

DNB Working Paper

No 871/ October 2026

Do Markets Dream of Artificial Intelligence? Macroeconomic Effects and International Spillovers of AI News Shocks

Maria Sole Pagliari

DeNederlandscheBank

EUROSYSTEEM

Do Markets Dream of Artificial Intelligence? Macroeconomic Effects and
International Spillovers of AI News Shocks
Maria Sole Pagliari*

* Views expressed are those of the author and do not necessarily reflect official positions of
De Nederlandsche Bank.

Working Paper No. 871

October 2026

De Nederlandsche Bank NV
P.O. Box 98
1000 AB AMSTERDAM
The Netherlands

Do Markets Dream of Artificial Intelligence?

Macroeconomic Effects and International Spillovers of AI News Shocks*

Maria Sole Pagliari[†]

Abstract

This paper identifies AI news shocks from financial markets and traces their macroeconomic effects in the United States and the euro area. We combine a daily equity factor based on differences in firms' AI exposures with sign, magnitude, and narrative restrictions in a Bayesian VAR to isolate revisions in AI's expected economic value. Monthly local projections reveal an investment-led adjustment in the US: capital-goods orders rise early and persistently, while aggregate activity responds more gradually and with greater uncertainty. Producer prices increase before consumer prices, and the expansion spreads from sectors closely connected to the AI production network to the broader economy. In the euro area, equity prices rise and the euro depreciates; import and energy costs feed into inflation before real activity strengthens. Inflation expectations and nominal rates also increase. The results highlight a monetary policy dilemma: financial spillovers and imported price pressures arrive before the broader productive gains from AI adoption.

Keywords: AI news shocks, macroeconomic dynamics, international spillovers, production networks

JEL codes: C32, E32, E44, F41, O33

*The author would like to thank Niccolò Battistini, Claus Brand, Kai Christoffel, Giacomo Mangiante, Falk Mazelis, Davide Porcellacchia, Annukka Ristiniemi and seminar participants at the ECB for useful comments and suggestions. The author also gratefully acknowledges excellent data support from Viola Maria Billerbeck, Sofia Caputo and Chiara Scenini. The views expressed in this paper are those of the author only and should not be attributed to the European Central Bank, De Nederlandsche Bank or the ESCB.

[†]European Central Bank, De Nederlandsche Bank & University of Amsterdam, e-mail: maria_sole.pagliari@ecb.europa.eu

1 Introduction

“The future is already here – it’s just not evenly distributed.”

— William Gibson

AI is increasingly viewed as a potentially transformative general-purpose technology, with implications for productivity, investment, labour markets, inflation, and the conduct of macroeconomic policy. Yet its aggregate effects remain difficult to quantify. A growing microeconomic literature documents significant, though heterogeneous, effects of AI adoption on worker productivity, firm growth, and innovation ([Brynjolfsson et al., 2025](#); [Babina et al., 2024](#)), while the magnitude and timing of the associated aggregate productivity gains remain debated ([Acemoglu, 2025](#)). This distinction is important for macroeconomic analysis: advances in the technological frontier do not necessarily translate immediately into measured productivity, as their effects may depend on adoption, complementary investment, organisational change, and the reallocation of resources across firms and sectors.

This paper approaches the measurement problem through financial markets. We use daily relative equity valuations to identify changes in expectations about AI’s economic value and trace their subsequent domestic and international macroeconomic effects.

Recent theoretical work has begun to investigate these propagation mechanisms. [Aldasoro et al. \(2024\)](#) show in a multi-sector model that AI-driven productivity gains raise output, consumption, and investment, while their inflationary implications depend on the interaction between supply and demand forces. [Cerutti et al. \(2025\)](#) further emphasise that the aggregate consequences of AI can differ across countries, reflecting differences in exposure, adoption, and economic structure. From a monetary policy perspective, [Hartmann and Maver \(2025\)](#) distinguish between indirect effects operating through changes in macroeconomic conditions and direct effects on the transmission of monetary policy, highlighting the roles of productivity, investment, labour markets, inflation, and financial intermediation. Empirically, however, identifying an AI-specific macroeconomic shock remains challenging. Technological progress is gradual, its effects are difficult to separate from broader innovation trends, and the information relevant for economic decisions may become available well before it materialises in conventional measures of productivity or technology adoption.

Existing empirical work has addressed this measurement challenge primarily through patent data. At the firm level, [David et al. \(2026\)](#) link AI patenting to higher market valuation and subsequent increases in firms’ capital, revenues, employment, and total factor productivity, while [Yang \(2022\)](#) identifies AI-related innovation from patent text and documents positive effects on productivity and employment. Patent information has also been used to measure the exposure of occupations and tasks to advances in AI ([Septiandri et al., 2024](#); [Webb, 2020](#)), and more recent work develops increasingly refined classifications of AI patents to study their organisation, diffusion, and economic value ([Fang et al., 2026](#)). These approaches provide a natural measure of realised technological invention and are particularly well suited to studying where AI innovation originates and how it diffuses across firms and workers.

The closest empirical contribution to our paper is [Gazzani and Natoli \(2026\)](#), who exploit changes in the AI content of US patents to identify variation in the AI intensity of innovation and document a delayed expansion in economic activity and productivity. Our paper takes a complementary approach by identifying AI news directly from financial markets. Patent-based measures capture realised invention, while asset prices capture revisions in the *expected economic value* of AI. The daily frequency of market data also provides sharper timing, allowing us to combine valuation movements with precisely dated announcements and narrative information. This distinction is particularly relevant for the recent generative AI episode, in which market expectations about adoption, commercialisation, regulation, and profitability can change well before their aggregate effects become observable.

This approach connects the empirical literature on AI to the macroeconomic literature on technology news. In that literature, forward-looking variables reveal changes in expectations about future productive opportunities before those changes appear in measured productivity or output ([Barsky and Sims, 2011](#); [Jaimovich and Rebelo, 2009](#); [Beaudry and Portier, 2006](#)). In the context of AI, [Christensen and Rudebusch \(2026\)](#) provide related evidence from bond markets, examining US Treasury yields and estimates of the natural interest rate around major generative AI model releases.

Our analysis focuses instead on relative equity valuations, exploiting cross-sectional differences in firms’ AI exposure to identify news about capabilities, adoption, complementary investment, regulation, and commercial value. We construct a daily equity factor that

captures the relative valuation performance of firms that stand to benefit more from advances in AI, building on the insight of [Eisfeldt et al. \(2023\)](#), who show that firms' valuations respond differentially to their exposure to generative AI. We combine these cross-sectional valuation movements with sign, magnitude, and narrative restrictions around precisely dated AI-related events to isolate news that is specific to AI rather than generic equity-market, technology-sector, or semiconductor shocks. The resulting disturbance is a high-frequency revision in the expected economic value of AI.

The paper's methodological contribution is to construct and identify a daily, financial-market-based measure of AI news, linking revisions in relative equity valuations to subsequent macroeconomic outcomes. Using this measure, we make three substantive contributions. First, we document its domestic macroeconomic transmission in the US. The clearest and earliest real-side response is a persistent increase in capital-goods orders, whereas aggregate output responds more modestly and with greater uncertainty. Capacity utilisation temporarily declines before recovering, employment expands at medium horizons, and average hours per worker fall. Price pressures build gradually: producer prices respond before headline and core consumer prices. This sequencing is consistent with AI valuation news stimulating investment and demand for complementary resources before any broader expansion in productive activity becomes clearly visible. It also distinguishes the shock from a conventional, immediately cost-reducing productivity innovation. The shock instead combines expected future productivity gains with near-term investment and resource demand.

Second, we show that these aggregate responses conceal substantial sectoral heterogeneity. Industries with greater upstream or downstream exposure to the AI production network initially record stronger relative production and more muted producer-price increases. Over time, this relative production advantage fades and the expansion rotates towards less-exposed industries, where price pressures are also larger. Employment rises particularly strongly among highly exposed upstream suppliers, consistent with demand for inputs and complementary capacity. The sectoral evidence therefore characterises the shock as a hybrid technology-adoption disturbance: favourable supply effects are most visible close to the AI ecosystem, while investment and resource-demand effects spread more broadly through the production network.

Third, we estimate sizeable spillovers from US-originated AI news shocks to the euro area. Import, energy, and headline consumer prices respond relatively early, while core inflation increases more gradually and real activity responds only with a delay. Euro area equity prices rise and the euro depreciates against the dollar; short-horizon inflation expectations and nominal rates also increase. The EA–US interest rate differential displays a phased adjustment, initially moving in favour of the US before reversing as euro area inflation becomes more persistent. These results show that AI-related news is transmitted internationally through financial prices and external costs well before its effects on euro area real activity are fully apparent. They point to a two-speed mechanism in which the global financial cycle carries the news rapidly across borders, while productive gains spread more slowly as the technology is adopted and integrated by domestic firms.

Viewed together, the three sets of results describe a common transition from expectations to adoption. Within the US, market news is followed by capital spending and demand for complementary inputs before its effects spread to aggregate production, employment, and consumer prices. Across industries, the initial supply-side gains close to the AI ecosystem coexist with broader resource pressures transmitted through production networks. Across borders, financial prices and external costs adjust more rapidly than productive capacity. The central empirical finding is therefore the gap between the speed at which the economic value of AI is repriced and the time required for that value to be realised throughout the economy.

The remainder of the paper is structured as follows. [Section 2](#) discusses the construction and validation of the AI news factor; [Section 3](#) presents the identification of the structural AI news shock; [Section 4](#) estimates its domestic, sectoral, and international macroeconomic effects; and [Section 5](#) concludes.

