marketing functions more broadly. If firms were broadly trimming commercial headcount or shedding non-essential positions, sales and non-sales marketing would weaken together; that sales instead expands while customer-service marketing and non-customer-service marketing decline points to a task-specific adjustment consistent with generative AI substituting for marketing work rather than a general reduction in commercial staffing.

5.3 Heterogeneity by Firm and Industry AI Exposure

To test whether the effects are stronger where generative AI is most relevant, we extend the baseline to a triple-differences (DDD) specification that interacts the post-

period indicator, the marketing-group indicator, and an above-median AI-exposure indicator at the firm or industry level. Online Appendix Table B6 reports the results across three exposure measures (GPT-β, FRS-AIOE, AEI) at both levels. The DDD coefficient is negative and statistically significant in all six specifications, implying additional relative declines of approximately 1.3% to 2.4% in high-exposure environments. In several specifications, marketing headcount does not decline relative to non-marketing in low-exposure settings; it is in high-exposure settings that the relative decline emerges. These results complement the occupation-level evidence in Brynjolfsson, Chandar, and Chen (2025) by showing a similar exposure gradient at the firm and industry level within marketing specifically.

6 Robustness Checks

We assess the sensitivity of the baseline result along several dimensions; we summarize these robustness checks in brief here, with full details reported in the online appendix.

Including tech employees in the control group yields a qualitatively similar, or slightly larger, estimate (Online Appendix Table B7), indicating that their exclusion is conservative. Restricting to U.S.-located employees yields a similar, highly significant coefficient, albeit of smaller magnitude (-0.0036; Online Appendix Table B8). The coefficient remains negative and significant under four alternative sample definitions (-0.0055 to -0.0101; Online Appendix Tables B9–B12), and extending the event window to 2021Q1–2025Q3 yields a virtually unchanged estimate of -0.00891 (Online Appendix Table B14).

7 Conclusion

We study how the release of ChatGPT affected marketing employment using more than 110 million online employment records from Revelio Labs covering 164,780 U.S.-headquartered firms over 2021–2025, comparing marketing to non-marketing employees within the same firm.

We find that marketing headcount declines by approximately 0.92% relative to non-marketing headcount within the same firm after ChatGPT's release, with the divergence growing to roughly 1.2–1.3% by the end of the sample and roughly doubling in more AI-exposed firms and industries. The relative decline is concentrated in non-digital marketing and customer-service marketing roles, while sales headcount is comparatively resilient and even expands slightly, pointing to task-specific adjustment rather than a general reduction in commercial staffing. Junior and senior marketing employees bear the largest relative shortfalls compared with same-seniority non-marketing peers, while mid-level employees remain closer to parity, a pattern consistent with the theoretical prediction that middle workers are retained for their coordinating role (Cheng, Dogan, and Yildirim, 2025).

That said, the analysis has several limitations. We observe headcounts but not the tasks, productivity, or wages of remaining employees, so we cannot determine whether the relative reduction reflects pure displacement or partial within-firm reallocation. Revelio Labs data are drawn from online profiles and resumes, which may underrepresent workers with limited online presence. We measure AI exposure from pre-period occupational composition rather than observing firms' actual adoption decisions, and the within-firm design, while absorbing firm-level shocks, cannot fully rule out marketing-specific shocks that coincide with ChatGPT's release. Nonetheless, the AI-exposure gradient in the triple-differences analysis helps mitigate this

last concern, since any confounding shock would need to be systematically stronger in more AI-exposed environments. Linking employment data to firm-level measures of AI tool adoption would allow researchers to test whether these patterns are driven by realized adoption, and extending the design to other high-exposure functions would clarify whether the dynamics documented here are specific to marketing.

Taken together, our findings suggest that the marketing function is undergoing a measurable but selective restructuring in the wake of generative AI, concentrated in roles where AI capabilities most directly overlap with existing tasks. As these tools continue to mature and diffuse, understanding how the marketing workforce reorganizes around them remains a first-order question for both marketing practice and labor market policy.

References

Acemoglu, Daron and David Autor (2011). "Skills, Tasks and Technologies: Implications for Employment and Earnings". In: *Handbook of Labor Economics* 4, pp. 1043–1171.

