Generative AI, Online Self-Presentation, and Labor Market Signaling: Evidence from LinkedIn Posting Behavior

Derya Oktay[*], Mona Birjandi[**]

[*] Derya Oktay, Mr. Derya Oktay is a Data Analyst on the Economic & Data Analytics team at Outten & Golden LLP. Before joining the firm, Mr. Oktay worked in economic consulting and developmental economics. He has provided consulting and analytical support on antitrust and labor & employment matters in a variety of industries including technology and financial services. He previously served as an Analyst at NERA, and a Summer Research Associate at the World Bank. He holds a B.S. in Economics from Duke University. He may be reached at doktay@outtengolden.com.

[**] Mona Birjandi, Dr. Mona Birjandi is the Principal Economist and Director of Data Analytics at Outten & Golden LLP. She leads the firm's economics and data analytics team, supporting employment litigation through statistical analysis, economic modeling, and data-driven insights. Before joining the firm, Dr. Birjandi worked in both economic consulting and corporate analytics. She has advised on matters involving antitrust, intellectual property, and labor and employment, and previously served as Principal Economist at Petco Health and Wellness. She holds a Ph.D. and M.A. in Economics from Yale University and an M.A. in Law and Diplomacy from Tufts University. She may be reached at mbirjandi@outtengolden.com.

I. Abstract

We analyze how the emergence of generative AI tools correlates with changes in LinkedIn employee posting behavior from 2019 through 2025. Using a novel dataset of daily LinkedIn posts by individual users worldwide, we document sharp increases in posting activity following major generative AI releases – notably the public release of OpenAI's ChatGPT on November 30, 2022 – and persistent growth through subsequent releases (GPT-4 launch in early 2023, Google's Gemini in late 2023, OpenAI's GPT-4o in mid-2024, and OpenAI o1 in late 2024). Descriptive statistics reveal an order-of-magnitude rise in the average posts per employee between 2019 and 2025.

To examine how communication itself changes, we analyze the textual content of posts from a random sample of employees at two large technology firms—Apple and Microsoft—using natural language processing methods. We find that while basic length measures (words, tokens, sentences) exhibit gradual upward trends that predate generative AI, features related to structure, formatting, and style show sharper and more discrete shifts after 2022. In particular, posts become more segmented, organized, and expressive, with substantial increases in paragraph breaks, bullet-point usage, and emojis. Lexical diversity rises more modestly and primarily during later waves of AI diffusion.

Regression evidence confirms that these post-2022 changes persist after controlling for individual characteristics, job roles, departments, seniority, company affiliation, and time effects. Overall, our findings suggest that generative AI has lowered the cost of professional self-presentation and altered the form and style of online labor market signaling. While these changes may democratize visibility and confidence in communication, they also raise questions about the informational content of professional signals. We conclude by discussing implications for labor market signaling and outlining directions for future research on the causal and distributional effects of AI-assisted communication.

II. Introduction

Advances in generative AI are transforming how professionals communicate and signal their skills online. In late 2022, OpenAI's ChatGPT was released, attracting over 100 million users within a few months of launch resulting in a wave of AI content being generated online.[1] Tools like ChatGPT allow users to instantly generate or refine written content, from emails to social media posts, with unprecedented ease. On LinkedIn – the world's largest professional networking platform – this technological shock has coincided with a surge in content creation.[2] One domain undergoing particularly notable change is the process of job search and career entry. Job seekers, especially early-career entrants, are increasingly turning to AI assistance to enhance their professional self-presentation.[3] Many now use tools like ChatGPT to polish résumés, craft cover letters, refine LinkedIn profiles, and even generate ideas for LinkedIn posts.[4,5] In 2025, LinkedIn reported a 45% year-over-year surge in job applications, a spike widely attributed to generative AI making it easier to create application materials.[6] These trends suggest that AI tools have rapidly permeated professional self-presentation, potentially altering the volume, style, and credibility of signals that workers send to the labor market. A candidate with modest writing skills or limited confidence can now produce a clear, well-structured cover letter or a compelling LinkedIn post by leveraging AI's help.

From a labor economics perspective, these developments raise important questions about signaling and skill communication. Classic job-market signaling theory posits that individuals convey their productivity or ability to employers through observable credentials or behaviors (e.g. education, certifications).[7] Active participation on a professional platform like LinkedIn – by posting industry insights, project updates, or thought leadership content – can likewise serve as a signal of one's skills, motivations, and social capital. However, the advent of AI-assisted writing complicates this signaling mechanism. On one hand, generative

AI can empower individuals to communicate with greater clarity and confidence, potentially leveling the playing field for those with weaker writing backgrounds or for non-native English speakers. By dramatically reducing the time and effort required to produce a polished post, AI tools enable more professionals to regularly broadcast their achievements and ideas. On the other hand, generative AI may weaken the informational content of professional communications. When chatbots can produce polished and articulate posts for a wide range of users, employers may find it more difficult to distinguish underlying ability, effort, and skill from AI-assisted presentation. In other words, generative AI tools such as ChatGPT reduce barriers by enabling clearer, more confident expression, but may also blur signals of underlying ability.

Empirical evidence on how generative AI is changing worker behavior remains limited. This paper provides early descriptive evidence by examining LinkedIn employee posting patterns from 2019 to 2025, a period encompassing the introduction of several major AI models. We leverage a novel dataset of daily LinkedIn posts worldwide to address two primary questions: (1) Have employees begun posting more frequently since the arrival of major generative AI tools (particularly after the release of ChatGPT in late 2022)? (2) How has the content of posts changed in the AI era, particularly with respect to linguistic complexity, elaboration, and other textual characteristics associated with professional self-presentation? We focus on several metrics – and track their evolution around key AI milestones: the launch of ChatGPT, GPT-4 launch (March 14, 2023), Google's Gemini model (Dec 6, 2023), and GPT-40 launch (May 13, 2024) and OpenAI o1 Launch (Sep 12, 2024) releases.

Our analysis yields three main findings. First, LinkedIn posting activity exhibits a strong upward trend over 2019–2025, with especially pronounced growth following late 2022. The data show a clear inflection after ChatGPT's introduction: daily post volumes in

2023–2024 are substantially higher than pre-ChatGPT baselines, and the average posts per user roughly doubled in the year after ChatGPT was launched (compared to modest growth in prior years). Further jumps or accelerations are observed around late 2023 and early 2025, suggesting that each new wave of AI capability kept fueling engagement.