2 The AI news factor

In this Section, we construct an equity-based AI news factor that captures changes in investors’ expectations about the future economic relevance of AI. Market prices respond immediately to information about technological capabilities, adoption, investment, competitive advantages, and regulatory constraints. This forward-looking dimension is especially useful in the recent

AI episode, where expectations can adjust well before the broader economic effects become visible. We identify a universe of companies relevant to AI and categorise them into a *long* portfolio, consisting of firms at the forefront of AI semiconductors, big-tech platforms, and software, and a *short* portfolio comprising legacy technology companies and non-AI peers.¹

The AI news factor is the difference between the returns on the long and short portfolios. The *long portfolio* consists of companies directly involved in AI development or positioned to benefit from its adoption. These include firms producing AI semiconductors and compute hardware (e.g., NVIDIA, AMD, and ASML), big-tech platforms investing heavily in AI capabilities (e.g., Microsoft, Alphabet, Amazon, and Meta), and software or enterprise companies providing AI-related solutions (e.g., Palantir, Salesforce, and Adobe).

The *short portfolio* comprises legacy technology companies and non-AI peers. It includes semiconductor firms with more limited AI exposure (e.g., Intel and Texas Instruments), traditional IT service providers (e.g., Accenture and Hewlett Packard Enterprise), and established enterprise technology firms (e.g., IBM and Cisco). This portfolio provides a benchmark against which AI-specific relative returns can be measured. Financial data for both portfolios, including adjusted closing prices, cover January 2015 to July 2026.

The long-short construction captures the relative performance of companies driving or benefiting from AI development and adoption. The portfolios use *time-varying market-capitalisation weights*, allowing changes in firms' market values and capital structures to enter the factor composition.²

To construct the time-varying weights, we obtain historical shares outstanding from SEC filings, adjust them for stock splits, and interpolate them to the daily frequency. This procedure incorporates stock splits, buybacks, and share issuances into daily market capitalisations. We rebalance the long and short portfolio weights quarterly.

Figure 1 examines the factor's ability to capture AI-related repricing rather than broad equity-market movements. Panel a shows the daily series alongside major AI-related events, while Panel b uses one-minute data around OpenAI's GPT-4o presentation on 13 May 2024 to examine the timing of relative valuation changes.

¹Section A of the Appendix provides additional details on the construction of the index, including the list of constituents.

²Section A.1.1 in the Appendix provides further details and descriptive statistics.

OpenAI had scheduled its Spring Update for 10:00 PT (13:00 ET), during US trading hours, without fully disclosing its content. The combination of a known release time and potentially unanticipated information about model capabilities and commercial availability provides a useful event window, paralleling the timing logic of high-frequency monetary-policy studies (Gürkaynak et al., 2005; Kuttner, 2001).

The AI portfolio initially underperforms during the product, access and API announcements, with NVIDIA and Microsoft accounting for the largest negative contributions. The differential begins to recover during the voice demonstrations, with positive contributions becoming more widespread during the vision and mathematics demonstrations, including Alphabet, Microsoft, Meta, Amazon and NVIDIA.³ By approximately 30 minutes after the event begins, the differential reaches around 20 basis points, while the S&P 500 remains below its pre-announcement level. The alignment of relative repricing with the presentation’s progression supports the interpretation that the factor captures market responses to AI-related information, complementing the daily-event evidence.⁴


To assess whether the factor also moves in the anticipated direction across a broader set of announcements, Table 1 compares expected signs with realised event-day returns for 31 AI-related events. We distinguish three categories: i) technological breakthroughs, i.e. moments when new AI capabilities become publicly known, such as the launch of large language models or the announcement of new chip architectures; ii) adoption and commercialisation signals (hereafter adoption events), including major corporate investment announcements and earnings releases from key AI infrastructure companies; iii) regulatory and industrial policy interventions, ranging from national bans on AI applications to executive orders and export restriction decisions. For each event, we assign an ex ante directional expectation: positive when the event is expected to raise the relative value of AI-intensive firms, negative when it signals adverse competitive or regulatory effects, and ambiguous when neither direction can

³Figure B.1 in the Appendix decomposes the intraday differential into stock-level contributions over four non-overlapping presentation phases: announcement, access and API economics ([0, 9) minutes), voice demonstrations ([9, 14)), vision and mathematics demonstrations ([14, 18)), and coding, translation and closing ([18, 26)). The phase markers are also shown in Panel b.

⁴In Section A.1.4, we compare the proposed factor with an independent, market-priced measure based on the first principal component of five AI-themed ETFs. We also examine alternative weighting, semiconductor-cycle exposure, stock selection, and the construction and coverage of the high-frequency GPT-4o exercise. Conclusions are unchanged.

be assigned with confidence. The realised factor return matches the expected direction for all 24 signed events; the remaining seven have ambiguous expected effects and are excluded from this comparison. This descriptive evidence motivates the event study below, which examines the factor’s response jointly with other financial variables while controlling for concurrent macro-financial conditions.

Figure 1



Sources: Yahoo Finance, Dukascopy, author’s computations.

Notes: Panel **a** reports the daily market-capitalisation-weighted AI factor; innovation, adoption and commercialisation, and policy events (vertical dashed lines) are listed in [Table 1](#). Panel **b** reports cumulative one-minute returns over the 60 minutes before and after the start of OpenAI’s Spring Update on 13 May 2024 at 13:00 ET. Intraday returns are normalised to zero immediately before the event. The black vertical dashed line marks the event start; grey vertical dotted lines mark the approximate starts of voice, vision, coding and translation demonstrations at minutes 9, 14, 18 and 22, respectively. “Differential” denotes the return on the AI portfolio minus that on the synthetic non-AI portfolio. The S&P 500 is shown only as a market benchmark and does not enter the construction of the differential. Details on the intraday data, portfolio coverage, and construction are reported in [Section B.1](#).

2.1 Event study

Following the literature⁵, we provide additional evidence that the AI factor captures AI-related surprises through an event study of the dates reported in [Table 1](#). Exploiting the ex ante sign classification of these events, we study the joint behaviour of the AI factor and a set of daily financial variables: broad equity returns (S&P 500), implied volatility (VIX), the US 2-year Treasury yield, the slope of the yield curve (the US 10-year minus 2-year yield spread), 2-year

⁵See [Gürkaynak et al. \(2005\)](#); [Kuttner \(2001\)](#); [MacKinlay \(1997\)](#)

Table 1: AI-related events, expected signs, and sign validation

Date	Event	Expected Sign	Validation	AI factor return
<i>Innovation Events (Technological Breakthroughs)</i>				
30 Nov 2022	OpenAI ChatGPT (GPT-3.5)	+	✓	+2.03%
14 Mar 2023	OpenAI GPT-4 release	+	✓	+1.84%
18 Jul 2023	Meta Llama 2	+	✓	+0.48%
06 Dec 2023	Google Gemini	−	✓	−0.47%
18 Mar 2024	NVIDIA Blackwell architecture	+	✓	+1.48%
18 Apr 2024	Meta Llama 3	−	✓	−0.30%
13 May 2024	OpenAI GPT-4o release	+	✓	+0.15%
27 Jan 2025	DeepSeek R1 release/mkt reass.	−	✓	−5.92%
24 Feb 2025	Anthropic Claude 3.7 Sonnet	?	−	−1.24%
04 Apr 2025	Meta Llama 4	?	−	+1.61%
22 May 2025	Anthropic Claude 4 (Opus and Sonnet)	+	✓	+1.36%
10 Jun 2025	OpenAI o3-pro release	−	✓	−0.51%
<i>Adoption and Commercialisation Events</i>				
23 Jan 2023	Microsoft \$10bn OpenAI investment	+	✓	+0.65%
24 May 2023	NVIDIA Q1 FY24 earnings	+	✓	+0.91%
23 Aug 2023	NVIDIA Q2 FY24 earnings	+	✓	+0.17%
25 Sep 2023	Amazon \$4bn Anthropic investment	+	✓	+0.64%
22 Feb 2024	NVIDIA FY24 Q4 earnings	+	✓	+3.22%
23 May 2024	NVIDIA Q1 FY25 earnings	+	✓	+2.74%
10 Jun 2024	Apple Intelligence announced	+	✓	+0.32%
29 Aug 2024	NVIDIA Q2 FY25 earnings	?	−	−2.11%
21 Jan 2025	Stargate Project announced	+	✓	+0.34%
27 Feb 2025	NVIDIA FY25 Q4 earnings	?	−	−3.17%
29 May 2025	NVIDIA Q1 FY26 earnings	+	✓	+0.78%
28 Aug 2025	NVIDIA Q2 FY26 earnings	+	✓	+0.12%
<i>Policy Events (Regulatory and Industrial Policy)</i>				
28 Mar 2023	Italy bans ChatGPT	−	✓	−0.76%
30 Oct 2023	Biden Executive Order on AI	?	−	+1.69%
08 Dec 2023	EU AI Act political agreement	?	−	+0.05%
23 Jul 2025	US AI Action Plan	+	✓	+2.88%
23 Jul 2025	Trump AI Executive Orders	+	✓	+2.88%
30 Jul 2025	US restrictions on advanced AI chip exports	?	−	+1.16%
12 Jun 2026	US export restrictions on Anthropic frontier models	−	✓	−2.13%
# of events				31
# of ambiguous events				7
# of signed events				24
# of matched signed events				24 (100%)

Notes: The table lists the AI-related events used in the analysis, grouped into innovation, adoption and commercialisation, and policy events. “Expected sign” denotes the ex ante anticipated direction of the market-cap-weighted AI factor (+ positive, − negative, ? ambiguous). For signed events, ✓ indicates that the realised event-day AI factor return matches the expected sign. The “AI factor return” column reports the realised market-cap-weighted AI factor return on the event day. For events with an ambiguous prior, the validation column is marked with − and only the empirical AI factor return is reported. Event returns are measured on the event date when it is a trading day and otherwise on the nearest trading day. Ambiguous events are excluded from the hit-rate calculation.

and 10-year inflation-linked swap (ILS) rates, and the broad US dollar nominal effective exchange rate (USDNEER). We estimate separate regressions for innovation, adoption and policy events, controlling for concurrent macro-financial conditions through the US economic growth surprise index (EGS).⁶ We also control for the US monetary policy surprise measure of [Jarociński and Karadi \(2020\)](#).