Acemoglu, Daron, David Autor, et al. (2022). "Artificial intelligence and jobs: Evidence from online vacancies". In: *Journal of Labor Economics* 40.S1, S293–S340.

Acemoglu, Daron and Pascual Restrepo (2018). "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment". In: *American Economic Review* 108.6, pp. 1488–1542.

Autor, David H., Frank Levy, and Richard J. Murnane (2003). "The Skill Content of Recent Technological Change: An Empirical Exploration". In: *The Quarterly Journal of Economics* 118.4, pp. 1279–1333.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen (2025). *Canaries in the Coal Mine?: Six Facts about the Recent Employment Effects of Artificial Intelligence*. Tech. rep. Stanford Institute for Economic Policy Research (SIEPR).

Cai, Wei et al. (2024). "Diversity targets". In: *Review of Accounting Studies* 29.3, pp. 2157–2208.

Chen, Jiafeng and Jonathan Roth (2024). "Logs with zeros? Some problems and solutions". In: *The Quarterly Journal of Economics* 139.2, pp. 891–936.

Cheng, Xienan, Mustafa Dogan, and Pinar Yildirim (2025). *Artificial Intelligence in Team Dynamics: Who Gets Replaced and Why?* Working Paper 34259. National Bureau of Economic Research.

Chui, Michael et al. (2023). *The Economic Potential of Generative AI: The Next Productivity Frontier*. McKinsey & Company.

Davenport, Thomas et al. (2020). "How Artificial Intelligence Will Change the Future of Marketing". In: *Journal of the Academy of Marketing Science* 48.1, pp. 24–42.

Dell'Acqua, Fabrizio et al. (2023). "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality". In: *Harvard Business School Working Paper* 24-013.

Demirci, Ozge, Jonas Hannane, and Xinrong Zhu (2025). "Who Is AI Replacing? The Impact of Generative AI on Online Freelancing Platforms". In: *Management Science*.

Eloundou, Tyna et al. (2024). "GPTs are GPTs: Labor market impact potential of LLMs". In: *Science* 384.6702, pp. 1306–1308.

Felten, Edward, Manav Raj, and Robert Seamans (2023). "Occupational Heterogeneity in Exposure to Generative AI". In: *SSRN Electronic Journal*.

Felten, Edward, Manav Raj, and Robert Channing Seamans (2019). "The effect of artificial intelligence on human labor: An ability-based approach". In: *Academy of Management Proceedings*. Vol. 2019. 1. Academy of Management Briarcliff Manor, NY 10510, p. 15784.

Handa, Kunal et al. (2025). "Which economic tasks are performed with ai? evidence from millions of claude conversations". In: *arXiv preprint arXiv:2503.04761*.

Hosseini Maasoum, Seyed Mahdi and Guy Lichtinger (2025). "Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data". In: *Working paper, Harvard University*. Available at SSRN: `https://ssrn.com/abstract=5425555`.

Huang, Ming-Hui and Roland T. Rust (2018). "Artificial Intelligence in Service". In: *Journal of Service Research* 21.2, pp. 155–172.

— (2021). "A Strategic Framework for Artificial Intelligence in Marketing". In: *Journal of the Academy of Marketing Science* 49.1, pp. 30–50.

Hui, Xiang, Oren Reshef, and Luofeng Zhou (2024). "The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market". In: *Organization Science* 35.6, pp. 1977–1989.

Humlum, Anders and Emilie Vestergaard (2025). *Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI*. Working Paper 33777. National Bureau of Economic Research.

Labaschin, Benjamin et al. (2025). "Extending "GPTs Are GPTs" to Firms". In: *AEA Papers and Proceedings* 115, pp. 51–55.

Li, Qin et al. (2022). "Employee turnover and firm performance: Large-sample archival evidence". In: *Management Science* 68.8, pp. 5667–5683.

Moorman, Christine and George S. Day (2016). "Organizing for Marketing Excellence". In: *Journal of Marketing* 80.6, pp. 6–35.

Noy, Shakked and Whitney Zhang (2023). "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence". In: *Science* 381.6654, pp. 187–192.

