

Is It AI or Is It the Fed?

Generative AI Exposure, Monetary Tightening and the Post-2022 White-Collar Slowdown

TMC Research

September 2026

Abstract

Since late 2022, employment in software, information, finance and professional services has stalled and entry-level hiring has weakened, a pattern widely attributed to generative AI. ChatGPT, however, arrived eight months into the fastest monetary tightening in four decades and at the end of a pandemic hiring overshoot. This paper confronts three predetermined, nearly orthogonal exposures across 205 detailed industries observed monthly from 2015 to 2026: language-model AI exposure, historical rate sensitivity estimated by local projections on identified monetary surprises, and pandemic displacement. In headcount, AI exposure predicts nothing: its coefficient is positive, imprecise, and of the size the design produces at fake base years. Rate sensitivity predicts relative losses of 0.9 log points per standard deviation during the policy hold and 2.0 during easing, with the growth gap narrowing once policy eases; the association survives sector clustering, but once the sampling uncertainty of the estimated sensitivities is propagated its interval includes zero, and a precisely measured cyclical beta predicts the same losses. Pandemic mean reversion is the largest force. Among 22 to 24 year olds, relative employment falls in both AI-exposed and rate-sensitive industries. The rate gradient breaks at liftoff. The AI gradient fell in two steps, one in the pandemic year and one from 2023, separated by a plateau; the data neither require nor exclude an AI contribution to the second.

JEL Codes: E24, E52, J23, J24, O33

Keywords: generative artificial intelligence, monetary policy transmission, industry employment, young workers, event study, discount-rate channel

1. Introduction

Since the end of 2022, payroll employment in the industries that employ most of the country’s programmers, analysts, accountants, lawyers and back-office professionals has stopped growing. Software publishing, computer systems design, finance and insurance, and the professional and technical services had added jobs at well above the economy-wide pace for a decade; after 2022 they flattened, and the hiring of recent graduates into these industries weakened noticeably. The popular explanation is generative artificial intelligence. Large language models write code, draft documents, summarize filings and answer customer queries, precisely the tasks that define these industries, and the public release of ChatGPT on 30 November 2022 lines up well with the turn in the data. Several recent studies find that early-career workers in the occupations most exposed to language models have lost ground since 2022 [Brynjolfsson et al., 2025a, Hosseini and Lichtinger, 2025], while others find no discernible disruption in aggregate employment or in the occupational mix [Humlum and Vestergaard, 2025, Gimbel et al., 2025, Eckhardt and Goldschlag, 2025].

There is, however, a second candidate explanation with an equally good alibi. ChatGPT arrived eight months into the fastest monetary tightening in forty years. The Federal Reserve raised its policy rate for the first time in March 2022, the two-year Treasury yield rose by 3.6 percentage points relative to its February 2022 level at the peak in late 2023, and the policy rate stayed at its terminal range for over a year before the first cut in September 2024. The industries that the AI narrative singles out are also industries whose labor demand is plausibly sensitive to the cost of capital: they hire for long-lived positions whose value is a capitalized stream of future surplus, they are disproportionately populated by young firms financed with equity, and their output is itself an investment good for their customers. A third force operated over the same window. These same industries had hired far ahead of trend in 2021 and early 2022, when remote work, low rates and pandemic-era demand for digital services combined into a hiring overshoot; some mean reversion was to be expected regardless of AI or the Federal Reserve.

This paper asks how much of the post-2022 slowdown in white-collar employment is attributable to generative AI, how much to monetary tightening and how much to the unwinding of the pandemic overshoot. The three explanations coincide in time and cannot be separated in aggregate data. They can be separated in the cross-section of industries, provided the characteristics that govern exposure to each shock differ from one another, and the paper’s first empirical claim is that they do: across 205 detailed private industries, exposure to language-model AI is essentially uncorrelated with historical sensitivity to monetary policy (correlation -0.06) and with the size of the pandemic hiring overshoot (correlation 0.07).

The three exposures are therefore run as a horse race, a continuous-exposure difference-in-differences design in event-study and pooled-regime forms in which each industry’s detrended employment after February 2022 is regressed on all three exposures at once.

Three features of the setting provide identification beyond the cross-sectional horse race itself. The first is timing. Liftoff preceded ChatGPT by eight months, so a rate effect should open during 2022 while an AI effect cannot. The second is the sign reversal that monetary policy, but not technology, implies. When the Federal Reserve began easing in September 2024, the growth disadvantage of rate-sensitive industries should have faded; an AI effect, if it reflects a technology that continues to diffuse, should have persisted or deepened. The third is the age gradient within industries. Both channels operate on the hiring margin, and young workers are disproportionately new hires, so both predict that young workers bear the adjustment; but the timing of the youth gradient relative to liftoff and relative to ChatGPT differs sharply between the two stories, and the Quarterly Workforce Indicators allow that timing to be observed.

The headcount results are informative about the rate channel and uninformative, in a way the paper is explicit about, on the AI channel. A one-standard-deviation higher historical rate sensitivity is associated with detrended employment 0.9 log points lower than the reference year during the tightening-and-hold regime (December 2022 to August 2024) and 2.0 log points lower during easing (September 2024 to July 2026), against nothing in the first nine months of tightening, consistent with the lags with which policy reaches employment. In growth form the disadvantage widens from -0.5 percentage points per year during the initial tightening to -0.8 during the hold and narrows to -0.3 under easing, a narrowing (95 percent interval -0.02 to 0.89) rather than an established reversal. How precise these estimates are depends on what is treated as random. Clustering by industry or by sector, and a wild-cluster bootstrap over the fourteen sectors, reject zero. But rate sensitivity is itself an estimate with a reliability ratio of 0.39, and once its sampling covariance is propagated through the design in the manner of [Murphy and Topel \[1985\]](#) the standard error of the easing coefficient doubles, from 0.66 to 1.35, and the interval includes zero; a cyclical beta measured with reliability 0.91 predicts the same relative losses (-4.0 , standard error 0.76). The robust statement is that industries with historically cyclical employment, including those whose cyclicity shows up in identified monetary shocks, lost ground during and after the tightening; how much of that is specifically monetary is less precisely determined than conventional standard errors suggest. The AI-exposure coefficient is positive in every regime (1.1 log points under easing, standard error 0.8) and does not break at ChatGPT (the change in its growth differential at the release is -0.4 points per year, interval -1.0 to $+0.2$). The largest force is neither: a one-standard-deviation larger pandemic overshoot is associated with 5.2 log points lower

detrended employment by the easing regime, as industries that had hired ahead of trend gave the excess back.

The youth results are where the AI narrative finds its support in public data, and where the question becomes one of timing. Within industries, employment of 22 to 24 year olds relative to 35 to 64 year olds is 4.9 log points lower per standard deviation of AI exposure and 3.2 log points lower per standard deviation of rate sensitivity in the easing regime, both precisely estimated under industry-robust and sector-clustered inference; the 25 to 34 year olds show nothing on either exposure. The canaries that [Brynjolfsson et al. \[2025a\]](#) identify in payroll microdata are therefore visible in the Census Bureau’s public industry-by-age series. The two gradients have different histories. The rate gradient is flat through 2015 to 2021, including the pandemic, and breaks at liftoff, falling 1.3 log points a year during 2022. The AI gradient fell in two steps: by 3.5 log points a year during the pandemic year (2020Q1 to 2021Q1), essentially not at all during a plateau from 2021Q1 to 2022Q4, and by 1.9 log points a year from 2022Q4 to 2025Q4. The post-ChatGPT pace equals the average pace over 2020Q1 to 2022Q1 (a difference of 0.03 log points a year, standard error 0.31) and is slower than the pandemic step (difference +1.6, standard error 0.6), but it is faster than the plateau that immediately preceded it (difference -1.6 , standard error 0.5), and that last comparison is what a reader who dates the decline from ChatGPT will see. A longer QWI panel back to 2005 shows that the same gradient also fell by one log point a year through the recovery from the Great Recession. The paper’s reading is that the youngest workers’ share of employment in AI-exposed industries has been falling in episodes since 2010, that the largest and fastest of those episodes began in 2020, and that the resumption from 2023 coincides with adoption, with the arrival of the tightening in the labor market, and with the unwinding of the hiring boom, so that the data neither require nor exclude an AI contribution to it. The worker flows for the 22 to 24 year olds themselves show more churn in AI-exposed industries relative to the reference year, higher hiring and higher separation rates, rather than a hiring freeze, while rate-sensitive industries show the freeze: fewer hires and a smaller share of them going to the youngest workers.

A stock-flow hiring model with a discount-rate channel in the spirit of [Hall \[2017\]](#) and [Kehoe et al. \[2023\]](#) organizes the estimates. Fitted to the post-2022 event-study coefficients, the model reproduces the rate path almost exactly ($R^2 = 0.99$) and the mean-reversion path well ($R^2 = 0.92$), while it cannot fit the AI path at all ($R^2 = -0.38$), because the AI coefficient rises before ChatGPT and never resembles an adoption curve. With standard errors from a cluster bootstrap over industries, the rate channel is estimated with a t -statistic of about three and the AI channel is not distinguishable from zero. Taken at face value, the rate channel implies that the value of a filled job in the most rate-sensitive quartile of

industries fell by roughly seven percent at the peak of tightening, in the range that [Hall \[2017\]](#) attributes to discount-rate movements, and a no-tightening counterfactual puts that quartile 2.9 log points (bootstrap standard error 0.8) higher relative to the mean industry in July 2026. These are model-dependent magnitudes conditional on the estimated exposures, and the paper labels them as such.

The paper is candid about what the headcount design can and cannot distinguish. The conventional 95 percent interval for the easing-regime AI coefficient runs from -0.5 to $+2.7$ log points per standard deviation. But industries sorted on any fixed characteristic drift over a four-year window for reasons that have nothing to do with the characteristic, and cross-sectional standard errors do not capture that drift. Replaying the exact design at nineteen fake base months between 1997 and 2015 produces AI coefficients with a standard deviation of 2.1 log points, and 58 percent of the placebo coefficients exceed the actual estimate in absolute value. A placebo event in December 2017 produces a significant negative AI coefficient. Treating the placebo dispersion as the scale of the design’s noise, an AI effect on headcount would have had to be about four log points per standard deviation, six across the interquartile range, to stand out from that noise; a decline of two or three percent in the most exposed quarter of industries relative to the least exposed would not. That is a limitation of the evidence, not a footnote to it.

The paper contributes to three literatures. To the literature on the labor market effects of generative AI [[Felten et al., 2023](#), [Eloundou et al., 2024](#), [Brynjolfsson et al., 2025a](#), [Humlum and Vestergaard, 2025](#), [Hosseini and Lichtinger, 2025](#)], it adds a design that confronts AI exposure with the two obvious competing explanations on the same units, and it shows that the entry-level pattern read as evidence of AI has antecedents that predate the technology. To the literature on the industry incidence of monetary policy [[Gertler and Gilchrist, 1994](#), [Bernanke and Gertler, 1995](#), [Dedola and Lippi, 2005](#), [Peersman and Smets, 2005](#)], it adds evidence from the 2022 to 2024 cycle that historically estimated rate sensitivities, cleaned of information effects [[Jarociński and Karadi, 2020](#), [Bauer and Swanson, 2023](#)], predict which industries lost employment and when their losses stopped widening, together with an accounting of how imprecise those historical sensitivities are. To the discount-rate literature on unemployment [[Hall, 2017](#), [Kehoe et al., 2023](#)], it offers a cross-industry estimate of the sensitivity of hiring to the cost of capital from the one episode since the 1980s in which the policy rate moved far enough to make that sensitivity visible. The design is an exposure design: the identifying assumption, stated in [Section 5](#), is that no other post-2022 shock loaded on industries in proportion to the same characteristics, and the paper treats that as an assumption to be probed rather than a fact.

The remainder of the paper proceeds as follows. [Section 2](#) describes the three shocks of

2022 and situates the paper in the literature. [Section 3](#) sets out the stock-flow hiring model and its predictions. [Section 4](#) describes the data and the construction of the three exposures. [Section 5](#) presents the empirical design and inference. [Section 6](#) reports the headcount horse race, the growth and sign-reversal tests and the continuous-stance specification. [Section 7](#) turns to the age gradient. [Section 8](#) estimates the model and reports counterfactuals. [Section 9](#) reports the robustness battery, the placebo failures and the permutation and placebo-in-time distributions. [Section 10](#) concludes.

2. Background and Related Literature

2.1 Three Shocks in 2022

Liftoff. The Federal Open Market Committee raised the target range for the federal funds rate from near zero in March 2022 and reached a terminal range of 5.25 to 5.50 percent in July 2023, where it held until September 2024. The two-year Treasury yield, the market rate that most closely tracks the expected path of policy over the horizon relevant to hiring decisions, rose by 3.63 percentage points relative to February 2022 at its peak in October 2023 and was still 2.8 percentage points above its February 2022 level in July 2026, after the cuts that began in September 2024. The paper treats February 2022, the last month before liftoff, as the base month of every design.

ChatGPT. OpenAI released ChatGPT on 30 November 2022. Adoption was rapid by the standards of general-purpose technologies: [Bick et al. \[2024\]](#) document that within two years a substantial minority of American workers used generative AI at work, with the highest rates among the young, the college-educated and workers in computing, management and business occupations. Experiments established large productivity gains in writing and customer-support tasks [[Noy and Zhang, 2023](#), [Brynjolfsson et al., 2025b](#)], and exposure measures constructed before the release [[Felten et al., 2021, 2023](#), [Eloundou et al., 2024](#)] identified legal, financial, administrative and software occupations as the most exposed. The paper dates the AI regime from December 2022 and, in the structural model, allows adoption to follow a logistic path with a midpoint eighteen months after release.

The pandemic hangover. Employment in the contact-intensive services collapsed in the spring of 2020 and recovered over the following two years, while employment in remote-capable industries [[Dingel and Neiman, 2020](#)] barely fell and then overshot its pre-pandemic trend as low rates, fiscal transfers and the shift to digital services raised demand for their output. By February 2022 the standard deviation across industries of the log deviation from the 2015 to 2019 trend was 8.6 log points, with restaurants and hotels still far below trend and warehousing, delivery, software and some professional services far above it. Industries

starting above trend would tend to fall back and industries below trend would tend to catch up, regardless of AI or policy; the paper treats this displacement as a third exposure and estimates its effect alongside the other two rather than assuming it away.

Section 174. A fourth event is harder to measure. The Tax Cuts and Jobs Act of 2017 required firms, for tax years beginning after 31 December 2021, to amortize research and experimental expenditures over five years (fifteen for foreign research) instead of deducting them immediately. The provision took effect in 2022, raised the effective tax cost of research-intensive activity, and was widely blamed in the software industry for layoffs of engineering staff. Legislation enacted in 2025 relaxed the requirement for domestic research. No public data source reports research intensity by detailed industry at monthly frequency, and the paper cannot measure this exposure directly. [Section 4](#) discusses the resulting confound and [Section 9](#) shows that dropping the research-intensive industries does not change the results.

2.2 AI and Jobs

The modern literature on technology and labor demand is organized around tasks. [Autor et al. \[2003\]](#) showed that computerization substituted for routine tasks and complemented non-routine analytical ones; [Autor \[2015\]](#) and [Acemoglu and Restrepo \[2019\]](#) developed the framework in which automation displaces labor from existing tasks while new tasks reinstate it, so that the net employment effect of a technology is an empirical question, and [Acemoglu and Restrepo \[2020\]](#) provided the leading example of a robust displacement effect, for industrial robots. For artificial intelligence, [Webb \[2020\]](#) and [Felten et al. \[2021\]](#) built occupation-level exposure measures from the overlap between AI capabilities and occupational task descriptions, and [Acemoglu et al. \[2022\]](#) found that exposed establishments changed the composition of their vacancy postings without measurable effects on employment. Large language models motivated a second generation of measures: [Felten et al. \[2023\]](#) reweighted their index toward language-modeling capabilities, and [Eloundou et al. \[2024\]](#) classified tasks by whether a language model could halve the time needed to perform them. Both find that the exposed occupations are white-collar, well paid and concentrated in finance, law, software and professional services, a reversal of the pattern of earlier automation.

Evidence on realized employment effects is recent and divided. [Brynjolfsson et al. \[2025a\]](#) use payroll microdata and report that employment of workers aged 22 to 25 in the most exposed occupations fell by about thirteen percent relative to less exposed occupations after late 2022, while older workers in the same occupations were unaffected. [Hosseini and Lichtinger \[2025\]](#) find in resume and posting data that firms adopting generative AI reduced junior hiring while retaining senior staff. Against this, [Humlum and Vestergaard \[2025\]](#) match chatbot adoption to Danish administrative records and find small effects on earnings and

hours; [Gimbel et al. \[2025\]](#) find that the occupational mix of employment has changed no faster since ChatGPT than during earlier technology waves; and [Eckhardt and Goldschlag \[2025\]](#) find no unusual unemployment in exposed occupations. At the firm level, [Babina et al. \[2024\]](#) find that firms investing in AI grew faster in employment before the generative wave, and [Eisfeldt et al. \[2023\]](#) find that the equity values of firms with high generative-AI exposure rose after the release of ChatGPT, both consistent with a productivity rather than a displacement reading in the short run; [Albanesi et al. \[2023\]](#) find that AI-enabled automation was associated with employment growth in European occupations over the 2010s, and [Acemoglu \[2025\]](#) argues that the macroeconomic effects of the current generation of models over a horizon of a few years should be modest. The estimands in this literature differ: occupation-level studies compare workers across occupations, firm-level studies compare adopters with non-adopters, and the present paper compares industries by the exposure of their occupational mix. An industry-level result neither validates nor overturns an occupation-level one when occupational composition changes within industries, and the paper reads its findings as complements rather than as tests of the microdata results. Its contribution is to confront AI exposure with the two competing explanations on the same units and over the same window, and to observe the timing of the young-worker gradient over a longer horizon than the post-2022 comparisons in the existing work.

2.3 Monetary Transmission across Industries

That monetary policy affects industries unequally is an old and robust finding. [Gertler and Gilchrist \[1994\]](#) showed that small manufacturing firms account for a disproportionate share of the decline in activity after tightening, and [Bernanke and Gertler \[1995\]](#) organized the mechanisms into the credit channels. [Dedola and Lippi \[2005\]](#) and [Peersman and Smets \[2005\]](#) estimated industry-level responses to monetary shocks across OECD and euro-area industries and related them to durability, capital intensity and financial structure; [Ottonello and Winberry \[2020\]](#) showed that the investment response is concentrated among firms with room to expand when the cost of capital falls; and [Fort et al. \[2013\]](#) showed that young and small firms are the most cyclically sensitive employers. These findings motivate a rate-sensitivity measure estimated directly from each industry’s employment response to identified policy shocks, with a cyclical beta and the young-firm employment share as alternatives.

The identification of monetary shocks has its own literature: [Romer and Romer \[2004\]](#) derived shocks from a policy rule fitted to the Federal Reserve’s own forecasts, [Gertler and Karadi \[2015\]](#) used high-frequency changes in interest-rate futures around announcements, and [Ramey \[2016\]](#) surveyed the resulting estimates and their instability. Two recent contributions bear directly on the design. [Jarociński and Karadi \[2020\]](#) showed that raw announcement

surprises mix policy shocks with central-bank information effects and separated the two by the co-movement of rates and stock prices; [Bauer and Swanson \[2023\]](#) showed that even these surprises are predictable from macroeconomic news available before the announcement, and orthogonalized them against it. The baseline uses the Bauer-Swanson surprises because, in the local projections that generate the measure, they produce the correct sign of the employment response in 79 percent of industries, whereas the Jarocinski-Karadi surprises produce the familiar employment puzzle. Local projections follow [Jordà \[2005\]](#); the empirical-Bayes shrinkage follows [Morris \[1983\]](#).