Let $y_{j,t}$ denote variable j on day t , where the event-study vector includes the AI equity factor, broad equity returns, changes in the VIX, changes in the 2-year yield, changes in the slope of the yield curve, changes in the 2-year and 10-year ILS rates, and changes in the USDNEER. All dependent variables are standardised before estimation. For each event category $c \in \{\text{innovation, adoption, policy}\}$, we define a signed event dummy

$$D_t^c = \sum_{e \in \mathcal{E}_c} s_e \mathbf{1}\{t = \tau_e\},$$

where $\mathcal{E}_c$ is the set of events in category c , τ_e is the corresponding trading day, and $s_e \in \{+1, -1, 0\}$ is the ex ante expected sign of the event. Events with ambiguous sign, $s_e = 0$, are retained in the event classification but do not contribute to the signed dummy. Positive events therefore enter with +1, negative events with -1, and ambiguous events with zero weight. For each variable j and event category c , we estimate

$$\tilde{y}_{j,t} = \alpha_j^c + \beta_{j,\text{EGS}}^c \widetilde{egs}_t + \beta_{j,\text{MP}}^c mp_t + \gamma_j^c D_t^c + u_{j,t}^c, \quad (1)$$

where $\tilde{y}_{j,t}$ denotes the standardised dependent variable, $\widetilde{egs}_t$ is the standardised change in the US economic growth surprise index, and mp_t denotes monetary policy surprises, demeaned over non-zero observations. The coefficient γ_j^c is the signed average abnormal response (SAAR). A positive value of γ_j^c indicates that variable j moves, on average, in the same direction as the event sign, while a negative coefficient indicates an average response in the

⁶US surprise indices are computed by Haver Analytics and are normalised, equally weighted standard deviations of data surprises based on Action Economics median forecasts versus their respective first-release indicator value or growth rate. A positive reading of the surprise index suggests that economic releases are beating the consensus. As a robustness check, we augment the baseline event-study specification with Haver's inflation surprise measure. The estimated responses are virtually unchanged, indicating that the results are not driven by omitted contemporaneous macroeconomic news. In additional checks, we substitute the EGS with Citigroup Economic Surprise Indices. The results are unchanged.

opposite direction. [Table 2](#) reports the estimates for the three event categories.⁷

Table 2: Event study results

Variable	Adoption and			
	Innovation	commercialisation	Policy	Sign Restriction
AI factor	+1.3392***	+0.8992***	+1.7969***	+
EQ returns	+0.7718***	+0.4510*	+0.2008	+
ΔVIX	-0.4544**	-0.0570	+0.2144	-
$\Delta 2Y$	+0.1560	+0.2910	-0.4554	.
$\Delta Slope$	+0.1260	-0.3538	+0.6577	.
$\Delta ILS2Y$	+0.3349*	-0.2883**	-0.3934	.
$\Delta ILS10Y$	+0.5032	-0.4695**	-0.9743***	-
$\Delta NEER$	-0.0406	-0.3168	+0.9414*	+ (additional)

Notes: Coefficients are in standardised units. *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$, based on Newey–West HAC standard errors, with bandwidth $L = [4(T/100)^{2/9}]$ ([Andrews, 1991](#)). This allows for heteroskedasticity and serial correlation in the daily regression residuals. All regressions control for the Haver US economic growth surprise index and the monetary policy surprise measure of [Jarociński and Karadi \(2020\)](#). The last column reports the sign restrictions suggested by the event-study evidence for the daily BVAR. The 2-year ILS response is left unrestricted because statistically significant innovation and adoption responses have opposite signs. By contrast, the 10-year ILS response is negative and statistically significant for both adoption and policy events, motivating a negative restriction on long-term inflation compensation. The positive $\Delta NEER$ restriction is supported by the policy event study and is used as an additional identifying dimension.

The AI equity factor responds positively and significantly across all three types of events. The strongest response is observed on policy event days, where the SAAR reaches 1.80 standard deviations, followed by innovation events at 1.34 and adoption events at 0.90. The uniform positive and highly significant response across categories provides direct validation of the factor as a high-frequency measure of AI-related news and strongly supports a positive sign restriction on the AI factor in the structural BVAR.

Broad equity returns also respond positively to AI events. The response is large and statistically significant for innovation events (0.77 standard deviations) and positive and marginally significant for adoption events (0.45), whereas the response to policy events is smaller and statistically insignificant (0.20). Importantly, the AI factor responds more strongly than the aggregate equity market in every category. The ratio of the AI-factor response to the broad-equity response is approximately 1.7 for innovation events and 2.0 for adoption events.⁸ This relative-response pattern is useful for identification: AI-related announcements generate a disproportionately large repricing of AI-exposed firms relative to the broad market, helping distinguish AI-specific news from generic risk-on or aggregate

⁷[Table B.5](#) in the Appendix reports the full set of results for each event in the sample. Policy events that occur on the same trading day are represented by a single event dummy in the regression.

⁸For policy events the ratio is approximately 9, although this partly reflects the small and statistically insignificant response of broad equities.

growth shocks. This interpretation is consistent with the literature on technology and news shocks, in which information about future productivity is capitalised immediately in asset prices (Barsky and Sims, 2011; Alexopoulos, 2011; Jaimovich and Rebelo, 2009; Beaudry and Portier, 2006).

The response of implied volatility provides an additional dimension of identification. For innovation events, the VIX falls by 0.45 standard deviations and the response is statistically significant at the 5% level. This is consistent with major capability breakthroughs being associated with higher equity valuations and lower aggregate market volatility. The VIX response is instead close to zero for adoption events and becomes positive, though statistically insignificant, for policy events. Cross-category Wald tests confirm that the VIX response is heterogeneous across event types: the difference between adoption and policy responses is statistically significant ($p = 0.006$), and the joint null of equal responses across the three categories is rejected ($p = 0.008$). We therefore retain a negative VIX restriction as an innovation-driven identifying restriction, while not interpreting it as a uniform feature of all AI-related announcements. This pattern is consistent with the broader literature linking uncertainty shocks, risk premia and financial volatility (Bekaert et al., 2013; Bloom, 2009).

The responses of Treasury yields and the yield-curve slope are considerably less uniform. The 2-year yield rises modestly after innovation and adoption events, but neither response is statistically significant, while it declines after policy events. The policy response differs significantly from both innovation and adoption responses, with pairwise Wald-test p -values of 0.001 and 0.002, respectively, and the joint equality test strongly rejects homogeneous responses across categories ($p < 0.001$). The slope of the yield curve also changes sign across event types and is never statistically significant in the signed specification. These results provide little support for imposing a contemporaneous restriction on either Treasury-rate variable.

Related evidence is provided by Christensen and Rudebusch (2026), who document declines in five- and ten-year US Treasury yields around major AI model releases. While our innovation events do not generate a significant decline in the two-year yield, they likewise provide no evidence of a significant increase. Despite differences in maturities and event windows, both studies caution against assuming that favourable AI news necessarily raises

US yields.

Inflation compensation provides a more informative additional dimension. The 2-year ILS response is positive after innovation events (0.33 standard deviations, or about 2.0 basis points) and negative after adoption events (-0.29 standard deviations, or about -1.7 basis points), with both responses statistically significant at conventional levels. The policy response is also negative but imprecisely estimated. The conflicting signs of the significant innovation and adoption responses therefore provide no robust basis for imposing a restriction on short-term inflation compensation. We accordingly leave the 2-year ILS rate unrestricted and exclude it from the baseline structural BVAR.

The 10-year ILS rate displays a clearer pattern among the event categories that are informative about its response. The innovation coefficient is positive but statistically insignificant (0.50 standard deviations; $p = 0.133$). By contrast, the response is negative and statistically significant for adoption events (-0.47 standard deviations, about -1.4 basis points) and strongly negative for policy events (-0.97 standard deviations, about -2.9 basis points). The adoption and policy responses also differ significantly from one another in magnitude ($p = 0.042$), but share the same sign. Thus, across the event categories that yield statistically informative responses for long-term inflation compensation, AI news is associated with a decline in the 10-year ILS rate. In augmented specifications that include long-term inflation compensation, we use this evidence to impose a negative contemporaneous restriction on ΔILS10Y . The baseline six-variable BVAR excludes the ILS rates for parsimony. The restriction in the augmented system is based on the concordant and statistically significant evidence from the adoption and policy studies; the insignificant innovation response is treated as uninformative for the sign decision.

The USDNEER response remains heterogeneous across categories. It is close to zero for innovation events (-0.04 standard deviations), negative and statistically insignificant for adoption events (-0.32), and positive and marginally significant for policy events ($+0.94$). The latter result is consistent with supportive US AI-policy announcements being accompanied by a stronger dollar, but the sign is not uniform across event categories and the cross-category equality tests do not reject common responses. We nevertheless use the positive NEER response as an additional identifying dimension in the baseline BVAR, while recognising

that its direct event-study support is weaker than for the AI factor, equity returns, VIX and long-term inflation compensation. This mixed exchange-rate response is consistent with the broader evidence that exchange rates jointly reflect interest rate expectations, risk premia and global financial conditions ([Miranda-Agrippino and Rey, 2020](#); [Engel, 2016](#)).