Second, the content of LinkedIn posts changes systematically in the AI era. While basic length measures such as words, tokens, and sentences exhibit gradual upward trends that predate generative AI, features related to post structure, formatting, and stylistic expression display sharper and more discrete changes following the introduction of ChatGPT and subsequent AI releases. In particular, posts become more segmented and organized, with marked increases in paragraph breaks, bullet-point usage, and expressive elements such as emojis, indicating a shift toward clearer, more polished, and more engagement-oriented professional communication.

Third, regression evidence confirms that these post-2022 changes persist even after controlling for individual characteristics, job roles, departments, and time effects. Posts written after ChatGPT's introduction are significantly longer, more structured, and more feature-rich, with higher usage of paragraphs, bullets, and emojis, while lexical diversity shows weaker and less immediate responses. These patterns are strongest for structural and stylistic features rather than raw length, suggesting that generative AI is more closely associated with how professionals write than with how much they write.

Taken together, these descriptive results, while not establishing causality, are consistent with the interpretation that generative AI has lowered the cost of professional self-presentation and reshaped the form and style of online labor market signaling, encouraging broader participation while altering the nature of the signals workers send.

The remainder of this paper is organized as follows. Section III provides a brief review of related literature on labor market signaling, communication technology, and the

integration of AI in professional writing. Section IV describes the data and key variables. Section V presents the descriptive analysis of LinkedIn posting trends and their alignment with AI release events. Section VI provides econometric analysis of LinkedIn posting trends and Section VII outlines study limitations and direction for future research, including econometric strategies to identify the causal impact of AI on communication behavior. Section IX concludes.

III. Literature Review

Labor Market Signaling and Communication: The concept of labor market signaling goes back to Spence (1973), who argued that individuals undertake certain costly actions (such as acquiring education) to signal their productivity to prospective employers under conditions of asymmetric information.[8] While formal credentials and observable characteristics have long been the focus of this theory, later work recognizes that many forms of communication and self-presentation can function as signals of underlying ability or ambition.[9] In modern labor markets, professional social media activity is increasingly viewed as part of one's signaling portfolio.[10] By sharing industry insights, accomplishments, or thoughtful commentary on platforms like LinkedIn, workers can signal their expertise, communication skills, and engagement to a broad audience of colleagues and recruiters. Empirical studies show that active LinkedIn users report greater informational benefits – such as timely access to job openings or career advice – than non-users.[11] This aligns with broader findings that communication and networking skills are growing in importance for labor market success.[12]

Generative AI as a Communication Tool: The emergence of large language models (LLMs) and AI writing assistants offers a potential game-changer in professional communication. Tools like ChatGPT can produce well-structured paragraphs, suggest persuasive phrasings, and correct grammatical errors, all within seconds. Early evidence

indicates that such AI assistance can significantly boost productivity on writing tasks: a controlled experiment by Noy and Zhang (2023) found that access to ChatGPT cut the time to complete business writing assignments by 40% and improved the quality by 18%, on average.[13] Notably, these improvements were particularly pronounced for participants with weaker writing skills, hinting that AI tools may help lower-ability or less-experienced workers close the gap in communication quality. In the context of LinkedIn, this means employees who previously struggled to write professional posts can now do so more readily, potentially increasing their visibility. In short, generative AI has dramatically lowered the cost of producing professional communications, effectively automating some aspects of personal branding and thought leadership.

Authenticity and Signal Dilution: A growing discourse accompanies the rise of AI-generated content, centered on whether these technological aids undermine authenticity in professional signaling.[14] Traditionally, a well-crafted LinkedIn post or blog could serve as a credible signal of a professional's insight and communication ability, precisely because it required effort and skill to produce. Now, however, the credibility of that signal is in question if a significant portion of content is written by AI. Potentially, AI might make LinkedIn posts sound formulaic and indistinguishable once users began relying on the same AI models for writing assistance, and therefore the content might lose its signaling value.[15] In economic terms, AI potentially can threaten the reliability of the signal, giving polished communication less informational value.[16] This concern echoes a key issue in signaling theory: for a signal to convey information, it must be easier for high-quality agents to produce than for low-quality agents.

Generative AI flips this dynamic by enabling even those with limited subject-matter knowledge or communication skill to generate articulate prose. Thus, while AI-aided posting may increase overall information sharing, it also poses a challenge for employers and peers

trying to discern genuine expertise and effort. The net effect on labor market efficiency is ambiguous. At the market-design level, a new study by Galdin and Silbert (2025) explicitly models and measures how LLMs can "make talk cheap," disrupting settings where writing previously served as a costly signal and changing equilibrium screening and selection.[17]

In summary, the literature and emerging evidence highlight a double-edged sword. Generative AI stands to democratize professional communication, giving a boost to those who historically struggled with self-expression, and potentially enhancing the diversity of voices and ideas shared in professional networks. At the same time, it risks diluting the informational content of communications as signals, complicating traditional hiring and networking processes. These considerations motivate a closer empirical look at how worker behavior on platforms like LinkedIn is evolving in the wake of AI's arrival. By focusing on descriptive patterns of posting behavior around AI milestone dates, this paper takes an initial step toward understanding the scope of these changes before moving on to identify their consequences

IV. Data

For our analysis, we compiled publicly available LinkedIn profile and post data from January 1, 2019, through mid-2025 (~500 million posts). Our primary dataset is constructed from LinkedIn's platform data and includes the daily count of posts made by individual users (employees) globally. We focus on employee posts, or content posted by individual member accounts (as opposed to company pages or paid advertisements). Based on this global data, for each day in the sample, we observe the total number of employee-generated posts. We also compute a normalized measure – the number of posts per employee –as a proxy for an individual's posting frequency on a given day.