2.4 The Discount-Rate Channel

The paper’s structural interpretation rests on the observation that a filled job is an asset. In the search-and-matching framework of [Mortensen and Pissarides \[1994\]](#) and [Pissarides \[2000\]](#), free entry equates the expected cost of recruiting a worker to the present value of the surplus the match will generate. [Shimer \[2005\]](#) showed that productivity shocks move this value too little to explain observed fluctuations in vacancies and unemployment; [Hall \[2017\]](#) argued that discount-rate movements do the work, and [Kehoe et al. \[2023\]](#) formalized the mechanism in a model in which the match is a long-duration asset. Monetary tightening is a direct movement in the discount rate, and cross-industry variation in the duration of matches implies cross-industry variation in the hiring response. [Section 3](#) builds this into a model with an AI channel and [Section 8](#) estimates it.

Two further strands inform the age analysis. The literature on graduating into a recession [[Kahn, 2010](#), [Oreopoulos et al., 2012](#)] established that entry-level workers bear a disproportionate and persistent share of cyclical adjustment, and [Forsythe \[2022\]](#) showed that this occurs because firms shift the composition of their reduced hiring toward experienced workers when labor is abundant. The Quarterly Workforce Indicators used here are described by [Abowd et al. \[2009\]](#).

3. Conceptual Framework

This section sets out a deliberately simple multi-industry stock-flow model of hiring in which monetary policy operates through the discount rate and generative AI operates through the surplus of a match. The model is used for three purposes: to state the identification levers of the empirical design as predictions, to interpret the estimated event-study paths, and, in [Section 8](#), to decompose the employment shortfall of exposed industries into rate, AI and mean-reversion components. [Section A](#) derives the equations.

3.1 The Value of a Filled Job

Consider industry i with a continuum of firms that hire in a frictional labor market. A filled job produces a flow surplus $p_i - w_i$, the gap between the worker’s marginal product and the wage, which grows at rate g_i over the life of the match. Matches dissolve at rate s_i , and firms discount at the rate r . The value of a filled job is the capitalized surplus,

$$J_i = \frac{p_i - w_i}{r + s_i - g_i}, \quad (1)$$

so that the semi-elasticity of the job value with respect to the discount rate is $-1/(r + s_i - g_i)$. Industries in which matches are long-lived (low s_i) and in which the surplus grows over the life of the match (high g_i) have job values that are more sensitive to r : they are long-duration assets. This is the discount-rate channel of [Hall \[2017\]](#) and [Kehoe et al. \[2023\]](#) applied across industries, and the industries emphasized in the AI narrative plausibly have low separation rates for professional staff, steep tenure-earnings profiles, and, for young firms, a financing structure that ties the discount rate to the equity market.

Generative AI enters through the surplus. Let a_i be industry i ’s exposure, the share of its tasks that language models can perform, and let $A_t \in [0, 1]$ be the economy-wide adoption of the technology at date t . Adoption changes the flow surplus of a match in exposed industries by an amount proportional to $a_i A_t$, with a sign that theory does not determine: if the technology mainly raises the productivity of the workers who use it, exposed industries hire more; if it mainly substitutes for tasks that workers performed, the surplus of a marginal match falls and they hire less [[Acemoglu and Restrepo, 2019](#)]. The sign of the AI channel is left to be estimated.

Combining the two channels in log-linear form, and writing x_t for the change in the relevant market interest rate since the base month and ρ_i for industry i ’s (standardized) discount-rate sensitivity,

$$\ln J_{it} = -\kappa_r \rho_i x_t + \kappa_a a_i A_t + \phi_t, \quad (2)$$

where ϕ_t collects everything that moves job values in all industries alike and κ_r and κ_a are the two structural semi-elasticities to be estimated. Because the empirical design compares industries at a point in time, ϕ_t is absorbed by the month effects and is not identified; the model is silent on aggregate effects, a limitation that [Nakamura and Steinsson \[2018\]](#) and [Chodorow-Reich \[2020\]](#) discuss for cross-sectional macro designs generally and that is taken up in [Section A](#).

3.2 Hiring and the Law of Motion

Firms post vacancies until the expected cost of filling one equals the value of a filled job. With a constant-returns matching function with elasticity α with respect to unemployment [Petrongolo and Pissarides, 2001], free entry implies that the flow of hires is proportional to J_{it}^η with $\eta = (1 - \alpha)/\alpha$. Measuring hires and employment relative to the industry’s steady-state employment, and normalizing so that hiring exactly replaces separations when $\ln J_{it} = 0$,

$$h_{it} = s \exp(\eta \ln J_{it}), \quad (3)$$

where h_{it} is hires as a fraction of steady-state (not current) employment, and employment relative to its steady-state level evolves as

$$n_{i,t+1} = (1 - s) n_{it} + h_{it}. \quad (4)$$

The pandemic overshoot enters as the initial condition: at the base month, $\ln n_{i,\text{base}} = d_i$, the log deviation of the industry’s employment from its pre-pandemic trend. With $\ln J = 0$, equation (4) implies that an initial displacement decays geometrically at rate s per month, with a half-life of $\ln(0.5)/\ln(1 - s)$ months; s therefore governs both the speed of mean reversion and the speed with which changes in the job value reach the stock of employment. Because measured separation rates include many short spells that do not correspond to the stock adjustment the model describes, s is treated as an adjustment speed to be estimated rather than as observed turnover; Section 8 reports a variant in which it is proportional to the industry’s observed separation rate. The baseline sets $\alpha = 0.5$, the center of the range that Petrongolo and Pissarides [2001] survey, so that $\eta = 1$.

3.3 Predictions

The model generates three predictions that correspond to the three identification levers of the empirical design.

Prediction 1 (Timing). *Employment in rate-sensitive industries begins to fall relative to other industries after liftoff, with a lag governed by s ; employment in AI-exposed industries cannot diverge before adoption begins.* In equation (2), x_t turns positive in March 2022 and A_t is zero until December 2022. Because hiring responds immediately but the stock adjusts at rate s , the rate gap opens gradually over 2022 and 2023; the AI gap, whatever its sign, opens only as A_t rises through 2023 and 2024.

Prediction 2 (Sign reversal). *When the interest rate falls, the growth disadvantage of*

rate-sensitive industries narrows and eventually reverses; the growth differential of AI-exposed industries does not reverse while adoption continues. The rate channel is proportional to x_t , a level, but the stock adjusts toward the level it implies only gradually: in [equation \(4\)](#) the employment gap keeps widening for as long as the stock is above the target that the current x_t implies, so a gap can continue to open after the yield stops rising, and it narrows in growth terms only once the yield has fallen enough to lift the target above the stock. The prediction is therefore about growth differentials, not levels, and it is a prediction of attenuation before reversal. The AI channel is proportional to A_t , which is monotone, so the AI growth differential keeps its sign throughout the adoption period.

Prediction 3 (Age gradient). *Both channels operate on the hiring margin, so both predict that young workers, who are disproportionately new hires, bear a larger share of the adjustment than incumbents; the channels differ in the timing of the youth gradient and in the pattern of flows.* Under the rate channel the youth gradient in a rate-sensitive industry should break at liftoff; under the AI channel it should break with adoption, after December 2022. Under a hiring-freeze mechanism, hiring rates fall and, as [Forsythe \[2022\]](#) shows, the composition of hires shifts toward experienced workers; under a reallocation mechanism, in which AI changes the tasks and the mix of workers that exposed firms want, both hiring and separation rates rise.

The model has an obvious limitation as a description of AI: it treats adoption as a single economy-wide curve and the effect as proportional to a fixed exposure. If the empirical AI path does not look like an adoption curve, the model will not fit it, which is what [Section 8](#) finds.

4. Data

4.1 The Industry Employment Panel

The outcome data are seasonally adjusted all-employee counts from the Current Employment Statistics (CES) survey of the Bureau of Labor Statistics, retrieved for each industry from the BLS public API. The panel covers 205 detailed private industries observed monthly from January 2015 to July 2026, 139 months, and the industries together employed 120.1 million workers in 2019. Detailed CES industries publish one month behind the headline release, so the panel ends in the last month in which at least 95 percent of series report; no industry has a gap within the sample and none is dropped.

The set of industries is defined by the AI exposure data. The language-model AI industry exposure index of [Felten et al. \[2023\]](#) is published for four-digit NAICS industries, and the

corresponding CES series were requested for every code on that list. The CES does not publish every four-digit industry; where no four-digit series exists, the three-digit parent stands in for its children, with the parent’s exposure the simple mean of the children’s. The final panel comprises 173 four-digit industries and 32 series at the three-digit level, of which 18 are industries with no published four-digit child (rail and pipeline transportation, warehousing and storage, and petroleum products are examples) and 14 are composite codes that the AIIE list carries as residuals, such as “Chemical manufacturing (3251, 3252, 3253 and 3259 only)”. For 8 of these composites the excluded children are in the panel in their own right, and to avoid counting any job twice the composite’s employment is the CES three-digit parent minus its listed children. The information and retail sectors were recoded from NAICS 2017 to NAICS 2022 in the CES in 2023, whereas the AIIE is on the 2017 basis; a hand crosswalk maps the changed codes, with many-to-one merges again taking the simple mean. [Section B](#) lists the residual composites and the crosswalk.

The outcome throughout is $100 \times \log$ employment net of the industry’s own linear trend estimated on January 2015 to December 2019. Detrending removes the large secular differences in growth between, for example, computer systems design and printing, which would otherwise be attributed to whichever exposure they happen to correlate with. The raw series are retained, and results without detrending and with a trend estimated on 2010 to 2019 are reported in [Section 9](#).

4.2 AI Exposure

The AI exposure measure is the language-modeling AI industry exposure (LM-AIIE) of [Felten et al. \[2023\]](#). Its construction starts from the AI occupational exposure index of [Felten et al. \[2021\]](#), which links AI capabilities to the abilities required by each occupation in the O*NET database, reweights the capabilities toward language modeling, and aggregates the occupational scores to industries using each industry’s occupational employment mix. The measure is predetermined in the sense that matters for the design: its inputs, the O*NET ability requirements, the capability weights and the occupational employment mix, all predate the release of ChatGPT and do not respond to post-2022 developments. The index itself was posted in March 2023, three months after the release and before almost all of the employment changes studied here, as was the index of [Eloundou et al. \[2024\]](#); the general AIIE of [Felten et al. \[2021\]](#) was published in 2021. The general AIIE is used in robustness; it is correlated 0.985 with the language-modeling version across the 205 industries, so the two cannot be distinguished empirically. Neither index is a literal share of automatable tasks; each is a standardized score of how closely an industry’s occupational abilities match what language models do, and the paper treats it as an ordering of industries by exposure rather than as a

quantity of tasks. [Table 12](#) in [Section B](#) lists the ten most and ten least exposed industries on each measure.

Both measures are standardized to mean zero and standard deviation one across the 205 industries, unweighted. The most exposed industries are legal services (2.2 standard deviations above the mean), securities and investment activities, insurance agencies, nondepository credit intermediation, insurance carriers, accounting, business support services and software publishers, all between 1.9 and 2.1. The least exposed are logging, foundation contractors, services to buildings, truck transportation, warehousing and animal slaughtering, all between -1.5 and -1.8 . The interquartile range of exposure is 1.55 standard deviations, a yardstick for the magnitudes reported below.

4.3 Monetary Surprises and Rate Sensitivity

The rate-sensitivity exposure is estimated, for each industry, from its own historical employment response to identified monetary policy shocks. The baseline shock is the orthogonalized monetary policy surprise of [Bauer and Swanson \[2023\]](#), the first principal component of high-frequency changes in short-term interest-rate futures around FOMC announcements, orthogonalized against macroeconomic and financial news available before each announcement. The alternative is the pure monetary policy shock of [Jarociński and Karadi \[2020\]](#), which separates policy shocks from information effects by the sign of the co-movement between rates and equity prices. Both are aggregated to monthly frequency and standardized to unit variance over the estimation sample.

For each industry i , the following local projection [[Jordà, 2005](#)] is estimated on monthly data from September 1990 to December 2016, so that the longest horizon ends in December 2019 and no pandemic month enters:

$$\frac{1}{5} \sum_{h \in \mathcal{H}} (y_{i,t+h} - y_{i,t-1}) = \alpha_i + \beta_i \varepsilon_t + \sum_{l=1}^6 \left(\gamma_{il} \Delta y_{i,t-l} + \delta_{il} \varepsilon_{t-l} + \theta_{il} \Delta y_{t-l}^{agg} \right) + u_{it}, \quad (5)$$

where y_{it} is $100 \times \log$ employment, ε_t is the standardized shock, y_t^{agg} is $100 \times \log$ total nonfarm employment, and $\mathcal{H} = \{12, 18, 24, 30, 36\}$. The dependent variable is the average cumulative response at horizons of one to three years, the window over which employment responses to policy are typically largest; averaging over five horizons reduces the noise of any single-horizon estimate. Standard errors use the [Newey and West \[1987\]](#) estimator with a bandwidth equal to the longest horizon. The raw rate sensitivity is $-\hat{\beta}_i$, so that a positive number means a larger employment loss after a contractionary surprise.

The industry-level estimates are noisy. Every industry has the same 316 usable months;

the Newey-West standard errors average 0.27 log points against a cross-industry standard deviation of the raw estimates of 0.41, and a block bootstrap over months (blocks of 24, the same blocks for every industry) gives standard errors averaging 0.31, so the HAC estimator is close to the resampling one. The estimates are also correlated across industries, since every industry is projected on the same shock months: the mean absolute correlation of the bootstrap errors across industry pairs is 0.36. The choice of horizons, bandwidth and lag length matters little for the ranking of industries: sensitivities built from horizons 12 to 24 only, from 24 to 36 only, with a bandwidth of 18 or with three lags are correlated 0.95, 0.98, 0.99 and 0.95 with the baseline after shrinkage. Local projections at long horizons on persistent series are the setting in which [Montiel Olea and Plagborg-Møller \[2021\]](#) show that lag augmentation makes conventional inference robust; the six lags of the industry’s own growth serve that purpose here.

Treating noisy estimates as data attenuates the coefficients on rate sensitivity in every subsequent regression and, separately, understates their uncertainty. The paper handles the two problems differently. Against attenuation, the raw estimates are shrunk toward the precision-weighted cross-industry mean by parametric empirical Bayes [[Morris, 1983](#)]: with $\hat{\tau}^2 = \max\{\text{Var}(\hat{\beta}) - \overline{\text{se}}^2, 0\}$ the estimated variance of true sensitivities, industry i ’s shrunk sensitivity is $\bar{\beta} + \lambda_i(\hat{\beta}_i - \bar{\beta})$ with $\lambda_i = \hat{\tau}^2/(\hat{\tau}^2 + \text{se}_i^2)$. The reliability ratio $\hat{\tau}^2/(\hat{\tau}^2 + \overline{\text{se}}^2)$, the share of the cross-industry variance of the raw estimates that is signal, is 0.39. Shrinkage with heterogeneous factors reorders industries toward their better-measured neighbors; it is a prediction device, not a guarantee of unbiasedness, and the paper makes no lower-bound claim on its account. Against understated uncertainty, [Section 5](#) propagates the sampling covariance of the 205 raw estimates through the shrinkage, the standardization and the second-stage regressions in the manner of [Murphy and Topel \[1985\]](#). A split-sample instrumental-variables estimate, in which the raw sensitivity estimated on shocks from September 1990 to June 2000 is instrumented by the one estimated on July 2003 to December 2016, with a 36-month gap so that no outcome month is shared between the two, is reported in [Section 9](#); its first stage turns out to be empty, which is itself informative about the stability of the measure.

The choice of the Bauer-Swanson surprises as the baseline was made on measurement diagnostics before any outcome regression was run. The local projections on the Jarocinski-Karadi shocks display the familiar employment puzzle of high-frequency surprises: the mean industry response to a contractionary shock is essentially zero (0.03 log points per standard-deviation shock) and only 54 percent of industries show the correct sign. The Bauer-Swanson surprises give a mean employment loss of 0.29 log points per standard-deviation shock, the correct sign in 79 percent of industries, and a slightly higher reliability ratio (0.39 against 0.35). A measure built from responses that are mostly of the wrong sign would not measure

rate sensitivity, so the Jarocinski-Karadi analogue is a robustness check rather than the baseline. The two shrunk measures are correlated 0.57 across industries.

Three further measures of exposure to the cost of capital are used in robustness: the cyclical beta, the slope of the industry’s twelve-month employment growth on aggregate twelve-month growth over 1990 to 2019, which depends on no identified shock but measures sensitivity to all aggregate shocks, and which is measured with a reliability ratio of 0.91; the young-firm share, the 2019 share of the industry’s employment at firms aged zero to five years from the Business Dynamics Statistics, a financial-conditions proxy in the spirit of [Fort et al. \[2013\]](#) that is not estimated from employment responses at all and is available for 190 industries; and the first principal component of the four measures, which accounts for 49 percent of their joint variance. All are standardized.

Rate sensitivity so measured is highest in residential building construction (2.9 standard deviations), employment services, newspaper and periodical publishing, logging, other transportation equipment and wood products, and lowest in grantmaking services, food manufacturing, psychiatric hospitals and petroleum refining. The interquartile range is 1.38 standard deviations.

4.4 Pandemic Displacement

Pandemic displacement is the log deviation of the industry’s February 2022 employment from its 2015 to 2019 linear trend: positive for an overshoot, negative for an unfinished rebound. Its mean across the 205 industries is -5.1 log points and its standard deviation 8.6, ranging from leisure and hospitality industries still 15 log points below trend on average to a few industries well above it. It is standardized like the other exposures. Because the outcome is measured relative to the same trend, a negative coefficient on displacement means that industries above trend at the base month subsequently lost ground relative to industries below it: mean reversion in both directions. Displacement is not a treatment in the sense the other two exposures are. It is the baseline level of the outcome itself, and a regression of subsequent changes on a baseline level picks up regression to the mean from any transitory component or measurement error in the baseline, as well as any anticipation of the tightening that was already in February 2022 employment. Its coefficient is therefore read as the pace at which industries returned toward trend, whatever the mix of reasons, and the paper checks that the rate and AI results do not depend on it: column (2) of [Table 2](#) omits it, and [Section 9](#) re-estimates the trend, and hence the displacement, on 2010 to 2019.

4.5 Correlation of the Exposures

The horse race is only informative if the exposures differ. [Figure 1](#) plots AI exposure against rate sensitivity for the 205 industries, with marker area proportional to 2019 employment. The two are nearly orthogonal: the correlation is -0.06 . AI exposure is correlated 0.07 with pandemic displacement and -0.30 with the cyclical beta, the latter because the most cyclical industries, construction and durable manufacturing, are among the least exposed to language models. Rate sensitivity and displacement are correlated -0.06 , and none of the three exposures is correlated with the 2015 to 2019 pre-trend (correlations of -0.01 , -0.10 and -0.02). The industries emphasized in the AI narrative are scattered across the rate-sensitivity distribution: software publishers are below the mean, legal services near it, and insurance carriers and nondepository credit well above it. The confound the paper set out to test, that AI-exposed industries are also the rate-sensitive ones, is not present in the historical employment sensitivities.




Figure 1: AI Exposure and Rate Sensitivity across 205 Industries

Notes: Each marker is one of the 205 detailed private industries; marker area is proportional to 2019 employment. The horizontal axis is the language-modeling AI industry exposure of [Felten et al. \[2023\]](#); the vertical axis is the shrunk local-projection estimate of the industry's employment response to [Bauer and Swanson \[2023\]](#) monetary surprises, both standardized across industries. The cross-industry correlation is -0.06 . Labeled industries are illustrative.