As an additional validation, we assess whether the AI factor responds systematically to salient technology news unrelated to AI. [Section B.2.1](#) in the Appendix reports a placebo event study based on 26 non-AI technology events, including generic big-tech product announcements, digital innovations, technology disruptions, semiconductor-cycle news and broad technology-policy announcements. The signed average abnormal response of the AI factor on placebo dates is close to zero and statistically insignificant, whereas the corresponding pooled response on AI-event dates exceeds one standard deviation. Moreover, a Wald test strongly rejects equality of the two responses. Importantly, the result survives orthogonalising the factor with respect to semiconductor returns, with the AI–placebo difference remaining significant. These robustness checks further indicate that the factor captures AI-specific information rather than generic mega-cap or semiconductor-market news.

Overall, the event-study evidence supports a sparse identification scheme. The most robust contemporaneous responses are a positive AI-factor response, a positive broad-equity response for innovation and adoption events, and a negative VIX response concentrated in innovation events. The 2-year Treasury yield and the slope of the yield curve are left unrestricted, while the positive NEER response provides an additional identifying dimension. The baseline structural BVAR does not include the ILS rates because their signed responses are not uniform across event categories.⁹ Together with the narrative classification of event dates and the relative magnitude of the AI-factor response, these results provide the empirical basis for the identification strategy used below, which combines sign, magnitude and narrative restrictions to isolate AI news shocks within the structural BVAR ([Arias et al., 2018a](#); [Antolín-Díaz and Rubio-Ramírez, 2018](#); [Uhlig, 2005](#)).

⁹In unreported analyses, we also estimate seven- and eight-variable BVARs that include the 10-year ILS rate and both the 2- and 10-year ILS rates, respectively. The results are qualitatively similar to the baseline but are estimated with greater uncertainty. We therefore retain the six-variable specification.

3 Identifying AI news shocks from financial markets

Our identification strategy exploits both the forward-looking content and the daily frequency of financial-market data. Relative equity valuations provide information about changes in the expected economic value of AI, while their joint movements with broader financial variables help distinguish AI-specific news from risk-on and monetary policy shocks. Narrative restrictions connect this variation to dated AI-related announcements.

We estimate a *daily* Bayesian vector autoregression (BVAR) with the following endogenous variables:

$$y_t = [AI_t \ EQ_t \ \Delta VIX_t \ \Delta 2Y_t \ \Delta Slope_t \ \Delta NEER_t].$$

Here, AI_t denotes the AI equity long–short factor, EQ_t broad equity returns, ΔVIX_t the daily change in implied volatility, $\Delta 2Y_t$ the daily change in the US 2-year yield, $\Delta Slope_t$ the change in the 10-year minus 2-year yield spread, and $\Delta NEER_t$ the daily change in the US dollar nominal effective exchange rate. All variables are standardised before estimation. The sample runs from January 2015 to July 2026. The reduced-form BVAR is

$$y_t = c + \sum_{l=1}^p A_l y_{t-l} + u_t, \quad u_t \sim \mathcal{N}(0, \Sigma_u), \quad (2)$$

where the Bayesian information criterion selects $p = 1$.¹⁰ We use a conjugate normal-inverse-Wishart prior for the autoregressive coefficients and the variance–covariance matrix. The structural representation is

$$A_0 Y_t = c^* + \sum_{l=1}^p A_l^* Y_{t-l} + \varepsilon_t, \quad (3)$$

where A_0 contains the contemporaneous structural relationships. The structural shocks satisfy $u_t = A_0^{-1} \varepsilon_t$, with $\varepsilon_t \sim \mathcal{N}(0, I_N)$, and are recovered as

$$\varepsilon_t = A_0 u_t, \quad (A_0^{-1})(A_0^{-1})' = \Sigma_u,$$

Admissible impact matrices A_0 are selected using sign, magnitude, and narrative restrictions,

¹⁰Re-estimating the BVAR with two lags produces results very similar to the baseline.

following [Arias et al. \(2018b\)](#); [Antolín-Díaz and Rubio-Ramírez \(2018\)](#). We identify three shocks: i) an AI news shock; ii) a risk-on shock, capturing general risk-on movements in financial markets; iii) a monetary policy shock. [Table 3](#) reports the main identification scheme, consisting of fourteen sign restrictions and one magnitude restriction.

Table 3: Identification Restrictions

Shock	AI Factor	EQ Returns	ΔVIX	$\Delta 2Y$	$\Delta Slope$	$\Delta NEER$
AI news shock	++	+	−	·	·	+
Risk-on shock	+	+	−	+	·	−
Monetary policy shock	+	+	−	−	·	−

Notes: + denotes a positive impact, − a negative impact, and · an unrestricted impact. ++ denotes a positive response whose magnitude exceeds that of broad equity returns.

The AI news shock raises both the AI factor and broad equity returns, with a larger response of the former, lowers market volatility, and appreciates the dollar. A risk-on shock raises both equity measures, lowers volatility, raises the 2-year yield, and depreciates the dollar. A monetary easing shock raises equity returns, lowers volatility and the 2-year yield, and depreciates the dollar.

We conduct the analysis using $R = 10,000$ draws from the posterior distribution. To distinguish the AI news shock further from generic technology, growth, or financial-market disturbances, we complement the sign and magnitude restrictions with a set of narrative restrictions based on the event-study evidence discussed in [Section 2.1](#). The identification proceeds in three steps. First, for each posterior draw of the reduced-form BVAR, we construct candidate structural decompositions and retain those satisfying the contemporaneous sign and magnitude restrictions reported in [Table 3](#). Second, we impose a set of narrative restrictions on the historical AI news shock realisations. Specifically, we focus on the events reported in [Table 4](#), corresponding to dates on which the AI factor experienced particularly large movements relative to its historical distribution (exceeding 1.5 standard deviations), and require the identified AI news shock to have the direction implied by the corresponding event. This second-stage restriction exploits information that is external to the BVAR dynamics and helps sharpen the economic interpretation of the identified disturbance. Third, in the spirit of [Antolín-Díaz and Rubio-Ramírez \(2018\)](#), we use the event-study information to assign differential posterior weight to the structural decompositions that survive the sign, magnitude, and narrative restrictions. In particular, on the selected AI-event dates we evaluate the

extent to which the unexpected movement in the AI factor is attributable to the identified AI news shock rather than to the other structural disturbances in the BVAR. Structural decompositions in which the AI news shock accounts for a larger share of the AI-factor innovation, and dominates competing shocks more clearly on these dates, receive a higher posterior weight. Importantly, this step does not impose an additional hard dominance restriction: decompositions in which other shocks also contribute materially to the observed market movement remain part of the posterior, but receive relatively less weight. The resulting event-informed posterior therefore exploits the information content of the event chronology while retaining uncertainty about the structural decomposition of individual announcement days.¹¹ We implement this re-weighting through importance resampling and use the resulting posterior distribution of AI news shocks throughout the subsequent analysis.

Table 4: Narrative Restrictions

Date	Event	Sign
30-Nov-2022	OpenAI ChatGPT launch	+
23-Jan-2023	Microsoft \$10bn OpenAI investment	+
14-Mar-2023	OpenAI GPT-4 release	+
28-Mar-2023	Italy bans OpenAI ChatGPT	−
24-May-2023	NVIDIA Q1 FY24 earnings	+
18-Jul-2023	Meta Llama 2	+
22-Feb-2024	NVIDIA FY24 Q4 earnings	+
18-Mar-2024	NVIDIA Blackwell architecture	+
18-Apr-2024	Meta Llama 3	−
23-May-2024	NVIDIA Q1 FY25 earnings	+
10-Jun-2024	Apple Intelligence announced	+
27-Jan-2025	DeepSeek R1 release	−
23-Jul-2025	US AI Action Plan	+
12-Jun-2026	US Anthropic ban	−

Notes: '+' : shock is positive; '-' : shock is negative.

3.1 Results

Figure 2 reports the posterior impact responses of the endogenous variables to AI news, risk-on, and monetary policy shocks. Results are economically coherent and distinguish the AI news shock from the two competing financial disturbances.

An AI news shock that raises the AI factor by 10% produces i) an increase in equity

¹¹See Section C.1 in the Appendix for details.

returns by around 4%; ii) a decrease in market volatility by around 3%; iii) an appreciation of the USD broad exchange rate by around 1.4%. These results are consistent with the view that AI-specific news operates as a broad, aggregate optimism shock rather than a purely sectoral or idiosyncratic tech shock. The response of the market-wide equity return is in line with the idea of AI as a general-purpose technology whose productivity implications extend well beyond the small set of firms directly involved in its development (Brynjolfsson et al., 2019; Acemoglu and Restrepo, 2019a; Aghion et al., 2017). The accompanying decline in volatility further suggests that positive AI news shocks resolve, rather than amplify, aggregate uncertainty, complementing the uncertainty-shock literature typically focused on the reverse case of negative shocks raising volatility (Bloom, 2009). Finally, the appreciation of the dollar following a positive AI news shock is consistent with the idea that AI-driven optimism is disproportionately concentrated in US firms and therefore attracts capital flows into US equities and, in turn, into the dollar, in line with the broader evidence on the link between US asset returns, capital flows and the dollar exchange rate (Miranda-Agrippino and Rey, 2020; Gourinchas and Rey, 2007).


