NBER WORKING PAPER SERIES

TIME TRAVEL ON PROFESSIONAL PROFILES

Nicholas Bloom
Gideon Moore
Lisa K. Simon
Caelan Wilkie-Rogers

Working Paper 35546
http://www.nber.org/papers/w35546

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2026

We thank Revelio Labs for access to the data and for assistance in constructing the monthly profile vintages. We thank Stanford Impact Labs for funding. The views expressed here are those of the authors and do not necessarily reflect the views of Revelio Labs, Stanford University, or the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w35546

Time Travel on Professional Profiles
Nicholas Bloom, Gideon Moore, Lisa K. Simon, and Caelan Wilkie-Rogers
NBER Working Paper No. 35546
July 2026
JEL No. J0

ABSTRACT

Economists increasingly use professional profile data to reconstruct employment histories and measure skill supply. We show that these records are not fixed historical snapshots, but mutable accounts that workers revise over time. Using monthly vintages of Revelio Labs data from 2020–2026, we document that 19.2 percent of U.S. LinkedIn users retroactively edit the title or description of a job they have already left. These “time-travel” edits are closely tied to labor market transitions: around such edits, workers are much more likely to change employers as compared to later-editing users. This mutability can bias historical measures of skills, but it also reveals workers’ beliefs about which skills are in demand. Retroactive edits show sharp post-2022 increases in AI-related language and recent reductions in work-from-home and DEI language. Finally, LLM-associated writing markers surge after ChatGPT, especially among less-educated groups and MBAs from lower-ranked programs, revealing heterogeneous AI-assisted profile editing.

Nicholas Bloom
Stanford University
Department of Economics
and NBER
nbloom@stanford.edu

Gideon Moore
Stanford University
Department of Economics
gsmoore@stanford.edu

Lisa K. Simon
Revelio Labs
lisa@reveliolabs.com

Caelan Wilkie-Rogers
Revelio Labs
caelan@reveliolabs.com

1 Introduction

Professional profile data has—rightly—become a linchpin of modern labor market research, particularly using modern NLP methods to analyze job descriptions (Babina et al. 2023; Babina et al. 2024; Hampole et al. 2025). While there is much to learn from these data sources, users should be aware that worker-provided accounts of their employment histories should not be taken as ground truth; workers can change their job descriptions over time in response to market forces. Analyses which neglect this fact can overstate the presence of skills desirable in the present day. We show first that retroactive editing of profiles is common, with one-in-five high-quality American LinkedIn users performing at least one retroactive edit from 2020-2026. Unsurprisingly, these retroactive edits are associated with job transitions. Workers are 66% more likely to change employers in the period they update their professional profile. Thus, we can *harness* these changes to better understand labor market skill demand—if workers are adding a skill to their resume, we take this as a signal they believe it to be desirable on the market. As a proof of concept, we track retroactive changes for three sets of skills—AI, work-from-home, and DEI—from 2021-present. Finally, we use these changes to measure AI adoption over this period harnessing a bag-of-words method from Kobak et al. (2025).

We lead with an illustration of when profile editing might matter. Suppose researchers wanted to estimate the rate AI skills penetrated the labor market. One reasonable first approach could be to plot the rate at which AI skills appear in workers' old job descriptions. However, if workers retroactively add these skills to old job descriptions to look more AI-ready, then the research team would overestimate the speed with which AI spread within the labor market. We have performed this exercise in Figure 1 to highlight the difference between these approaches, tracking the presence of "AI," "gpt," "llm," and "artificial intelligence" over time.[1] We compute an "uncorrected" estimate of AI presence in jobs each year using the 2026 vintage of LinkedIn and compare it to how those jobs appeared on LinkedIn in the relevant year (e.g., 2022 job descriptions for 2022).[2] By our estimates, the uncorrected approach overestimates AI skill presence pre-ChatGPT by around 30%.

[1] Words require break characters on each side; this means "**ai** specialist" will be picked up while "h**ai**l" will not.

[2] Our updates file only goes back to 2020; thus, consider this a lower bound in the disparity as we neglect any retroactive edits from 2019 or earlier.

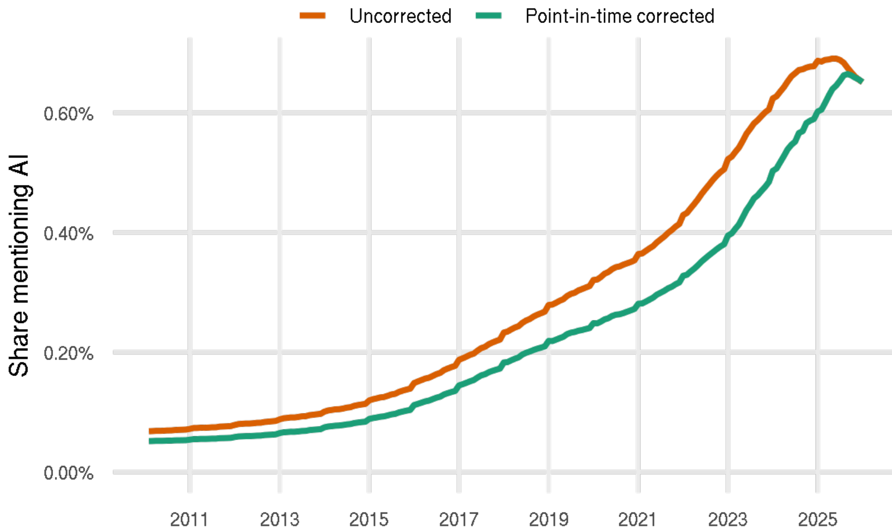

(a) Shares of jobs

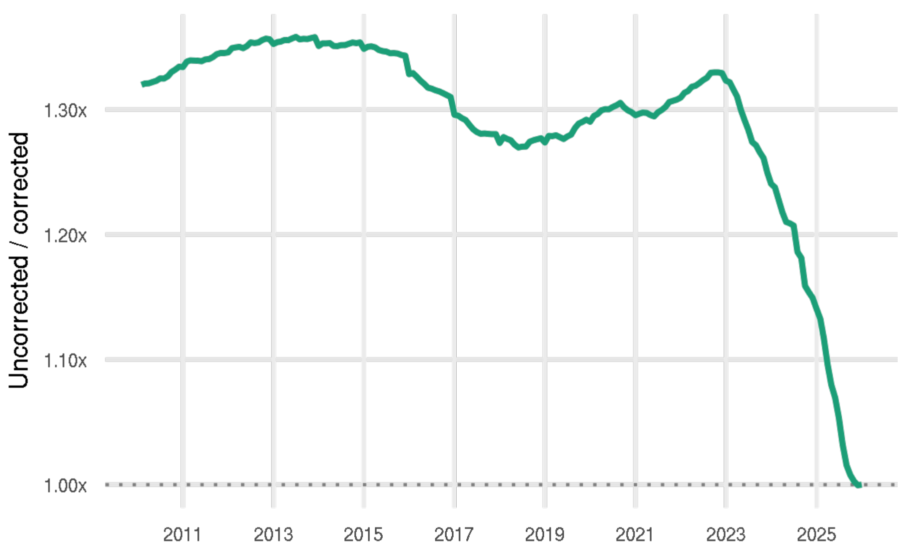

(b) Ratio of shares

Figure 1: Rates of "AI word" appearance on LinkedIn job descriptions. The orange line computes rates for each year based on LinkedIn *as it exists in 2026*; i.e., it counts up workers who said they worked in an AI job in 2022. The green line computes rates based on LinkedIn *as it was*; i.e., did you say you worked in an AI job in 2022? The ratio of these two lines is in the bottom figure. This is mechanically 1 today, as there are no updates to correct for. Uncorrected estimates based only on the 2026 snapshot would overstate "AI skills" in 2022 relative to the contemporaneous baseline by around 30%. Note the updates file begins in 2020, so prior to 2020 the green line is not LinkedIn "as it was" but rather LinkedIn "as it was in 2020."

Some of these changes are "contemporaneous," with people editing their profile to reflect new responsibilities; however, we document that substantial shares of LinkedIn users (19.2%) update their job descriptions *after they have already finished the job*, with most of these changes coming *years* after they have left the job in question. The gap between workers leaving their jobs and when they update their job description leads us to call those who do this "time travelers," reflecting the idea these workers are altering their pasts. By identifying and tracking time travelers, researchers can turn the mutable nature of professional profiles from a weakness into a strength. Workers retroactively adjust their profiles when preparing for a labor market transition; thus, we can consider the changes they make to reflect their understanding of skills "in demand" in the market at the time of their job search. We use these adjustments to track demand for three skill sets—artificial intelligence, work from home, and DEI—from 2021 through 2026. Finally, we can use the introduction of "AI signature" words from Kobak et al. (2025) to measure AI use across demographic groups based on profile updates.