Santos Silva, J. M. C. and Silvana Tenreyro (2006). "The Log of Gravity". In: *The Review of Economics and Statistics* 88.4, pp. 641–658.

Verhoef, Peter C. and Peter S. H. Leeflang (2009). "Understanding the Marketing Department's Influence Within the Firm". In: *Journal of Marketing* 73.2, pp. 14–37.

Wooldridge, Jeffrey M. (2023). "Simple Approaches to Nonlinear Difference-in-Differences with Panel Data". In: *The Econometrics Journal* 26.3, pp. C31–C66.

A Data Construction Details

A.1 O*NET Pipeline and Subgroup Classification

We use the U.S. Department of Labor's O*NET (O*NET-SOC) database to obtain standardized, occupation-level descriptors of work content (tasks, work activities, skills, and related requirements) at the SOC level. O*NET provides a stable and externally maintained characterization of what workers do on the job, which we treat as predetermined relative to firms' subsequent adoption of generative AI. In our pipeline, O*NET enters in three steps.

First, to construct marketing employees, we begin with all positions assigned to the job role "Sales and Marketing" in the dataset and use occupation-level SOC/O*NET information to distinguish sales from non-sales roles. Based on the O*NET database, we code occupations in the Sales and Related Occupations major group (SOC 41-) as sales. Our baseline marketing measure is therefore the non-sales subset of the "Sales and Marketing" role, consisting of customer-service marketing and non-customer-service marketing roles but excluding SOC 41 sales occupations. Sales occupations, which fall within the "Sales and Marketing" role but are excluded from the marketing group, do not enter either group in the baseline and are examined separately in Section 5.2. For the decomposition in Section 5.2, Customer Service Representatives (SOC 43-4051) are coded separately as customer-service marketing.

Second, for subgroup construction, we classify occupations as digital or non-digital using an occupation-level digital score constructed from O*NET database. Beginning with O*NET-SOC occupations, we combine information from the Technology Skills, Tools Used, Task Statements, Task Ratings, Skills, Knowledge, and Work Activities files. The measure is built in two stages. In the first stage, we identify digital content

in technology-skill examples and tool descriptions using a curated dictionary of software, computing, digital-business, and digital-marketing terms. Matches in the Technology Skills file receive additional weight when the underlying technology is flagged by O*NET as "hot" or "in demand." In the second stage, we augment this text-based measure with task statements and structured descriptors. Task-based digital content is weighted using O*NET importance ratings when available, and otherwise by expected frequency constructed from the Task Ratings file. The structured component is based on selected Skills, Knowledge, and Work Activities dimensions, including programming, systems analysis, technology design, systems evaluation, computers and electronics, telecommunications, and interacting with computers. Each component is converted to a percentile rank and then aggregated into a composite score, with weights of 0.10 for Technology Skills, 0.10 for Tools Used, 0.35 for Tasks, and 0.45 for the structured descriptors. The resulting occupation-level score is merged to job positions using exact O*NET-SOC codes. In the empirical analysis, occupations with digital scores above the sample median are classified as digital, and the remainder are classified as non-digital.

Third, we use the O*NET database to classify firms and industries into high versus low AI exposure groups. We map each worker's job position in our workforce data to a SOC code using the provider's occupation crosswalk, merge the O*NET-based occupation measures to each spell by SOC, and then aggregate position-level measures to the firm level (and, separately, to industries identified by four-digit NAICS codes) using only the pre-period to form predetermined exposure measures used in the difference-in-differences analyses.

A.2 Aggregation to Firm- and Industry-Level Baselines

The O*NET-based occupation measures are merged to job positions and aggregated to firm-level baselines by taking headcount-weighted averages across occupations within each firm in the pre-period (2018Q1–2022Q3). The weighted average uses employment duration within firm to emphasize occupations that account for more worker-time in the firm's pre-period workforce; these baselines remain fixed over time and are merged to the firm–quarter panel used for estimation.