4.6 The Quarterly Workforce Indicators

The age analysis uses the Quarterly Workforce Indicators (QWI) of the Census Bureau’s Longitudinal Employer-Household Dynamics program [[Abowd et al., 2009](#)], which report employment, hires and separations by detailed industry and worker age group from linked employer-employee records. The QWI are published by state with no national endpoint, so national series were built by summing the state-by-industry cells, which weights every state by its employment by construction. Data were retrieved for private ownership at the four-digit NAICS level for the age groups 22 to 24, 25 to 34, 35 to 44, 45 to 54 and 55 to 64 for every quarter from 2015Q1 to 2025Q4. A state enters only if it reports from the first quarter through the modal last quarter, which excludes Alaska and Michigan and leaves 49 state-level jurisdictions including the District of Columbia; within those, a state-by-industry cell is kept only if employment is reported for all five age groups in every quarter, so that the aggregates are not moved by cells entering and leaving. The balanced cells hold 99.6 percent of the employment reported in the kept states, and 204 of the 205 industries are matched. The QWI are not seasonally adjusted, so each outcome is purged of industry-specific quarter-of-year effects and of an industry-specific linear trend, both estimated on 2015 to 2019.

The outcomes constructed for each industry and quarter are: $100 \times$ the log of employment aged 22 to 24, 25 to 34, and 22 to 34, each relative to employment aged 35 to 64; $100 \times$ log employment of each age group and of ages 22 to 64; the hiring rate, hires during the quarter as a percent of employment, and the separation rate, each for all ages and separately for ages 22 to 24 and 25 to 34; and the share of all hires accounted for by each age group. In 2019 the mean hiring rate across industries was 14.2 percent of employment per quarter for all ages and 29.3 percent for 22 to 24 year olds, the mean separation rates 14.1 and 25.7 percent, and the mean share of hires aged 22 to 24 was 12.8 percent. Because a log ratio moves when either age group’s employment moves, and because a cohort hired at 23 ages out of the 22 to 24 group within two years, the levels of the two age groups and the age-specific flows are reported alongside the ratio. The base quarter for the youth event study is 2022Q1, the quarter containing the base month of the monthly design.

4.7 Other Series

The two-year Treasury yield and the federal funds rate are monthly averages from FRED. The two-year yield enters the continuous-stance specification of [Section 5](#) as the twelve-month change lagged six, twelve or eighteen months, and enters the model of [Section 8](#) as the change since February 2022. The young-firm employment shares come from the 2019 Business Dynamics Statistics by four-digit NAICS.

4.8 Section 174 as an Unmeasured Confounder

The one shock the paper cannot measure is the 2022 change in the tax treatment of research expenditures. Research intensity is plausibly correlated with AI exposure: software publishers, computer systems design and data processing are both research-intensive and among the most exposed industries. If the amortization requirement reduced employment in research-intensive industries after 2022, it would appear in the horse race as a negative AI effect. No public source reports research spending by detailed industry at a level of detail that could be merged with the panel; the National Science Foundation’s business research surveys cover broad industry groups with a lag of several years. The paper therefore treats Section 174 as a stated limitation and reports a robustness check that drops the thirteen industries that account for the bulk of business research spending in the NSF surveys (pharmaceuticals, computer and electronic products, motor vehicles, aerospace, software publishers, computer systems design, scientific research services and computing infrastructure). Because the estimated AI effect on headcount is positive, an unmeasured negative Section 174 effect would, if anything, imply that the true AI coefficient is more positive than reported.

4.9 Summary Statistics

[Table 1](#) reports the composition of the panel by CES sector, with sector cells the unweighted means of the standardized exposures. Financial activities (1.35), information (1.22), education and health (0.87) and professional and business services (0.55) have the highest average AI exposure; construction (−1.04), mining (−1.05) and nondurable manufacturing (−0.87) the lowest. Rate sensitivity is highest on average in construction (0.63), financial activities (0.57) and professional and business services (0.34) and lowest in utilities (−1.39). The sector means show why the two exposures are uncorrelated at the industry level: the sectors that are high in one are a mixture in the other, and there is wide dispersion within sectors. Pandemic displacement in February 2022 was most negative in leisure and hospitality (−15.0 log points) and construction and mining (−8.0), and positive only in utilities. The 2015 to 2019 pre-trend averaged 1.37 percent per year with a standard deviation of 2.52, the variation that detrending removes.

Table 1: Industry Panel: Summary Statistics by Sector

Sector	N	Emp. 2019 (mil.)	AI exposure	Rate sensitivity	Pandemic displacement	Pre-trend (% p.a.)
Education and health services	24	24.2	0.87	-0.39	-6.3	2.27
Professional and business services	20	18.9	0.55	0.34	-3.9	2.15
Leisure and hospitality	13	16.6	-0.19	0.02	-15.0	2.39
Retail trade	18	12.0	0.17	-0.10	-2.7	0.05
Financial activities	10	8.8	1.35	0.57	-4.1	1.57
Durable goods	33	8.0	-0.39	0.30	-3.4	0.56
Construction	10	7.5	-1.04	0.63	-8.0	3.24
Wholesale trade	8	5.9	0.30	-0.14	-0.3	-0.04
Transportation and warehousing	19	5.7	-0.51	0.03	-6.2	2.74
Nondurable goods	21	4.4	-0.87	-0.47	-3.3	0.44
Other services	12	4.2	-0.05	-0.34	-6.2	1.35
Information	8	2.9	1.22	0.07	-3.0	2.01
Mining and logging	6	0.7	-1.05	0.05	-8.0	-1.75
Utilities	3	0.5	-0.03	-1.39	1.1	0.33
All industries	205	120.1	0.00	0.00	-5.1	1.37
standard deviation			1.00	1.00	8.6	2.52

Notes: The panel has 205 detailed private industries (CES, seasonally adjusted, 2015m1–2026m07) covering 120.1 million jobs in 2019. AI exposure is the language-modeling AIIE of [Felten et al. \[2023\]](#); rate sensitivity is minus the average 12–36 month employment response to a one-s.d. [Bauer and Swanson \[2023\]](#) monetary surprise, estimated by local projections on 1990–2019 and shrunk by empirical Bayes. Both are standardized across industries (sector cells are unweighted means). Pandemic displacement is the log-point deviation of February 2022 employment from the industry’s 2015–2019 trend. Pre-trend is the 2015–2019 annual log growth.

5. Empirical Strategy

5.1 Event Study

Let $\tilde{y}_{it}$ denote $100 \times$ log employment of industry i in month t net of the industry’s 2015 to 2019 linear trend, and let t_0 be February 2022. The event study estimates, for each month t in the panel, the cross-sectional regression

$$\tilde{y}_{it} - \tilde{y}_{it_0} = \alpha_t + \beta_t^{AI} AI_i + \beta_t^{RS} RS_i + \beta_t^D D_i + e_{it}, \quad (6)$$

where AI_i , RS_i and D_i are the standardized AI exposure, rate sensitivity and pandemic displacement. The coefficient β_t^{AI} is the difference in detrended employment growth since February 2022 between an industry one standard deviation more AI-exposed and an otherwise identical industry, holding the other two exposures fixed. In a balanced panel the sequence of cross-sectional long-difference coefficients is numerically identical to the coefficients on exposure-by-month interactions in a two-way fixed effects regression with industry and

month effects; the replication code verifies the identity for the pooled version below to machine precision. Because each month’s regression is a cross-section of 205 industries, heteroskedasticity-robust standard errors are the industry-clustered standard errors of the panel version.

The months before February 2022 provide the pre-trend check. Coefficients for 2015 to 2019 are mechanically close to zero because the outcome is detrended on that window, so the informative pre-period is the year before liftoff, March 2021 to January 2022. [Roth \[2022\]](#) and [Rambachan and Roth \[2023\]](#) caution that passing a pre-trend test does not certify parallel trends and that conditioning on passing biases the post-period estimates; the paper reports the pre-period coefficients as descriptive evidence rather than as a gate, and relies on the placebo-in-time distribution described below as its main statement of what the design’s null looks like.

5.2 Pooled Regime Difference-in-Differences

The pooled specification groups the months after February 2022 into three policy regimes: tightening before ChatGPT (March to November 2022), tightening and hold after ChatGPT (December 2022 to August 2024), and easing (September 2024 to July 2026). With R_{kt} an indicator for regime k and the reference period the year before liftoff (March 2021 to February 2022),

$$\tilde{y}_{it} = \mu_i + \lambda_t + \sum_{k=1}^3 R_{kt} \left(\beta_k^{AI} AI_i + \beta_k^{RS} RS_i + \beta_k^D D_i \right) + e_{it}, \quad (7)$$

estimated on the 65 months of the reference period and the three regimes, 13,325 industry-month observations, with standard errors clustered by industry. Each β_k is the average of the event-study coefficients over the months of regime k , less their average over the reference year. The regime split serves the timing lever directly: the first regime is tightening without AI, the second is tightening and hold with AI, the third is easing with AI.

The specification is a difference-in-differences with a continuous treatment, and the identifying assumption is the strong parallel trends assumption that [Callaway et al. \[2024\]](#) show is required to interpret the coefficients on a continuous dose as average causal responses: absent the shocks, industries with different exposures would have continued on their own pre-2020 trends, conditional on the other exposures, and the path of an industry with exposure d is the path it would have followed at any other dose. The assumption is stronger than that of a binary design because it rules out selection into dose on the basis of the treatment effect. Here the doses are predetermined characteristics, so the concern is not selection but the possibility that industries sorted on a characteristic drift for other reasons; the placebo-in-time exercise measures that drift. The design resembles a shift-share regression in which the shocks

(the policy path and adoption) are common to all industries and the identifying variation is in the exposure shares, so that, as [Goldsmith-Pinkham et al. \[2020\]](#) emphasize, the exogeneity of the shares is what matters, rather than the exogeneity of the shocks that [Borusyak et al. \[2022\]](#) exploit. The problems with staggered adoption and heterogeneous effects in two-way fixed effects designs [[Goodman-Bacon, 2021](#), [de Chaisemartin and D’Haultfœuille, 2020](#)] do not arise because every industry is treated at the same date.

5.3 Growth Form and the Sign-Reversal Test

The sign-reversal lever concerns growth rates rather than levels. For each regime k , including the reference year, the annualized detrended growth of industry i over the regime is

$$g_{ik} = \frac{12}{m_k} \left(\tilde{y}_{i,\text{end}(k)} - \tilde{y}_{i,\text{start}(k)-1} \right), \quad (8)$$

with m_k the number of months, and it is regressed on the three exposures across the 205 industries; the growth coefficients are the slopes of the event-study paths within each regime. The sign-reversal statistic is the coefficient on each exposure in a regression of $g_{i,\text{easing}} - g_{i,\text{hold}}$ on the exposures: a positive value for rate sensitivity means the growth disadvantage of sensitive industries narrowed once policy eased. The ChatGPT-break statistic is the analogous coefficient for $g_{i,\text{hold}} - g_{i,\text{tight,pre}}$: a negative value for AI exposure would mean that the growth of exposed industries deteriorated when the technology arrived.

5.4 Continuous Policy Stance

Regimes are coarse. A continuous version replaces the regime indicators by the policy stance itself. Let $\Delta^{12}r_{t-\ell}$ be the twelve-month change in the two-year Treasury yield, in percentage points, lagged ℓ months, and let $\Delta\tilde{y}_{it}$ be the annualized monthly change in detrended log employment. The specification is

$$\begin{aligned} \Delta\tilde{y}_{it} = & \mu_i + \lambda_t + \gamma_r RS_i \Delta^{12}r_{t-\ell} + \gamma_a AI_i \mathbb{I}[t \geq \text{Dec 2022}] \\ & + \gamma_{ar} AI_i \Delta^{12}r_{t-\ell} + \gamma_d D_i \mathbb{I}[t \geq \text{Mar 2022}] + e_{it}, \end{aligned} \quad (9)$$

estimated on the full panel from February 2015 with industry and month effects and industry-clustered errors, for $\ell \in \{6, 12, 18\}$. The coefficient γ_r is the effect of a one-percentage-point rise in the two-year yield over the preceding year on the annualized growth of an industry one standard deviation more rate-sensitive, ℓ months later. The AI-by-stance interaction γ_{ar} tests whether AI-exposed industries respond to rates over and above what their measured rate sensitivity implies, the form the confound would take if the AIIE were secretly a rate-sensitivity

measure.

5.5 The Age Gradient

The QWI design mirrors the monthly one at quarterly frequency. For each of the seven outcomes w_{iq} defined in [Section 4](#), net of industry-specific seasonal effects and trends, the quarterly event study is

$$w_{iq} - w_{i,2022Q1} = \alpha_q + \beta_q^{AI} AI_i + \beta_q^{RS} RS_i + \beta_q^D D_i + e_{iq}, \quad (10)$$

and the pooled version regresses the regime mean minus the mean of the reference year 2021Q2 to 2022Q1 on the three exposures, with the hold regime 2023Q1 to 2024Q3 and the easing regime 2024Q4 onward. For the youth ratio, the timing analysis regresses the annualized change in the outcome over four windows, 2015Q1 to 2019Q4, 2020Q1 to 2022Q1, 2022Q1 to 2022Q4 and 2022Q4 to 2025Q4, on the exposures, and tests whether the slope after ChatGPT differs from the slope during 2020 and 2021. Robustness adds the industry’s own detrended youth ratio at 2022Q1 as a control (so that an industry that had over-hired young workers in the pandemic is not credited to AI), adds sector fixed effects, replaces rate sensitivity by the cyclical beta, drops the tech core and weights by 2019 employment.

5.6 Inference

Standard errors are clustered by industry in the panel specifications and heteroskedasticity-robust in the cross-sectional ones, which in a balanced panel is the same thing up to degrees-of-freedom corrections. Industry clustering accommodates serial correlation within an industry; it does not accommodate disturbances that are correlated across industries with similar exposure structures, the concern that [Adão et al. \[2019\]](#) raise for shift-share designs and that applies here because industries in the same sector share demand, occupational composition and financial conditions. Two devices address cross-industry dependence. Standard errors are re-computed with clustering by the fourteen CES sectors, and, because fourteen clusters are too few for the asymptotics of the cluster-robust estimator, wild-cluster bootstrap p -values with Rademacher weights and the null imposed [[Cameron et al., 2008](#)] are reported for every headline coefficient; sector-by-regime effects are examined as a specification. All event-study bands in the figures are pointwise; the pooled regime coefficients are the objects on which the paper’s tests rest.

The second issue is that rate sensitivity is estimated. Its sampling covariance across the 205 industries, obtained from the block bootstrap of the local projections described in [Section 4](#), is propagated through the shrinkage, the standardization and the cross-sectional

regressions by the delta method: the second-stage coefficients are differentiated numerically with respect to the 205 raw sensitivities, and the resulting Jacobian, applied to the first-stage covariance, gives the additional variance that the two-step structure adds to the clustered variance [Murphy and Topel, 1985]. The paper reports both the clustered and the two-step (Murphy-Topel) standard errors. A naive alternative, a bootstrap that re-estimates the sensitivities on resampled months and re-runs the second stage, is not centered on the estimate, because resampling adds a second dose of measurement error to every draw and attenuates each toward zero; it is computed and reported in the replication files but not used for inference.

Three further devices are descriptive rather than inferential, and are labeled as such. First, minimum detectable effects: for each coefficient, the effect that a test at the five percent level would detect with 80 percent power, 2.80 times the standard error; for AI exposure in the easing regime this is 2.32 log points per standard deviation with the clustered standard error. An MDE is a statement about power under the sampling model, not a confidence bound, and the paper reports intervals alongside it. Second, permutation reference distributions [Young, 2019]: the AI exposure vector is reshuffled across industries 2,000 times and the pooled regression re-estimated on each draw, holding rate sensitivity and displacement in place, with the p -value computed as $(k + 1)/(B + 1)$; a second version reshuffles only within CES sectors. A fixed industry characteristic is not randomly assigned, so these are reference distributions under exchangeability, not exact tests. Third, a placebo-in-time distribution. The entire design, with the same window geometry (a 60-month trend window ending 26 months before the base month, a twelve-month reference period ending at the base month, a hold analogue at months 10 to 30 after the base and an easing analogue at months 31 to 53), is replayed at nineteen fake base months, February 1997 through February 2015, using the same industries and the same AI and rate-sensitivity exposures, with displacement recomputed at each fake base. The windows overlap, so the nineteen coefficients are not independent draws, and the business-cycle exposure they reveal is a potential omitted variable rather than pure noise; what the distribution offers is a scale for how large a coefficient on a fixed industry characteristic arises over a four-year window in the absence of the shock. Because rate sensitivity is estimated on data overlapping these windows, its placebo distribution is not a null at all, but it does say how large a rate-sensitivity gap is typical of an ordinary cycle. The same replay is applied to the youth-ratio design at five fake base quarters in Section 7.

6. Main Results

6.1 The Event Study

Figure 2 is the paper’s central figure. It plots the monthly event-study coefficients on AI exposure and on rate sensitivity from equation (6), relative to February 2022, with pointwise 95 percent bands and vertical markers at liftoff (March 2022), ChatGPT (November 2022) and the first cut (September 2024). The displacement coefficient is omitted from the figure for scale; it is reported in the tables.




Figure 2: Employment since Liftoff: the Horse Race between AI Exposure and Rate Sensitivity

Notes: Each point is the coefficient on the standardized exposure in a cross-sectional regression of the change in detrended log employment ($\times 100$) since February 2022 on AI exposure, rate sensitivity and pandemic displacement across 205 industries (equation (6)). Solid: AI exposure. Dashed: rate sensitivity. Shaded areas are pointwise 95 percent confidence bands from heteroskedasticity-robust standard errors. Vertical lines mark liftoff (March 2022), the release of ChatGPT (November 2022) and the first rate cut (September 2024).

Three features stand out. First, neither exposure predicts anything in the year before liftoff: the mean of the AI coefficients over March 2021 to January 2022 is -0.18 log points and the mean of the rate-sensitivity coefficients is -0.23 , with pointwise bands that comfortably include zero. Second, the rate-sensitivity path is flat through the summer of 2022, turns down in the autumn, declines steadily through 2023 and 2024 to about -2 log points around the time of the first cut, drifts down more slowly thereafter, and stands at -2.44 log points

(industry-robust standard error 0.71) in July 2026, an endpoint gap that is larger than the easing-regime average reported below because the average spans the whole regime. This is the shape that [equation \(4\)](#) predicts for a shock to the job value that arrives with liftoff and accumulates in the stock of employment at the turnover rate. Third, the AI path is positive throughout, rises during 2022 before ChatGPT was released, and ends at +1.18 log points (standard error 0.88) in July 2026, with a band that includes zero in every month; there is no break at November 2022 and no acceleration as adoption spread. The displacement path (not shown) falls monotonically to -4.17 log points (standard error 1.08) by July 2026.

[Figure 3](#) shows the same information without regression, as the average detrended employment path of industries in the bottom, middle and top terciles of each exposure. The rate-sensitivity terciles fan out after 2022 in order, with the most sensitive tercile lowest; the AI terciles do not separate in any consistent order. The figure also shows the pandemic dynamics that the regressions control for: all terciles rise from the 2020 trough through 2021 and flatten in 2022 as the recovery completed.




Figure 3: Detrended Employment by Tercile of AI Exposure and of Rate Sensitivity
Notes: Unweighted mean of $100 \times \log$ employment net of the industry's 2015 to 2019 linear trend, normalized to zero in February 2022, for industries in the bottom, middle and top tercile of AI exposure (left) and of rate sensitivity (right). Vertical lines mark liftoff, the release of ChatGPT and the first rate cut.

6.2 The Pooled Horse Race

[Table 2](#) reports the pooled regime estimates from [equation \(7\)](#). Column (1) includes AI exposure alone, column (2) adds rate sensitivity, and column (3) adds pandemic displacement, the baseline.