After outlining our data in Section 2, in Section 3 we show time travel is commonplace on professional profiles. This creates the potential to bias naive estimates of historical skill supply as illustrated in Figure 1. Time travel is most common in knowledge work, particularly tech. On the other hand, it's relatively rare in manufacturing and services. Thus, time travel can bias different trends *different amounts,* complicating any potential cross-industry comparisons.

Second, Section 4 shows that worker time travel is strongly associated with labor market transitions. Relative to the not-yet-treated (in the tradition of Callaway and Sant'Anna (2021)), time travelers experience a jump in their probability of changing employers. Following their time travel event, transition probabilities fall until workers approach the baseline rate of job transitions in the broader population. This could be selection—workers polish their resumes when preparing to job search—or treatment—resume polish improves the probability of finding a job, as shown in Wiles, Munyikwa, and Horton (2025). In either case, we can be confident workers are adjusting their resumes with the goal of being *more attractive* to employers.

Third, Section 5 examines what types of updates workers make to their profiles to capture their second-order beliefs about skill demand: what do workers believe firms want? Tracking normal job descriptions confounds two forces: the skills that firms *demand*, and the skills that workers can *supply*. Using time travel edits allows us to hold skills fixed and isolate only variation coming from firm demand. We show three applications where time travelers are informative about changing market conditions. First, we document explosive increases in demand for AI skills since late 2022; however, these increases may have tempered in recent months. Second, we show softening demand for work-from-home since 2024. Finally, "DEI" positions have seen significantly weakening demand since 2024—particularly following a series of executive orders in early 2025.

Finally, Section 6 estimates AI tool adoption across demographic groups by measuring use of distinctive "AI words" documented in Kobak et al. (2025). We find use of these AI words explodes in the first half of 2023. This is part of a broader tradition of detecting AI-assisted writing in free text data pioneered by M. Shin, Kim, and J. Shin (2025), Liang et al. (2024), and Liang et al. (2026). Liang et al. (2025) in particular assess LLM assistance in job *postings*; we complement their "demand side" estimates with "supply side" measures. We show this adoption is heterogeneous across education levels. Generally, using AI tools to edit your resume declines with education; however, MBAs exhibit significantly higher rates of AI use relative to other master's degree holders. This association with MBAs is driven largely by relatively low-ranked programs; Among the "magnificent 7" (M7) MBA programs, use of Kobak style words is roughly half the MBA base rate.

2 Revelio Labs Vintages, 2020-Present

Using novel data from Revelio, we can reconstruct LinkedIn "as it was" each month from August 2020 to January 2026. To understand how profiles change over time, we need multiple "vintages" of professional profile data. In addition to full data snapshots from September 2022 and January 2026, we also use the *updates* collected by Revelio starting in August 2020: each row represents a timestamped profile change Revelio detected and merged into their primary dataset. These include character-by-character changes to the raw text of each profile's titles and job descriptions. By "winding back" updates from the full snapshot, we can approximate LinkedIn as it was in any month over this period.

For our analysis, we limit only to "high quality" US profiles. Given our text-as-data methods below, we limit to US profiles given the dominance of English as the professional language—minimizing the need to standardize metrics across languages. We also require profiles to have at least 100 connections by the end of our sample period in order to limit the influence of fake and inactive profiles. Finally, we turn off labor market entry by limiting only to workers present in our September 2022 snapshot.

We define a "time traveler" to be someone who edits a job title or description on their profile *after* they have already left the position. This allows us to isolate "branding" changes from true responsibility changes. If I update my job description while still employed at a company, that could reflect *either* a change in how I describe my job *or* true changes in my responsibilities. If a worker has already left a job, their responsibilities *cannot* be changing, as the job is in the past. Thus, adjustments can *only* reflect a change in branding—the job is the same, but the way they talk about it has changed. Thus, we call these workers time travelers: they go "back in time" to adjust job responsibilities even *after* they have already left. We enforce a 60-day gap between job exit and profile edit to be considered "time travel" to ensure the worker has truly left.

Because scrapes happen at discrete points we cannot perfectly time updates. For example, a profile may be scraped on May 10 and then again on May 24. If the profile changed at some point between those two points we observe that change but not the exact day it happened on. To proxy the date of a time travel event we take the midpoint of the day the new profile is observed and the last day the *old* profile was observed. In the example above this would be May 18. Because scrapes occur irregularly, we cannot know the exact date of the event. However, under an assumed Poisson arrival rate of edit events and conditioning on one edit within the interval, this midpoint estimator recovers the expected date of the change. To enforce an upper bound on our measurement error, we drop the 9.9% of updates where old and new snapshots of the data are more than four months apart; thus, the attributed date can be no more than sixty days from the true edit date.[3]

[3] The median scrape gap overall is 60 days; after applying the four month maximum, this falls to 49 days in our analysis sample.

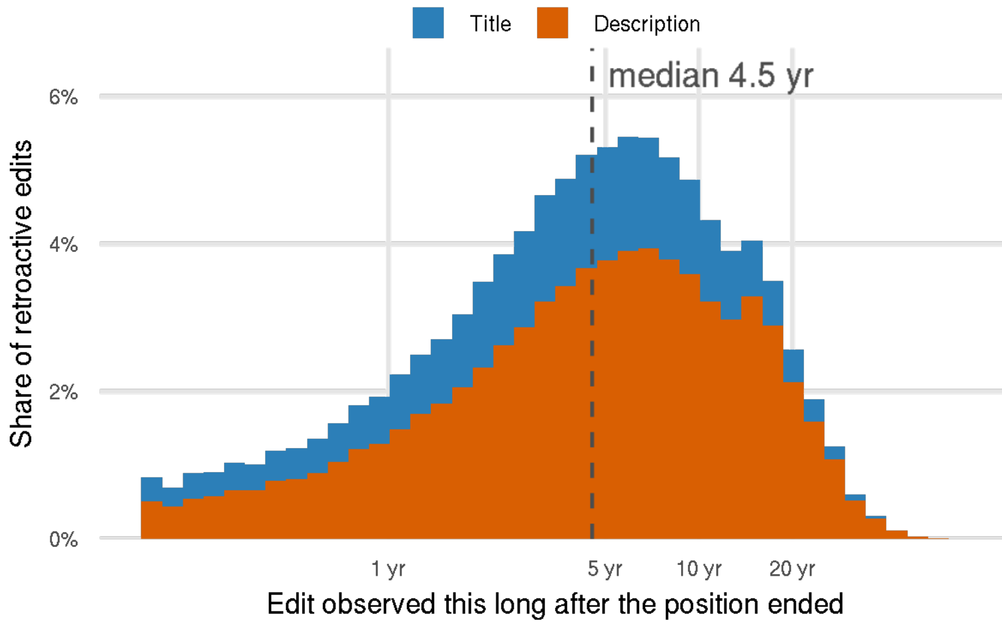

Figure 2: Time between position end date and time travel edit (log scale). If a title and description are changed at the same time, those are counted as two separate changes.

3 Time Travel is Common

We find that time travel is widespread; among American profiles present 2022-2026 with 100 connections by the end of the period, we find one in five (19.2%) had time traveled at least once since 2020. This represents a lower bound, as it is likely some people time traveled before that who we cannot observe. These revisions are not coming shortly after leaving jobs: Figure 2 shows that the median retroactive edit took place more than four years after leaving the job—well after responsibilities had been finalized and people have moved on to newer employment.

The rate of time travel is strongly heterogeneous by industry, as seen in Figure 3. Technology and information has distinctively the highest rate, with nearly one-in-three workers time traveling by the end of our sample. In contrast, construction and real estate both show distinctively low levels. This translates to strong geographic heterogeneity as seen in Figure 4. West coast tech hubs like San Francisco, Seattle, and San Jose have the highest rates, while time travel is less common in Southwestern cities like Houston, Phoenix, and Dallas.