A risk-on shock that raises broad equity returns by 10% produces i) an increase in the AI factor of around 5%; ii) a decrease in market volatility by around 18%; iii) an increase in the US 2-year yield of around 52 basis points; iv) a depreciation of the USD broad exchange rate by around 3%. The risk-on shock generates a smaller response of the AI factor relative to the broad market than does the AI news shock. This contrast supports the identification of an AI-specific impulse despite the two shocks' likely co-movement in the raw data. The joint rise in equities and short-term yields, together with the decline in volatility, is characteristic of the risk-on and risk-off dynamics documented in the global financial-cycle and exchange-rate literature (Miranda-Agrippino and Rey, 2020; Ready et al., 2017). The dollar *depreciates* following the risk-on shock, in contrast to its appreciation after AI news. This pattern is consistent with the safe-haven dollar literature (Jiang et al., 2021; Engel, 2016): generic risk-on episodes shift capital towards riskier global assets, whereas the AI news shock contains a more specifically US component.

Finally, a monetary easing shock that lowers the 2-year yield by 10 basis points produces i) an increase in the AI factor of around 1%; ii) an increase in equity returns of around

1.7%; iii) a decrease in market volatility by around 3%; iv) a depreciation of the USD broad exchange rate by around 0.6%. The decline in volatility and the depreciation of the dollar are standard features of monetary policy transmission and are consistent with the extensive literature on high-frequency identified monetary policy shocks (Jarociński and Karadi, 2020; Gürkaynak et al., 2005; Kuttner, 2001). The response of equity returns is also consistent with the duration channel of monetary policy, as growth assets are typically more sensitive to changes in the risk-free rate (Gormsen and Lazarus, 2023).

Overall, the impact responses distinguish AI news most clearly through the joint behaviour of relative equity valuations and the dollar. AI-related assets also behave as long-duration, growth-sensitive assets following monetary easing, although their response relative to the broader market is less distinctive. The dollar appreciates after positive AI news but depreciates following risk-on and monetary easing shocks. This contrast points to a US-specific component of AI-related optimism and complements evidence on the global financial cycle (Miranda-Agrippino and Rey, 2020) and the safe-haven role of the dollar (Jiang et al., 2021).

Figure 2: Impact of shocks at horizon 0



Notes: Impact responses of the endogenous variables to scaled structural shocks. The AI news shock is scaled to raise the AI factor by 10%; the risk-on shock to raise broad equity returns by 10%; and the monetary easing shock to lower the US 2-year yield by 10 basis points, all on impact. Whiskers indicate 68% highest-density posterior sets.

Having distinguished the AI news shock from risk-on and monetary policy shocks, we next examine its daily history. For each day t , we obtain the posterior distribution of ε_t^{AI} across accepted rotations and use its cross-rotation median as the point estimate. Panel [a](#) plots this series together with the events subject to narrative sign restrictions. These include positive developments, such as major model releases, chip-architecture announcements, and large investment deals, as well as adverse developments, such as regulatory bans and export restrictions. The largest shock realisations align closely with the narrative dates, providing an additional check on the economic plausibility of the identification. As an additional diagnostic, we examine the cumulative identified AI shocks in a $[-5, +5]$ trading-day window around the events used for the narrative restrictions ([Table 4](#)). For each admissible structural draw, shocks are cumulated within each event window and then averaged separately across positive and negative events; the resulting distributions are summarized using the event-weighted posterior. Results are reported in [Figure 3](#). The two groups display a marked divergence around the event date, with cumulative shocks rising sharply for positive events and falling for negative events, while remaining clearly separated in the subsequent days. The shaded 68% highest-density posterior sets indicate that this pattern is robust across admissible structural rotations..^{[12](#)}

4 The macroeconomic impact of AI news shocks


We assess the macroeconomic implications of US-originated AI news shocks by first aggregating the daily structural shocks to the monthly frequency. The aggregation is performed separately for each structural draw m ,

$$\varepsilon_{q,m}^{AI} = \sum_{t \in q} \varepsilon_{t,m}^{AI}, \quad (4)$$

where q indexes calendar months and the sum runs over all trading days t in month q . Before aggregation, the daily shocks are expressed in the original units of the AI factor using its daily standard deviation. The posterior weights attached to the structural draws are preserved, and posterior summaries are computed only after aggregation. This procedure retains the

¹²Additional checks appear in the Appendix: variance and historical decompositions ([Figures C.2 and C.3](#)), predictability tests ([Section C.4](#)), and validation against the release dates of [Christensen and Rudebusch \(2026\)](#), conditional on the raw-factor salience screen ([Section C.2](#)).

Figure 3



Notes: In Panel a, shocks are standardised and vertical lines mark the event days reported in Table 4. Panel b reports cumulative identified AI shocks in a $[-5, +5]$ trading-day window around the events used for the narrative restrictions. For each admissible structural rotation, daily AI shocks are cumulated from $t - 5$ through each relative trading day and then averaged separately across positive and negative narrative events. Solid lines show the event-weighted posterior medians across admissible rotations, while shaded areas denote 68% highest-posterior-density sets. The vertical line marks the narrative-event date $t = 0$. Since the underlying data are measured at the daily close, the movement between $t - 1$ and t captures the contemporaneous market reaction to the event. Positive and negative events are classified according to the sign imposed in the narrative restrictions.

full distribution of monthly shock histories instead of replacing the structural posterior with a single point estimate.

For each outcome variable y and horizon $h = 0, 1, \dots, H$, we estimate a cumulative local projection following Jordà (2005):

$$y_{t+h} - y_{t-1} = \alpha_h + \beta_h \varepsilon_t^{AI} + \sum_{j=1}^p \gamma_{h,j} y_{t-j} + \sum_{j=1}^p \delta_{h,j} \varepsilon_{t-j}^{AI} + \theta'_h \mathbf{X}_t + u_{t+h}, \quad (5)$$

where β_h is the response of y at horizon h to an AI news shock in month t , $\mathbf{X}_t$ is a vector of controls,¹³ p is the lag length, and u_{t+h} is the projection residual. The sequence $\{\beta_h\}_{h=0}^H$ traces the *cumulative* response of y_{t+h} , relative to the pre-shock month, to a monthly AI news shock at time t .

¹³To account for the macroeconomic disruption associated with the COVID-19 pandemic, the baseline includes an indicator equal to one whenever the cumulative response window $[t - 1, t + h]$ overlaps March–June 2020 (Lenza and Primiceri, 2022).

To propagate the event-informed BVAR posterior through the local projections stage while retaining the existing simulation-based inference procedure, we construct an equally weighted Monte Carlo representation of Equation (C.28) through systematic importance resampling. BVAR shock histories are sampled with replacement according to their probabilities ω_m , and each resampled monthly history is treated as an equal-weight draw in the local projection estimation.¹⁴ We use systematic rather than independent multinomial resampling to reduce Monte Carlo noise. The normalisation to a one-standard-deviation AI news shock is computed directly from the weighted source distribution rather than from the resampled approximation.

Rather than estimating each horizon independently, we regularise the sequence of impulse-response coefficients across horizons using a hierarchical second-order random-walk prior. In particular, letting $\beta = (\beta_0, \dots, \beta_H)'$, the prior penalises changes in the slope of the impulse response,

$$\Delta^2 \beta_h = \beta_h - 2\beta_{h-1} + \beta_{h-2}, \quad (6)$$

while leaving the overall shape and sign of the response unrestricted. The strength of the smoothing is not fixed ex ante, but is governed by a variance parameter that is itself estimated hierarchically from the data. Hence, the procedure borrows information across neighbouring horizons while allowing for non-monotonic and hump-shaped responses.¹⁵

An additional complication is that $\varepsilon_t^{\text{AI}}$ is a *generated regressor*: it is inferred from the daily BVAR rather than directly observed. Conditioning the local projections on a single estimated shock series would therefore discard the first-stage uncertainty associated with the identification of the structural shock and would generally understate uncertainty in the impulse responses (Pagan, 1984). We address this problem by integrating the local projection posterior over the posterior distribution of the identified AI news shock.

More specifically, let $\{\varepsilon_t^{\text{AI},(r)}\}_{r=1}^R$ denote draws from the event-informed posterior distribution of the structural AI news shock. The daily draws are first aggregated to the monthly frequency draw by draw, preserving their posterior weights. We then use systematic impor-

¹⁴Repeating the local projections with equal weights on the original admissible BVAR shock histories, rather than event-informed weights, yields broadly similar median responses but wider posterior intervals that generally include zero at the reported credibility levels.

¹⁵Details of the Bayesian smooth local projection estimator and the hierarchical smoothing prior are provided in Section D.1.

tance resampling to obtain R shock histories distributed according to the event-weighted posterior. Conditional on each history r , we estimate the Bayesian smooth local projection and draw B realisations from the conditional posterior of the impulse response,

$$\beta_h^{(r,b)} \sim p\left(\beta_h \mid \varepsilon_{1:T}^{\text{AI},(r)}, \mathbf{Y}\right), \quad r = 1, \dots, R, \quad b = 1, \dots, B. \quad (7)$$

The reported posterior distribution is therefore the mixture

$$p(\beta_h \mid \mathbf{Y}) \simeq \frac{1}{R} \sum_{r=1}^R p\left(\beta_h \mid \varepsilon_{1:T}^{\text{AI},(r)}, \mathbf{Y}\right), \quad (8)$$

Following importance resampling, the histories entering the local projections carry equal weights, while their frequency reflects the original event-informed probabilities. Posterior medians and credible intervals are computed from the resulting pooled distribution

$$\left\{ \beta_h^{(r,b)} : r = 1, \dots, R, \quad b = 1, \dots, B \right\}. \quad (9)$$

This procedure propagates both layers of uncertainty: variation across r captures uncertainty about the structural shock generated in the first-stage BVAR, while variation across b conditional on a given r captures estimation uncertainty in the local projection. In this way, the generated AI news shock is not treated as an observed regressor, and inference incorporates rather than conditions away first-stage identification uncertainty.