Table 2: The Horse Race: Employment Effects by Policy Regime

	(1)	(2)	(3)
<i>Tightening, pre-ChatGPT (2022m3–2022m11)</i>			
AI exposure $\times$ regime	0.328 (0.326)	0.335 (0.331)	0.446 (0.319)
Rate sensitivity $\times$ regime		0.118 (0.276)	0.030 (0.277)
Pandemic displacement $\times$ regime			−1.637*** (0.480)
<i>Tightening and hold, post-ChatGPT (2022m12–2024m8)</i>			
AI exposure $\times$ regime	0.530 (0.610)	0.490 (0.614)	0.746 (0.556)
Rate sensitivity $\times$ regime		−0.700 (0.477)	−0.906** (0.448)
Pandemic displacement $\times$ regime			−3.799*** (0.672)
<i>Easing (2024m9–2026m07)</i>			
AI exposure $\times$ regime	0.858 (0.904)	0.759 (0.897)	1.108 (0.827)
Rate sensitivity $\times$ regime		−1.732** (0.681)	−2.011*** (0.657)
Pandemic displacement $\times$ regime			−5.169*** (0.997)
Industry and month fixed effects	Yes	Yes	Yes
Industries	205	205	205
Observations	13,325	13,325	13,325

Notes: Dependent variable: $100 \times$ log employment net of the industry's 2015–2019 linear trend. Each exposure (standardized) is interacted with the three regime indicators; the reference period is the year before liftoff (2021m3–2022m2). Standard errors clustered by industry in parentheses. The minimum detectable effect for AI exposure in the easing regime (80% power, 5% two-sided) is 2.32 log points per s.d. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The rate channel opens with a lag and grows with the duration of the tightening. During the first nine months of tightening, before ChatGPT, the coefficient on rate sensitivity is 0.030 log points (standard error 0.277): nothing had happened yet. During the tightening-and-hold regime it is −0.906 (0.448), significant at five percent, and during easing it is −2.011 (0.657), significant at one percent. The easing coefficient is a level relative to the reference year, and its being larger in absolute value than the hold coefficient does not contradict the sign-reversal prediction: a gap that opened during the hold remains open, in levels, until sensitive industries grow faster than others for long enough to close it; the growth-rate results below are where

the reversal is tested. Column (2) shows the same pattern without the displacement control (0.118, -0.700 and -1.732), so the rate effect does not depend on controlling for the pandemic overshoot, though the control sharpens it.

The AI channel is positive in every regime and every column. In the baseline the coefficients are 0.446 (0.319) before ChatGPT, 0.746 (0.556) during the hold and 1.108 (0.827) during easing; with AI exposure alone in column (1) they are 0.328, 0.530 and 0.858. None is distinguishable from zero, and the sign is the opposite of a displacement story. Whatever generative AI has done to the employment of exposed industries between December 2022 and July 2026, it has not produced a decline in headcount that this design can detect, and the point estimates lean toward exposed industries having done slightly better than others once their rate sensitivity and pandemic position are held fixed. The minimum detectable effect for the easing coefficient is 2.32 log points per standard deviation, so a decline of about 3.6 log points across the interquartile range of exposure would have been detected with 80 percent power; smaller declines would not.

Pandemic mean reversion dominates both. A one-standard-deviation (8.6 log point) larger overshoot at February 2022 is associated with detrended employment 1.6 log points lower before ChatGPT, 3.8 lower during the hold and 5.2 lower during easing, all significant at one percent. Industries that had hired ahead of trend gave up about 60 percent of the excess over the following four and a half years, which the model of [Section 8](#) translates into a half-life of a little over three years. Part of this is ordinary regression to the mean in a baseline level, as [Section 4](#) notes; whatever its composition, any reading of post-2022 industry employment that does not account for where each industry stood relative to trend when tightening began will attribute this reversion to whichever shock is under discussion.

6.3 How Precise Is the Rate Effect?

[Table 3](#) re-examines the baseline coefficients under the two forms of dependence that industry clustering ignores.

Table 3: Inference for the Baseline Horse Race under Cross-Industry Dependence and First-Stage Uncertainty

	Estimate	Industry cluster	Sector cluster	Wild bootstrap p	First-stage component	Murphy–Topel SE	p (MT)
<i>Tightening, pre-ChatGPT (2022m3–2022m11)</i>							
AI exposure	0.446	(0.319)	(0.354)	0.256	(0.021)	(0.320)	0.164
Rate sensitivity	0.030	(0.277)	(0.265)	0.916	(0.388)	(0.476)	0.951
Pandemic displacement	−1.637	(0.480)	(0.801)	0.184	(0.021)	(0.481)	0.001
<i>Tightening and hold, post-ChatGPT (2022m12–2024m8)</i>							
AI exposure	0.746	(0.556)	(0.689)	0.340	(0.163)	(0.580)	0.198
Rate sensitivity	−0.906	(0.448)	(0.445)	0.044	(0.736)	(0.862)	0.293
Pandemic displacement	−3.799	(0.672)	(1.194)	0.025	(0.048)	(0.674)	0.000
<i>Easing (2024m9–2026m07)</i>							
AI exposure	1.108	(0.827)	(1.150)	0.411	(0.355)	(0.900)	0.218
Rate sensitivity	−2.011	(0.657)	(0.648)	0.004	(1.185)	(1.355)	0.138
Pandemic displacement	−5.169	(0.997)	(1.834)	0.020	(0.097)	(1.002)	0.000

Notes: Column (1) repeats the column (3) estimates of Table 2. Column (2): standard errors clustered by industry. Column (3): clustered by the 14 CES sectors (CR1). Column (4): wild-cluster bootstrap p -value with sector clusters (Rademacher weights, null imposed, 1,999 draws). Column (5): the first-stage component of the standard error, the delta-method propagation of the sampling covariance of the 205 raw rate sensitivities (block bootstrap over months of the local-projection sample, blocks of 24, 200 draws; mean absolute cross-industry error correlation 0.36) through the shrinkage, standardization and cross-sectional regressions [Murphy and Topel, 1985]. Column (6): the Murphy–Topel standard error, the square root of the sum of the squared industry-clustered and first-stage components; column (7) its normal p -value. AI exposure and displacement are not generated regressors, but their coefficients change when the rate-sensitivity vector changes, which is why they carry a first-stage component too.

Cross-industry dependence within sectors changes little. The sector-clustered standard errors are close to the industry-clustered ones (0.65 against 0.66 for the easing-regime rate coefficient, 1.15 against 0.83 for AI exposure), and the wild-cluster bootstrap, which is the appropriate test with fourteen clusters, gives p -values of 0.044 for the rate coefficient in the hold regime and 0.004 in the easing regime, 0.02 to 0.03 for displacement, and 0.26 to 0.41 for AI exposure in every regime. Sector-by-regime effects, reported as the sector fixed effects row of Table 9, leave a within-sector rate effect of −1.2 (0.68).

First-stage uncertainty changes a great deal. The first-stage component of the standard error of the easing-regime rate coefficient is 1.18 log points, larger than the clustered component of 0.66, and the Murphy–Topel standard error is 1.35, with a p -value of 0.14; for the hold regime the corresponding figures are 0.86 and 0.29. The reason is not that any single industry’s sensitivity is badly measured but that all of them are measured from the same 316 months of shocks: the first-stage errors are correlated across industries (a mean absolute correlation of 0.36), so the effective number of independent observations behind the cross-industry pattern of rate sensitivity is closer to the number of independent blocks of shock months than to the number of industries. Standardization removes a common shift in all estimates but not a

correlated reordering of them. The AI and displacement coefficients acquire small first-stage components (0.36 and 0.10 in the easing regime) because they are estimated jointly with rate sensitivity, and their conclusions are unchanged.

The reading the paper adopts is the following. Conditional on the estimated sensitivities, the rate effect is precisely estimated and robust to how industries are grouped. Unconditionally, the cross-industry pattern of monetary sensitivity that the local projections deliver is uncertain enough that the specifically monetary attribution carries a p -value of 0.14 in the easing regime. The pattern itself is not in doubt: the cyclical beta, measured with reliability 0.91 from thirty years of aggregate cycles and requiring no identified shock, gives -4.0 (0.76) in the same regression (Table 9), and the two measures are correlated 0.57. Industries whose employment has historically moved with the aggregate cycle, and with identified monetary shocks in particular, lost ground after liftoff. Whether the loss is the monetary transmission mechanism specifically, or the broader cyclical sensitivity with which monetary sensitivity is entangled, the design cannot say with the precision that the clustered standard errors alone would suggest.

6.4 Growth Differentials and the Sign-Reversal Test

Table 4 Panel A reports the annualized growth differentials by regime from equation (8), and Figure 4 plots the AI and rate-sensitivity coefficients with 95 percent intervals.

Table 4: Growth Differentials, the Sign-Reversal Test and the Policy Stance

	AI exposure	Rate sensitivity	Displacement
<i>Panel A. Annualized growth differential by regime (percentage points per s.d.)</i>			
Year before liftoff	0.518 (0.470)	0.435 (0.502)	-2.492*** (0.872)
Tightening, pre-ChatGPT	0.546 (0.359)	-0.538** (0.263)	-1.917*** (0.455)
Tightening and hold, post-ChatGPT	0.147 (0.243)	-0.784*** (0.219)	-1.265*** (0.302)
Easing	0.267 (0.198)	-0.348* (0.185)	-0.268 (0.258)
<i>Panel B. Tests of changes in the growth differential</i>			
Post-ChatGPT minus pre-ChatGPT	-0.399 (0.314)	-0.245 (0.239)	0.651 (0.424)
Easing minus hold (sign-reversal test)	0.120 (0.194)	0.436* (0.234)	0.997*** (0.267)
<i>Panel C. Continuous stance: monthly growth (annualized), industry and month fixed effects</i>			
Stance lagged 6 months	0.163 (0.179)	-0.123 (0.081)	-2.145*** (0.239)
Stance lagged 12 months	0.112 (0.206)	-0.331*** (0.091)	-2.156*** (0.240)
Stance lagged 18 months	0.153 (0.282)	-0.344** (0.139)	-2.154*** (0.240)

Notes: Panels A and B: cross-sectional regressions of annualized detrended employment growth within each regime on the three standardized exposures ($N = 205$ industries), heteroskedasticity-robust standard errors. Panel C: the AI column is AI exposure $\times$ post-ChatGPT; the rate column is rate sensitivity $\times$ the 12-month change in the 2-year Treasury yield (percentage points) at the stated lag; the displacement column is displacement $\times$ post-liftoff; an AI exposure $\times$ stance term is included and is never distinguishable from zero. Standard errors clustered by industry. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.




Figure 4: Growth Differentials by Policy Regime

Notes: Coefficients on standardized AI exposure (circles) and rate sensitivity (squares) in cross-sectional regressions of annualized detrended employment growth within each regime on the three exposures (Table 4, Panel A), with pointwise 95 percent intervals from heteroskedasticity-robust standard errors. Regimes: the year before liftoff (March 2021 to February 2022), tightening before ChatGPT (March to November 2022), tightening and hold after ChatGPT (December 2022 to August 2024) and easing (September 2024 to July 2026).

The growth differential on rate sensitivity follows the policy cycle. In the year before liftoff it is +0.44 percentage points per year (standard error 0.50): sensitive industries were, if anything, growing slightly faster than others as rates sat at zero. In the first nine months of tightening it is -0.54 (0.26), during the hold -0.78 (0.22), and during easing -0.35 (0.19). Panel B reports the differences. The sign-reversal statistic, easing minus hold, is +0.44 percentage points per year (standard error 0.23, 95 percent interval -0.02 to 0.89, $p = 0.065$): the growth disadvantage of rate-sensitive industries narrowed by more than half once the Federal Reserve began cutting, but it did not close and the interval does not exclude no change. The easing regime was also not a period of uniform easing. The federal funds rate fell from 5.33 percent in August 2024 to 3.63 percent by early 2026, but the two-year yield, after falling from 4.9 percent at its 2023 peak to 3.5 percent in February 2026, rose again to 4.2 percent by August 2026, and remained 2.8 percentage points above its February 2022 level in July 2026. The result is therefore an attenuation of the growth disadvantage that is consistent with Prediction 2 and with the level of rates, not merely their direction, being

what matters; it is not an established reversal.

The growth differential on AI exposure does none of these things. It is +0.52 in the year before liftoff, +0.55 in the tightening months before ChatGPT, +0.15 during the hold and +0.27 during easing, with no coefficient distinguishable from zero. The ChatGPT-break statistic, hold minus pre-ChatGPT tightening, is -0.40 percentage points per year (standard error 0.31, 95 percent interval -1.01 to $+0.22$): the growth of exposed industries slowed somewhat when the technology arrived, from a positive differential, and the interval is consistent with a slowdown of up to one percentage point a year as well as with none. The easing-minus-hold difference for AI is $+0.12$ ($p = 0.54$), neither the persistence nor the deepening that a continuing technology shock predicts. A failure to reject a break is not evidence that the break is negligible; the interval, not the p -value, is the statement. The displacement differential is large and negative in every regime and shrinks steadily, from -2.5 percentage points per year before liftoff to -0.3 under easing; its easing-minus-hold difference is $+1.0$ ($p < 0.001$), the signature of mean reversion running its course.

6.5 The Continuous Policy Stance

Panel C of [Table 4](#) reports the continuous specification of [equation \(9\)](#). With the twelve-month change in the two-year yield lagged twelve months, a one-percentage-point rise in the yield lowers annualized detrended growth in an industry one standard deviation more rate-sensitive by 0.331 percentage points (standard error 0.091), significant at one percent. The effect is smaller and insignificant at a six-month lag (-0.123 , standard error 0.081) and essentially the same at an eighteen-month lag (-0.344 , standard error 0.139). The lag structure agrees with the regime results, in which the rate effect was absent in the first nine months of tightening: monetary policy reaches industry employment with a lag of about a year. Because the two-year yield rose by 3.6 percentage points at the peak, the estimate implies that at the height of the tightening, rate-sensitive industries were growing about 1.2 percentage points per year more slowly per standard deviation of sensitivity than they otherwise would have, and across the interquartile range of rate sensitivity about 1.6 percentage points per year.

The AI terms in the same regression are uninformative in the same way as before. The coefficient on AI exposure interacted with the post-ChatGPT indicator is 0.112 (0.206) at the twelve-month lag and similar at the others. The coefficient on AI exposure interacted with the policy stance is -0.005 (0.084): AI-exposed industries do not respond to rate changes over and above what their own measured rate sensitivity implies. This is the most direct test of the confound the paper began with, and it finds none. The displacement term, interacted with the post-liftoff indicator, is -2.156 (0.240), the mean-reversion result again. Two caveats apply to the continuous specification. The two-year yield is an endogenous object that reflects

expected policy, expected activity and term premia, so the interaction identifies an association between the yield path and the relative growth of sensitive industries, not the causal effect of a policy surprise; and its standard errors are clustered by industry and do not carry the first-stage component of [Table 3](#), which would inflate them by a similar proportion.

6.6 Magnitudes

The rate effect is economically meaningful. The interquartile range of rate sensitivity is 1.38 standard deviations, so the easing-regime coefficient of -2.0 log points per standard deviation implies that an industry at the 75th percentile of rate sensitivity had, by 2025 and 2026, detrended employment about 2.8 log points lower relative to its reference-year level than an industry at the 25th percentile, and about 3.4 log points lower by July 2026 using the event-study endpoint. For the industries most often named in the AI narrative the implications differ: software publishers, with below-average rate sensitivity, are predicted by the rate channel to have done slightly better than the mean industry, while nondepository credit and insurance carriers, well above average, are predicted to have done worse.

7. Who Bears It: the Age Gradient

The headcount results say that rates, not AI, moved the employment of exposed industries. The most influential evidence for an AI effect, however, is not about headcount but about who is hired: [Brynjolfsson et al. \[2025a\]](#) document a sharp relative decline in the employment of the youngest workers in exposed occupations, and [Hosseini and Lichtinger \[2025\]](#) document a decline in junior hiring at adopting firms. This section asks whether the same pattern is visible in public industry data, whether it belongs to AI or to rates, and, decisively, when it began.

7.1 Levels

[Table 5](#) reports the pooled QWI estimates for the hold and easing regimes.

Table 5: Who Bears It: the Age Gradient and Worker Flows

Outcome	AI exposure		Rate sensitivity		2019
	Hold	Easing	Hold	Easing	mean
<i>Panel A. Relative employment (100× log ratio)</i>					
Ages 22–24 relative to 35–64	−2.32*** (0.73)	−4.88*** (1.08)	−2.23*** (0.61)	−3.24*** (0.90)	−252.5
Ages 25–34 relative to 35–64	0.13 (0.33)	−0.23 (0.56)	0.14 (0.32)	0.04 (0.48)	−103.9
Ages 22–34 relative to 35–64	−0.25 (0.37)	−0.97 (0.60)	−0.25 (0.33)	−0.51 (0.50)	−82.8
<i>Panel B. Employment levels (100× log)</i>					
Ages 22–24	−0.87 (1.14)	−2.54 (1.98)	−2.89*** (0.88)	−4.85*** (1.25)	946.5
Ages 35–64	1.45** (0.72)	2.34* (1.30)	−0.65 (0.45)	−1.61*** (0.60)	1199.0
Ages 22–64	1.37* (0.75)	2.02 (1.37)	−0.74 (0.50)	−1.80*** (0.67)	1236.4
<i>Panel C. Flows, ages 22–24 (% of own employment; % of all hires)</i>					
Hiring rate, ages 22–24	1.43 (0.87)	2.00*** (0.47)	−0.58 (0.45)	−0.15 (0.48)	29.3
Separation rate, ages 22–24	1.45*** (0.44)	2.46*** (0.42)	−0.34 (0.26)	−0.57 (0.50)	25.7
Share of hires aged 22–24	−0.13** (0.06)	−0.08 (0.09)	−0.14*** (0.05)	−0.20*** (0.07)	12.8
<i>Panel D. Flows, ages 25–34 and all ages</i>					
Hiring rate, ages 25–34	1.14** (0.46)	1.38*** (0.28)	−0.72** (0.29)	−0.52 (0.32)	18.5
Separation rate, ages 25–34	0.91*** (0.21)	1.74*** (0.28)	−0.34* (0.18)	−0.57 (0.38)	17.9
Share of hires aged 25–34	0.06 (0.08)	0.21 (0.13)	−0.02 (0.06)	−0.00 (0.10)	32.9
Hiring rate, all ages	0.71** (0.32)	0.76*** (0.23)	−0.46* (0.25)	−0.31 (0.28)	14.2
Separation rate, all ages	0.56*** (0.15)	1.04*** (0.23)	−0.19 (0.15)	−0.39 (0.32)	14.1

Notes: Census QWI, private ownership, national aggregates (sums over balanced state-by-industry cells) over 49 jurisdictions with complete coverage through 2025-Q4 (99.6% of reported employment), $N = 204$ industries in every row. Outcomes are net of industry-specific quarter-of-year effects and linear trends estimated on 2015–2019. Each row is a pair of cross-sectional regressions of the regime mean minus the 2021Q2–2022Q1 mean on AI exposure, rate sensitivity and pandemic displacement; ‘hold’ is 2023Q1–2024Q3, ‘easing’ is 2024Q4 onward. The last column is the 2019 mean level of the outcome. Heteroskedasticity-robust standard errors in parentheses. For the easing-regime coefficients in the first row, sector-clustered standard errors are 1.43 (AI) and 0.77 (rate sensitivity), with wild-cluster bootstrap p -values of 0.015 and 0.001; the permutation p -value for the AI coefficient is 0.0005. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

For the youngest workers, both exposures matter, and with similar magnitudes (Panel A). Employment of 22 to 24 year olds relative to 35 to 64 year olds within industry is 2.3 log points lower per standard deviation of AI exposure during the hold (robust standard error 0.73) and 4.9 log points lower during easing (1.08); it is 2.2 log points lower per standard deviation of rate sensitivity during the hold (0.61) and 3.2 lower during easing (0.90). Clustering by sector gives standard errors of 1.43 and 0.77 for the two easing coefficients, with wild-cluster

bootstrap p -values of 0.015 and 0.001, and the permutation reference distribution places the AI coefficient outside all 2,000 reshuffles ($p = 0.0005$). Across the interquartile range of AI exposure, the youth ratio fell by about 7.6 log points more in exposed industries; across the interquartile range of rate sensitivity, by about 4.5 log points more. The canaries exist in public data.