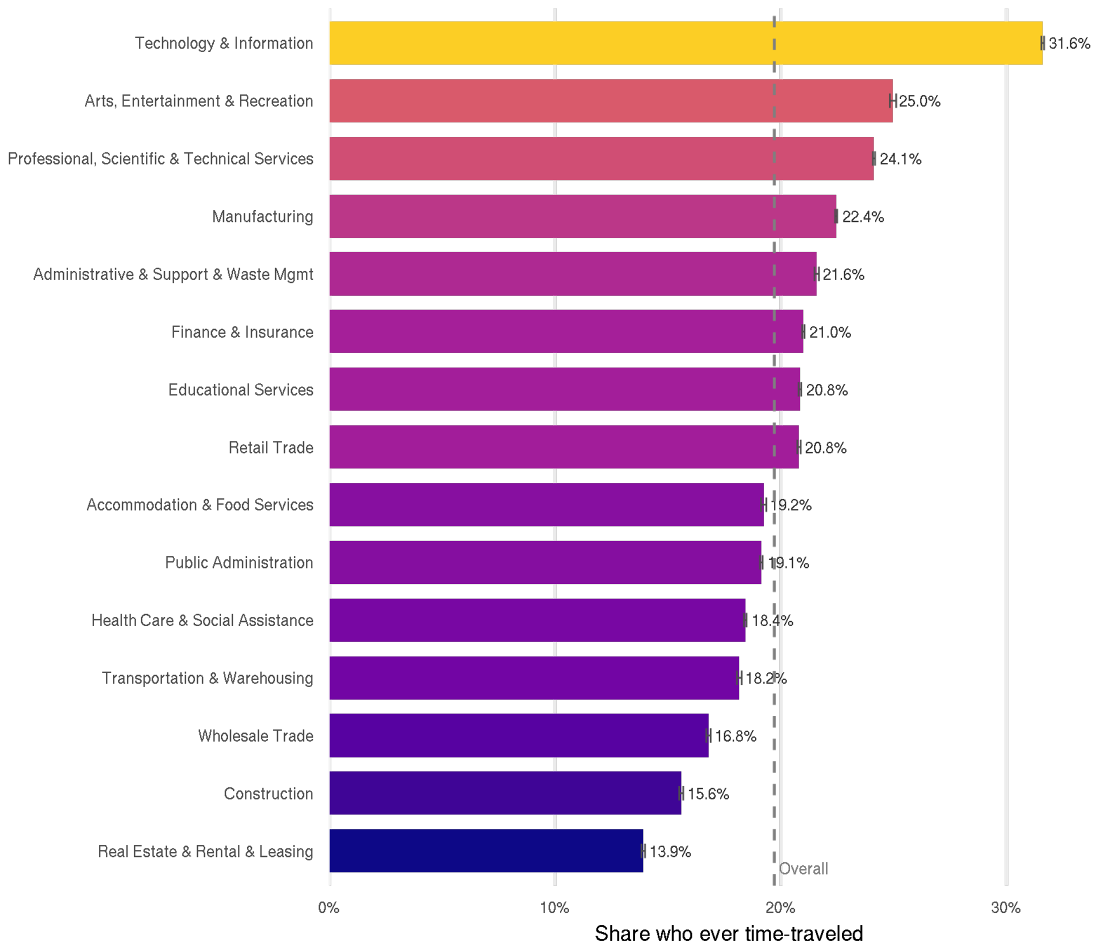

Figure 3: Time travel rates for 15 most common NAICS 2 codes on LinkedIn. Overall time travel rate in the sample given by the vertical dashed line. Excludes NAICS codes containing "Other," e.g. "Other Services."

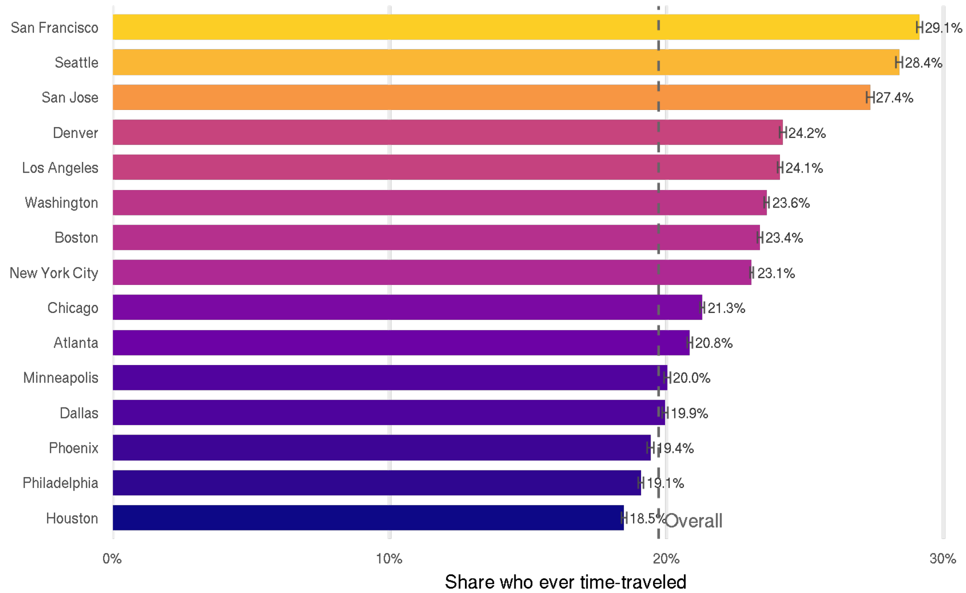

Figure 4: Time travel rates in the 15 most common American metro areas on LinkedIn. Overall time travel rate in the sample given by the vertical dashed line.

Time travel is most common among the most-connected users. Figure 5 shows that time travel is nearly three times more common among profiles with more than 500 connections relative to profiles with between 100 and 200 connections. This could reflect multiple channels: for example, profiles with more connections may be more active users, or have had a longer lifespan creating more opportunities to time travel. However, this means limiting only to profiles with a large number of connections—usually done to improve sample quality—actually *increases* exposure to the possibility of worker time travel.

Further subgroup heterogeneity is available in Table 1. In the first panel we see that time travel is most common among master's degree holders–especially MBAs–and least common among high school and associate's degree holders. Time travel exhibits a strong age gradient, with under-30 workers time traveling at rates more then four times higher than over-60 workers. Unsurprisingly, the high-time-travel professions are common in high-time-travel industries:[4] product manager, UX researcher, and the like. Similarly, low-time-travel professions match low-time-travel industries:

[4] We limit to professions with at least 1,000 profiles to minimize the influence of sampling variation.

Rank	Subgroup	Profiles	Base rate	95% CI
	Overall	29,448,462	19.7%	[19.7, 19.7]
By education				
1	MBA	1,822,779	29.3%	[29.2, 29.4]
2	Master's	4,597,254	26.2%	[26.2, 26.3]
3	Bachelor's	10,964,677	23.8%	[23.8, 23.8]
4	Doctorate	2,216,420	20.2%	[20.2, 20.3]
5	High school	437,857	18.6%	[18.5, 18.7]
6	Associate's	1,139,700	16.4%	[16.3, 16.5]
	Missing	8,269,775	8.9%	[8.9, 9.0]
By gender (predicted)				
1	Female	13,093,264	20.2%	[20.2, 20.2]
2	Male	16,151,050	19.4%	[19.4, 19.4]
	Missing	204,148	13.7%	[13.6, 13.9]
By age (estimated from education)				
	< 30	3,008,166	38.1%	[38.1, 38.2]
	30s	6,357,379	27.4%	[27.3, 27.4]
	40s	4,059,796	21.0%	[20.9, 21.0]
	50s	2,440,250	17.1%	[17.1, 17.2]
	≥ 60	2,163,018	9.8%	[9.8, 9.9]
	Missing	11,419,853	12.6%	[12.6, 12.6]
By profession				
1	AI product manager	1,321	70.6%	[68.1, 73.0]
2	Platform Product Owner	4,542	60.3%	[58.9, 61.7]
3	UX Research	3,847	59.2%	[57.7, 60.8]
4	Design and Qualitative Research	1,361	56.6%	[54.0, 59.3]
5	AI Leadership Roles	3,315	56.6%	[54.9, 58.3]
⋮	⋮	⋮	⋮	⋮
5019	Chiropractor	12,051	3.3%	[3.0, 3.6]
5020	Dental Health Specialist	3,365	3.1%	[2.5, 3.6]
5021	Vehicle Rental Sales	91,170	2.5%	[2.4, 2.6]
5022	Orthodontist	2,435	2.4%	[1.8, 3.0]
5023	Dental Surgeon	2,538	2.4%	[1.8, 3.0]

Table 1: Heterogeneous rates of time travel among users in January 2026. A user is counted as a time traveler if they made at least one retroactive edit to a job they already left from August 2020 to January 2026. Standardized education and occupation codes are taken from Revelio's cleaned versions of free-text LinkedIn entries.