We incorporate three complementary AI exposure indices that measure how strongly different occupations (and, by aggregation, firms and industries) are exposed to generative AI and related AI capabilities. The indices are merged to worker job positions by O*NET-SOC codes and then aggregated to firm-level, predetermined pre-period baselines, which are subsequently attached to the firm–quarter panel and used for high-versus-low splits (median split) at the firm level. We also construct industry-level exposure measures at the four-digit NAICS level by taking pre-period headcount-weighted averages of firm baselines within each four-digit NAICS industry, which allows parallel firm- and industry-exposure heterogeneity analyses while preserving the predetermined nature of the exposure variables.

A.3 AI Exposure Index Descriptions

The GPT-β index, developed by Eloundou et al. (2024), is an occupation-level measure designed to capture the extent to which the tasks performed in an occupation can be assisted or substituted by large language models. (See also Brynjolfsson, Chandar, and Chen, 2025.) In the underlying framework, each O*NET task receives a β rating equal to 1 if the task can be completed at least twice as fast by an LLM alone, 0.5 if achieving that speedup requires additional software or tools built on top of the LLM,

and 0 otherwise. Occupation-level exposure is then obtained by aggregating task-level ratings within occupation, with O*NET core tasks receiving twice the weight of supplemental tasks. We merge this occupation-level score to worker job positions using O*NET-SOC codes and compute each firm's baseline GPT-β exposure as the pre-period headcount-weighted average across its occupational mix. This baseline serves as a predetermined measure of the firm's exposure to generative-AI capabilities at the time of the ChatGPT shock.

The FRS-AIOE index, developed by Felten, Raj, and Seamans (2019), measures occupational exposure to AI using a capability-based framework that links common AI applications to occupational abilities in O*NET. (See also Acemoglu, Autor, et al., 2022.) The dataset provides AI Occupational Exposure (AIOE) scores by six-digit SOC occupation and AI Industry Exposure (AIIE) scores by four-digit NAICS industry, and the construction proceeds from an application–ability relatedness matrix and standardized ability-level exposure measures. As with GPT-β, we merge the SOC-level FRS-AIOE score to job positions and compute a firm-level predetermined baseline using pre-period headcount weights. We also construct an industry-level analogue by aggregating firm baselines to NAICS4 industries and, where appropriate, use the provider's AIIE series as a direct industry-level benchmark.

The Anthropic Economic Index (AEI) captures a different concept: realized AI use rather than pure technological feasibility. The AEI is built from millions of anonymized Claude interactions and maps observed usage to O*NET task statements and occupational categories. It is designed to measure which theoretically feasible tasks are actually seeing meaningful AI use in professional settings, with greater weight placed on tasks that are work-related, more automated rather than merely augmentative. We merge the occupation-level AEI measure to worker job positions by O*NET-SOC code and compute firm-level predetermined baselines in the same

manner as for the other indices, using pre-period headcount-weighted averages across the firm's occupational mix.

Table A1: Summary statistics for employee headcounts

	Before ChatGPT			After ChatGPT		
Variable	Mean	SD	Median	Mean	SD	Median
Marketing employee headcount	42.66	539.87	11.00	43.88	585.13	11.50
Non-marketing employee headcount	258.10	2,307.78	41.00	265.68	2,283.82	42.50
Tech employee headcount	71.87	1,330.50	7.50	75.63	1,364.57	8.00
Total employee headcount	397.23	3,770.91	76.00	410.58	3,834.84	79.00
Marketing share of firm employment	0.189	0.141	0.155	0.185	0.141	0.150
Pre: Obs=659,120, Firms=164,780, Quarters=4 (2021-10-01 to 2022-07-01)						
Post: Obs=1,977,360, Firms=164,780, Quarters=12 (2022-10-01 to 2025-07-01)						
Overall: Obs=2,636,480, Firms=164,780, Quarters=16 (2021-10-01 to 2025-07-01)						

Note: Statistics are computed on the firm-quarter panel restricted to firms with at least 25 total employees and at least 5 marketing employees and 5 non-marketing employees at some point during the pre-period, to ensure representativeness and meaningful variation.

A.4 Summary Statistics

Tables A1 and A2 report summary statistics for the estimation sample.