For the 25 to 34 year olds nothing happens on either exposure: the coefficients are 0.13 and -0.23 for AI and 0.14 and 0.04 for rate sensitivity, all small relative to their standard errors. The gradient is confined to the youngest group, the workers within two years of entering the labor market, and the combined 22 to 34 ratio is correspondingly muted (-0.97 and -0.51 under easing, neither significant at five percent). The concentration at the entry point is what both channels predict under Prediction 3, and it matches the age profile in [Brynjolfsson et al. \[2025a\]](#), where the decline is concentrated among workers aged 22 to 25 and absent for older workers.

Panel B decomposes the ratio into its two levels, and the decomposition differs across the two channels. In AI-exposed industries the ratio fell because employment of 22 to 24 year olds fell somewhat (-2.5 log points per standard deviation under easing, imprecisely estimated with a standard error of 2.0) while employment of 35 to 64 year olds rose ($+2.3$, standard error 1.3): the exposed industries grew their experienced workforce and did not grow their entry-level one. In rate-sensitive industries both levels fell, the young by more (-4.8 , standard error 1.2) than the experienced (-1.6 , standard error 0.6): a contraction that fell hardest on the youngest. The QWI levels also reproduce the CES headcount results, which is reassuring about the comparability of the two sources: log employment of 22 to 64 year olds is higher per standard deviation of AI exposure (1.4 during the hold, 2.0 during easing, imprecise) and lower per standard deviation of rate sensitivity (0.7 and 1.8, the latter significant at one percent), as in the CES.

7.2 Flows

Panels C and D of [Table 5](#) distinguish a hiring freeze from reallocation, using the flows of the 22 to 24 year olds themselves, since that is the group whose employment moved. Relative to the reference year, which was the year of the Great Resignation and of the highest hiring and quit rates on record, the hiring rate of 22 to 24 year olds in AI-exposed industries (hires as a percent of their own employment) is 1.4 percentage points higher per standard deviation during the hold and 2.0 higher during easing, and their separation rate is 1.5 and 2.5 percentage points higher; all but the first are significant at one percent. The share of all hires accounted for by 22 to 24 year olds does not fall with AI exposure (-0.13 percentage points during the hold, significant at five percent against a 2019 mean of 12.8 percent, and -0.08

during easing, insignificant). The youngest workers in AI-exposed industries are therefore being hired and are leaving at higher rates than before, not being frozen out at the hiring stage; the stock fell because separations rose by more than hires. The 25 to 34 year olds in the same industries show the same elevated churn (Panel D) with no change in their stock, so the churn is a feature of exposed industries generally and the net loss is confined to the youngest.

Rate-sensitive industries look different. The hiring rate of their 22 to 24 year olds is lower rather than higher (-0.6 during the hold, -0.2 during easing, neither significant), their separation rate is lower (-0.3 and -0.6 , insignificant), and the share of all hires going to 22 to 24 year olds falls (-0.14 during the hold, -0.20 during easing, both significant at one percent). This is the hiring-freeze pattern of Forsythe [2022]: fewer hires, stable separations, and a shift in the composition of hires away from the young. The two channels produce a similar youth gradient in levels through different flows, which is evidence that the gradient on AI exposure is not the rate channel wearing a different label. Two limitations of the flow evidence should be stated. A rate expressed relative to a shrinking stock rises mechanically, so part of the higher hiring rate of the young in exposed industries reflects the smaller denominator; and industry-level flows cannot separate a reallocation of tasks within firms, a change in the composition of firms, and a change in the labor supply of young workers to these industries.

7.3 Timing

Figure 5 plots the quarterly event-study coefficients on AI exposure and on rate sensitivity for the 22 to 24 ratio from 2015Q1, with 2022Q1 as the base, and Table 6 Panel A reports the annualized slope of the gradient over four windows.




Figure 5: Young Workers: the AI Gradient Falls from 2020, the Rate Gradient from Liftoff
Notes: Coefficients on standardized AI exposure (solid) and rate sensitivity (dashed) in quarterly cross-sectional regressions of the change since 2022Q1 in $100 \times \log$ employment of ages 22 to 24 relative to ages 35 to 64 within industry on the three exposures (equation (10)), Census QWI, 204 industries, with pointwise 95 percent bands. Outcomes are net of industry-specific quarter-of-year effects and linear trends estimated on 2015 to 2019. Vertical lines mark liftoff, the release of ChatGPT and the first rate cut.

Table 6: The Young-Worker Gradient: Timing

	AI exposure	Rate sensitivity
<i>Panel A. Short panel (49 jurisdictions), detrended: annual change by window</i>		
2015Q1–2019Q4 (detrending window)	0.12*** (0.04)	0.00 (0.04)
2020Q1–2022Q1	−1.98*** (0.28)	0.45 (0.29)
2022Q1–2022Q4	−0.15 (0.53)	−1.33*** (0.42)
2022Q4–2025Q4 (post-ChatGPT)	−1.86*** (0.34)	−0.78** (0.30)
Post-ChatGPT minus 2020Q1–2022Q1	0.12 (0.31)	−1.23*** (0.39)
<i>Panel B. Long panel (47 jurisdictions), not detrended: annual change by window</i>		
2005Q1–2009Q4	0.21 (0.22)	−1.62*** (0.22)
2010Q1–2014Q4	−1.04*** (0.22)	0.24 (0.19)
2015Q1–2019Q4	0.24 (0.20)	−0.04 (0.20)
2019Q4–2022Q4	−1.19*** (0.23)	−0.07 (0.24)
2022Q4–2025Q4 (post-ChatGPT)	−2.00*** (0.31)	−0.84*** (0.28)
Post-ChatGPT minus 2015–2019	−2.24*** (0.37)	−0.80** (0.33)
Post-ChatGPT minus 2019Q4–2022Q4	−0.82** (0.35)	−0.76** (0.37)
<i>Panel C. Common sample (47 jurisdictions, detrended as in Panel A): segments and break dates</i>		
Pandemic step, 2020Q1–2021Q1	−3.46*** (0.59)	0.77 (0.68)
Plateau, 2021Q1–2022Q4	−0.28 (0.45)	−0.44 (0.39)
Post-ChatGPT, 2022Q4–2025Q4	−1.86*** (0.34)	−0.83*** (0.30)
2020Q1–2022Q1 (Panel A window)	−1.90*** (0.27)	0.50* (0.29)
Post minus pandemic step	1.59*** (0.59)	−1.60** (0.68)
Post minus plateau	−1.58*** (0.52)	−0.39 (0.51)
Post minus 2020Q1–2022Q1	0.03 (0.31)	−1.33*** (0.41)
Post slope with break at 2022Q3 (post minus plateau)	−1.74 (−1.46, 0.50)	−0.90
Post slope with break at 2022Q4 (post minus plateau)	−1.86 (−1.58, 0.52)	−0.83
Post slope with break at 2023Q1 (post minus plateau)	−1.85 (−1.35, 0.51)	−0.75
Post slope with break at 2023Q2 (post minus plateau)	−1.73 (−0.95, 0.50)	−0.69

Notes: Outcome: $100 \times$ log employment of ages 22–24 relative to ages 35–64 within industry (QWI); coefficients on the standardized exposure in cross-sectional regressions of the annualized change over the window on AI exposure, rate sensitivity and pandemic displacement (Panel B: the two exposures only). Panel A: the 2015–2025 panel net of 2015–2019 seasonal effects and trends, so the first row is close to zero by construction. Panel B: the 2005–2025 panel with seasonal effects removed and no detrending. Panel C holds the sample (the 47 jurisdictions of the long panel) and the detrending convention of Panel A fixed and splits 2020–2025 into a pandemic step, a plateau and a post-ChatGPT segment; the break-date rows report the post-ChatGPT AI slope and, in parentheses, its difference from the plateau slope with its standard error, for four alternative start dates of the post segment. Robust standard errors. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The AI coefficient sits between +4 and +5.6 log points throughout 2015 to 2019, at +4.4

in 2020Q1, and then falls: to +1.5 by 2020Q3, +0.7 by 2021Q1, zero at the 2022Q1 base by construction, -0.9 in 2023Q1, -3.0 in 2024Q1, -4.6 in 2025Q1 and -5.6 in 2025Q4. Panel A quantifies the slopes on the short panel. Over 2015 to 2019 the gradient was flat by construction (0.12 log points per year, the residual of detrending). Over 2020Q1 to 2022Q1 it fell at 1.98 log points per year (standard error 0.28). Over 2022Q1 to 2022Q4, the tightening months before ChatGPT, it was flat (-0.15 , standard error 0.53). Over 2022Q4 to 2025Q4, after ChatGPT, it fell at 1.86 log points per year (standard error 0.34). The difference between the post-ChatGPT slope and the 2020Q1 to 2022Q1 slope is +0.12 log points per year (standard error 0.31).

That comparison averages over two very different sub-periods, and Panel C separates them on a common sample. It uses the 47 jurisdictions of the long panel and the detrending convention of Panel A, so that the only thing that changes across rows is the window. The pandemic step, 2020Q1 to 2021Q1, was steep: -3.46 log points per year (standard error 0.59). The plateau, 2021Q1 to 2022Q4, was flat: -0.28 (0.45). The post-ChatGPT segment, 2022Q4 to 2025Q4, fell at -1.86 (0.34), slower than the pandemic step by 1.59 (0.59) and faster than the plateau by 1.58 (0.52), both differences significant at one percent; against the 2020Q1 to 2022Q1 average the difference is 0.03 (0.31). The break-date rows show that the post-ChatGPT slope is between -1.7 and -1.9 whether the segment is started in 2022Q3, 2022Q4, 2023Q1 or 2023Q2, and that its difference from the plateau is between -0.95 and -1.58 , so the resumption is not an artifact of where the break is placed. The youngest workers' relative employment in AI-exposed industries thus fell in two steps: a large one in the pandemic year, then a plateau of seven quarters, then a second decline from about the turn of 2023 at roughly half the pace of the first step and continuing to the end of the sample.

The rate-sensitivity coefficient behaves as Prediction 3 says a rate channel should. It is within ± 0.6 log points of zero throughout 2015 to 2022Q2, including the pandemic years, during which its slope is +0.45 per year (standard error 0.29, insignificant). It then breaks at liftoff: -0.85 by 2022Q4, -2.28 by 2023Q4, -2.76 by 2024Q4 and -3.06 by 2025Q4. The slope over 2022Q1 to 2022Q4 is -1.33 log points per year (standard error 0.42), significant at one percent, and over the post-ChatGPT window -0.78 (0.30). On the common sample of Panel C the rate gradient is flat in the pandemic step (+0.77, 0.68) and the plateau (-0.44 , 0.39) and falls at -0.83 (0.30) after 2022Q4; its post-ChatGPT slope differs from the 2020Q1 to 2022Q1 slope by -1.33 (0.41). That difference is a break at liftoff, not at ChatGPT: the rate gradient had been falling for three quarters when ChatGPT was released, and its steepest decline was in 2022.

This is the paper's key fact about the age gradient, and it should be stated with both of its edges. Young workers' relative employment has fallen much more in AI-exposed industries

than elsewhere, by a large amount that survives every robustness check, and in that sense the canaries are real and visible in public data. The decline began in 2020, in the same industries and at a faster pace than anything since, when the technology did not exist. It then paused, and it resumed from about 2023Q1, one quarter after the release of ChatGPT, at a pace that equals the 2020 to 2021 average and is well below the pandemic step but well above the plateau. A reader who dates the decline from ChatGPT will see the resumption and read it as the technology’s effect; a reader who dates it from 2020 will see a two-step decline whose first step could not have been the technology and whose second step coincides with three things at once: adoption, the arrival of the tightening in the labor market (the rate gradient’s own break is in 2022 and its decline continued through 2023 and 2024), and the unwinding of the 2021 to 2022 hiring boom. The design cannot assign the second step among the three. What it can say is that the first step, which was larger, had a cause other than generative AI, that the same gradient fell at one log point a year through 2010 to 2014 (the next subsection), and that an account of the youngest workers’ position in these industries that starts in November 2022 starts too late. Candidate causes of the pre-2023 decline that the data cannot test include the shift to remote work in these industries, which [Dingel and Neiman \[2020\]](#) show are the most remote-capable and which [Barrero et al. \[2023\]](#) document persisted after the pandemic, and which may have reduced the value of on-site junior staff; the composition of pandemic-era hiring; and an earlier wave of automation of routine professional tasks.

7.4 The Long Run

The detrended 2015 to 2025 panel cannot say whether the decline that began in 2020 was itself the continuation of an older trend or a genuine break, because the trend is removed on 2015 to 2019 by construction. [Figure 6](#) extends the youth ratio back to 2005 using the longer QWI history, with 2019Q4 as the base, seasonal effects removed but no detrending, so that the AI coefficient on the youth ratio can be seen over two full business cycles.




Figure 6: Two Decades of the Young-Worker Gradient in AI-Exposed Industries

Notes: Coefficient on standardized AI exposure in quarterly cross-sectional regressions of the change since 2019Q4 in $100 \times \log$ employment of ages 22 to 24 relative to ages 35 to 64 within industry on AI exposure and rate sensitivity, Census QWI, 2005Q1 onward, with pointwise 95 percent bands. Outcomes are net of industry-specific quarter-of-year effects but are not detrended. Shaded areas are NBER recessions; the vertical line marks the release of ChatGPT.

Two decades of data put the post-2020 decline in proportion. The AI gradient in the youth ratio was about +3.7 log points per standard deviation in 2005, relative to its 2019Q4 level, rose to a peak of +5.1 in 2009 as the Great Recession destroyed young workers' jobs in the low-exposure goods industries faster than in the exposed professional ones, and then declined through the 2010s: Panel B of Table 6 puts the slope at -1.04 log points per year over 2010 to 2014 (standard error 0.22) and flat over 2015 to 2019 ($+0.24$, standard error 0.20). A decline of about one log point a year is therefore not unprecedented; it happened in the recovery from the last recession, when no language model existed. The path since 2019 has three segments. In 2020 the gradient fell by more than three log points in a single year, from -0.3 in 2020Q1 to -3.4 in 2021Q1. It then plateaued between -3.2 and -4.1 from 2021Q1 through 2022Q4, the period of the pandemic hiring boom and the first three quarters of tightening. From 2023Q1 it declined again, at 2.0 log points a year (standard error 0.31), to -9.6 by 2025Q4. Panel B reports the slope over the three-year window 2019Q4 to 2022Q4, which averages the 2020 drop and the plateau, as -1.19 per year, and the post-ChatGPT slope as -2.00 ; the difference of -0.82 (standard error 0.35) is significant at five percent. The rate gradient in the same panel fell at 1.6 log points a year over 2005 to 2009, was flat

from 2010 through 2022Q4, and has fallen at 0.84 a year since.

The long and short panels are consistent with each other and with Panel C. The renewed decline since 2023 is real, faster than the average pace of the preceding three years, faster than the plateau that immediately preceded it, and faster than anything in 2015 to 2019. It is not the first such decline, and it is slower than the decline of 2020, which began before the technology existed. The detrended comparison of Panel A, which sets the 2020Q1 to 2022Q1 average against the post-ChatGPT pace, finds them equal; the segment comparison of Panel C finds the post-ChatGPT pace between the pandemic step and the plateau; the undetrended comparison of Panel B, which averages the 2020 step with the plateau, finds the later decline steeper. None of the three finds a break at ChatGPT in an otherwise stable gradient, because the gradient was not stable before ChatGPT; and none of them excludes a contribution of the technology to the second step.

7.5 Robustness of the Age Gradient

Panel A of [Table 7](#) re-estimates the easing-regime coefficients on the 22 to 24 ratio under alternative specifications. Adding the industry’s own detrended youth ratio at 2022Q1 as a control, so that industries which had accumulated an unusually young workforce during the pandemic are not credited to AI, reduces the AI coefficient from -4.88 to -3.55 (standard error 1.01) and leaves the rate coefficient at -3.42 (0.89). Adding sector fixed effects, so that only within-sector variation in exposure is used, gives -4.19 (1.37) and -2.19 (0.98). Both together give -3.33 (1.62) and -2.42 (1.00). Replacing rate sensitivity by the cyclical beta strengthens both: -6.43 (1.12) for AI and -5.71 (1.22) for the beta, which says that the youth ratio in AI-exposed industries fell by more once the comparison is made against industries of similar cyclicalities. Dropping the tech core (software publishers, computer systems design, computing infrastructure, web search and portals, and streaming and social networks) gives -3.74 (0.98) and -3.53 (0.87): the AI gradient is not a tech-industry phenomenon but is present across the exposed professional and financial industries. Weighting by 2019 employment gives -4.93 (1.06) and -4.48 (0.89). In every specification both coefficients remain negative and significant at five percent or better; the youth gradient on AI exposure is a robust feature of the data.

Table 7: The Young-Worker Gradient: Robustness and Placebo-in-Time

	AI exposure	Rate sensitivity
<i>Panel A. Easing-regime level effect, alternative specifications</i>		
Baseline	−4.88*** (1.08)	−3.24*** (0.90)
Plus own-outcome displacement	−3.55*** (1.01)	−3.42*** (0.89)
Plus sector fixed effects	−4.19*** (1.37)	−2.19** (0.98)
Own displacement and sector FE	−3.33** (1.62)	−2.42** (0.99)
RS: cyclical beta	−6.43*** (1.12)	−5.71*** (1.22)
Drop tech core	−3.74*** (0.98)	−3.53*** (0.87)
Weighted by 2019 employment	−4.93*** (1.06)	−4.48*** (0.88)
<i>Panel B. Placebo-in-time: easing-analogue coefficient at fake base quarters</i>		
Base 2012Q1	−1.08 (1.27)	7.54 (1.18)
Base 2013Q1	2.75 (1.61)	3.49 (1.41)
Base 2014Q1	4.47 (1.56)	1.00 (1.60)
Base 2015Q1	5.19 (1.28)	−1.26 (1.53)
Base 2016Q1	3.67 (1.17)	−1.61 (1.19)
Actual, base 2022Q1	−4.88	−3.24

Notes: Outcome as in Table 6. Panel A repeats the easing-regime regression of Table 5, first row, changing one element; ‘own-outcome displacement’ adds the detrended level of the outcome itself in 2022Q1. Panel B replays the design on the 2005–2025 panel at five fake base quarters with its own geometry (seasonal and trend window from 28 to 9 quarters before the base, reference year ending at the base, easing analogue 11 to 15 quarters after it, CES displacement recomputed at the fake base month). The placebo windows overlap and rate sensitivity is estimated on data that overlap them, so Panel B is a descriptive reference distribution, not a null distribution. Robust standard errors. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Panel B applies the placebo-in-time replay to the youth-ratio design itself, at the five fake base quarters, 2012Q1 to 2016Q1, for which the long panel supplies the design’s full geometry without touching the pandemic. The easing-analogue AI coefficients are −1.1, +2.7, +4.5, +5.2 and +3.7 log points per standard deviation (mean +3.0, standard deviation 2.5), against the actual −4.9: every placebo has the opposite sign, and the actual estimate lies outside their range. The placebos are large, which says that the youth-ratio design drifts as much as

the headcount design does (the mid-2010s were a period in which the youngest workers' share in exposed industries was recovering from its 2010 to 2014 decline), and with five overlapping windows they are a scale rather than a test. By that scale the post-2022 movement is unusual in size and unusual in sign, which is the most that a descriptive reference distribution can say.

7.6 Reconciling the Canaries

The results reconcile the two sides of the recent literature more completely than either side alone. [Brynjolfsson et al. \[2025a\]](#) are right that the employment of the youngest workers in the most AI-exposed parts of the economy has fallen sharply relative to older workers, and the pattern appears here in Census data covering essentially all private employment, not only in one payroll provider's clients. [Humlum and Vestergaard \[2025\]](#), [Gimbel et al. \[2025\]](#) and [Eckhardt and Goldschlag \[2025\]](#) are right that there is no detectable effect on aggregate employment, on headcount in exposed industries, or on the occupational mix. What the longer window adds is that the youth pattern is not new. The gradient fell faster in 2020, when generative AI was a research curiosity, than it has since 2023, and it fell at half the recent pace through the recovery from the last recession. A researcher who observes the youth ratio in exposed industries only from 2022 onward, as post-2022 comparisons necessarily do, sees a decline that coincides with ChatGPT; a researcher who observes it from 2015 sees a decline that coincides with the pandemic, pauses, and resumes; a researcher who observes it from 2005 sees a gradient that has been falling, with interruptions, since 2010. The canaries were singing before the technology arrived, and they have sung louder since.