In the baseline specification we set $H = 60$ and $p_y = p_\varepsilon = 12$. We use $R = 10,000$ event-weighted shock histories and draw $B = 50$ posterior realisations of the local projection coefficients for each history. The hierarchical smoothing parameter is estimated using the procedure described in [Section D.1](#) of the Appendix.

4.1 US economy

Aggregates. We begin with the effects of AI news shocks on the US economy. The outcomes include headline and core CPI, core PCE prices, the headline PPI, real GDP, industrial production, nonfarm employment, average weekly hours, hourly earnings, core capital-goods orders (CAPEX), and total capacity utilisation. [Figure 4](#) reports their monthly responses to

a one-standard-deviation shock over a five-year horizon.¹⁶

The response of prices points to a gradual build-up of inflationary pressures following the AI news shock, with producer prices responding earlier than consumer prices. PPI starts increasing within the first year and is already significantly positive between months 8 and 11. Following a brief loss of precision, the response becomes persistently positive from month 19 to month 44 and reaches a peak of approximately 0.84% around month 34. The relatively early response of producer prices is consistent with stronger demand for intermediate inputs and investment goods during the initial phase of AI adoption, as well as with input-cost pressures associated with the expansion of productive capacity (Cerra and Swanson, 2017; Carvalho and Gomes, 2017; Gordon, 2010).

Consumer-price pressures emerge more gradually. Headline CPI becomes significantly positive around month 29 and remains so through approximately month 45, reaching a peak of about 0.68% around month 42. Core CPI displays an even more delayed response, becoming significantly positive from around month 34 and peaking at approximately 0.53% around month 43. Core PCE inflation follows a closely related pattern, with the cumulative price response becoming significant around month 31 and reaching approximately 0.49% at its peak around month 42. The sequencing of the responses—from producer prices to headline and subsequently core consumer prices—is consistent with a gradual pass-through of upstream price pressures (Blanchard, 2010; Galí, 2008). The positive price response differs from the disinflationary effects found by Gazzani and Natoli (2026) for patent-based AI technology shocks. One possible interpretation is that the two exercises capture different types of information about AI. While patent-based shocks are more directly associated with technological improvements and therefore display the characteristics of a positive supply shock, our market-based news measure captures revisions in expectations about the adoption, commercialisation, and economic value of AI. Such news may stimulate investment and resource demand before the associated expansion in productive capacity is fully realised, generating upward price pressures during the transition.

Turning to real activity, the response is considerably more heterogeneous. Real GDP

¹⁶Section D.2.1 shortens the lag length to six, while Section D.2.2 removes smoothing across horizons. The results for both the US and the euro area are similar to the baseline.

increases modestly in the short to medium run, with a peak of about 0.22% around month 16 and a 68% credible interval excluding zero at months 15–16. Thereafter, the posterior median remains mostly positive, but the response is imprecisely estimated. Industrial production displays a similar pattern: after a small initial decline, it rises significantly around months 16–18, reaching approximately 0.18%, before temporarily weakening around months 22–24. The posterior median subsequently turns positive again, but the posterior sets remain relatively wide over most of the medium-run horizon. These results therefore provide some evidence of an expansion in aggregate activity, although considerably less precise than the response of investment.

Capacity utilisation exhibits a more pronounced temporary contraction. The response falls significantly below zero between months 21 and 24, reaching a trough of approximately -0.34% around month 23, before recovering rapidly and turning positive thereafter. While the posterior median reaches around 0.33% near month 49, the positive medium-run response is estimated with substantial uncertainty. The temporary decline in capacity utilisation is consistent with adjustment costs and reallocation frictions during the initial phase of technology adoption ([Acemoglu, 2025](#); [Caballero, 2019](#); [Fujita, 2014](#)). Together with the earlier rise in investment, the subsequent recovery suggests that firms initially undertake costly reorganisation and capacity-building before the productive effects of the new technology become fully visible.

The labour market displays a delayed but economically meaningful expansion. Non-farm employment rises significantly over several medium-run windows, first around months 17–19 and more persistently from roughly month 29 onwards, with particularly strong posterior support between months 38 and 44. The cumulative response peaks at approximately 0.44% around month 43. Hourly earnings also increase in the medium run and reach a peak of about 0.30% around months 42–43, although the response is estimated less precisely and the 68% credible interval excludes zero only around months 41–42. The positive responses of employment and earnings are qualitatively consistent with the labour-market effects documented by [Gazzani and Natoli \(2026\)](#), suggesting that the expansionary general-equilibrium effects associated with AI adoption can outweigh direct labour displacement effects over the medium run.

Average weekly hours move in the opposite direction. The response turns significantly negative from approximately month 37 to month 46 and again towards the end of the estimation horizon, with a cumulative decline of about -0.12% around months 42–44 and a somewhat larger decline at the longest horizons. Thus, the increase in employment is accompanied by a reduction in average hours per worker. This differs from the positive response of hours reported by [Gazzani and Natoli \(2026\)](#) and suggests that labour-market adjustment in our sample operates more through the extensive than the intensive margin. Firms expand headcount while reducing hours per worker, a pattern that may reflect changes in the organisation and composition of tasks as AI adoption progresses. More generally, the divergence between employment and hours indicates that aggregate measures of labour input may conceal substantial adjustment across margins ([Acemoglu, 2025](#); [Fujita, 2014](#)).

The most pronounced real-side response is found for investment. Core capital goods orders increase rapidly from the first year after the shock: the response becomes significantly positive around month 8, reaches approximately 0.73% by month 11, and continues to strengthen thereafter. CAPEX peaks at about 1.18% around month 29, remaining significantly positive through month 35. Following a temporary moderation around months 36–42, the response strengthens again and remains significantly positive from month 43 through the end of the horizon, although the second increase is smaller than the earlier peak, reaching approximately 1.03% around month 49. The strength and persistence of the investment response are consistent with investment-led propagation of news about productivity-enhancing technologies ([Acemoglu, 2025](#); [Mokyr, 2015](#); [Jaimovich and Rebelo, 2012](#)). In particular, financial markets may react immediately to news about future AI-related productivity, while firms subsequently translate these expectations into a gradual expansion of the capital stock.

Overall, the results point to a propagation mechanism centred on an early and persistent investment response rather than on an immediate increase in aggregate output. Following the AI news shock, firms rapidly expand capital expenditure and producer prices begin to rise, while aggregate activity responds only moderately and capacity utilisation temporarily declines, consistent with adjustment and reallocation costs. Over the subsequent two to four years, employment strengthens and producer-price pressures gradually pass through to headline and core consumer prices, while average hours per worker decline. The resulting pattern

is therefore consistent with a technology-news shock that initially stimulates investment and resource demand before the associated productive-capacity effects fully materialise.

This transmission mechanism shares with [Gazzani and Natoli \(2026\)](#) the idea that the macroeconomic effects associated with AI emerge only gradually, but differs importantly in the response of prices. Whereas their patent-based shocks resemble conventional positive technology shocks, the market-based news shock considered here generates positive producer- and consumer-price responses. This difference is consistent with our measure capturing not only expected future productivity gains but also revisions in expectations about the commercial value and diffusion of AI, which can trigger capital deepening, resource reallocation, and strong investment demand well before the corresponding increase in productive capacity is completed ([Acemoglu, 2025](#); [Carvalho and Gomes, 2017](#); [Gordon, 2010](#)). The evidence therefore suggests that the macroeconomic transmission of AI news is characterised by a relatively rapid investment response but a much more gradual adjustment of output, employment, and prices.


Sectoral effects. We next examine sectoral prices, production, and employment to assess whether the aggregate responses conceal heterogeneous adjustment across industries with different degrees of AI exposure.

To construct the sectoral aggregates, we map individual statistical series into a common set of BEA industries and assign each industry upstream and downstream AI-network exposure measures.¹⁷ [Figure 5](#) illustrates the production linkages underlying these measures. Upstream exposure is highest in industries such as software publishing, broadcasting and telecommunications, motion pictures and sound recording, and computer and electronic products. Downstream exposure is also high in software publishing and computer and electronic products, as well as in broadcasting and telecommunications, data processing, and internet services. AI-related activity is therefore embedded in a broad production network through which news shocks can propagate.

For each economic domain, we then split sectors into high- and low-exposure groups using the corresponding exposure distribution and construct fixed-composition aggregate


¹⁷[Section A.2.3](#) provides details.

Figure 4: Impact of AI news shocks on US aggregates



Notes: Cumulative impulse response functions (IRFs) to a one-standard-deviation AI news shock. Variables are expressed as percent deviations from the baseline ($t-1$). Darker (lighter) shaded areas represent 68% (90%) highest-density posterior sets.

Figure 5: Sectoral exposures to AI



Sources: BEA and author's computations.

Notes: Vertical dashed lines indicate the 75th percentile.

indices for each group. The same sectoral classification is kept throughout the sample, so changes in the aggregates reflect within-sector dynamics rather than time variation in group membership. We separately construct upstream and downstream aggregates, and for each dimension consider both the responses of the high- and low-exposure groups. This yields a direct comparison of AI-related shocks' effects between sectors that are more and less exposed to AI-related activity through the production network.

Figures 6 and 7 show that the aggregate two-stage pattern documented above has a clear sectoral counterpart. The differential responses in Figure 7 reveal a distinctive dynamic reallocation across the production network. During the first two years following a positive AI news shock, production tends to expand relatively more in sectors with greater AI-network exposure, both upstream and downstream. At the same time, these sectors experience systematically weaker producer-price pressures than their less exposed counterparts. As the adjustment proceeds, however, the relative production advantage of highly exposed industries disappears and eventually reverses, with sectors that are less directly connected to AI accounting for an increasing share of the expansion. The evidence therefore points to an initial concentration of the shock around the AI ecosystem followed by a gradual diffusion towards the broader economy.