Table A2: Summary statistics for employee flows

	Before ChatGPT			After ChatGPT		
Variable	Mean	SD	Median	Mean	SD	Median
Marketing hires	2.89	35.33	0.00	1.83	20.07	0.00
Marketing exits	2.25	19.73	0.00	1.91	17.68	0.00
Non-marketing hires	15.98	166.69	2.00	10.94	103.25	1.00
Non-marketing exits	13.21	143.26	2.00	10.60	105.57	2.00
Tech hires	5.09	103.49	0.00	3.06	53.64	0.00
Tech exits	3.73	72.38	0.00	2.93	51.10	0.00
Total employee hires	25.40	282.04	5.00	16.82	163.78	3.00
Total employee exits	20.27	215.18	4.00	16.37	161.67	3.00
Pre: Obs=659,120, Firms=164,780, Quarters=4 (2021-10-01 to 2022-07-01)						
Post: Obs=1,977,360, Firms=164,780, Quarters=12 (2022-10-01 to 2025-07-01)						
Overall: Obs=2,636,480, Firms=164,780, Quarters=16 (2021-10-01 to 2025-07-01)						

Note: Extended summary statistics for employee flows before and after ChatGPT release. Statistics are computed on the firm-quarter panel restricted to firms with at least 25 total employees and at least 5 marketing employees and 5 non-marketing employees at some point during the pre-period.

Table B1: Marketing versus non-marketing employee exits and hires after GPT release

	Exits	Hires
Post × Marketing	0.0422***	-0.0168**
	(0.00406)	(0.0068)
Firm-quarter observations	2,471,577	2,471,577
Firm × group FE	✓	✓
Firm × time FE	✓	✓
Pseudo R^2	0.936	0.941

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Note: PPML difference-in-differences estimates for employee exits and hires. Marketing defined as non-sales marketing employees. Sample restricted to firms with at least 25 total employees and at least 5 marketing and 5 non-marketing employees in pre-period. Standard errors clustered at firm level.

B Extended Empirical Analyses

B.1 Exits versus Hires Decomposition

Table B2: Seniority heterogeneity in marketing versus non-marketing employee exits and hires

	Exits	Hires
Junior marketing	0.0256***	-0.0645***
	(0.00343)	(0.00752)
Mid-level marketing	0.062***	0.0505***
	(0.00399)	(0.00777)
Senior marketing	0.071***	0.0373***
	(0.00698)	(0.00918)
Firm-quarter observations	2,464,330	2,464,330
Firm × group FE	✓	✓
Firm × time FE	✓	✓
Pseudo R^2	0.913	0.917

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Note: PPML difference-in-differences estimates for each marketing seniority subgroup relative to corresponding non-marketing subgroup of same seniority. Standard errors clustered at firm level.

Table B3: Digital heterogeneity in marketing versus non-marketing employee exits and hires

	Exits	Hires
Digital marketing	0.0592***	0.0115*
	(0.00327)	(0.00653)
Non-digital marketing	0.0438***	-0.217***
	(0.0083)	(0.014)
Firm-quarter observations	2,466,655	2,466,655
Firm × group FE	✓	✓
Firm × time FE	✓	✓
Pseudo R^2	0.917	0.92

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Note: PPML difference-in-differences estimates for digital and non-digital marketing employees relative to corresponding non-marketing subgroup. Standard errors clustered at firm level.

Table B4: Digital heterogeneity in marketing versus non-marketing employee headcount

	Headcount
Digital marketing	-0.0079***
	(0.00131)
Non-digital marketing	-0.0318***
	(0.0039)
Firm-quarter observations	2,636,480
Firm × group FE	✓
Firm × time FE	✓
Pseudo R^2	0.990

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Note: PPML difference-in-differences estimates for digital and non-digital marketing employees relative to corresponding non-marketing subgroup. Standard errors clustered at firm level.

B.2 Digital versus Non-Digital Marketing

Table B4 decomposes the baseline headcount effect by digital intensity, as discussed in Section 5.

Table B5: Headcount changes for sales, customer-service marketing, and non-customer-service marketing roles

	Relative post-ChatGPT change (vs. non-marketing employees)
Sales	0.00587**
	(0.00213)
Customer-service marketing	-0.0258***
	(0.00601)
Non-customer-service marketing	-0.0108***
	(0.00151)
Firm-quarter observations	2,636,480
Firm × group FE	✓
Firm × time FE	✓
Pseudo R^2	0.994

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Note: PPML difference-in-differences estimates for employee headcount. Each coefficient captures subgroup-specific post-ChatGPT change relative to non-marketing control group. Customer-service marketing and non-customer-service marketing together constitute the baseline marketing group; sales, excluded from the baseline, is reintroduced here. Standard errors clustered at firm level.