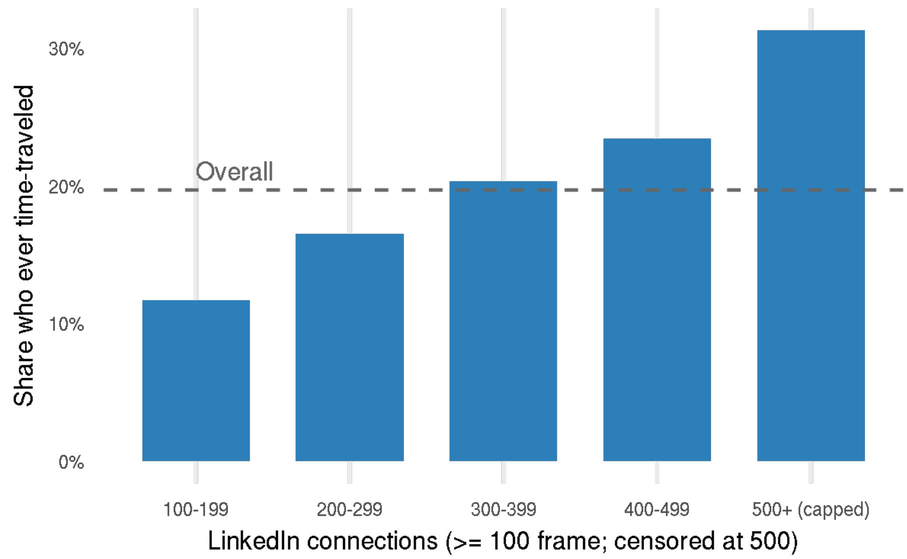

Figure 5: Rate of time travel by profile connections in 2026. The x-axis shows the number of connections belonging to a profile; the y-axis shows, conditional on having that many connections, the probability of having time traveled over our sample period. Dashed gray line shows the overall population rate of time travel. Data is censored at 500+ connections.

construction workers and dentists generally don't retroactively update their resumes.

4 Time Travel is Associated with Employment Transitions

Employment transitions are positively associated with time travel: starting six months prior to an employment transition, workers see increasing rates of time travel. Following their employment transition, time travel rates return to baseline within three months. This reaffirms our argument that time-travel changes capture workers' beliefs about what employers are looking for: they appear to time travel just as they prepare to change employers.

We estimate the pattern of time travel around employment transitions using a classic difference-in-differences design, using workers who never change employers over the course of our sample period as a control group. We assign each non-mover a placebo move date drawn from the distribution of realized move dates. For tractability, we limit to a 500,000 person subsample of the data. We do not wish to compare observations too far from each other in time; the unrestricted estimator

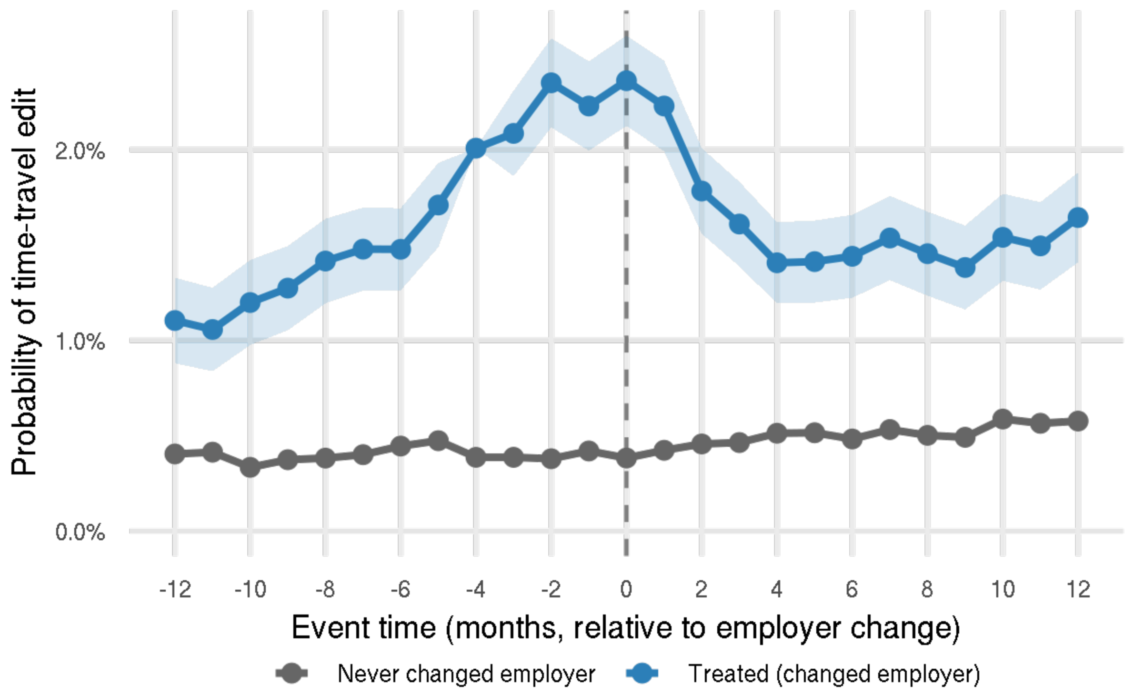

Figure 6: Probability of time traveling around an employer change. Employer change occurs at $t = 0$, blue line represents path of moving workers, gray line represents path of non-movers. Sample limited to American LinkedIn users with at least 100 connections by 2026 who were present on the platform in 2022. Blue ribbon is the 95% confidence interval around treated group mean.

could assign a 2026 move date to workers in 2020 even if the 2020 workers had not yet entered the labor force. Thus, we limit to moves occurring in January 2023 or later. Combined with our sample restrictions, this yields 27,310 movers and 117,900 non-movers. We downsample non-movers to 27,310 as well.

Figure 6 shows that job transitions are strongly associated with time travel. The blue line shows rates of time travel edits among workers who change employers around their move date. The gray line shows the time travel rate among non-movers around their *placebo* move date. Note, as described above, that time travel events are imperfectly measured; thus, the elevated levels of time travel after move could reflect workers *actually* time traveling after moving, *or* pre-move time travel incorrectly attributed to after the move. Relative to non-movers, workers who change employers are more than twice as likely to time travel even a year prior to their move. Starting six months ahead of their move, this rate elevates even further—eventually reaching double their rate from twelve months prior (2.5 percent vs. 1.2 percent). Time travel propensity falls rapidly in the three months following the move, returning roughly to baseline levels.

5 Tracking Skill Demand Using Time Travel

The previous section argues that time traveling is associated with employment transitions; it is reasonable to think, then, that these updates could reflect workers' second-order beliefs about firms' preferences. That is: workers modify their resume to reflect what they believe employers want to hear. By examining the content of these edits, we can recover how these beliefs change over time.

By looking specifically at time travelers, we can disentangle changing *demand* for skills from changing *supply.* A naive estimate would count, for example, the number of profiles generally containing the word "AI" and use this as a proxy for AI demand. However, this is an equilibrium object; "AI" could appear in more profiles either because demand for AI skills is increasing *or* because supply of AI skills is increasing. Time travel holds the underlying job fixed, making changes in wording likely to reflect changing perceived demand rather than new experiences. The same worker is describing the same job they held previously, where they used the same skills both before and after changing the description. The only thing that can change, then, is the *way* the worker chooses to describe the job: a reflection of changing demand conditions.

Below we trace out three specific cases where we track changing demand over time: artificial intelligence, work-from-home, and DEI. Because time-travel is generally resume polishing, words are almost always net added. Note a data collection error in June and July 2024 creates coincident artificial "removals" across all our series, so we interpolate over that period as marked by the dashed lines. Details, along with full, uninterpolated figures, are available in Appendix B.