Three cautions bound the interpretation. First, the estimands differ. The microdata results concern workers in exposed occupations or at adopting firms; the present results concern industries ranked by the exposure of their occupational mix, and the seniority-biased pattern of [Hosseini and Lichtinger \[2025\]](#) at adopting firms is a within-firm result that the industry-level design cannot test. Second, an industry-level youth ratio is consistent with labor-supply as well as labor-demand explanations: young graduates may have chosen other industries, delayed entry, or entered these industries through channels that the QWI classify elsewhere, and the flows in [Table 5](#) cannot distinguish a firm that stops replacing junior staff from junior staff who leave for better offers. Third, the second step of the decline coincides with adoption, and the paper does not claim otherwise. The evidence here does not permit the inference, which the timing of the popular narrative invites, that the youngest professionals began losing ground because of large language models; it is equally clear that it does not permit the inference that the technology has had no effect on them. They began losing ground in 2020, in the same industries, faster than since; what has happened since 2023 is consistent with a continuation, with a technology effect, or with both.

8. Model Estimation and Counterfactuals

8.1 Estimation

The model of [Section 3](#) has three free parameters: the monthly turnover rate s , the rate semi-elasticity κ_r and the AI semi-elasticity κ_a . The interest-rate path x_t is the change in the two-year Treasury yield since February 2022, which peaks at 3.63 percentage points in October 2023; the adoption path A_t is zero before December 2022 and thereafter a logistic curve with midpoint eighteen months after release and scale six months, reaching one half in mid-2024 and 0.98 by mid-2026; the initial condition d_i is the industry’s pandemic displacement; and $\eta = 1$. Because the empirical objects are cross-sectional regression coefficients, the model is fitted to them directly: for a candidate parameter vector, the model is simulated for all 205 industries, the event study of [equation \(6\)](#) is run on the simulated paths, and the model-implied coefficients on the three exposures in every month after February 2022 are compared to the estimated ones. The objective is the sum of squared differences weighted by the inverse variance of each estimated coefficient, over $3 \times 53 = 159$ moments, minimized by bounded L-BFGS-B from several starting values. The 159 moments are estimated on the same industries and are correlated across months and exposures, so the curvature of the objective is not a sampling covariance. Standard errors are instead from a cluster bootstrap over industries: industries are resampled with replacement, the event-study coefficient paths are recomputed on the resampled set, and the model is refit, 200 times. The bootstrap holds the estimated exposures fixed, so it inherits the conditional character of the second-stage inference discussed in [Section 6](#); the first-stage component of [Table 3](#) would widen the interval for κ_r further.

[Table 8](#) reports the results and [Figure 7](#) plots the fitted against the estimated paths.

Table 8: A Stock-Flow Hiring Model: Parameters, Fit and Decomposition

	Estimate	(s.e.)
<i>Panel A. Parameters (cluster-bootstrap standard errors over industries, 200 draws)</i>		
Adjustment speed s (per month)	0.0178	(0.0047)
Rate channel κ_r (per s.d., per pp of the 2-year yield)	0.0155	(0.0056)
AI channel κ_a (per s.d., at full adoption)	0.0465	(0.0343)
Hiring elasticity η (fixed)	1.0	
Implied half-life of employment adjustment (months)	38.6	
<i>Panel B. Fit: R^2 of model against estimated event-study path</i>		
Rate sensitivity	0.99	
Pandemic displacement	0.92	
AI exposure	-0.38	
<i>Panel C. Relative employment change, February 2022 to 2026-07 (log points)</i>		
	Top quartile, AI	Top quartile, RS
Rate channel	0.12	-2.84
AI channel	2.21	-0.30
Pandemic mean reversion	-1.00	0.76
Total (model)	1.35	-2.43
No-tightening counterfactual: gain for the top RS quartile		2.92 (0.83)
<i>Panel D. Heterogeneous adjustment speed: $s_i = \phi \times 2019$ QWI separation rate</i>		
ϕ (implied mean s_i 0.0143, s.d. 0.0081)	0.303	
$\kappa_r; \kappa_a$	0.0182	0.0536
Objective (baseline 69.5); R^2 rate, displacement	62.4	0.99, 0.94
Corr. of rate sensitivity with s_i ; with $1/(r + s_i)$	0.15	-0.21

Notes: Three parameters fitted by bounded L-BFGS-B to 159 inverse-variance-weighted event-study coefficients (three exposures, every month after February 2022). Standard errors are from a cluster bootstrap over industries (resample industries, recompute the event-study moments, refit), which respects the correlation of the moments across months and exposures but holds the estimated exposures fixed; the curvature-based standard errors of the weighted objective are 0.0007, 0.0008 and 0.0072 and are not sampling standard errors. Panel C reports model-implied changes relative to the mean industry; effects common to all industries are not identified by the cross-section and are excluded. Panel D replaces the common s by an industry-specific speed proportional to the industry's 2019 QWI quarterly separation rate converted to monthly.




Figure 7: Model Fit to the Estimated Event-Study Paths

Notes: Solid lines with shaded pointwise 95 percent bands are the estimated event-study coefficients of [equation \(6\)](#) on rate sensitivity (left), AI exposure (center) and pandemic displacement (right), February 2022 to July 2026. Dashed lines are the coefficients implied by the stock-flow hiring model at the fitted parameters of [Table 8](#), obtained by running the same cross-sectional regressions on simulated industry employment paths. The R^2 in each panel title is one minus the ratio of the squared fit error to the variance of the estimated path.

The adjustment speed is estimated at $s = 0.0178$ per month (bootstrap standard error 0.0047), which implies that the stock of employment closes half of any gap between itself and the level implied by the current job value in 38.6 months. This is the speed of pandemic mean reversion in the displacement path, and it is also the speed with which the rate gap accumulates. It is well below the 4.7 percent monthly separation rate that the QWI report for these industries in 2019, because QWI separations include the many short spells that turn over without moving the stock; s is an adjustment parameter, not observed turnover, and Panel D of [Table 8](#) shows that letting it vary across industries in proportion to observed turnover changes little (the proportionality factor is 0.30, the fit improves slightly, and κ_r moves from 0.0155 to 0.0182). The rate semi-elasticity is $\kappa_r = 0.0155$ (0.0056): a one-percentage-point rise in the two-year yield lowers the log value of a filled job by 1.55 percent in an industry one standard deviation more rate-sensitive than the mean. The AI semi-elasticity is $\kappa_a = 0.0465$ (0.0343), positive but not distinguishable from zero.

Panel D also reports a check on the microfoundation. The discount-rate channel of [equation \(12\)](#) says that long-duration matches, those with low separation rates, should be the most rate-sensitive. In the data the estimated rate sensitivity is positively correlated with the observed separation rate (0.15) and negatively with the implied duration $1/(r + s_i)$ (-0.21): the industries that respond most to identified monetary shocks (residential construction,

employment services, transportation equipment) are high-turnover industries whose sensitivity runs through the demand for their output, not through the duration of their matches. The historical sensitivity that the model takes as ρ_i is therefore a reduced-form loading on monetary shocks that bundles the demand, credit and investment channels, and the model's discount-rate language should be read as one interpretation of that loading rather than as an identified mechanism.

8.2 Fit

The model fits the rate path almost perfectly ($R^2 = 0.99$): the gradual opening of the gap during 2022 and 2023, the flattening as the yield stopped rising, and the continued slow drift as the yield remained elevated are all reproduced by a job value proportional to the yield change and a stock that adjusts at rate s . It fits the displacement path well ($R^2 = 0.92$): geometric decay from the initial overshoot with a three-year half-life is close to what the data show. It cannot fit the AI path ($R^2 = -0.38$, worse than a constant at the mean). The reason is visible in [Figure 7](#): the estimated AI coefficient rises during 2022, before adoption begins, and then meanders around one log point with no acceleration during the adoption period, whereas the model's AI path is by construction zero until December 2022 and then follows the logistic curve. The best the model can do is to set κ_a positive so that the fitted path rises to meet the estimated one by 2025, at the cost of missing the pre-ChatGPT rise entirely. The failure is informative: whatever generates the small positive AI coefficient in the headcount data does not have the time profile of the diffusion of a technology released in November 2022.

8.3 Decomposition and Counterfactuals

Panel C of [Table 8](#) decomposes the model-implied change in employment between February 2022 and July 2026, relative to the mean industry, for the top quartile of AI-exposed industries and the top quartile of rate-sensitive industries. The top AI quartile is 1.35 log points above the mean industry by July 2026 in the model, the sum of +0.12 from rates (these industries are on average slightly less rate-sensitive than the mean), +2.21 from the AI channel at its fitted positive value, and -1.00 from mean reversion (they were on average above trend at the base). The top rate-sensitivity quartile is 2.43 log points below the mean industry, the sum of -2.84 from rates, -0.30 from the AI channel and +0.76 from mean reversion (they were on average below trend). The rate channel accounts for more than the whole of the shortfall of rate-sensitive industries, part of which was offset by their catching up from below trend. A no-tightening counterfactual, in which x_t is held at zero throughout, puts the top

rate-sensitivity quartile 0.45 log points above the mean industry in July 2026 rather than 2.43 below it; the difference, 2.9 log points with a bootstrap standard error of 0.8, is the model’s estimate of what the 2022 to 2024 cycle cost the most rate-sensitive quarter of industries in relative employment. It is a model-dependent quantity: it takes the reduced-form loading ρ_i as a structural sensitivity, holds the estimated exposures fixed, and, because only relative effects are identified, says nothing about the aggregate employment effect of tightening.

The fitted κ_r has a direct interpretation in terms of the job value of [equation \(1\)](#). The mean standardized rate sensitivity in the top quartile is about 1.2, so the implied semi-elasticity of the job value with respect to the two-year yield in these industries is about 1.8 percent per percentage point. At the peak change of 3.63 percentage points, the value of a filled job in a top-quartile industry fell by roughly seven percent relative to the mean industry. [Hall \[2017\]](#) attributes movements of this order in the value of a match to discount-rate variation over the business cycle and finds them sufficient, through free entry, to generate the observed volatility of hiring. The cross-industry estimate here is in the same range and was obtained from an episode in which the discount-rate movement was a deliberate policy action rather than an inferred component of asset prices.

9. Robustness

9.1 The Battery

[Table 9](#) re-estimates the easing-regime coefficients on AI exposure and rate sensitivity under sixteen alternative specifications, and [Figure 8](#) plots them.

Table 9: Robustness of the Headcount Results (Easing Regime)

Specification	AI exposure	Rate sensitivity	<i>N</i>
Baseline	1.108 (0.833)	−2.011*** (0.661)	205
AI only, no controls	0.858 (0.906)		205
General AIIIE	1.063 (0.868)	−2.019*** (0.660)	205
RS: Jarocinski-Karadi	1.315 (0.807)	−3.474*** (0.687)	205
RS: cyclical beta	0.014 (0.870)	−3.984*** (0.761)	205
RS: young-firm share	1.187 (0.899)	−0.295 (1.085)	190
RS: first principal component	0.750 (0.791)	−3.694*** (0.665)	205
RS: no shrinkage	1.021 (0.839)	−1.701* (0.951)	205
No detrending	0.929 (0.767)	−3.030*** (0.699)	205
Pre-trend from 2010-2019	0.905 (0.789)	−2.858*** (0.694)	205
Weighted by 2019 employment	1.128 (0.765)	−2.601*** (0.629)	205
Sector fixed effects	2.780** (1.282)	−1.248* (0.676)	205
Drop tech core	1.771** (0.776)	−2.202*** (0.645)	200
Drop leisure and hospitality	1.068 (0.828)	−2.188*** (0.666)	192
Drop R&D-intensive (Section 174)	1.193 (0.852)	−2.047*** (0.686)	192
Services only	0.188 (1.194)	−1.989** (0.973)	135
Stable industry definitions only	1.829** (0.802)	−2.186*** (0.674)	188

Notes: Each row re-estimates the easing-regime coefficients of Table 2, column (3), changing one element; the heteroskedasticity-robust cross-sectional standard errors differ from the industry-clustered panel standard errors of Table 2 only through degrees-of-freedom corrections. The rate sensitivity (RS) rows replace the [Bauer and Swanson \[2023\]](#) measure by the [Jarociński and Karadi \[2020\]](#) analogue, by the industry’s cyclical beta, by the 2019 employment share of firms aged 0–5 (BDS), or by the first principal component of all four. ‘Stable industry definitions’ drops every industry whose code changed between NAICS 2017 and NAICS 2022. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.




Figure 8: Robustness of the Easing-Regime Coefficients

Notes: Easing-regime coefficients on standardized AI exposure (left) and rate sensitivity (right) from the specifications of Table 9, with 95 percent intervals. Each row changes one element of the baseline (Table 2, column 3). The rate-sensitivity panel is empty in the rows where an alternative measure replaces it and in the AI-only row.

The rate coefficient is negative and significant at five percent or better in every specification that uses an employment-response measure of sensitivity: -2.0 in the baseline and with the general AIIE, -3.5 with the Jarocinski-Karadi measure, -4.0 with the cyclical beta, -3.7 with the first principal component, -3.0 without detrending, -2.9 with a 2010 to 2019 pre-trend, -2.6 weighted by employment, -2.2 dropping the tech core or leisure and hospitality, -2.0 dropping the research-intensive industries and -2.0 in services only. Without shrinkage it is -1.7 with a standard error of 0.95, significant at ten percent, the noisier regressor attenuating as expected. With sector fixed effects it is -1.2 (0.68, significant at ten percent), so part of the rate effect operates between sectors, as one would expect when construction and durable manufacturing are among the most sensitive; a within-sector effect remains. The one measure on which the rate effect disappears is the young-firm share (-0.3 , standard error 1.09). That measure is nearly uncorrelated with the others (0.13 with the baseline) and, as Fort et al. [2013] note, young-firm employment shares reflect industry life cycles as much as financial fragility; it is reported because it is the one measure not estimated from employment responses, and it fails.

The AI coefficient is positive in every specification and significant in three. It ranges from 0.01 (with the cyclical beta as the rate measure, which absorbs the negative correlation between AI exposure and cyclicalities) to 2.78 (with sector fixed effects, standard error 1.28,

significant at five percent). The sector-fixed-effects result deserves a word, because within-sector permutation p -values of 0.05 to 0.07 point the same way. Within information, finance and professional services, the less exposed industries are newspaper and periodical publishing, broadcasting, telecommunications carriers, depository branches and printing-adjacent services, industries in secular decline whose 2015 to 2019 trend understates their post-2022 losses; comparing the exposed industries with those neighbors flatters the exposed ones. The between-sector comparison, which the baseline uses, does not have this feature, and the paper does not read the within-sector coefficient as an AI employment gain. Dropping the tech core raises the coefficient to 1.77 (0.78): the five industries at the center of the AI narrative did worse than the other AI-exposed industries, so their removal makes the rest look better. Restricting the panel to the 188 industries whose NAICS codes did not change between the 2017 and 2022 vintages, which removes the information and retail industries whose crosswalk is approximate ([Section B](#)), gives 1.83 (0.80) for AI and -2.19 (0.67) for rate sensitivity. Across the fourteen leave-one-sector-out estimates in [Section C](#) it ranges from 0.17 (dropping transportation and warehousing) to 1.87 (dropping information). No specification produces a negative AI coefficient, and none produces a positive one that a skeptic would take as evidence of an AI employment gain; the range is that of a noisy null.

9.2 The Placebo Event of December 2017

[Table 10](#) collects the checks discussed in the rest of this section.

Table 10: Additional Checks on the Headcount Results

Check	AI exposure	Rate sensitivity
Split-sample IV for rate sensitivity	4.15	157.5
first-stage $F = 0.02$; Anderson–Rubin set unbounded	(25.02)	(1286.3)
Leave-one-sector-out range, 14 sectors	[0.17, 1.87]	
Placebo event, December 2017 (2015–2019 data)	−1.188***	−0.478
	(0.441)	(0.365)
Placebo-in-time, 19 base years: mean	−0.43	−1.61
standard deviation	2.13	3.91
range	[−4.59, 3.04]	[−8.09, 6.56]
Permutation p -value, unrestricted	0.144	
Permutation p -value, within sector	0.063	

Notes: Easing-regime coefficients on the two exposures (log points per s.d.) unless stated. The split-sample IV instruments the raw sensitivity estimated on shocks of 1990m9–2000m6 by the one estimated on 2003m7–2016m12; a 36-month gap separates the two samples so that no outcome month is shared; the Anderson–Rubin set is the set of values of the rate coefficient not rejected at 5% by a robust test of the instrument in the reduced form. The placebo-in-time rows replay the full design (trend window, reference year, horizon) at 19 fake base months, February 1997 to February 2015; the windows overlap and rate sensitivity is estimated on data that overlap them, so the distribution is descriptive. Permutation p -values are $(k + 1)/(B + 1)$ with $B = 2,000$ reshuffles of the AI exposure vector, unrestricted or within CES sectors. Robust standard errors. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A placebo assigns a fake ChatGPT date of December 2017, uses data from January 2015 to December 2019 only, estimates the industry trend on the months through November 2017, and regresses the change in detrended employment between the year before the fake event (December 2016 to November 2017) and the second year after it (December 2018 to December 2019) on AI exposure and rate sensitivity. The rate coefficient is -0.48 (standard error 0.37), unremarkable. The AI coefficient is -1.19 (0.44), significant at one percent. The placebo fails for AI.

The failure deserves to be read carefully rather than explained away. It means that, relative to a trend fitted on a 35-month window and relative to the year before the fake event, industries sorted on AI exposure drifted by about 1.2 log points per standard deviation over a two-year window in a period in which nothing happened to them. The drift is negative, the opposite sign of the actual post-2022 estimate, so it cannot be the source of the positive coefficient; if anything, drift of this sign in 2022 to 2026 would have made the actual estimate more negative than the truth. What the placebo establishes is that a coefficient of about one log point in absolute value on a fixed characteristic over a two-year window is ordinary for this design, whether one calls it noise or the influence of whatever else moves industries sorted on this characteristic. The actual easing coefficient on AI is 1.1 log points. It is the size of a placebo.

9.3 The Placebo-in-Time Distribution

The 2017 placebo is one draw. [Figure 9](#) reports nineteen, obtained by replaying the exact design, trend window, reference year, hold and easing horizons and all, at fake base months of February 1997 through February 2015, alongside the permutation distribution described below.




Figure 9: What Is a Normal Coefficient? Placebo-in-Time and Permutation Distributions

Notes: Left: the easing-analogue coefficient on standardized AI exposure (months 31 to 53 after the base month, relative to the twelve months ending at the base month, with the industry trend estimated on a 60-month window ending 26 months before the base and displacement recomputed at the base) when the full design is replayed at nineteen fake base months, February 1997 to February 2015 (circles), and the actual estimate at the February 2022 base (diamond). Right: the distribution of the easing-regime AI coefficient across 2,000 random reassignments of the AI exposure vector across industries, holding rate sensitivity and displacement fixed; the vertical line is the actual estimate.

The nineteen placebo AI coefficients have a mean of -0.43 log points, a standard deviation of 2.13 and a range from about -4.6 (base February 2011) to $+3.0$ (base February 2007). Eleven of the nineteen, 58 percent, exceed the actual estimate of 1.11 in absolute value. The distribution has a visible business-cycle pattern: the AI coefficient is positive for base years in the mid-2000s expansion, when the horizon window falls in 2008 to 2010 and the exposed white-collar industries lost less employment than the goods industries, and negative for base years around 2010, when the horizon falls in the recovery of construction and manufacturing. The windows overlap, so these are not nineteen independent draws, and the pattern is not noise in the statistical sense: it is the influence of the business cycle on industries sorted

by AI exposure, an omitted variable in any single window. That is the point. Industries sorted on AI exposure move over four-year windows, for reasons unrelated to AI, by amounts that dwarf the cross-sectional standard error of any single window, because the exposure is correlated with the sectoral composition of the business cycle even though it is uncorrelated with the measured rate sensitivity.