B.3 Sales, Customer-Service Marketing, and Non-Customer-Service Marketing

Table B5 reports estimates for sales, customer-service marketing, and non-customer-service marketing roles, as discussed in Section 5.2. Customer-service marketing and non-customer-service marketing together form the baseline marketing group; sales, excluded from the baseline, is reintroduced for this decomposition.

Table B6: Triple-differences estimates for marketing headcount by AI exposure

	Firm-level exposure			Industry-level exposure		
	GPT-β	FRS	AEI	GPT-β	FRS	AEI
Post × Marketing	0.00344*	0.00225	-0.000489	0.00789***	0.00427	0.0024
	(0.00137)	(0.00166)	(0.00146)	(0.00164)	(0.00288)	(0.00201)
Post × Marketing × High AI exposure	-0.0225***	-0.0245***	-0.0135***	-0.024***	-0.0182***	-0.015***
	(0.00275)	(0.00274)	(0.00255)	(0.00255)	(0.00327)	(0.00268)
Firm-quarter observations	2,636,461	2,636,461	2,636,461	2,636,461	2,636,461	2,636,461
Firm × group FE	✓	✓	✓	✓	✓	✓
Firm × time FE	✓	✓	✓	✓	✓	✓
Pseudo R^2	0.993	0.993	0.993	0.993	0.993	0.993

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Note: PPML triple-differences estimates for employee headcount. Columns 1–3 use above-/below-median firm-level AI exposure; columns 4–6 use industry-level exposure. Post × High AI exposure absorbed by Firm × time FE. Standard errors clustered at firm level.

B.4 Triple-Differences by AI Exposure

Table B6 reports triple-differences estimates across three AI exposure measures at both the firm and industry level, as discussed in Section 5.

B.5 Including Tech Employees in the Control Group

As discussed in Section 3, we exclude employees in "Software Engineer" and "Project and IT Specialist" roles from the baseline control group because their labor-market dynamics may differ from those of other non-marketing employees. Table B7 expands the control group to include tech employees. The estimated post-ChatGPT coefficient for marketing is -0.0148, larger in magnitude than the baseline estimate, confirming that the main result is not driven by the exclusion of tech workers.

Table B7: Robustness to including tech employees in the non-marketing control group

	Relative post-ChatGPT change (marketing vs. non-marketing + tech employees)
Marketing	-0.0148***
	(0.00141)
Firm-quarter observations	2,636,480
Firm × group FE	✓
Firm × time FE	✓
Pseudo R^2	0.995

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Note: PPML difference-in-differences estimate for employee headcount. Control group combines non-marketing and tech employees. Standard errors clustered at firm level.

B.6 Restricting to U.S.-Located Employees

Table B8 restricts the sample to U.S.-located employees only. The PPML estimate remains negative and statistically significant, but is smaller in magnitude than in the full sample, implying an approximately 0.36% lower expected marketing headcount relative to non-marketing headcount after ChatGPT's release. This is materially smaller than the baseline full-sample estimate of 0.92%. Figure B1 shows that the U.S.-only quarterly event study follows the same qualitative pattern as the baseline but with smaller post-period declines. One possible interpretation is that the full-sample estimate reflects both domestic employment changes and shifts in the geographic composition of marketing work; because the full-sample estimate is more negative, the results are consistent with a relative decline in marketing headcount that is at least as strong, and likely stronger, outside the United States.