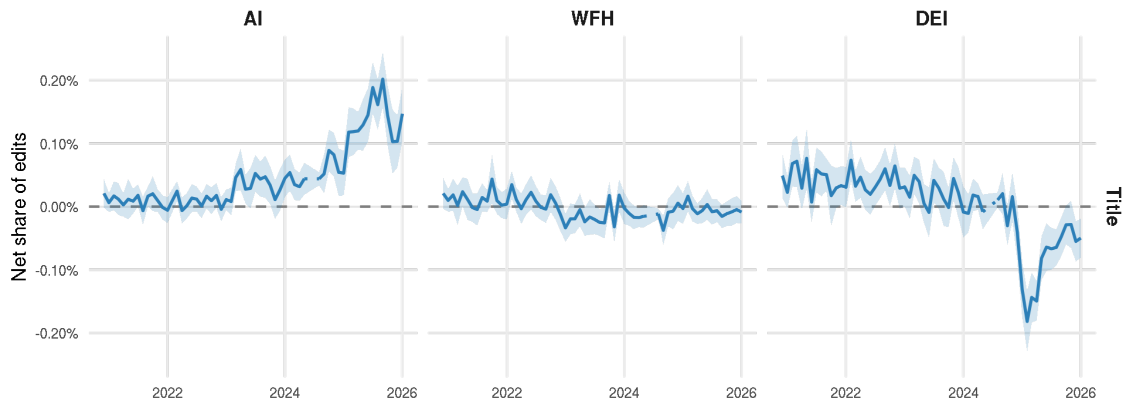

(a) Titles

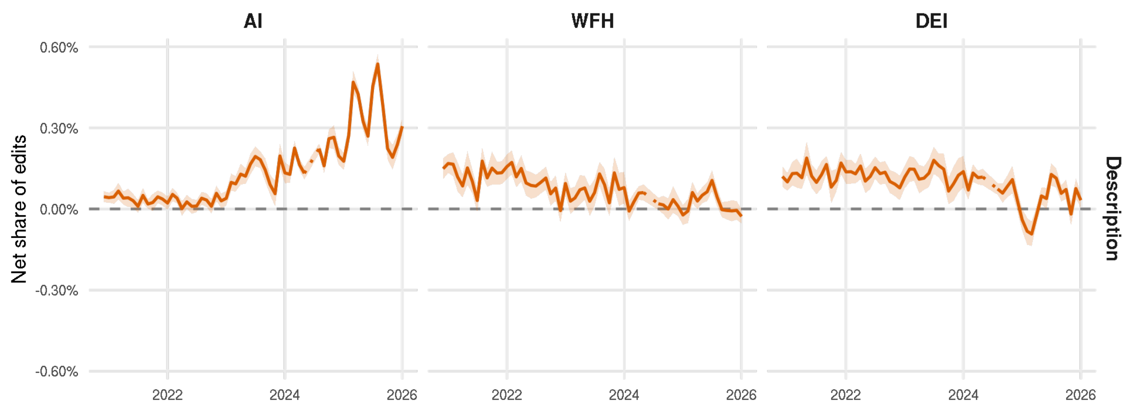

(b) Descriptions

Figure 7: Each panel shows how frequently words associated with the given skill are net added in retroactive edits. Blue lines indicate changes in the job title, orange lines indicate changes in the job description. Dashed lines cover an error in scraping that leads to more perceived removals; see Appendix B for details.

5.1 Artificial Intelligence

Over the course of our sample, expertise in artificial intelligence becomes increasingly common to retroactively add to resumes. The left panels of Figure 7 show the net rate of time travel edits adding one of "llm," "gpt," "artificial intelligence," or "ai."[5] Note this is a net rate, so edits *removing* a word count negatively. AI is relatively flat until the end of 2022—the introduction of ChatGPT. After this point, retroactive edits adding AI explode more than 5x. However, AI-related additions have declined in recent months, suggesting some softening in the salience or perceived value of AI language.

5.2 Work-from-Home

Though work-from-home is added to profiles steadily throughout our sample, its addition generally becomes less common each year. The middle panels of Figure 7 show the net rate of time-travel edits adding one of "remote," "remotely," "work from home," "wfh," or "telework." Descriptions generally see workers adding remote work to their profiles throughout our sample period. However, this declines gradually until around the end of 2024. Although the trend recovers in the first half of 2025, it falls again before the end of the year, coincident with return-to-office mandates. By the end of the sample period, work-from-home is no longer being net added to profiles.

5.3 Diversity, Equity, and Inclusion

Similarly to work-from-home, workers seem to be moving their resumes away from "DEI" in the latter half of our sample. The right panels of Figure 7 show the net rate of time-travel edits adding one of "diversity," "equity," "inclusion," or "dei." Though these are added throughout the first half of the decade, there is a sharp decline at the beginning of 2025—coinciding with Executive Orders 14151 and 14173 targeting DEI programs as discriminatory and disqualifying for federal contracts. This stands counter to the trend seen in Baker et al. (2024), which finds that firms increased discussion of diversity fivefold from 2008 to 2021–far outpacing any real increases in

[5] Words require break characters on each side; this means "ai specialist" will be picked up while "hail" will not.

workplace diversity. Here we find workers *removing* diversity from their job descriptions—without any changes in the underlying jobs being described.

6 Measuring AI-Assisted Profile Changes

Starting with the advent of ChatGPT, we document an explosion in "LLM-signature" words appearing in time-travel revisions—suggesting workers are using LLMs to update their professional profiles. Kobak et al. (2025) document words with significant "excess usage" in 2023 relative to prior years. Separating "content" words from "style" words, they argue these style words are hallmarks of LLM-assisted writing; for example, they document a 30-fold increase in prevalence of the word "delve" relative to prior years. We count appearances of the top-25 highest-signal words in time travel edits as a metric of AI use. By this measure, less educated workers generally use AI to revise their profiles more often. We find MBAs adopt at higher rates than other Master's degree holders. Exploring this relationship further, we find AI tool use is strongly negatively associated with MBA program prestige—alumni of "magnificent 7" programs are half as likely to use a Kobak word as a randomly selected MBA.

6.1 Take-Up is Heterogeneous by Education

Figure 8 shows the frequency with which high-signal Kobak et al. (2025) words appear in time travel edits each month. Pre-ChatGPT, these words are present in only a third of a percent of time travel edits. By August 2023, at least one of these style words (e.g., delve, underscore, showcase) is present in roughly 3% of time-travel edits: a ten-fold increase. We include the "content" words in blue as a placebo—these words reflect *topics* with increased discussion (e.g., COVID, LLM, transgender) rather than changes in style. The stylistic changes are more than three times as large as the content changes.

Adoption is heterogeneous by education: PhD holders use these words significantly less than bachelor's holders, while high school and associate's holders use these words significantly more. Note this is opposite the conclusion of Bick, Blandin, and Deming (2024); while their team finds AI

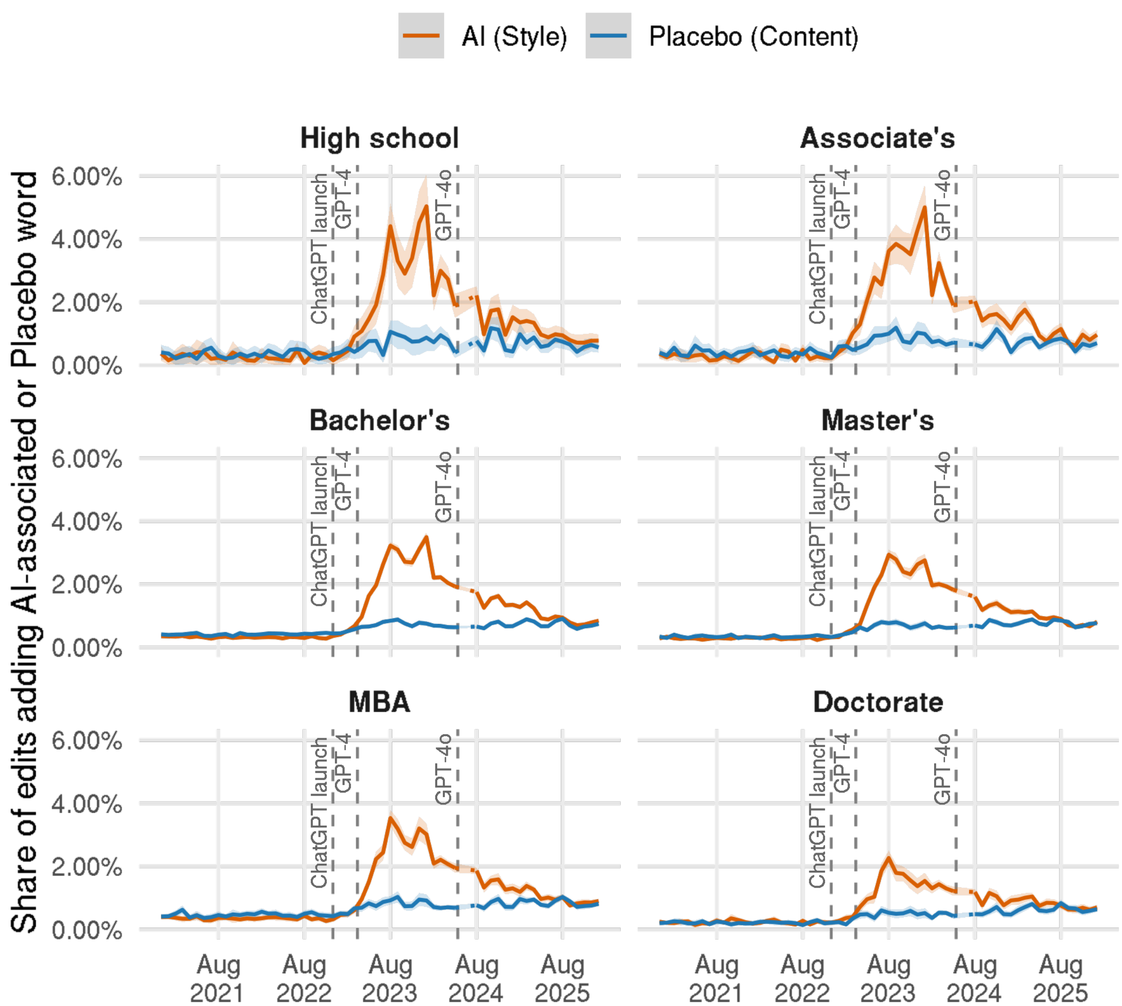

Figure 8: Frequency of Kobak et al. (2025) words in time-travel edits by education level. Counts the share of time-travel edits introducing either a top-25 signal style word (in orange) or a top-25 signal content word (in blue) according to Kobak et al. (2025). Gray vertical lines show major ChatGPT releases over this period. Dashed lines in the series cover an error in scraping; see Appendix B for details.