We focus on five major generative AI release events that represent notable shocks to professional communication costs and content production. Specifically, we mark:

- ChatGPT's release on Nov 30, 2022 is treated as a major external shock in our analysis timeline.[18]
- GPT-4 on March 14, 2023,[19]
- Google Gemini launch on December 6, 2023, [20]
- GPT-4o on May 13, 2024, [21]
- OpenAI o1 launch on September 12, 2024[22]

These events were selected because they represent widely publicized inflection points in the accessibility, capability, and adoption of large language models, and each coincided with substantial increases in public and professional use of generative AI tools.[1]

In addition to analyzing aggregate LinkedIn posting activity, we evaluate changes in post content using a detailed subsample drawn from two large technology firms, Apple and Microsoft. From each company, we randomly selected 10,000 current U.S.-based employees, yielding a combined sample of 20,000 individuals. All sampled individuals were employees of their respective employees as of the end of 2025, ensuring that we capture recent activity and that each individual has an active LinkedIn profile with up-to-date information. For each employee, we collected comprehensive profile attributes—including education history, employment history, listed skills, spoken languages, and total years of experience—as reported on LinkedIn. We then obtained the complete history of publicly available LinkedIn posts made by these employees between January 2017 and December 2025. This resulted in a longitudinal panel containing 8,896 posts for Apple employees and 17,433 posts for Microsoft employees (approximately one post per employee on average). For each post, we also recorded metadata on the posting date and engagement outcomes, including the number of likes and comments received.

[1] LinkedIn itself introduced an AI-assisted post drafting feature for its Premium users in late 2023 (around the same time as Gemini's launch).

Taken together, these data provide both a high-level view of global LinkedIn posting dynamics over a six-year period and a detailed longitudinal record of individual posting behavior before and after the widespread adoption of generative AI. By aligning these time-series data with key AI innovation dates, we can document how the frequency and linguistic complexity of professional activity on online labor market platforms such as LinkedIn evolved in the age of generative AI.

V. Descriptive Analysis

1) Changes in the Volume of Professional Self-Presentation on LinkedIn

We begin by plotting the daily volume of LinkedIn employee posts from 2019 through the end of 2024, highlighting both the trend over time and the day-of-week variation in posts. We then examine the normalized metric of posts per employee per day to understand changes in individual posting propensity. Key generative AI release dates are indicated on the graphs to assess temporal alignment between those milestones and any shifts in posting behavior.

Figure 1 shows the total number of LinkedIn employee posts per day. The time series exhibits a strong upward trajectory over the sample period, but the growth is characterized by distinct accelerations rather than a smooth linear trend. In particular, the series displays clear step increases around major generative AI release dates (highlighted by the red ovals). LinkedIn posting was growing gradually, roughly tracking the platform's expanding user base. Immediately after ChatGPT's public release, we observe a noticeable jump in the level of daily posts. Following the public launch of ChatGPT in late 2022, daily posting levels shift upward sharply, with post counts remaining persistently higher than in the pre-ChatGPT period. A second pronounced level increase occurs in late 2023 to early 2024, coinciding with the public launch of Google's Gemini model and LinkedIn's rollout of AI-assisted posting

tools. A further acceleration is visible in late 2024 and early 2025, around subsequent OpenAI model releases, marked by another discrete upward shift in posting activity. By 2025, daily employee post volumes are approximately five to ten times higher than their pre-ChatGPT levels. Overall, the figure suggests that the diffusion of generative AI is associated not only with faster growth but with repeated upward level shifts in professional content creation.

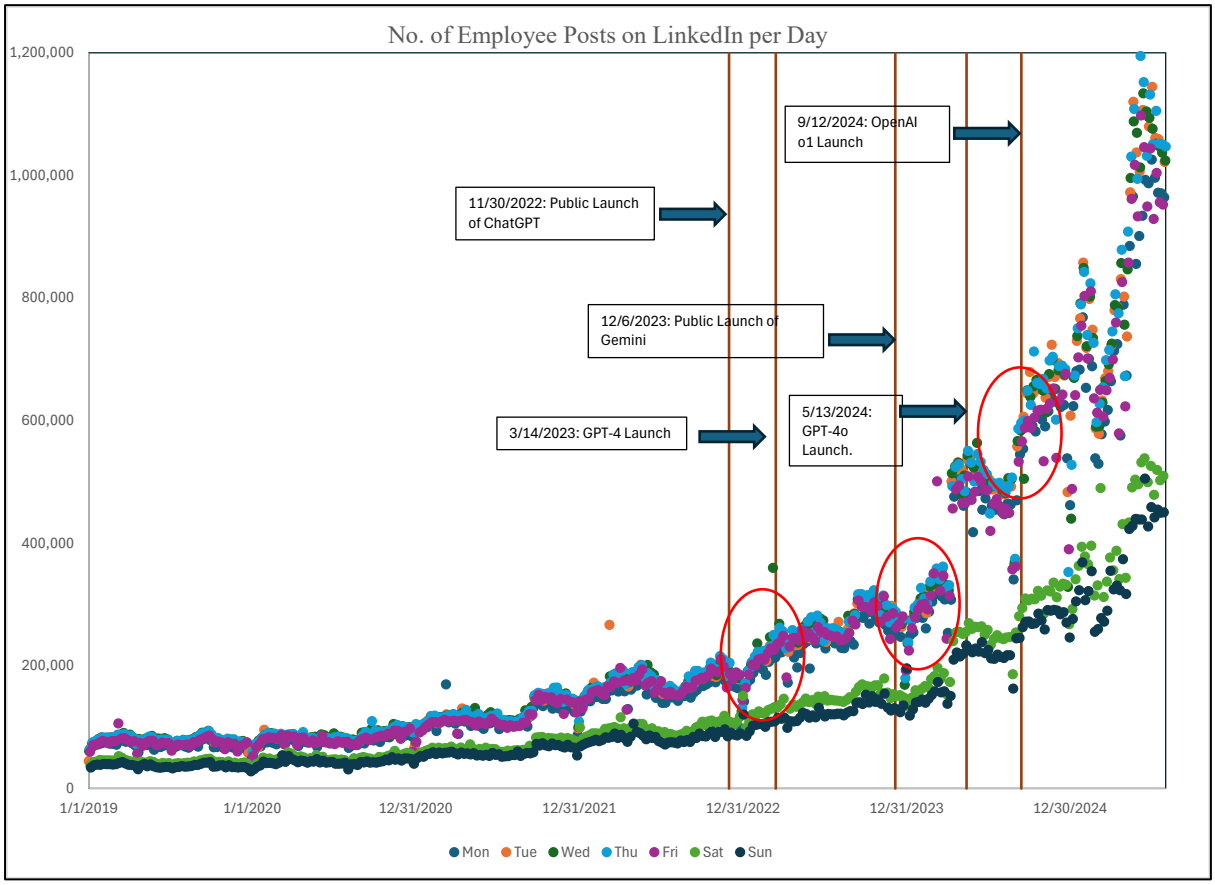


Figure 1-Total Number of LinkedIn Employee Posts per Day, 2019–2025 (color-coded by weekday). Four vertical reference lines indicate the dates of major generative AI releases.