The rate-sensitivity placebos have a mean of -1.61 , a standard deviation of 3.91 and a range from about -8.1 (base February 2007, with the horizon in the Great Recession) to $+6.6$ (base February 2012, with the horizon in the recovery); 53 percent are at or below the actual estimate of -2.01 . Because rate sensitivity is estimated on 1990 to 2019 data that overlap these windows, this is not a null distribution: a measure estimated from an industry’s response to past cycles will of course predict its performance in those cycles. What the distribution does say is that the gap that opened between rate-sensitive and insensitive industries over 2022 to 2026 is of ordinary size for a cycle, no larger than a typical episode of the past three decades.

9.4 Permutation Inference

The right panel of [Figure 9](#) shows the permutation reference distribution of the easing-regime AI coefficient. The two-sided permutation p -value, computed as $(k + 1)/(B + 1)$, is 0.143 for the easing regime, 0.144 for the hold and 0.173 for the pre-ChatGPT tightening months. Restricting the reshuffling to within CES sectors, so that the sectoral composition of exposure is preserved and only the within-sector allocation varies, gives 0.062, 0.045 and 0.067. The within-sector values are lower because the within-sector comparison is the one in which AI-exposed industries did better, as the sector fixed effects row of [Table 9](#) also shows and for the reason given there. A fixed industry characteristic is not randomly assigned, so these are reference distributions under an exchangeability assumption rather than exact tests. A reader inclined to see a positive AI effect on headcount in these numbers should weigh them against the placebo-in-time distribution; a reader inclined to see a negative one has nothing to point to.

9.5 The Generated Regressor

Rate sensitivity is estimated, not observed. Two consequences follow and are handled separately: its measurement error attenuates its coefficient, and its sampling error is not in the clustered standard errors. On the second, [Table 3](#) and the discussion in [Section 6](#) propagate the first-stage covariance through the design; the Murphy-Topel standard error of the easing-regime rate coefficient is 1.35 against a clustered 0.66, and its p -value is 0.14. On the

first, shrinkage reorders industries toward better-measured neighbors, and the no-shrinkage row of [Table 9](#) shows that the unshrunk measure gives a smaller coefficient with a larger standard error, as attenuation predicts; but shrinkage does not remove errors-in-variables bias in the second stage, and the paper makes no lower-bound claim.

The split-sample instrumental-variables estimate was designed to address attenuation directly. It instruments the raw sensitivity estimated on shocks from September 1990 to June 2000 by the raw sensitivity estimated on shocks from July 2003 to December 2016; the 36-month gap between the two samples ensures that the early sample’s dependent variables, which reach three years past each shock, share no outcome month with the late sample’s shocks. The exercise fails. The first-stage F statistic is 0.02, the Anderson-Rubin confidence set for the rate coefficient is the whole real line, and the two raw sensitivities are correlated only 0.01 across industries. If the true sensitivities were stable across the two halves of the sample, the split-sample reliabilities of 0.28 and 0.31 would imply a correlation between the two raw estimates of about 0.29, so the observed correlation of 0.01 says either that the ranking of industries by monetary sensitivity changed between the 1990s and the 2000s, or that the first-stage errors are larger and more structured than the Newey-West standard errors indicate. Either reading reinforces the conclusion of [Table 3](#): the cross-industry pattern of monetary sensitivity is measured with substantial uncertainty, and the paper’s rate estimates are conditional on the full-sample pattern. The AI coefficient is unaffected by any treatment of rate sensitivity, as the near-zero correlation between the two exposures implies.

9.6 What the Headcount Design Can and Cannot Distinguish

The results of this section should be stated as a limitation of the evidence on AI, and stated plainly, with the distinction between a confidence bound and a power statement kept clear. Under the sampling model of the clustered standard errors, the 95 percent interval for the easing-regime AI coefficient is -0.5 to $+2.7$ log points per standard deviation, and the minimum detectable effect at 80 percent power is 2.3. But industries sorted on any fixed characteristic drift over a four-year window for reasons the sampling model does not contain, and the placebo-in-time standard deviation of 2.1 log points is a descriptive scale for that drift, not a standard error. Used as if it were one, it would give an interval of about -3.1 to $+5.3$ around the point estimate and a minimum detectable effect of about six log points per standard deviation. The paper’s statement is therefore the weaker one: an AI effect on headcount would have to be roughly four log points per standard deviation, six across the interquartile range, to stand out from the ordinary movement of industries sorted on this characteristic; an effect that reduced employment in the most exposed quarter of industries by two or three percent over four years relative to the least exposed quarter would not stand

out, and the positive point estimate does not exclude it. What the design does establish is that no effect of the magnitude implied by the popular narrative, in which generative AI has visibly reduced employment in exposed industries, is present in the headcount data, while an association between historical cyclical and monetary sensitivity and post-2022 employment losses, of the predicted sign and timing and with an attenuation under easing, is present, robust to the grouping of industries, and imprecise once the estimation of the sensitivities is taken into account.

10. Conclusion

The post-2022 slowdown in white-collar employment coincided with three events: the fastest monetary tightening in four decades, the arrival of generative AI, and the end of a pandemic hiring overshoot. This paper separates them in the cross-section of 205 detailed industries using three predetermined exposures that turn out to be nearly orthogonal, and it uses the timing of the events, the sign reversal that policy but not technology implies, and the age composition of employment within industries to tell the channels apart.

In headcount, the evidence points to the pandemic hangover and to cyclical sensitivity, including monetary sensitivity, and not to AI. Industries with a history of losing employment after identified monetary shocks lost about two log points of detrended employment per standard deviation of that sensitivity relative to the year before liftoff; the loss opened with the roughly one-year lag familiar from the transmission literature, was absent in the months of tightening before ChatGPT and present in the months after, and the growth disadvantage narrowed by more than half once the Federal Reserve began cutting, without closing. A continuous specification associates a one-percentage-point rise in the two-year yield with 0.33 percentage points lower annualized growth twelve months later per standard deviation of sensitivity. These associations survive clustering by sector and a wild-cluster bootstrap. They are less precise than the clustered standard errors suggest, because the sensitivities are estimated from 26 years of shocks with a reliability ratio of 0.39 and errors that are correlated across industries; once that uncertainty is propagated, the easing-regime rate coefficient has a p -value of 0.14, and a precisely measured cyclical beta predicts the same losses. Industries that had overshoot their trend during the pandemic gave back about 60 percent of the excess, the largest single force in the data. AI-exposed industries, holding the other two exposures fixed, did slightly and insignificantly better than others, showed no break at ChatGPT within an interval that allows a slowdown of up to one point a year, showed no reversal under easing, and responded to rates exactly as their measured sensitivity implied and no more. A stock-flow hiring model with a discount-rate channel reproduces the rate and mean-reversion

paths and cannot reproduce the AI path, because the AI path does not look like the diffusion of a technology.

Among young workers the picture is genuinely mixed and the paper’s contribution is about timing. The relative employment of 22 to 24 year olds has fallen sharply in AI-exposed industries, by nearly five log points per standard deviation of exposure by 2025, a result that survives every robustness check, lies outside the range of the same design at five fake base quarters, and vindicates the canaries reported in payroll microdata. It has fallen by a comparable amount in rate-sensitive industries, through the hiring freeze that tightenings impose on the young. The two gradients differ in their histories. The rate gradient was flat until liftoff and then broke. The AI gradient fell in two steps: a large one in the pandemic year, before the technology existed, then a plateau, then a second decline from about the start of 2023 at half the pace of the first, faster than the plateau and equal to the 2020 to 2021 average. The first step had a cause other than generative AI; the second coincides with adoption, with the arrival of the tightening in the labor market and with the unwinding of the hiring boom, and the data do not assign it. An account of the youngest professionals’ position in these industries that begins in November 2022 begins too late; an account that ends there would be equally incomplete.

Three limitations bound these conclusions. The design is an exposure design, and its identifying assumption, that no other post-2022 shock loaded on industries in proportion to AI exposure, historical rate sensitivity or pandemic displacement, is an assumption; the paper probes it with placebos and alternative measures rather than proving it. The headcount design’s power against modest AI effects is limited, not by conventional standard errors but by the drift of industries sorted on any characteristic over multi-year windows; the placebo distributions put the effect that would stand out at roughly four log points per standard deviation. And the design is cross-sectional: it identifies relative effects across industries, its rate estimates are conditional on an imprecisely estimated pattern of sensitivity, and it is silent on the aggregate employment effects of either tightening or AI. All three argue for observing the same exposures at the level of firms and workers, where the adoption of generative AI can be measured directly rather than proxied by occupational exposure.

For the policy debate the message is a modest one. The evidence that generative AI has reduced employment in exposed industries is, as of mid-2026, not present in public headcount data at any magnitude those data could reveal. The evidence from young workers is real, and it has antecedents that predate the technology by three years and a decade. The evidence that the industries which historically respond to monetary policy lost ground after the 2022 liftoff, including several of the white-collar industries at the center of the AI narrative, is present and correctly timed, and its growth disadvantage eased when policy did. When the

next round of layoffs in software or finance is attributed to a language model, the two-year yield eighteen months earlier, and the hiring of two years earlier, deserve at least as much attention.

References

- John M. Abowd, Bryce E. Stephens, Lars Vilhuber, Fredrik Andersson, Kevin L. McKinney, Marc Roemer, and Simon Woodcock. The LEHD infrastructure files and the creation of the quarterly workforce indicators. In Timothy Dunne, J. Bradford Jensen, and Mark J. Roberts, editors, *Producer Dynamics: New Evidence from Micro Data*, pages 149–230. University of Chicago Press, 2009.
- Daron Acemoglu. The simple macroeconomics of AI. *Economic Policy*, 40(121):13–58, 2025.
- Daron Acemoglu and Pascual Restrepo. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30, 2019.
- Daron Acemoglu and Pascual Restrepo. Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6):2188–2244, 2020.
- Daron Acemoglu, David Autor, Jonathon Hazell, and Pascual Restrepo. Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1):S293–S340, 2022.
- Rodrigo Adão, Michal Kolesár, and Eduardo Morales. Shift-share designs: Theory and inference. *Quarterly Journal of Economics*, 134(4):1949–2010, 2019.
- Stefania Albanesi, António Dias da Silva, Juan F. Jimeno, Ana Lamo, and Alena Wabitsch. New technologies and jobs in Europe. Working Paper 31357, National Bureau of Economic Research, 2023.
- David H. Autor. Why are there still so many jobs? the history and future of workplace automation. *Journal of Economic Perspectives*, 29(3):3–30, 2015.
- David H. Autor, Frank Levy, and Richard J. Murnane. The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4):1279–1333, 2003.
- Tania Babina, Anastassia Fedyk, Alex He, and James Hodson. Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745, 2024.
- Jose Maria Barrero, Nicholas Bloom, and Steven J. Davis. The evolution of work from home. *Journal of Economic Perspectives*, 37(4):23–50, 2023.
- Michael D. Bauer and Eric T. Swanson. A reassessment of monetary policy surprises and high-frequency identification. *NBER Macroeconomics Annual*, 37:87–155, 2023.

- Ben S. Bernanke and Mark Gertler. Inside the black box: The credit channel of monetary policy transmission. *Journal of Economic Perspectives*, 9(4):27–48, 1995.
- Alexander Bick, Adam Blandin, and David J. Deming. The rapid adoption of generative AI. Working Paper 32966, National Bureau of Economic Research, 2024.
- Kirill Borusyak, Peter Hull, and Xavier Jaravel. Quasi-experimental shift-share research designs. *Review of Economic Studies*, 89(1):181–213, 2022.
- Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen. Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence. Working paper, Stanford Digital Economy Lab, 2025a.
- Erik Brynjolfsson, Danielle Li, and Lindsey Raymond. Generative AI at work. *Quarterly Journal of Economics*, 140(2):889–942, 2025b.
- Brantly Callaway, Andrew Goodman-Bacon, and Pedro H. C. Sant’Anna. Difference-in-differences with a continuous treatment. Working Paper 32117, National Bureau of Economic Research, 2024.
- A. Colin Cameron, Jonah B. Gelbach, and Douglas L. Miller. Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3):414–427, 2008.
- Gabriel Chodorow-Reich. Regional data in macroeconomics: Some advice for practitioners. *Journal of Economic Dynamics and Control*, 115:103875, 2020.
- Clément de Chaisemartin and Xavier D’Haultfoeulle. Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–2996, 2020.
- Luca Dedola and Francesco Lippi. The monetary transmission mechanism: Evidence from the industries of five OECD countries. *European Economic Review*, 49(6):1543–1569, 2005.
- Jonathan I. Dingel and Brent Neiman. How many jobs can be done at home? *Journal of Public Economics*, 189:104235, 2020.
- Sarah Eckhardt and Nathan Goldschlag. AI and jobs: The final word (until the next one). Report, Economic Innovation Group, 2025.
- Andrea L. Eisfeldt, Gregor Schubert, and Miao Ben Zhang. Generative AI and firm values. Working Paper 31222, National Bureau of Economic Research, 2023.

- Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702):1306–1308, 2024.
- Edward Felten, Manav Raj, and Robert Seamans. Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12):2195–2217, 2021.
- Edward Felten, Manav Raj, and Robert Seamans. How will language modelers like ChatGPT affect occupations and industries? Technical Report 2303.01157, arXiv, 2023.
- Eliza Forsythe. Why don’t firms hire young workers during recessions? *Economic Journal*, 132(645):1765–1789, 2022.
- Teresa C. Fort, John Haltiwanger, Ron S. Jarmin, and Javier Miranda. How firms respond to business cycles: The role of firm age and firm size. *IMF Economic Review*, 61(3):520–559, 2013.
- Mark Gertler and Simon Gilchrist. Monetary policy, business cycles, and the behavior of small manufacturing firms. *Quarterly Journal of Economics*, 109(2):309–340, 1994.
- Mark Gertler and Peter Karadi. Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, 7(1):44–76, 2015.
- Martha Gimbel, Molly Kinder, Joshua Kendall, and Maddie Lee. Evaluating the impact of AI on the labor market: Current state of affairs. Report, The Budget Lab at Yale, 2025.
- Paul Goldsmith-Pinkham, Isaac Sorkin, and Henry Swift. Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624, 2020.
- Andrew Goodman-Bacon. Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277, 2021.
- Robert E. Hall. High discounts and high unemployment. *American Economic Review*, 107(2):305–330, 2017.
- Seyed M. Hosseini and Guy Lichtinger. Generative AI as seniority-biased technological change: Evidence from U.S. résumé and job posting data. Working paper, SSRN, 2025.
- Anders Humlum and Emilie Vestergaard. Large language models, small labor market effects. Working Paper 33777, National Bureau of Economic Research, 2025.
- Marek Jarociński and Peter Karadi. Deconstructing monetary policy surprises: The role of information shocks. *American Economic Journal: Macroeconomics*, 12(2):1–43, 2020.

- Òscar Jordà. Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182, 2005.
- Lisa B. Kahn. The long-term labor market consequences of graduating from college in a bad economy. *Labour Economics*, 17(2):303–316, 2010.
- Patrick J. Kehoe, Pierlauro Lopez, Virgiliu Midrigan, and Elena Pastorino. Asset prices and unemployment fluctuations: A resolution of the unemployment volatility puzzle. *Review of Economic Studies*, 90(3):1304–1357, 2023.
- José Luis Montiel Olea and Mikkel Plagborg-Møller. Local projection inference is simpler and more robust than you think. *Econometrica*, 89(4):1789–1823, 2021.
- Carl N. Morris. Parametric empirical Bayes inference: Theory and applications. *Journal of the American Statistical Association*, 78(381):47–55, 1983.
- Dale T. Mortensen and Christopher A. Pissarides. Job creation and job destruction in the theory of unemployment. *Review of Economic Studies*, 61(3):397–415, 1994.
- Kevin M. Murphy and Robert H. Topel. Estimation and inference in two-step econometric models. *Journal of Business & Economic Statistics*, 3(4):370–379, 1985.
- Emi Nakamura and Jón Steinsson. Identification in macroeconomics. *Journal of Economic Perspectives*, 32(3):59–86, 2018.
- Whitney K. Newey and Kenneth D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987.
- Shakked Noy and Whitney Zhang. Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654):187–192, 2023.
- Philip Oreopoulos, Till von Wachter, and Andrew Heisz. The short- and long-term career effects of graduating in a recession. *American Economic Journal: Applied Economics*, 4(1):1–29, 2012.
- Pablo Ottonello and Thomas Winberry. Financial heterogeneity and the investment channel of monetary policy. *Econometrica*, 88(6):2473–2502, 2020.
- Gert Peersman and Frank Smets. The industry effects of monetary policy in the euro area. *Economic Journal*, 115(503):319–342, 2005.
- Barbara Petrongolo and Christopher A. Pissarides. Looking into the black box: A survey of the matching function. *Journal of Economic Literature*, 39(2):390–431, 2001.

- Christopher A. Pissarides. *Equilibrium Unemployment Theory*. MIT Press, Cambridge, MA, 2nd edition, 2000.
- Ashesh Rambachan and Jonathan Roth. A more credible approach to parallel trends. *Review of Economic Studies*, 90(5):2555–2591, 2023.
- Valerie A. Ramey. Macroeconomic shocks and their propagation. In John B. Taylor and Harald Uhlig, editors, *Handbook of Macroeconomics*, volume 2, pages 71–162. Elsevier, 2016.
- Christina D. Romer and David H. Romer. A new measure of monetary shocks: Derivation and implications. *American Economic Review*, 94(4):1055–1084, 2004.
- Jonathan Roth. Pretest with caution: Event-study estimates after testing for parallel trends. *American Economic Review: Insights*, 4(3):305–322, 2022.
- Robert Shimer. The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, 95(1):25–49, 2005.
- Michael Webb. The impact of artificial intelligence on the labor market. Working paper, Stanford University, 2020.
- Alwyn Young. Channeling Fisher: Randomization tests and the statistical insignificance of seemingly significant experimental results. *Quarterly Journal of Economics*, 134(2): 557–598, 2019.

Appendix Overview

A	Model Derivation	p. 66
B	Data Appendix	p. 69
C	Additional Results	p. 74

A. Model Derivation

This appendix derives the equations of [Section 3](#), explains why only relative effects are identified, and records the parameter restrictions used in estimation.

A.1 The Job Value as Capitalized Surplus

Consider a filled job in industry i that produces a flow surplus $\pi_i(\tau) = (p_i - w_i)e^{g_i\tau}$ at tenure τ , where p_i is the worker's marginal product, w_i the wage, and g_i the growth rate of the surplus over the life of the match, which can reflect learning on the job, the accumulation of match-specific capital, or a wage profile that is back-loaded relative to productivity. The match dissolves at the Poisson rate s_i and the firm discounts at r . The value of the job at the moment it is filled is

$$J_i = \int_0^\infty e^{-(r+s_i)\tau} (p_i - w_i) e^{g_i\tau} d\tau = \frac{p_i - w_i}{r + s_i - g_i}, \quad (11)$$

which requires $r + s_i > g_i$. The semi-elasticity of the job value with respect to the discount rate is

$$\frac{\partial \ln J_i}{\partial r} = -\frac{1}{r + s_i - g_i}, \quad (12)$$

which is the (Macaulay) duration of the surplus stream. A match with low separation risk and a rising surplus is a long-duration asset, and its value is sensitive to r in the same way that the price of a long bond or a growth stock is. Writing $r = \bar{r} + x_t$ for the change in the discount rate since the base month and taking a first-order approximation around $x_t = 0$,

$$\ln J_{it} \approx \ln J_i(\bar{r}) - \frac{x_t}{\bar{r} + s_i - g_i}. \quad (13)$$

The coefficient on x_t varies across industries with s_i and g_i , neither of which is observed. The empirical strategy replaces it with the estimated rate sensitivity: it sets $1/(\bar{r} + s_i - g_i) = \bar{\kappa} + \kappa_r \rho_i$, where ρ_i is the standardized local-projection estimate of the industry's employment response to identified monetary shocks and $\bar{\kappa}$ is the mean duration. The paper's κ_r is therefore the

cross-industry slope of the job-value semi-elasticity in the estimated rate sensitivity, in log points per percentage point per standard deviation.