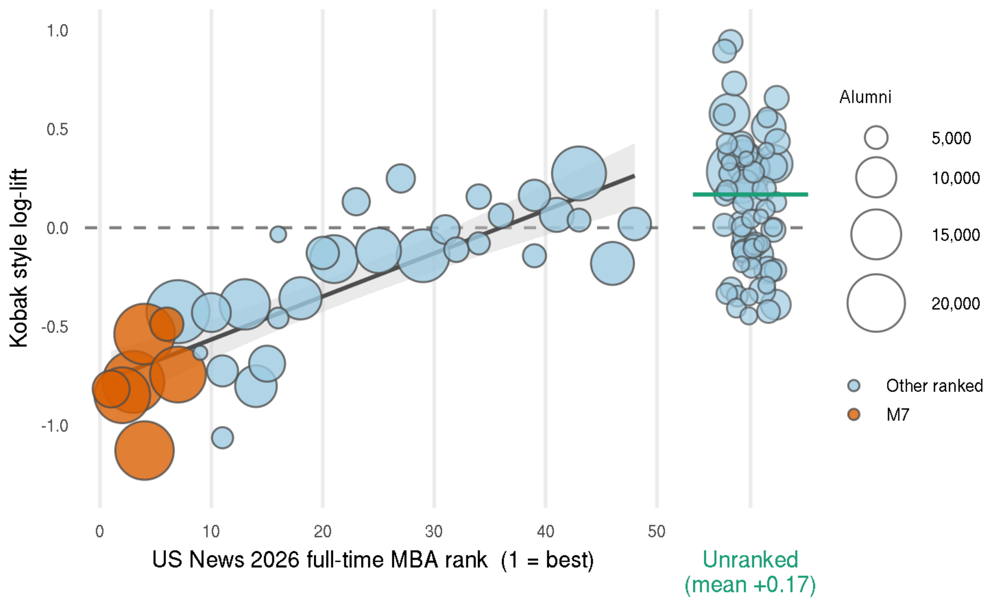

Figure 9: Kobak style log-lift by US News MBA program ranking. Points are sized by the number of alumni observed. A greater log lift means that program is more associated with Kobak style (AI) words. Unranked programs grouped on the right side with jitter.

use *increasing* in education level, we broadly find AI adoption—for this particular task—*declining* in education.

It is difficult to explain the explosive rise of these Kobak words in the first half of 2023 as anything other than widespread LLM adoption; however, their slow decline in subsequent years confounds several forces. It is tempting to conclude we are "past peak LLM;" however, we believe this conclusion is too hasty. It is possible the stock of "unpolished" pre-LLM profiles has been exhausted; everyone who was going to polish their profile has done so, and new profile creators use LLMs for initial construction so time travel is no longer necessary. Moreover, models have certainly improved over this period; if the Kobak et al. (2025) "tics" are specific to that vintage of model, then the decline in the prevalence of these words could simply reflect obsolescence of older models.

6.2 LLM Language Negatively Correlated with MBA Program Prestige

Given the distinctive patterns of MBAs relative to other master's degree holders, we dig further into program heterogeneity. For each of the top 100 most common programs in our data, we compute the "log lift:" how seeing a Kobak style word in a time travel edit impacts the probability the editing worker attended the school. A positive log lift suggests alumni of school s are more likely than average to use Kobak style words; and, symmetrically, that when you see a Kobak style word, it increases your posterior estimate a worker attended the given school. Formally:

$$\ln Lift(school) = \ln \frac{P(kobak|school)}{P(kobak|mba)}$$

Using AI assistance to update your profile is less common among alumni of top-ranked programs, as seen in Figure 9. We plot US News 2026 ranking on the x axis, the school's log lift on the y axis, and the number of alumni we observe in the data as the size of the point. Note since rank "1" is the best, top-ranked programs are on the left. Unranked programs are plotted separately at right with horizontal jitter. Knowing a student attended an M7 program lowers the probability of using a Kobak word when they time travel by around half relative to the baseline MBA probability.

7 Conclusion

Professional profile data is a powerful new tool for labor economics—but should be understood not as ground truth, but as a mutable equilibrium object. We document that one in five high-quality American profiles retroactively edit a job description after they have already left the position, reflecting a change not in the actual job content but rather a change in the way the job content is *described*. We then develop a method to use these retroactive edits to assess second-order beliefs regarding skill demand, showing that worker editing tracks intuitive patterns for three skill sets. Finally, we extend the literature on LLM adoption, showing that workers use AI to edit their profiles heterogeneously by educational background.

Although we apply it here to track LinkedIn profile discussion of AI, work-from-home, and

DEI, we believe this approach has broad applicability across different types of skills and different types of worker profiles so long as multiple vintages are available. For example, a researcher who had multiple resumés for a worker over time could apply a very similar approach. Similarly, it is applicable to a variety of different types of skill sets: do workers highlight social skills more with the advent of automation? How do workers respond to softening labor markets more generally?

Our method requires only that workers *believe* the skills on their profile influence their hiring probability; we have nothing to say about the efficacy of these revisions in job search. Wiles, Munyikwa, and Horton (2025) show that writing assistance improves job search outcomes; we think extending this to include updating skills is a promising avenue for future research.

References

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson (Jan. 1, 2024). "Artificial intelligence, firm growth, and product innovation". In: *Journal of Financial Economics* 151, p. 103745. ISSN: 0304-405X. DOI: 10.1016/j.jfineco.2023.103745. URL: https://www.sciencedirect.com/science/article/pii/S0304405X2300185X (visited on 06/23/2026).

Babina, Tania, Anastassia Fedyk, Alex X. He, and James Hodson (June 2023). *Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition*. DOI: 10.3386/w31325. URL: https://www.nber.org/papers/w31325 (visited on 06/23/2026).

Baker, Andrew C, David F Larcker, Charles G McClure, Durgesh Saraph, and Edward M Watts (2024). "Diversity washing". In: *Journal of Accounting Research* 62.5, pp. 1661–1709.

Bick, Alexander, Adam Blandin, and David J. Deming (Sept. 2024). *The Rapid Adoption of Generative AI*. DOI: 10.3386/w32966. URL: https://www.nber.org/papers/w32966 (visited on 06/22/2026).

Callaway, Brantly and Pedro H. C. Sant'Anna (Dec. 1, 2021). "Difference-in-Differences with multiple time periods". In: *Journal of Econometrics*. Themed Issue: Treatment Effect 1 225.2, pp. 200–230. ISSN: 0304-4076. DOI: 10.1016/j.jeconom.2020.12.001. URL: https://www.sciencedirect.com/science/article/pii/S0304407620303948 (visited on 06/15/2025).

Hampole, Menaka, Dimitris Papanikolaou, Lawrence D.W. Schmidt, and Bryan Seegmiller (Feb. 2025). *Artificial Intelligence and the Labor Market*. DOI: 10.3386/w33509. URL: https://www.nber.org/papers/w33509 (visited on 06/23/2026).

Kobak, Dmitry, Rita González-Márquez, Emőke-Ágnes Horvát, and Jan Lause (July 2, 2025). "Delving into LLM-assisted writing in biomedical publications through excess vocabulary". In: *Science Advances* 11.27, eadt3813. DOI: 10.1126/sciadv.adt3813. URL: https://www.science.org/doi/10.1126/sciadv.adt3813 (visited on 02/23/2026).