Several features of the data are worth noting. First, the LinkedIn network grew substantially during 2019–2025: the platform's membership roughly doubled from about 643 million users in 2019 to 1.2 billion by 2025.[23] Monthly active user counts also rose, reaching an estimated 310 million active users in 2025.[24] This expansion in the user base is an important backdrop for interpreting raw post counts – increasing total posts could simply

reflect more people on the platform. Therefore, we also examine posts *per user*, to separate pure membership growth from changes in individual behavior. Figure 2 below normalizes the posting volume by LinkedIn's user base to track how the *per user* posting rate has changed. We observe that posts per employee per day also exhibit a strong upward trend, with much of this increase occuring after 2022. In particular, the figure shows discrete upward level shifts in per-user posting intensity following the public release of ChatGPT, with posting rates remaining persistently higher thereafter. Additional accelerations are visible around late 2023 and again in late 2024 to early 2025, coinciding with subsequent major generative AI releases. These patterns mirror those observed Figure 1 for aggregate post counts and suggest that the rise in LinkedIn activity is not solely driven by platform growth. Instead, even after accounting for the expanding user base, individual LinkedIn members became substantially more active content creators with introduction of generative AI tools.[2]

[2] Note that the observed increase in posting activity following major generative AI releases could in principle reflect short-run attention or discussion effects (e.g., employees posting about the new technology itself). However, if this were the primary mechanism, one would expect posting levels to revert toward pre-release trends as interest faded. Instead, posting activity remains persistently elevated after each major AI release, consistent with a sustained change in posting behavior rather than a transitory topic-driven effect.

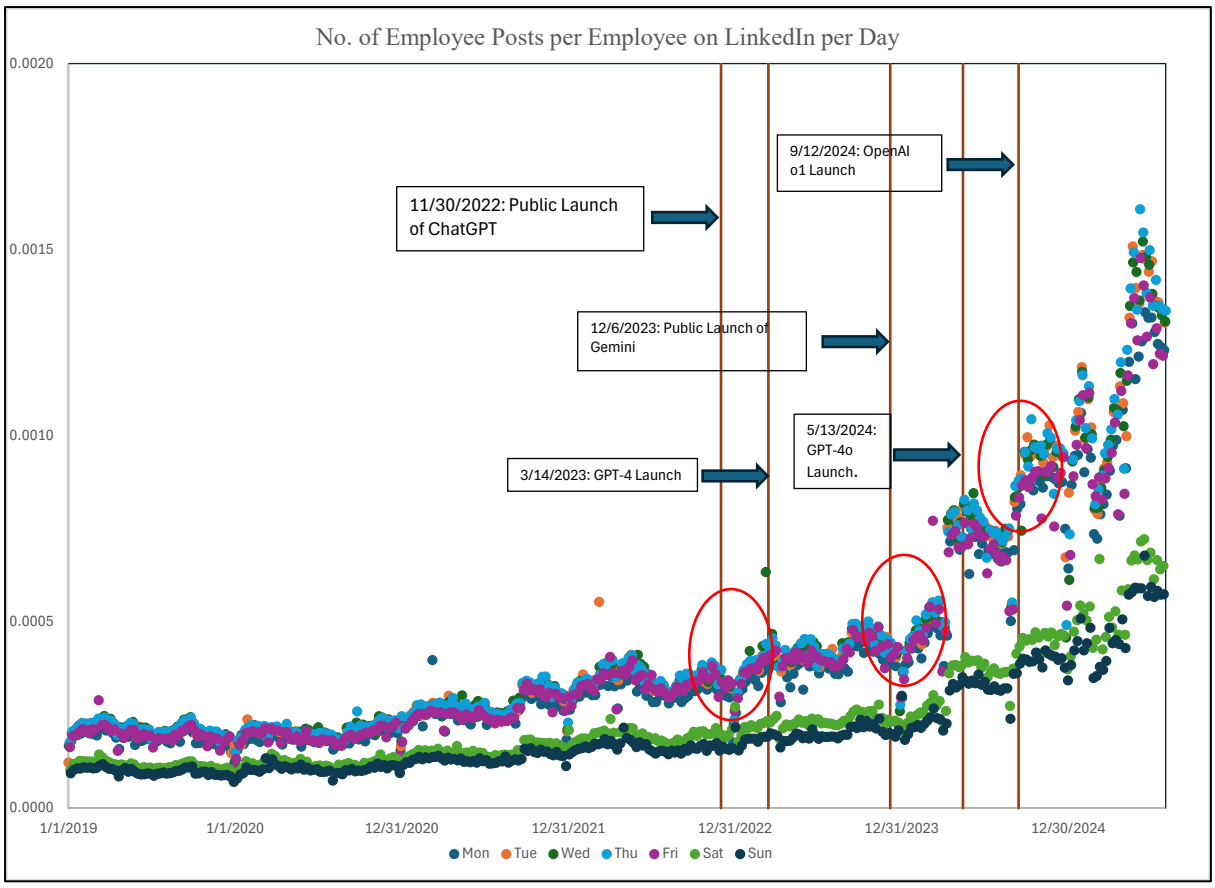


Figure 2- Average Number of LinkedIn Posts Per Employee Per Day, 2019 - 2025 (color-coded by weekday). Four vertical reference lines indicate the dates of major generative AI releases.