For producer prices, [Figure 6](#) shows positive and persistent responses across both upstream and downstream groups, but with considerably larger increases among less-exposed industries. Upstream, producer prices peak at around 1.2% among highly exposed sectors, compared with about 1.6% among less-exposed sectors. The contrast is still sharper downstream, where the corresponding peaks are approximately 0.8% and 1.6%. The differential responses in [Figure 7](#) make this pattern particularly clear. The price differential is predominantly negative along both network dimensions, and especially persistent downstream, where it reaches roughly -0.8 percentage points at longer horizons. Thus, conditional on the positive aggregate response of producer prices, sectors more closely connected to AI activity experience *smaller* price increases than sectors further away from the AI network.

This relative-price pattern is informative about the nature of the transmission mechanism. The inflationary consequences of the AI news shock are not concentrated in the industries most directly exposed to AI-related activity. Instead, the industries closest to the shock combine comparatively stronger early quantity responses with comparatively weaker price pressures, a configuration consistent with relatively favourable supply or productivity effects within the directly exposed part of the production network. As the shock propagates outward, price pressures become more pronounced among less-exposed industries. The positive aggregate PPI response documented above therefore reflects not only pressures within AI-intensive sectors themselves, but importantly the demand for inputs, capacity and complementary resources generated elsewhere in the economy. Such a pattern is consistent with the production-network literature, in which sector-specific disturbances propagate beyond their point of origin through intermediate-input demand, complementarities, and general-equilibrium reallocation ([Baqae and Farhi, 2018](#); [Carvalho, 2014](#)). Existing network relationships may similarly shape the gradual diffusion of new inputs and technologies across industries ([Carvalho and Voigtländer, 2014](#)).

The production responses reinforce this interpretation and reveal a particularly clear change in the relative incidence of the shock over time. At short and intermediate horizons, industrial production performs relatively better among more-exposed sectors. The differential production responses in [Figure 7](#) become positive over the first two years and peak at around 0.4–0.5 percentage points, consistent with AI-related investment and adoption initially

stimulating activity relatively close to the AI production network. This relative advantage subsequently dissipates. After roughly two to three years the differential turns negative along both network dimensions and becomes particularly pronounced at longer horizons. Upstream, the production differential falls to about -0.5 percentage points around month 50, while a broadly similar reversal emerges downstream.

The timing of this reversal closely coincides with the late expansion in aggregate US industrial production. This suggests that the aggregate recovery is not primarily accounted for by a persistent expansion of the industries most closely linked to AI. Rather, once the adjustment reaches its second phase, production increasingly spreads towards sectors with weaker direct exposure. The network evidence therefore points to a form of sectoral catch-up: activity initially rises disproportionately close to the AI ecosystem, but the relative advantage of these industries subsequently declines as the effects of the shock are transmitted to the rest of the economy. The production network thus appears to transmit the initial impulse outward from AI-related activity, consistent with models in which local productivity or demand shocks are amplified and redistributed through input–output linkages and factor reallocation ([Baqae and Farhi, 2019](#); [Carvalho, 2014](#)).

Employment adds an important distinction between the upstream and downstream propagation of the shock. Among upstream suppliers, employment expands much more strongly in sectors with high AI exposure. The response peaks at around 0.9% among highly exposed upstream industries, approximately twice the peak response of the corresponding low-exposure group. Consistently, the upstream differential employment response in [Figure 7](#) is positive over much of the medium-run horizon and peaks at more than 0.5 percentage points around month 33. This suggests that industries supplying the AI ecosystem respond to the initial investment and implementation boom partly through an expansion of headcount. The result is particularly notable because it indicates that the upstream propagation of the AI news shock is labour-demanding rather than simply substituting capital for labour.


Downstream, the employment pattern is different. Both high- and low-exposure groups eventually expand, but employment increases more strongly among less-exposed sectors, with peak responses of around 0.3% and 0.5% , respectively. Accordingly, the employment differential becomes negative at longer horizons. The contrast between strong relative

employment creation among highly exposed upstream suppliers and the eventual predominance of less-exposed sectors downstream is consistent with the broader literature on automation and AI, which stresses that technological change can simultaneously generate displacement effects in directly affected tasks and capital-deepening and reinstatement effects that raise labour demand in complementary activities ([Acemoglu et al., 2022](#); [Acemoglu and Restrepo, 2019a](#)).

Overall, the sectoral evidence points to a propagation mechanism characterised by both *concentration* and subsequent *diffusion*. In its initial phase, the AI news shock disproportionately stimulates activity in sectors closely connected to the AI ecosystem. These industries expand production relative to less-exposed sectors while facing comparatively muted price pressures, and upstream suppliers also experience a sizeable relative increase in employment. At longer horizons, however, the incidence of the shock rotates towards the broader economy. Production increasingly shifts towards less-exposed industries, downstream employment follows a similar pattern, and price pressures are systematically larger outside the most AI-exposed sectors. The production network therefore operates not simply as an amplification mechanism, but as a diffusion mechanism through which an initially concentrated AI-related impulse progressively spreads across sectors and becomes a broader macroeconomic expansion.¹⁸


¹⁸As a robustness exercise, we reconstruct the aggregates using the 10th/90th percentiles of the empirical exposure distribution, rather than the 25th/75th percentiles used in the baseline, and re-estimate the local projections. The results, reported in [Section D.2.4](#), are qualitatively similar.

Figure 6: Impact of AI news shock on sectors



Notes: Cumulative impulse responses to a one-standard-deviation AI news shock, shown separately by downstream and upstream input–output exposure to AI-related sectors. Darker (lighter) shaded areas represent 68% (90%) highest-density posterior sets. Variables are expressed as percentage deviations from the baseline ($t - 1$).

Figure 7: *Differential* impact of AI news shock on sectors



Notes: Differential impulse responses to a one-standard-deviation AI news shock for industries with high AI-network exposure relative to industries with low exposure. Darker (lighter) shaded areas represent 68% (90%) highest-density posterior sets. Responses are expressed in percentage points.

AI news shocks versus conventional productivity shocks. The sectoral evidence also helps distinguish the identified AI news shock from a conventional aggregate productivity innovation. A favourable TFP shock is typically interpreted as an immediate improvement in productive efficiency and therefore operates primarily as a supply-side disturbance. By contrast, the responses documented here point to a more gradual adoption process. The AI news shock initially stimulates investment and generates heterogeneous adjustments across sectors linked to AI-intensive activity, while the broader expansion in production and employment materialises only gradually as the shock diffuses through the production network. Moreover, producer prices increase at the aggregate level rather than fall. This combination is difficult to reconcile with a purely cost-reducing productivity shock and instead suggests that AI news generates substantial demand for complementary capital, intermediate inputs, energy and production capacity before its full productivity benefits are realised.

At the same time, the differential sectoral responses reveal an important supply-side component within this broader adjustment. Sectors that are more directly exposed to the AI production network initially exhibit stronger relative production together with weaker relative price pressures. In this sense, the industries closest to AI activity display the quantity–price configuration that one would expect from a favourable technology or productivity effect, even though this signature does not characterise the aggregate economy. As the shock propagates, the initial relative production advantage fades and eventually reverses, whereas price pressures increasingly concentrate among less-exposed industries. The evidence therefore suggests that productivity gains and the resource costs of adoption operate simultaneously but at different points of the production network.

The production-network evidence reinforces this interpretation. Standard network models imply that supply-side productivity disturbances propagate predominantly downstream, whereas demand disturbances transmit upstream towards input suppliers ([Acemoglu et al., 2016](#)). In our estimates, AI news shocks generate economically relevant responses along both dimensions. The initially stronger production response and lower relative prices among highly exposed sectors are consistent with a local supply-side improvement, while the particularly strong employment response among upstream suppliers points to substantial demand for complementary inputs and implementation capacity. At longer horizons, production and

price effects spread more broadly towards less directly exposed industries. The resulting pattern is therefore consistent with a hybrid technology-adoption shock that combines an eventual productivity improvement with an initial investment and implementation impulse. In this respect, AI differs from a generic TFP innovation because the technology must first be embodied in new capital, reorganised production processes, and complementary tasks before its productivity effects are fully realised, generating potentially important displacement, capital-deepening, and reinstatement effects along the transition ([Acemoglu and Restrepo, 2019b,a](#)).

4.2 Spillovers to the Euro Area

We next examine how US-originated AI news shocks propagate to the euro area. [Figure 8](#) reports the responses of real and price variables, while [Figure 9](#) presents the corresponding financial responses, including two- and ten-year EA–US interest rate differentials.¹⁹ The estimates reveal a clear difference in timing: financial prices, external costs, and inflation respond before domestic activity. This sequencing is consistent with evidence that US technology and uncertainty shocks travel internationally through financial, trade, and expectations channels (e.g., [De Gasperi and Ricco, 2023](#); [Miranda-Agrippino and Rey, 2020](#); [Jordà et al., 2019](#)).

Real and price responses. The response of consumer prices is positive and persistent. Headline HICP increases relatively quickly, following a hump-shaped path that peaks at about 1.6% after 18 months and then gradually declines. Core HICP rises more slowly but persistently, reaching a peak of about 1.3% after 30 months. Producer prices also increase markedly. The ordering of these responses—producer and headline prices first, core inflation later—suggests that upstream cost pressures gradually pass through the pricing chain and become embedded in underlying inflation, as in the energy- and commodity-price spillover literature ([De Gasperi and Ricco, 2023](#); [Känzig, 2021](#)).