Liang, Weixin, Zachary Izzo, Yaohui Zhang, Haley Lepp, Hancheng Cao, Xuandong Zhao, Lingjiao Chen, Haotian Ye, Sheng Liu, Zhi Huang, Daniel A. McFarland, and James Y. Zou (May 19, 2026). *Monitoring AI-Modified Content at Scale: A Case Study on the Impact of ChatGPT on AI Conference Peer Reviews*. DOI: 10.48550/arXiv.2403.07183. arXiv: 2403.07183[cs.CL]. URL: http://arxiv.org/abs/2403.07183 (visited on 06/23/2026).

Liang, Weixin, Yaohui Zhang, Mihai Codreanu, Jiayu Wang, Hancheng Cao, and James Zou (Feb. 17, 2025). *The Widespread Adoption of Large Language Model-Assisted Writing Across Society*. DOI: 10.48550/arXiv.2502.09747. arXiv: 2502.09747[cs.CL]. URL: http://arxiv.org/abs/2502.09747 (visited on 06/23/2026).

Liang, Weixin, Yaohui Zhang, Zhengxuan Wu, Haley Lepp, Wenlong Ji, Xuandong Zhao, Hancheng Cao, Sheng Liu, Siyu He, Zhi Huang, Diyi Yang, Christopher Potts, Christopher D. Manning, and James Y. Zou (Apr. 1, 2024). *Mapping the Increasing Use of LLMs in Scientific Papers*. DOI: 10.48550/arXiv.2404.01268. arXiv: 2404.01268[cs.CL]. URL: http://arxiv.org/abs/2404.01268 (visited on 06/23/2026).

Shin, Minkyu, Jin Kim, and Jiwoong Shin (Feb. 11, 2025). *The Adoption and Efficacy of Large Language Models: Evidence From Consumer Complaints in the Financial Industry*. Rochester, NY. DOI: 10.2139/ssrn.5004194. URL: https://papers.ssrn.com/abstract=5004194 (visited on 06/23/2026).

Wiles, Emma, Zanele Munyikwa, and John Horton (Dec. 2025). "Algorithmic Writing Assistance on Jobseekers' Resumes Increases Hires". In: *Management Science* 71.12, pp. 10144–10164. ISSN:

0025-1909. DOI: 10.1287/mnsc.2024.04528. URL: https://pubsonline.informs.org/doi/10.1287/mnsc.2024.04528 (visited on 06/23/2026).

A Traditional Difference-in-Difference

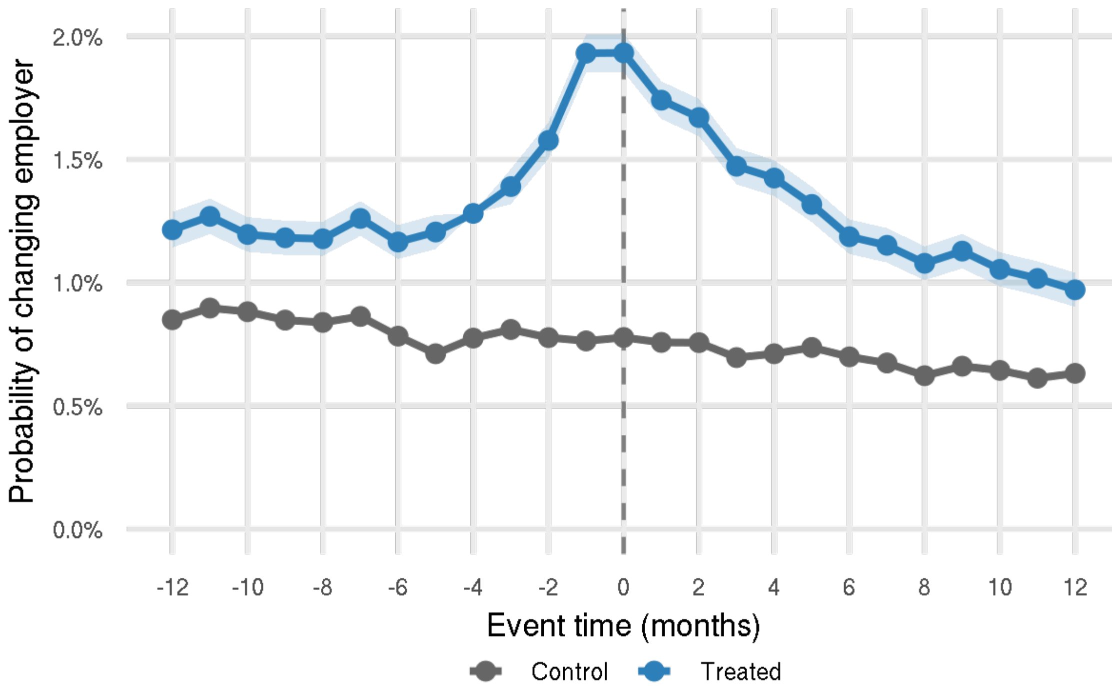

Figure A.1: Probability of changing employer around a time-travel event. Time travel estimated to occur at $t = 0$. Blue line represents path of treated group, gray line represents path of non-time traveling LinkedIn users around a placebo treatment date drawn from the distribution of treatment dates. Analogous to Figure 6, but compares against 100+ connection American LinkedIn updaters who *never* time travel, rather than Callaway and Sant'Anna (2021) not-yet-treated units.

B Scraping Error

In June and July of 2024 the profile scraper experienced an error where job descriptions were truncated. For example, instead of "William Eberle Professor of Economics at Stanford University, a Senior Fellow of SIEPR, and the Co-Director of the Productivity, Innovation and Entrepreneurship program at the National Bureau of Economic Research," the scraper would record "William Eberle Professor of [See More]" for a significant fraction of profiles. This creates an artificially large number of word removals that *appear* to be time-travel edits.

Figure B.1 shows diagnostics for our data in and around the error window, affirming that the jump is a data artifact. The period with the scraping error—June and July 2024—is shaded. We see coincident drops in *all* our skill series precisely in this window. Time travel edits occurring in this window on average make profiles shorter, whereas in nearly all other periods—particularly on either side of this window—time travel makes profiles significantly longer. Looking specifically at edits estimated to occur in July, we see significant heterogeneity in length changes by the month the scrape occurred. For edits occurring in July which were *scraped* in July, descriptions became 22 characters shorter on average. For edits occurring in July which were *scraped* in August or September, descriptions became 17 or 31 characters *longer*, respectively.

In the text, we linearly interpolate over this period. Otherwise, *all* our series in Figures 7 and 8 show a decrease at exactly the same point. See Figures B.2 and B.3 for the full series.