2) Changes in Length, Structure, and Style of Professional Communication

In addition to the raw post counts and engagement metrics, we analyze the textual content of historic LinkedIn posts from a random sample of 10,000 employees each from two large technology firms— Apple and Microsoft—to extract a range of linguistic and structural features using natural language processing (NLP) methods. In particular, we consider the following linguistic measures: number of words, number of sentences, number of paragraphs, number of tokens, whether the post contains bullets, number of emojis, and MTLDin a given post.[3]

[3] As a part of this analysis, Python packages Pandas, Spacy, Textstat, Matplotlib, Scikit-learn, and Statsmodels were used.

Figure 3 synthesizes the evolution of multiple linguistic and structural characteristics of LinkedIn employee posts for Apple and Microsoft employees between 2017 and 2026. Taken together, the panels document a broad and systematic transformation in how professionals communicate on LinkedIn.

Across all length-based measures—words, tokens,[4] and sentences—posts exhibit a clear upward trend over the full sample period. These increases begin well before the introduction of widely accessible generative AI tools, indicating a gradual, long-run shift toward more verbose and elaborated professional communication. Following the public release of ChatGPT in late 2022, these metrics remain elevated and continue to rise, with additional incremental increases around subsequent AI model releases in 2024 and 2025. However, for these raw length measures, the post-2022 changes largely build on pre-existing trends rather than reflecting sharp structural breaks, suggesting that generative AI may have reinforced or accelerated an ongoing evolution toward longer posts rather than initiating it outright.

In contrast, metrics capturing post structure and presentation—notably paragraph counts and bullet usage—display more pronounced and discrete changes around major AI milestones. While paragraph counts rise gradually prior to 2022, they increase sharply and persistently in the post-ChatGPT period, with especially large step increases during later AI release waves. A similar pattern emerges for bullet-point usage, which rises markedly after main AI waves, remaining at substantially higher levels through 2024–2025. These shifts suggest a meaningful reorganization of posts toward more segmented, list-like, and visually

[4] In natural language processing, a word is a linguistic unit as conventionally defined in human language (typically delimited by spaces or punctuation). A token is a unit produced by a model's tokenization process and may correspond to a full word, part of a word (subword fragment), punctuation, or special symbols. Token boundaries are model- and tokenizer-dependent and do not necessarily align with word boundaries. For example, the single word "unbelievable" may be represented as multiple tokens such as "un", "believ", "able" under a subword-based tokenizer.

structured formats, consistent with AI-assisted drafting tools that emphasize clarity and professional polish.

Stylistic features exhibit similarly distinct post-2022 changes. Emoji usage—largely limited prior to 2020—rises gradually through the early 2020s and then increases sharply following the diffusion of generative AI, with additional discrete jumps during later model releases. By 2024–2025, emoji usage is several times higher than pre-2022 levels, particularly among Microsoft employees. This pattern points to a shift in tone toward more expressive, conversational, and engagement-oriented communication, even within a professional context.

Finally, lexical diversity, measured by MTLD,[5] shows a more nuanced evolution. While vocabulary richness fluctuates over time and trends upward modestly over the long run, the clearest increases occur during later waves of generative AI diffusion rather than immediately after ChatGPT's release. In 2024–2025, MTLD reaches persistently higher levels, especially for Microsoft employees, indicating less repetitive and more varied word choice. This suggests that improvements in AI model sophistication may be increasingly reflected not just in post length or structure, but in the composition and diversity of language itself.

Taken together, these patterns reveal a consistent hierarchy of change. Length-based metrics evolve gradually over time, while structure, formatting, expressiveness, and lexical diversity exhibit sharper and more temporally aligned shifts following the diffusion of generative AI tools. The evidence points to a broad transformation in professional self-

[5] MTLD is a metric of lexical diversity designed to be relatively robust to text length. A text is processed sequentially and divided into segments ("factors") such that the type–token ratio (TTR) of each segment reaches a fixed threshold τ (commonly $\tau = 0.72$). MTLD is defined as

$$\text{MTLD} = \frac{N}{F},$$

where N is the total number of tokens in the text and F is the total number of factors, including any partial final factor (counted proportionally). Higher MTLD values indicate greater lexical diversity.

presentation: employees are not only posting more frequently and writing longer posts, but are increasingly producing content that is more structured, stylistically polished, expressive, and lexically varied. These changes are most pronounced in the post-2022 period and intensify with later AI model releases, suggesting that generative AI is more strongly associated with *how* professionals write than with *how much* they write.