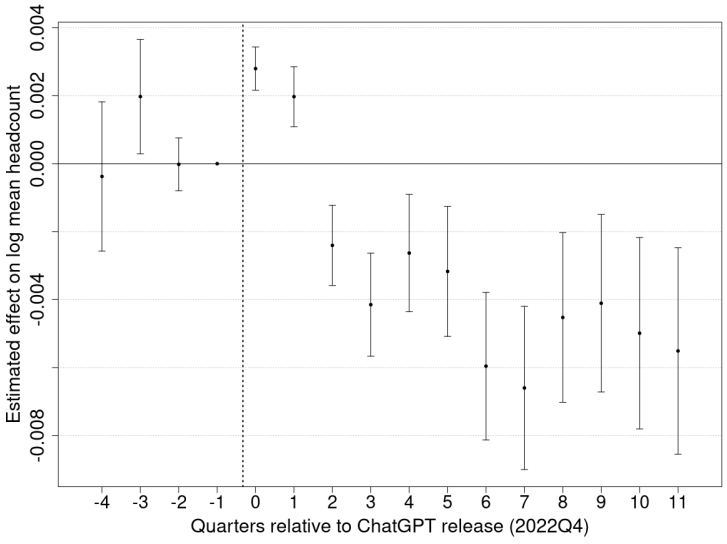
Table B8: Robustness to US employees only

	Headcount, PPML log link
Post × Marketing	-0.00364**
	(0.0012)
Firm-quarter observations	2,146,576
Firm × group FE	✓
Firm × time FE	✓
Pseudo R^2	0.992

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Note: PPML difference-in-differences estimate for employee headcount. Sample restricted to US employees only. Standard errors clustered at firm level.

Figure B1: Quarterly event study using US employees only



Note: Each point plots estimated coefficient $\hat{\beta}_s$ from specification (2) using US employees only. Vertical bars show 95% confidence intervals. Omitted reference quarter: 2022Q3 ($s = -1$). Standard errors clustered at firm level.

B.7 Alternative Sample Definitions

Tables B9–B12 re-estimate the baseline PPML headcount specification under four alternative sample definitions, as discussed in Section 6. Table B9 loosens the pre-period eligibility threshold to 25 total employees. Table B10 requires the 25/5/5 threshold in all pre-period quarters rather than at any point. Table B11 tightens the threshold to 25/10/10. Table B12 drops the transition quarter 2022Q4. The estimated coefficient is negative and statistically significant in all four cases, with magnitudes ranging from −0.00554 to −0.0101.

Table B9: Robustness to loosening pre-period employment thresholds

	Relative post-ChatGPT change (marketing vs. non-marketing employees)
Post × Marketing	-0.00554***
	(0.00137)
Firm-quarter observations	4,990,656
Firm × group FE	✓
Firm × time FE	✓
Pseudo R^2	0.991

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Note: PPML difference-in-differences estimate for employee headcount. Sample restriction loosened to firms with at least 25 total employees at some point during pre-period. Standard errors clustered at firm level.

Table B10: Robustness to restricting the sample to firm-quarters meeting the 25/5/5 employment thresholds

	Relative post-ChatGPT change (marketing vs. non-marketing employees)
Post × Marketing	-0.00887***
	(0.00148)
Firm-quarter observations	2,220,080
Firm × group FE	✓
Firm × time FE	✓
Pseudo R^2	0.994

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Note: PPML difference-in-differences estimate for employee headcount. Sample restricted to firms with at least 25 total employees, including at least 5 marketing and 5 non-marketing employees for all pre-event quarters. Standard errors clustered at firm level.

Table B11: Robustness to stricter pre-period firm eligibility thresholds

	Relative post-ChatGPT change (marketing vs. non-marketing employees)
Post × Marketing	-0.00949***
	(0.00156)
Firm-quarter observations	1,491,136
Firm × group FE	✓
Firm × time FE	✓
Pseudo R^2	0.995

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Note: PPML difference-in-differences estimate for employee headcount. Sample restricted to firms with at least 25 total employees and at least 10 marketing and 10 non-marketing employees at some point during pre-period. Standard errors clustered at firm level.

Table B12: Robustness to excluding 2022Q4 from the sample

	Relative post-ChatGPT change (marketing vs. non-marketing employees)
Post × Marketing	-0.0101***
	(0.00152)
Observations	2,471,700
Firm × group FE	✓
Firm × time FE	✓
Pseudo R^2	0.993

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Note: PPML difference-in-differences estimate for employee headcount. Transition quarter 2022Q4 excluded from sample; post period begins in 2023Q1. Standard errors clustered at firm level.