The AI channel enters the surplus. Let adoption A_t change the flow surplus of a match in industry i by a factor $\exp(\kappa_a a_i A_t)$, with a_i the standardized exposure. Under the productivity interpretation, $\kappa_a > 0$: the technology raises the output of an exposed worker by more than it raises the wage. Under the displacement interpretation, $\kappa_a < 0$: the tasks an exposed worker performed can be done without a worker, so the surplus of a marginal match falls. [Acemoglu and Restrepo \[2019\]](#) discuss why the net sign is an empirical matter. Since $\ln J$ is linear in $\ln(p_i - w_i)$, the combined job value is [equation \(2\)](#).

A.2 Free Entry and the Hiring Rate

Firms in industry i post vacancies v at a flow cost c_i per vacancy and fill them at the rate $q(\theta_i)$, where θ_i is the ratio of vacancies to job seekers in the industry's labor market and q is decreasing in θ_i . Free entry sets the expected cost of a hire equal to the value of the job:

$$\frac{c_i}{q(\theta_i)} = J_i. \quad (14)$$

With a Cobb-Douglas matching function $m = \mu u^\alpha v^{1-\alpha}$, $q(\theta) = \mu \theta^{-\alpha}$, and [equation \(14\)](#) gives $\theta_i = (\mu J_i / c_i)^{1/\alpha}$. Hires per period are $m = \theta q(\theta) u = \mu \theta^{1-\alpha} u$, so, holding the pool of job seekers fixed, hires are proportional to $J_i^{(1-\alpha)/\alpha}$. Defining $\eta = (1 - \alpha)/\alpha$ and normalizing the level of hiring so that it equals separations when $\ln J_i = 0$ gives [equation \(3\)](#). With $\alpha = 0.5$, the center of the range surveyed by [Petrongolo and Pissarides \[2001\]](#), $\eta = 1$ and the hiring rate is proportional to the job value.

Holding the pool of job seekers fixed is the assumption that makes the model partial equilibrium. In a general-equilibrium version, a fall in job values across industries would raise unemployment and lower $c_i/q(\theta_i)$ for every industry, an aggregate response that the cross-sectional design does not identify and the model does not attempt to capture.

A.3 The Law of Motion and Mean Reversion

Employment in industry i relative to its steady-state level, n_{it} , evolves according to [equation \(4\)](#). Setting $\ln J_{it} = 0$ gives $n_{i,t+1} - 1 = (1 - s)(n_{it} - 1)$: any initial deviation decays geometrically with monthly persistence $1 - s$ and half-life $\ln(0.5)/\ln(1 - s)$. At the estimated $s = 0.0178$ the half-life is 38.6 months. With $\ln J_{it} \neq 0$ the steady-state level toward which employment converges is $\exp(\eta \ln J_{it})$, so a permanent reduction in the job value of $\kappa_r \rho_i x$ lowers the industry's long-run employment by $\eta \kappa_r \rho_i x$ log points, approached at the same

geometric rate. In the model, the pandemic displacement d_i and the post-2022 movements in job values therefore share the same adjustment speed, which the data confirm in the sense that a single s fits both the displacement path and the rate path.

A.4 Why Only Relative Effects Are Identified

Every month effect λ_t in the empirical specifications absorbs whatever moves employment in all industries alike. In the model, this is the term ϕ_t in [equation \(2\)](#), which contains the aggregate effect of the change in discount rates on the mean industry, the aggregate effect of adoption, and any general-equilibrium feedback from the pool of job seekers. The event-study coefficient on ρ_i identifies $\kappa_r x_t$ times the cross-industry slope of employment in ρ_i , which does not depend on ϕ_t ; it does not identify the level effect. This is the missing-intercept problem of cross-sectional macroeconomic designs, discussed by [Nakamura and Steinsson \[2018\]](#) and, for regional data, by [Chodorow-Reich \[2020\]](#): the design delivers the relative effect of the shock on more-exposed versus less-exposed units, and translating this into an aggregate effect requires a model of the aggregate response that the cross-section cannot supply. The paper accordingly reports every model quantity, including the decomposition and the counterfactual, relative to the mean industry, and makes no claim about the aggregate employment cost of the 2022 to 2024 tightening or the aggregate employment effect of generative AI.

A.5 Estimation Details

The simulated event study is computed as follows. For a candidate (s, κ_r, κ_a) , employment paths are simulated for all 205 industries from February 2022 to July 2026 using [equations \(2\) to \(4\)](#), with x_t the observed change in the two-year yield since February 2022, A_t the logistic adoption path, and $\ln n_{i,\text{Feb } 2022} = d_i$. The change in $100 \times \ln n_{it}$ since the base month is regressed cross-sectionally, month by month, on a constant and the three standardized exposures, exactly as in [equation \(6\)](#); because the simulated data contain no other variation, the simulated coefficients are deterministic functions of the parameters. The objective is $\sum_c \sum_{t>t_0} \left(\hat{\beta}_t^c - \beta_t^c(s, \kappa_r, \kappa_a) \right)^2 / \widehat{\text{se}}(\hat{\beta}_t^c)^2$, summed over the three exposures c and the 53 post-base months. Bounds are $s \in [0.002, 0.25]$, $\kappa_r \in [-0.2, 0.2]$ and $\kappa_a \in [-0.5, 0.5]$; six starting values are used and the minimum is unique. Standard errors are from a cluster bootstrap over industries: in each of 200 replications the 205 industries are resampled with replacement, the 159 target coefficients are recomputed by regressing the resampled industries' long differences on their (fixed) standardized exposures, and the model is refit from the point estimate; the reported standard errors and the interval for the counterfactual gap are the standard deviation and percentiles across replications. The curvature-based standard errors (the square roots of

the diagonal of twice the inverse of the numerical Hessian of the objective) are reported in the notes to [Table 8](#) for comparison only; they describe the sensitivity of the objective to the parameters, not sampling uncertainty, because the 159 moments are correlated across months and exposures. Neither set of standard errors carries the first-stage uncertainty of the rate-sensitivity exposure.

B. Data Appendix

B.1 CES Series Construction

Candidate CES series identifiers were built for every four-digit NAICS code in the [Felten et al. \[2023\]](#) industry list as CES followed by the two-digit CES supersector code, the NAICS code padded to six digits with zeros, and the data-type code 01 (all employees, thousands, seasonally adjusted). The supersector code was assigned from the NAICS prefix (for example, 10 for mining and logging, 31 and 32 for durable and nondurable manufacturing, 50 for information, 60 for professional and business services). The BLS API was queried in two windows (1990 to 2009 and 2010 to 2026, within the API’s twenty-year limit per request) for each candidate, and 187 four-digit series were found to exist. For the remaining codes, the three-digit parent series was requested and used where it exists and where no four-digit child of the same parent is in the panel in its own right, giving 18 three-digit series. Fourteen further codes on the AIIE list are residual composites written with a trailing zero (for example, 3250 for “Chemical manufacturing (3251, 3252, 3253 and 3259 only)”); in the CES a trailing-zero code is the full three-digit industry, so for the 8 composites whose excluded children are in the panel, the composite’s employment is the parent minus the children, as [Table 11](#) lists. Total nonfarm employment (CES0000000001) is used as the aggregate control in the local projections and to form the cyclical beta.

Table 11: The Eight Residual Composites

AIIE code	CES parent	Children subtracted (in the panel in their own right)
3250	325 Chemical manufacturing	3254 Pharmaceutical and medicine manufacturing
3320	332 Fabricated metal products	3327 Machine shops and turned products; 3328 Coating, engraving and heat treating
3330	333 Machinery manufacturing	3333 Commercial and service industry machinery; 3335 Metalworking machinery; 3336 Engine, turbine and power transmission equipment
4230	423 Merchant wholesalers, durable goods	4231 Motor vehicle and parts; 4234 Professional and commercial equipment; 4238 Machinery, equipment and supplies
4240	424 Merchant wholesalers, nondurable goods	4243 Apparel, piece goods and notions; 4245 Farm product raw materials
4450	445 Food and beverage retailers	4453 Beer, wine and liquor retailers
5220	522 Credit intermediation	5222 Nondepository credit intermediation
5320	532 Rental and leasing services	5321 Automotive equipment rental and leasing

Notes: The AIIE list of [Felten et al. \[2023\]](#) carries these codes as composites of the four-digit industries not listed separately. The composite’s monthly employment is the CES three-digit parent minus the listed children, so that no job is counted twice; the residual is positive in every month for all eight. Six further trailing-zero composites on the list (3270, 3370, 4840, 5170, 5230 and 5310) have no children in the panel and use the parent series as published.

B.2 The Most and Least Exposed Industries

[Table 12](#) lists the ten highest and ten lowest industries on each of the two main exposures.

Table 12: The Most and Least Exposed Industries

Industry	s.d.	Industry	s.d.
<i>AI exposure: ten highest (left) and ten lowest (right)</i>			
Legal Services	2.19	Logging	-1.78
Securities, Commodity Contracts...	2.10	Foundation, Structure, and Buildi...	-1.66
Agencies, Brokerages, and Other I...	2.08	Services to Buildings and Dwellings	-1.65
Nondepository Credit Intermediation	2.04	Truck Transportation	-1.57
Insurance Carriers	1.97	Warehousing and Storage	-1.56
Accounting, Tax Preparation, Book...	1.90	Animal Slaughtering and Processing	-1.55
Business Support Services	1.87	Sawmills and Wood Preservation	-1.53
Software Publishers	1.87	Support Activities for Water Tran...	-1.48
Lessors of Nonfinancial Intangibl...	1.86	Building Finishing Contractors	-1.47
Business Schools and Computer and...	1.81	Waste Collection	-1.46
<i>Rate sensitivity: ten highest (left) and ten lowest (right)</i>			
Residential Building Construction	2.94	Grantmaking and Giving Services	-3.16
Employment Services	2.62	Fruit and Vegetable Preserving an...	-2.40
Newspaper, Periodical, Book, and...	1.98	Psychiatric and Substance Abuse H...	-2.36
Logging	1.98	Dairy Product Manufacturing	-2.27
Other Transportation Equipment Ma...	1.93	Petroleum and Coal Products Manuf...	-2.21
Other Wood Product Manufacturing	1.91	Colleges, Universities, and Profe...	-2.11
Furniture and Related Product Man...	1.83	Natural Gas Distribution	-2.06
Vocational Rehabilitation Services	1.78	Animal Food Manufacturing	-2.04
Sawmills and Wood Preservation	1.33	Other Food Manufacturing	-2.02
Nondepository Credit Intermediation	1.33	Electric Power Generation, Transm...	-1.83

Notes: Standardized exposures (mean zero, standard deviation one across the 205 industries). Titles are the AIIE industry titles of [Felten et al. \[2023\]](#), abbreviated; residual composites carry the parent’s title. Rate sensitivity is the shrunk [Bauer and Swanson \[2023\]](#) local-projection estimate.

B.3 NAICS 2017 to NAICS 2022 Crosswalk

The CES adopted the 2022 NAICS in 2023, and the reclassification restated the information and retail sectors; the AIIE is on the 2017 basis. [Table 13](#) lists the codes that changed and appear in the AIIE list. Where two 2017 codes map to one 2022 code, the 2022 industry’s exposure is the simple mean of the two 2017 exposures, since no pre-period employment weights exist on the old basis in the CES for the merged industries.

Table 13: NAICS 2017 to NAICS 2022 Codes Used

NAICS 2017	NAICS 2022	NAICS 2022 title
<i>Information</i>		
5111	5131	Newspaper, periodical, book and directory publishers
5112	5132	Software publishers
5151	5161	Radio and television broadcasting stations
5152	5162	Media streaming distribution services, social networks and other content providers
5173	5171	Wired and wireless telecommunications carriers (except satellite)
5179	5178	All other telecommunications
5191	5192	Web search portals, libraries, archives and other information services
<i>Retail trade</i>		
4421, 4422	4491	Furniture and home furnishings retailers
4431	4492	Electronics and appliance retailers
4461	4561	Health and personal care retailers
4471	4571	Gasoline stations and fuel dealers
4481	4581	Clothing and clothing accessories retailers
4482	4582	Shoe retailers
4483	4583	Jewelry, luggage and leather goods retailers
4511	4591	Sporting goods, hobby and musical instrument retailers
4512	4592	Book retailers and news dealers
4522	4551	Department stores
4523	4552	Warehouse clubs, supercenters and other general merchandise retailers
4531	4593	Florists
4532	4594	Office supplies, stationery and gift retailers
4533	4595	Used merchandise retailers
4539	4599	Other miscellaneous retailers

Notes: Only codes that changed between the 2017 and 2022 vintages and that appear on the AIIE industry list are shown. For the pair 4421 and 4422, the 2022 industry's AI exposure is the simple mean of the two 2017 exposures. The 2022 titles are abbreviated. The mapping is approximate where the 2022 revision redistributed rather than renamed an industry: the activities of 2017 code 5191 (other information services, including internet publishing) were split among 5162 (media streaming and social networks), 5192 (web search portals, libraries and archives) and the publishing codes, and the 2017 nonstore retailers (4541 to 4543) were distributed across the 2022 retail codes and are not on the panel. The 'stable industry definitions' row of [Table 9](#) drops every industry in this table and the three-digit information and retail parents affected by the revision.

B.4 The Local-Projection Specification

For each industry i and each shock series, [equation \(5\)](#) is estimated by ordinary least squares on monthly data from September 1990 to December 2016, with the dependent variable the

average over $h \in \{12, 18, 24, 30, 36\}$ of $y_{i,t+h} - y_{i,t-1}$, six lags each of the industry’s own monthly log employment change, of the shock and of the change in log total nonfarm employment as controls, and Newey and West [1987] standard errors with a bandwidth of 36 months. The shock is standardized to unit variance over the estimation sample. Every industry has 316 usable months in the full sample; the raw sensitivity is $-\hat{\beta}_i$. The split-sample estimates use shocks from September 1990 to June 2000 (118 usable months, with the minimum lowered to 100 for this purpose) and from July 2003 to December 2016 (162 months); the 36-month gap equals the longest horizon, so that the early sample’s dependent variables never extend into the late sample. Shrinkage is applied separately to each set of estimates as described in Section 4; the reliability ratios are 0.39 for the Bauer-Swanson full sample and 0.35 for the Jarocinski-Karadi full sample, and the split-sample reliabilities are 0.28 (Bauer-Swanson early), 0.31 (late), 0.24 (Jarocinski-Karadi early) and 0.03 (late). Variants of the baseline built from horizons 12 to 24 only, 24 to 36 only, a bandwidth of 18 and three lags have reliabilities of 0.20, 0.47, 0.30 and 0.32 and are correlated 0.95, 0.98, 0.99 and 0.95 with the baseline after shrinkage. The first-stage sampling covariance used in Table 3 is the covariance across 200 block-bootstrap replications (blocks of 24 months, drawn once per replication and applied to every industry) of the 205 raw estimates; its diagonal agrees closely with the Newey-West variances (correlation 0.97 across industries, means 0.31 and 0.27).

B.5 The QWI Balanced-Cell Rule

The QWI were retrieved from the Census API endpoint `timeseries/qwi/sa` for private ownership (`ownercode = A05`), all firm sizes and ages, at the four-digit NAICS level, for the age groups A03 (22 to 24), A04 (25 to 34), A05 (35 to 44), A06 (45 to 54) and A07 (55 to 64), for every quarter from 2015Q1 onward, with one request per state and block of age groups (the API has no national endpoint and no state wildcard) and results cached per state. The indicators used are beginning-of-quarter employment (**Emp**), hires (**HirA**) and separations (**Sep**). A state is retained if its earliest reported quarter is 2015Q1 or earlier and its latest reported quarter is at least the modal latest quarter across states, 2025Q4; this excludes Alaska and Michigan and retains 49 jurisdictions including the District of Columbia. Within retained states, a state-by-industry cell is retained only if **Emp** is reported for all five age groups in all 44 quarters; the retained cells account for 99.6 percent of the employment reported in the retained states. National aggregates are sums over retained cells. The QWI four-digit codes, on the 2022 NAICS basis, are matched to the CES panel’s industries through the same code list, with three-digit and residual-composite series handled as in the CES; 204 of the 205 industries are matched, with mean 2019 employment of 98.0 million in the matched cells. Each outcome is regressed, industry by industry, on four quarter-of-year indicators and

a linear trend estimated on 2015Q1 to 2019Q4, and the residual is used throughout.

B.6 Other Series

The Bauer-Swanson orthogonalized surprises (`MPS_ORTH`) were obtained from the Federal Reserve Bank of San Francisco’s replication files and the Jarocinski-Karadi monthly pure policy shocks (`MP_median`) from the authors’ public repository. The two-year Treasury yield (`GS2`), the ten-year yield (`GS10`) and the effective federal funds rate (`FEDFUNDS`) are monthly averages from FRED. Young-firm employment shares are from the Business Dynamics Statistics API for 2019 at the four-digit NAICS level, computed as employment at firms aged zero to five years over total employment, summed over the AII codes of each panel industry; 190 industries have a value. The tech core used in robustness comprises software publishers (5132), computer systems design (5415), computing infrastructure and data processing (518), and web search portals, streaming and social networks (519 and 5162). The research-intensive group comprises pharmaceuticals (3254), computer and electronic products (3341, 3342, 3344, 3345), motor vehicles and parts (3361, 3362, 3363), aerospace (3364), software publishers, computer systems design, scientific research and development services (5417) and computing infrastructure (518); it covers 13 panel industries.

C. Additional Results

C.1 Leave-One-Sector-Out Estimates

[Table 14](#) reports the coefficient on AI exposure in the baseline pooled specification of [Table 2](#), column (3), when each of the fourteen CES sectors is dropped in turn, for the tightening-and-hold and easing regimes.

Table 14: Leave-One-Sector-Out: the AI Exposure Coefficient

Sector dropped	Tightening and hold	Easing
None (baseline)	0.746	1.108
Mining and logging	0.843	1.189
Construction	0.495	0.803
Durable goods	0.791	1.114
Nondurable goods	0.421	0.782
Wholesale trade	0.734	1.149
Retail trade	0.834	1.147
Transportation and warehousing	0.267	0.174
Utilities	0.755	1.125
Information	1.114	1.867
Financial activities	1.075	1.479
Professional and business services	0.714	1.327
Education and health services	0.802	0.920
Leisure and hospitality	0.727	1.068
Other services	0.791	1.162
Range	[0.267, 1.114]	[0.174, 1.867]

Notes: Each row re-estimates the coefficient on standardized AI exposure (log points $\times 100$ per standard deviation, relative to the reference year March 2021 to February 2022) in the baseline specification with rate sensitivity and pandemic displacement as controls, dropping all industries in the named CES sector. Standard errors are of the same order as in the baseline (0.56 for the hold regime and 0.83 for easing) and no estimate is significant at five percent.

The AI coefficient remains positive in every case. The smallest easing-regime estimate, 0.17, arises when transportation and warehousing is dropped: the low-exposure warehousing, delivery and trucking industries, which overshot strongly during the pandemic and then contracted, contribute to the positive AI coefficient in the full sample through their low exposure and poor post-2022 performance, and the displacement control does not absorb all of it. The largest, 1.87, arises when information is dropped: the information sector contains both the most exposed software and data industries and the poorly performing publishing, broadcasting and telecommunications industries, and removing it makes the remaining exposed industries, mostly in finance and professional services, look better relative to the rest. Neither extreme is distinguishable from zero.