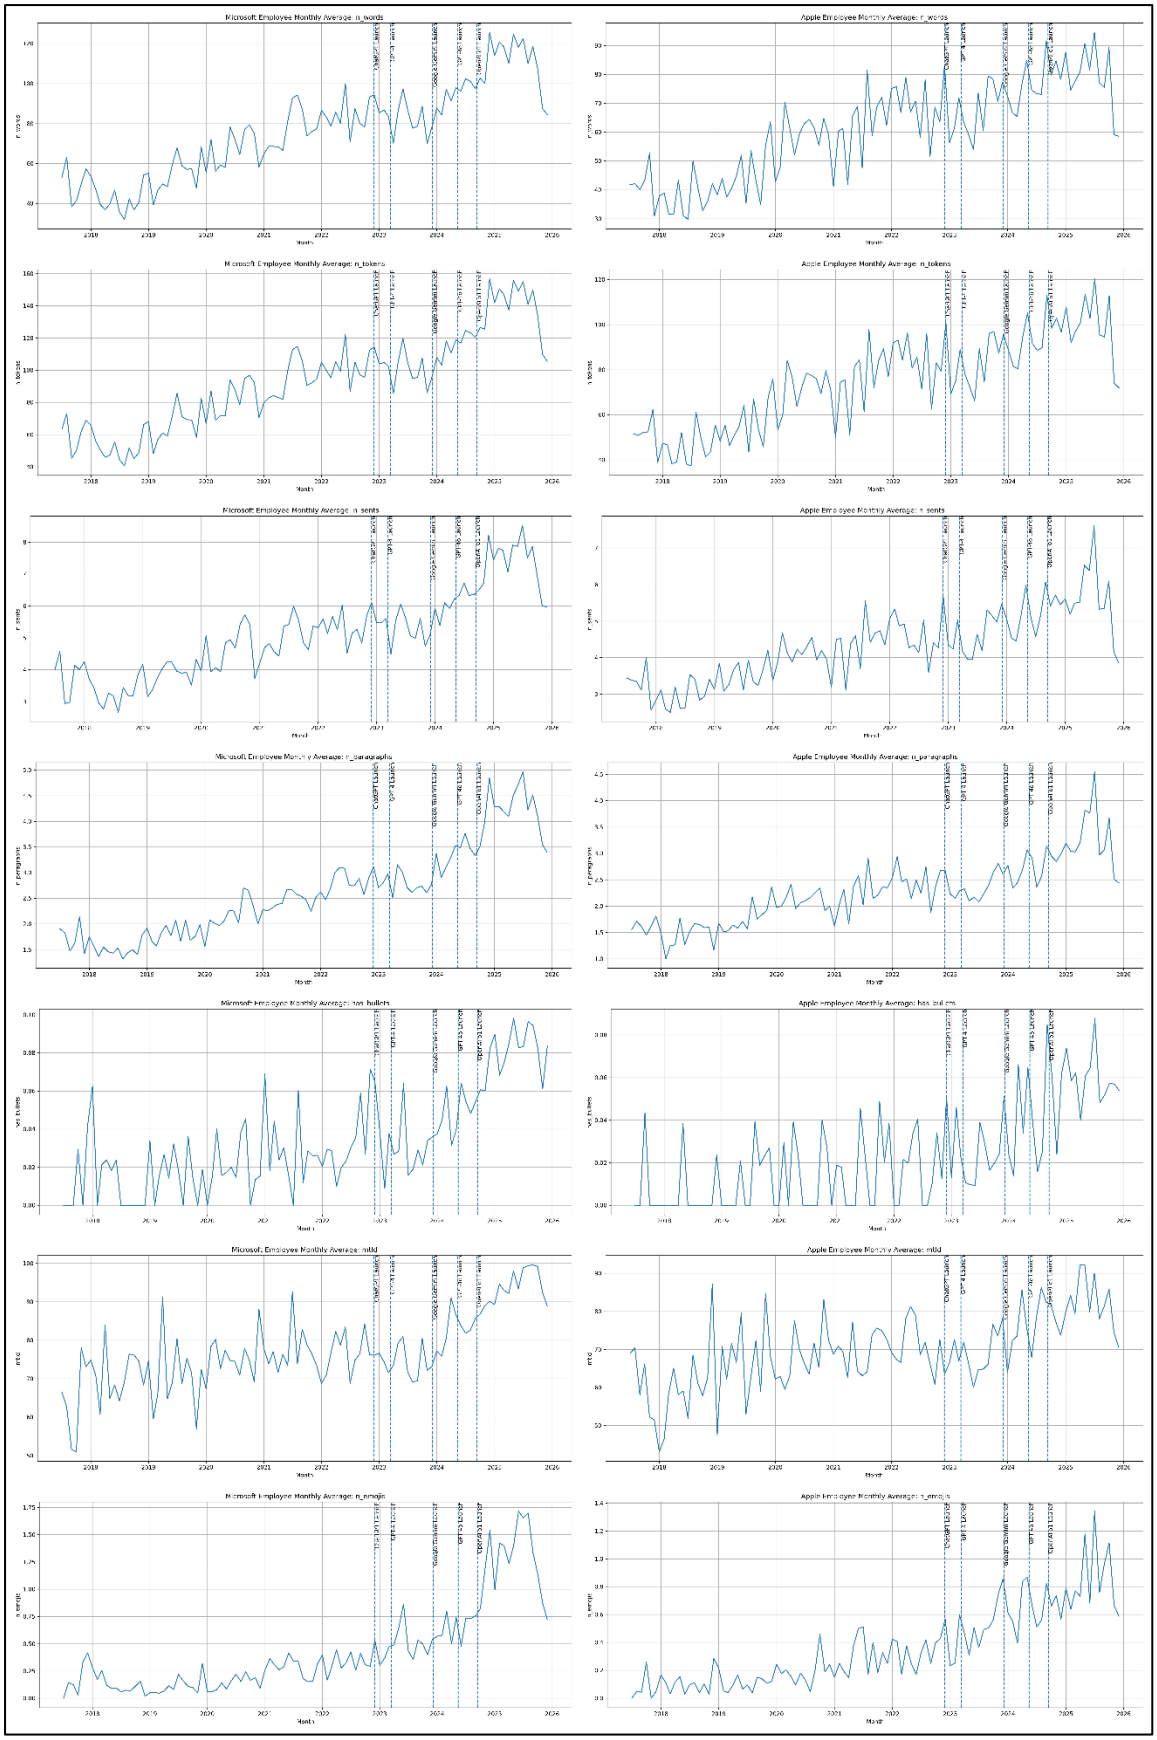


Figure 3 - Monthly Content Metrics in LinkedIn employee postings for Apple (Panel A) and Microsoft (Panel B), 2017-2026. Vertical dashed lines indicate major AI model launch dates

VI. Econometric Analysis

We estimated a series of OLS regressions to examine how LinkedIn post characteristics changed after the introduction of ChatGPT. Each column of the regression table corresponds to a different textual outcome variable: (1) MTLD, (2) number of words, (3) number of tokens, (4) number of sentences, (5) number of paragraphs, (6) number of emojis, and (7) a binary indicator for whether the post contains bullet points. The key independent variable is a dummy ChatGPT introduced that equals 1 for posts made after ChatGPT's introduction on November 30, 2022 (and 0 for posts before). We also include the poster's total work experience in months (total_experience_duration_months) as a control for seniority. All regressions incorporate a rich set of fixed effects: day-of-week fixed effects, year-month fixed effects, department fixed effects (based on LinkedIn-defined departments such as Engineering, Marketing, etc.), management level fixed effects (e.g. Specialist, Manager, Director, C-level, etc.), and company fixed effects (indicators for Microsoft vs. Apple). We use the same sample of posts from the previous analysis, with 20,073 observations in total after some observations are dropped due to missing values for certain variables.