B.8 Finer Seniority Classification

The main-text seniority analysis groups employees into three broad tiers (junior, mid-level, senior). Table B13 reports PPML difference-in-differences estimates for each of the seven Revelio seniority levels separately, comparing marketing employees at each level with non-marketing employees at the same level. The finer decomposition reveals that within the junior tier, the relative decline is largest among associate-level employees (-4.29%), followed by junior (-3.65%) and entry (-2.49%). Within the senior tier, the effect is driven by executive-level employees (-2.92%), while senior executives show no significant relative change. Manager and director coefficients are small, confirming the near-parity finding for mid-level employees in the main analysis.

Table B13: Finer seniority heterogeneity in marketing versus non-marketing employee headcount

	Relative post-ChatGPT change (marketing vs. corresponding non-marketing subgroup)
Entry (S1)	-0.0252***
	(0.00308)
Junior (S2)	-0.0372***
	(0.00353)
Associate (S3)	-0.0438***
	(0.00271)
Manager (S4)	-0.00368
	(0.00197)
Director (S5)	0.00308
	(0.00165)
Executive (S6)	-0.0296***
	(0.00248)
Senior executive (S7)	-0.00135
	(0.00287)
Firm-quarter observations	2,636,461
Firm × group FE	✓
Firm × time FE	✓
Pseudo R^2	0.978

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Note: Each coefficient is a PPML difference-in-differences estimate for a single seniority level, comparing marketing employees at that level with non-marketing employees at same level. Standard errors clustered at firm level.

B.9 Longer Event Window

The baseline analysis uses a 2021Q4–2025Q3 event window with four pre-treatment quarters, starting in late 2021 to reduce the influence of COVID-era labor-market disruptions on the pre-period. To assess sensitivity to this choice, we re-estimate the baseline specification over a longer 2021Q1–2025Q3 window, which adds three earlier pre-period quarters, as used by Brynjolfsson, Chandar, and Chen (2025). Results are presented in Table B14 and Figure B2. The point estimate (−0.00891) is virtually unchanged from the baseline (−0.00923), and the extended event study shows a qualitatively similar pattern, confirming that the results are not sensitive to the choice of pre-period length.

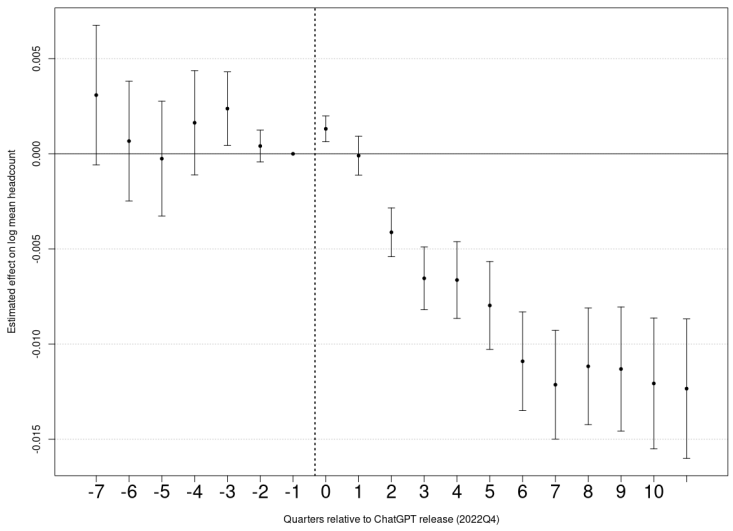
Table B14: Robustness checks for a longer event window (2021Q1-2025Q3)

	Log headcount
Post × Marketing	-0.00891***
	(0.00175)
Firm-quarter observations	3,261,141
Firm × group FE	✓
Firm × time FE	✓
Pseudo R^2	0.993

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Note: PPML difference-in-differences estimate for employee headcount using longer event window (2021Q1–2025Q3). Standard errors clustered at firm level.

Figure B2: Quarterly event study for 2021Q1-2025Q3



Note: Each point plots estimated coefficient $\hat{\beta}_s$ from specification (2). Vertical bars show 95% confidence intervals. Omitted reference quarter: 2022Q3 ($s = -1$). Standard errors clustered at firm level.