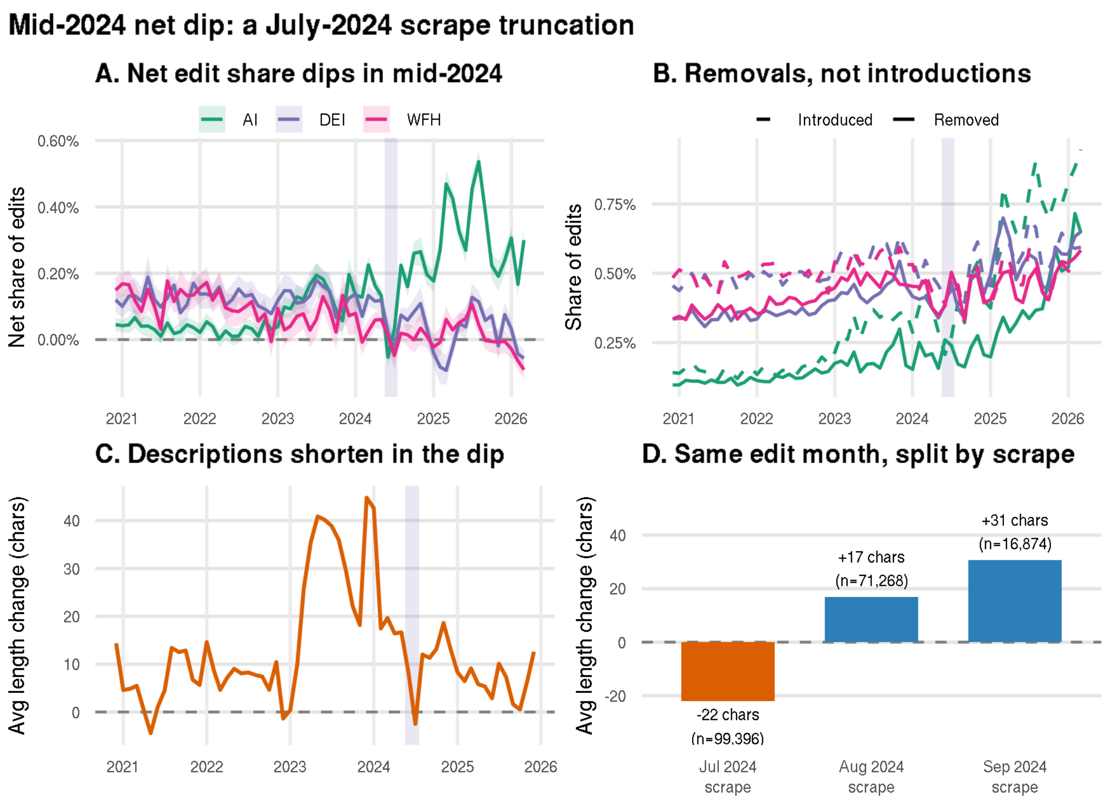

Figure B.1: Diagnostics identifying the scraping error. Panel A reproduces the description changes from Figure 7, shading the region with the error. All three series show a clear, discontinuous negative jump, followed by a return to trend. Panel B breaks the net change down into additions and removals; all series experience both a jump in removals and a fall in additions. Panel C shows the *change in length* of profiles experiencing time travel. In nearly all periods, time travel profiles get longer. In this window in particular, they get somewhat shorter, then jump back up to being net additions. To cement that this is a scraping artifact and not true behavior, Panel D shows profile edits *all* inferred to occur in July of 2024, but broken out by the month they were scraped. Profiles scraped in July became significantly shorter; those scraped in other months became longer—even though the *edits* are estimated to occur at the same time.

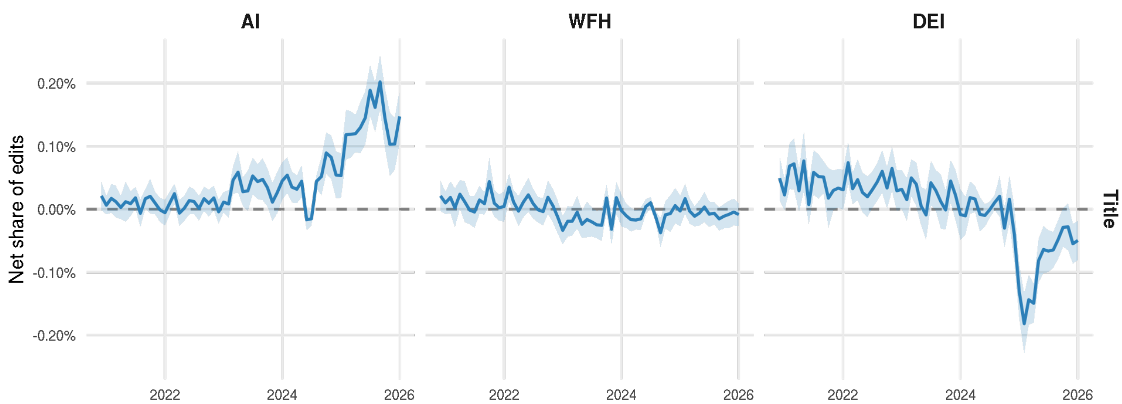

(a) Titles

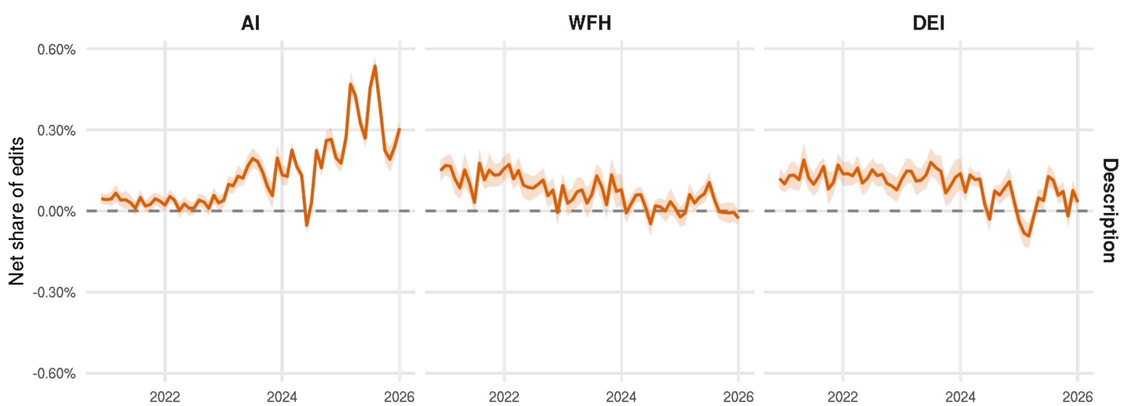

Figure B.2: Each panel shows how frequently words associated with the given skill are net added in retroactive edits. Orange lines indicate changes in the job description, blue lines indicate changes in the job title. Analogous to Figure 7, but without the interpolation over June and July 2024.

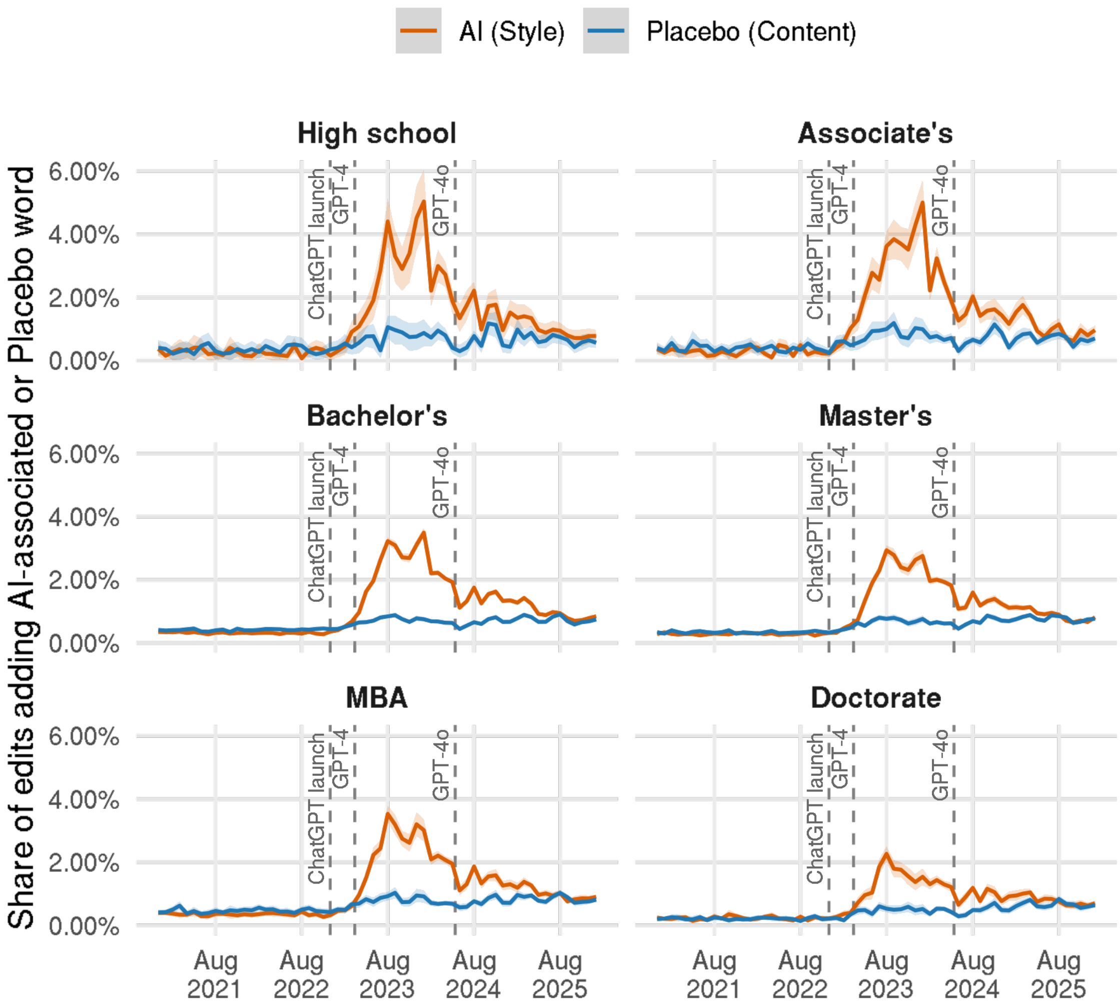

Figure B.3: Frequency of Kobak et al. (2025) words in time-travel edits by education level. Counts the share of time-travel edits introducing a top-25 signal style word (in orange) or a top-25 signal content word (in blue) according to Kobak et al. (2025). Gray vertical lines show major ChatGPT releases over this period. Analogous to Figure 8, but without the interpolation over June and July 2024.