The regression results show notable differences in post content after ChatGPT's introduction, even after controlling for all the above factors. In particular, the coefficient on ChatGPT time dummy is positive for all seven outcome variables and statistically significant for most of them. Posts written after ChatGPT's release tend to be longer and more detailed: they contain roughly 31 more words and 40 more tokens on average than posts before ChatGPT (both effects significant at the 5% level). These posts also have about 1.63 additional sentences (significant at the 10% level) and 1.16 more paragraphs ($p < 0.05$), suggesting that post-ChatGPT content is broken into more sentences and paragraphs, potentially indicating richer or more complex narratives. We also observe a substantial

increase in the use of expressive or structured elements: after ChatGPT's introduction, posts include about 0.59 more emojis on average ($p < 0.01$) and are more likely to contain bullet lists (the probability of having bullet points is higher by roughly 4.6 percentage points, $p < 0.01$). The coefficient on MTLD is positive but not statistically distinguishable from zero, implying no strong evidence of a lexical diversity change. Meanwhile, the poster's work experience in months is negatively associated with all these post characteristics. Although the magnitudes are small, the coefficients on total_experience_duration_months are significantly negative ($p < 0.01$) in every column, indicating that more experienced individuals tend to write slightly shorter, more succinct posts with fewer emojis and bullet points.

In summary, after accounting for fixed effects, our regressions suggest that the introduction of ChatGPT correlates with longer and more feature-rich LinkedIn posts, consistent with the evidence we documented using descriptive analytics. Importantly, this "post-ChatGPT" effect should not be interpreted as the effect of ChatGPT alone: because later model launches and broader diffusion (e.g., subsequent OpenAI releases and competing systems) occur within the post-period, the estimated shift is most plausibly the cumulative effect of multiple generative AI introductions and adoption waves rather than a single product launch.

Table 1 - Regression Analysis of Content Metrics in LinkedIn employee postings for Apple and Microsoft, 2017-2026.

	Model 1: *MTLD*	**Model 2:** *No. of Words*	**Model 3:** *No. of Tokens*	**Model 4:** *No. of Sentences*	**Model 5:** *No. of Paragraphs*	**Model 6:** *No. of Emojis*	**Model 7:** *Has Bullets*
ChatGPT Time Dummy	6.87	31.34**	39.81**	1.63*	1.16**	0.59***	0.05***
	(9.62)	(14.42)	(17.50)	(0.87)	(0.46)	(0.06)	(0.01)
Experience in Months	-0.0222***	-0.0817***	-0.1008***	-0.0052***	-0.0034***	-0.0015***	-0.0001***
	(0.0039)	(0.0059)	(0.0072)	(0.0004)	(0.0002)	(0.0001)	(0.0000)
Day FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Year-Month FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Department FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Management level FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Company FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
N	20,073	20,073	20,073	20,073	20,073	20,073	20,073
R^2	0.05	0.103	0.107	0.099	0.117	0.07	0.026

1.The regression data contains all historic LinkedIn posts going back to July 24, 2017 for 10,000 employees at Microsoft, and 10,000 employees at Apple.
2.Experience in Month is defined as total work experience (across all companies worked) for the individual poster.
3.Department Fixed Effects control for the poster's active department at their respective company, defined as one of Engineering and Technical, Customer Service, Other, Research, C-Suite, Sales, Design, Finance & Accounting, Project Management, Product, Operations, Administrative, Marketing, Education, Consulting, Legal, Human Resources, Medical, Real Estate, or Trades
4.Management level Fixed Effects control for the poster's active management level at their respective company, defined as one of Specialist, Senior, Manager, President/Vice President, Director, C-Level, Founder, Partner, Intern, Owner, or Head.
5.Company Fixed Effects control for the poster's current company, either Apple or Microsoft.

VII. Limitations and Future Directions

This study has several limitations that are important for interpreting the results and for guiding future research. First, our analysis relies on publicly available LinkedIn data, which captures only observable online self-presentation and not underlying skill, effort, or offline job-search behavior. Posting behavior is inherently selective: linguistic outcomes are observed only for individuals who choose to post, and generative AI may affect both the likelihood of posting (the extensive margin) and the content of posts (the intensive margin). While we partially address this by analyzing posting frequency and conditioning on posts, future work could more explicitly model selection into posting.

Second, for illustrative purposes, the detailed content analysis focuses on a random sample of employees from two large technology firms, Apple and Microsoft. While this allows for rich individual-level controls and longitudinal analysis, these firms are not representative of the broader labor market. Employees in large technology firms may be earlier adopters of AI tools, more digitally engaged, and more active on LinkedIn than workers in other industries. As a result, the magnitude and nature of the observed changes may differ in non-technology sectors, smaller firms, or less digitally intensive occupations.

Third, our econometric analysis is descriptive and does not establish causal effects of generative AI adoption. Although the timing of changes in posting behavior and content aligns closely with major AI releases, other concurrent factors may confound these patterns. In particular, LinkedIn introduced its own AI-assisted writing features, algorithmic feed changes, and user-interface updates during the study period, any of which could independently affect posting behavior or content structure. Broader macroeconomic trends, shifts toward remote work, and evolving norms around professional self-presentation may also contribute to the observed changes. While our regressions include rich fixed effects to

absorb many confounders, unobserved platform-level or time-varying shocks remain a limitation.

These limitations point to several promising directions for future research. First, expanding the analysis across geographies would allow researchers to exploit cross-country variation in AI adoption, language, and labor market institutions to better isolate AI-related effects. Second, extending the sample to additional industries and occupations would help assess whether AI-driven changes in professional communication are concentrated in knowledge-intensive or communication-heavy roles, or whether they generalize across the workforce. Third, future work could examine heterogeneity by worker characteristics such as gender, native language, education, or career stage, shedding light on whether AI tools reduce or exacerbate existing communication and visibility gaps. Finally, linking changes in online self-presentation to downstream outcomes—such as engagement, recruiter contact, job transitions, or wage growth—would be crucial for understanding the welfare and labor market implications of AI-assisted communication.

Overall, this paper provides an initial empirical foundation for studying how generative AI reshapes professional self-presentation. Addressing these limitations through richer data and causal designs represents an important next step for understanding the broader labor market consequences of AI-driven communication technologies.

VIII. Concluding Remarks

This paper provides early empirical evidence on how generative AI is reshaping professional self-presentation and communication in online labor market platforms. Using global LinkedIn posting data and detailed content analysis of employee posts from two large technology firms, we document a pronounced increase in both the frequency and form of professional communication following the introduction of widely accessible generative AI tools.

Our findings highlight a clear distinction between how much professionals write and how they write. Measures of post length increase gradually over time and appear to reflect longer-run trends in online professional communication. In contrast, structural and stylistic features—such as paragraph segmentation, bullet-point usage, emojis, and, to a lesser extent, lexical diversity—exhibit sharper and more temporally aligned shifts following major AI releases. Regression results confirm that these changes persist after accounting for individual heterogeneity, job characteristics, and time effects. Taken together, the evidence suggests that generative AI is more strongly associated with changes in the organization, tone, and presentation of professional content than with raw verbosity alone.

These patterns have important implications for labor market signaling. By lowering the cost of producing polished, structured, and expressive communication, generative AI may enable broader participation in professional self-presentation, particularly for individuals with weaker writing skills or lower confidence. At the same time, the widespread adoption of AI-assisted communication may reduce the informational content of professional signals, complicating employers' ability to infer underlying ability or effort from observable online behavior. In this sense, generative AI both democratizes access to visibility and challenges traditional signaling mechanisms.

While our analysis is descriptive and does not establish causal effects, it provides a foundation for understanding how AI technologies are transforming the early stages of job search, networking, and professional identity formation. Future work that exploits variation in AI adoption across countries, industries, occupations, and worker characteristics—and that links changes in online communication to downstream labor market outcomes—will be essential for assessing the broader economic consequences of AI-assisted self-presentation. As generative AI continues to evolve, understanding how it reshapes communication-based

signals will be critical for workers, employers, and policymakers navigating the future of work.

---

[1] OpenAI, *ChatGPT Usage and Adoption Patterns at Work* (PDF), https://cdn.openai.com/pdf/3c7f7e1b-36c4-446b-916c-11183e4266b7/chatgpt-usage-and-adoption-patterns-at-work.pdf.

[2] Singh, Jaspreet, LinkedIn Deepens Video Ad Push, Taps More Publishers and Creators to Spur Growth, Reuters (Aug. 25, 2025), https://www.reuters.com/business/media-telecom/linkedin-deepens-video-ad-push-taps-more-publishers-creators-spur-growth-2025-08-25.

[3] Anne Griffin, *How Job Seekers Should Use AI To Get A New Job In 2026*, *Forbes* (Jan. 12, 2026), https://www.forbes.com/sites/annegriffin/2026/01/12/how-job-seekers-should-use-ai-to-get-a-new-job-in-2026.

[4] *Id.*

[5] https://www.nber.org/system/files/working_papers/w30886/w30886.pdf

[6] Sarah Kessler, *Employers Are Buried in A.I.-Generated Job Applications*, *N.Y. Times* (June 21, 2025), https://www.nytimes.com/2025/06/21/business/dealbook/ai-job-applications.html.

[7] Michael Spence, Job Market Signaling, 87 Q.J. Econ. 355 (1973), https://doi.org/10.2307/1882010.

[8] *Id.*

[9] https://www.aeaweb.org/articles?id=10.1257%2Fjep.9.4.133.

[10] Matthew Baird et al., *Skill Signals in a Digital Job Search Market and Duration in Employment Gaps*, *J. of Labor Research* (Oct. 16, 2024), https://doi.org/10.1007/s12122-024-09363-y (finding that listed LinkedIn skills act as signals of competencies that affect labor market outcomes).

[11] Sonja Utz & Johannes Breuer, *The Relationship Between Networking, LinkedIn Use, and Retrieving Informational Benefits*, 22 *Cyberpsychol. Behav. & Soc. Network* 180 (2019), doi:10.1089/cyber.2018.0294.

[12] *The Quarterly Journal of Economics*, Volume 132, Issue 4, November 2017, Pages 1593–1640, https://doi.org/10.1093/qje/qjx022.

[13] Shakked Noy & Whitney Zhang, *Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence*, 381 *Science* 187 (2023), https://doi.org/10.1126/science.adh2586.

[14] C. P. Kirk and J. Givi. *The AI-authorship effect: Understanding authenticity, moral disgust, and consumer responses to AI-generated marketing communications*. Journal of Business Research, 186:114984, 2025.

[15] Originality.AI, *Over ½ of Long Posts on LinkedIn are Likely AI-Generated Since ChatGPT Launched*, Originality.AI (Oct. 28, 2025), https://originality.ai/blog/ai-content-published-linkedin.

[16] F. Draxler, A. Werner, F. Lehmann, M. Hoppe, A. Schmidt, D. Buschek, and R. Welsch. The AI Ghostwriter Effect: *When users do not perceive ownership of AI-generated text but self-declare as authors*. arXiv preprint arXiv:2303.03283, 2023.

[17] Anaïs Galdin & Jesse Silbert, *Making Talk Cheap: Generative AI and Labor Market Signaling*, *arXiv* (Nov. 11, 2025) (working paper), https://arxiv.org/abs/2511.08785.

[18] OpenAI, *Introducing ChatGPT* (Nov. 30, 2022), https://openai.com/index/chatgpt/.

[19] OpenAI, GPT-4 (Mar. 14, 2023) (describing the launch and capabilities of OpenAI's fourth-generation foundation model), https://openai.com/index/gpt-4-research/.

[20] *Shiona McCallum & Zoe Kleinman, Google Claims New Gemini AI "Thinks More Carefully,"* BBC News: Technology (Dec. 6, 2023), https://www.bbc.com/news/technology-67630454.

[21] *Kyle Wiggers, OpenAI Debuts GPT-4o 'Omni' Model Now Powering ChatGPT,* TechCrunch (May 13, 2024), https://techcrunch.com/2024/05/13/openais-newest-model-is-gpt-4o/.

[22] *Cade Metz, OpenAI Unveils New ChatGPT That Can Reason Through Math and Science,* N.Y. Times (Sept. 12, 2024), https://www.nytimes.com/2024/09/12/technology/openai-chatgpt-math.html.

[23] *50 LinkedIn Statistics 2026 (Active Users Data)*, Demand Sage (Jan. 12, 2026), https://www.demandsage.com/linkedin-statistics/ (showing LinkedIn's membership grew from about 643 million in 2019 to approximately 1.20 billion in 2025).

[24] *Id.*