Import prices remain above baseline for a prolonged period, while energy and gas prices react especially strongly. Gas prices peak at about 24% after only ten months, making them among the fastest-moving variables in the exercise. This response may partly reflect the energy intensity of data centres and AI-related computing infrastructure, consistent with

¹⁹See [Section A.2.2](#) for the variables and data sources.

evidence linking expanding data-centre capacity to electricity and natural-gas demand (IEA, 2025). Together, the import- and energy-price responses help account for the relatively early rise in headline HICP. The subsequent persistence of core inflation indicates that the pressure does not remain confined to imported or energy-intensive components, echoing established terms-of-trade and imported-inflation channels (De Gasperi and Ricco, 2023; Forbes, 2019; Auer et al., 2019).

Real activity responds considerably later. Industrial production moves little at short horizons and expands more clearly only thereafter; hours worked follow a similarly delayed path, while the employment response is muted. The real spillover therefore emerges after asset prices, external costs, and inflation have already begun to adjust. This timing mirrors the gradual US response in Section 4.1 and is consistent with evidence that the real effects of foreign technology and financial shocks unfold slowly (De Gasperi and Ricco, 2023; Jordà et al., 2019).

Financial spillovers. Euro area equity prices increase and remain above baseline over much of the horizon, peaking at about 4.6% after 47 months. Thus, although the shock originates in the US, the associated repricing extends to European assets. At the same time, the euro depreciates against the dollar. Higher equity prices and a weaker euro can coexist if favourable US AI news raises the value of technology-sensitive assets globally while strengthening the relative appeal of US assets and the dollar. The combination is therefore consistent with the global comovement of equity markets and the dollar’s role in international financial conditions (Miranda-Agrippino and Rey, 2020; Rey, 2015).

Inflation expectations and nominal rates. Inflation expectations increase, with two-year inflation-linked swap rates peaking by about 36 basis points, compared with roughly 16 basis points at the ten-year maturity. The stronger short-maturity response indicates a sizeable near- to medium-term inflationary impulse without a comparable de-anchoring of long-run expectations. At the same time, the persistent response of core HICP suggests that the shock is not merely a transitory energy-price disturbance: imported cost pressure gradually feeds into underlying inflation even as its effect on longer-term expectations remains

comparatively contained ([Känzig, 2021](#)).

Both two- and ten-year OIS rates increase, with a particularly pronounced response at the shorter maturity. Their movements exceed those of the corresponding inflation-linked swap rates, suggesting that the adjustment also contains a real-rate component and, especially at longer maturities, changes in term premia.

Interest rate differentials and the policy dilemma. The EA–US interest rate differential follows a phased path. Initially, US rates rise by more than euro area rates at both maturities, moving the differential in favour of the US and accompanying the euro depreciation. The differential later reverses as domestic rates respond more strongly, broadly when inflation spreads from external and energy prices to headline and then core HICP. This pattern suggests that the relative monetary and financial stance adjusts as the composition of the spillover changes.

The interpretation should not be that the interest rate differential identifies a mechanical policy rule. Rather, it illustrates the dilemma created by global financial spillovers ([Rey, 2015](#)). Leaning immediately against a US-led rise in rates and the associated currency depreciation may restrain activity before the real benefits of adoption arrive; allowing those movements to pass through can instead amplify imported inflation. As price pressure becomes broader and more persistent, the balance shifts towards a stronger domestic rate response.

These results provide an empirical counterpart to the conceptual assessment of [Hartmann and Maver \(2025\)](#), who emphasise that AI can affect the monetary policy stance indirectly through its joint effects on productivity, investment, aggregate demand, labour markets, energy costs, and inflation, with the balance between supply and demand determining whether the overall impulse is inflationary or disinflationary. Here, the near- to medium-run configuration is inflationary: external and energy prices rise early, short-horizon inflation expectations and nominal rates increase, and core inflation strengthens before real activity responds clearly. Together with the early US investment response, this indicates that demand for inputs and implementation capacity can generate cost pressure before the broader productive gains emerge.

Taken together, the results point to a two-speed open-economy transmission mechanism.


At short horizons, the shock is transmitted through the global financial cycle: US AI news is rapidly reflected in euro area equity prices, the exchange rate, inflation expectations, and nominal rates, while the euro depreciation and higher external input costs contribute to the initial inflationary response. These financial and nominal spillovers arise before a clear expansion in domestic activity and elicit a monetary policy response as price pressures become broader and more persistent ([Jordà et al., 2019](#); [Auer et al., 2019](#); [Rey, 2015](#)).

At longer horizons, the distinction between AI innovator and AI adopter helps explain the slower real adjustment. In the present episode, the US is more directly exposed to the development and commercialisation of frontier AI and therefore experiences an early response in valuations, capital spending, and demand for complementary inputs. The productive benefits for the euro area depend more heavily on the subsequent adoption and integration of the new technology by domestic firms. Related firm-level evidence is provided by [Aldasoro et al. \(2026\)](#), who estimate that AI adoption raises labour productivity in EU firms by around 4%, with no significant employment effect. The gains are larger among firms undertaking complementary investments in software, data and workforce training. These findings support the relevance of adoption and complementary investment for realising productivity gains in Europe, although they do not establish the timing of the aggregate response to AI news documented here.

This pattern echoes the history of earlier general-purpose technologies: the gains from steam, electricity, and ICT emerged gradually as firms accumulated complementary capital and reorganised production, while international diffusion lagged invention and varied with countries' capacity to absorb and deploy new technologies ([Crafts, 2021](#); [Brynjolfsson et al., 2019](#); [Comin and Hobijn, 2010](#); [Keller, 2004](#); [Eaton and Kortum, 1999](#)). The global-financial-cycle and innovator–adopter interpretations are therefore complementary rather than competing. The former explains the rapid transmission of valuations, relative prices, and financial conditions; the latter accounts for the slower diffusion of real gains. Monetary and financial conditions adjust to the inflationary pressure that arises between these two stages.²⁰


²⁰The results excluding the 2020–2022 period are reported in [Section D.2.3](#). For the US, the short-run inflationary response remains visible, but investment and several medium-run real and labour-market responses are sensitive to the sample restriction. For the euro area, the short-run price responses are broadly preserved,

Figure 8: Impact of AI news shocks on euro area real and price aggregates



Notes: Cumulative impulse response functions (IRFs) to a one-standard-deviation AI news shock. Variables are expressed as percent deviation from the baseline ($t - 1$). Darker (lighter) shaded areas represent 68% (90%) highest-density posterior sets.

Figure 9: Impact of AI news shocks on euro area *financial* aggregates



5 Conclusions

This paper develops a financial-market-based approach to measuring and identifying AI news shocks. By linking daily relative equity valuations to monthly macroeconomic outcomes, it characterises the adjustment following changes in the expected economic value of AI, rather than changes in realised AI adoption or measured productivity. The results reveal a gradual and uneven adjustment rather than an immediate productivity-led expansion. News about AI first stimulates investment and demand for the resources needed to adopt the technology; production, employment, and consumer prices respond only as this process unfolds.

In the US, core capital-goods orders provide the clearest and most persistent real-side response. Producer prices also rise relatively early, whereas GDP and industrial production expand more modestly and with greater uncertainty. Capacity utilisation temporarily declines before recovering, consistent with firms building capacity and reorganising production while existing resources are being reallocated. The labour-market response emerges later: employment and, less precisely, hourly earnings increase, but average hours per worker fall. Price pressures follow a similarly ordered path, passing from producer prices to headline and then core consumer prices. The LPs therefore describe an investment-led transition in which adoption costs and resource demand precede the broader productive effects of AI.

The production-network results show why this aggregate adjustment contains both supply and demand forces. Industries with greater upstream or downstream exposure to AI initially record stronger relative production and smaller producer-price increases, revealing a favourable supply-side effect close to the AI ecosystem. At the same time, the pronounced employment response among highly exposed upstream suppliers points to demand for complementary inputs and implementation capacity. As the shock diffuses, the initial production advantage of highly exposed industries fades and the effects spread towards less-exposed sectors. AI news thus combines localised technological gains with an investment and resource-demand impulse that propagates through the wider economy.

International transmission follows the same logic but operates at two speeds. In the euro area, equity prices, the exchange rate, import costs, energy prices, inflation expectations, and

while the real responses remain imprecise and the financial responses are less stable. Owing to the reduced number of usable observations, the exercise is reported over the first 24 months.

nominal rates respond before industrial production and hours worked. Headline inflation rises ahead of core inflation, indicating that external price pressures gradually pass through to underlying prices. The EA–US interest rate differential initially moves in favour of the US and later reverses as domestic inflation broadens and financial conditions tighten. The rapid financial and nominal adjustment is consistent with transmission through the global financial cycle, while the delayed real response reflects the time required for an adopter economy to integrate a technology developed and commercialised primarily abroad.

Taken together, the findings reconcile rapid market repricing with the slower macroeconomic effects of AI. Expected future gains can generate substantial capital spending, input demand, employment adjustment, and inflation before their full effect on productive capacity is visible. They can also create cross-border price and financial pressures before productive benefits diffuse internationally. For monetary policy, the central issue is therefore not only the eventual supply potential of AI, but also the timing and composition of the demand, cost, and financial effects generated along the transition.

References

- Acemoglu, D. “The simple macroeconomics of AI”. *Economic Policy*, 40(121):13–58, 01 2025.
- Acemoglu, D. and Restrepo, P. “Artificial Intelligence, Automation, and Work”. In Agrawal, A., Gans, J., and Goldfarb, A., editors, *The Economics of Artificial Intelligence: An Agenda*, pages 197–236. University of Chicago Press, 2019a.
- Acemoglu, D. and Restrepo, P. “Automation and New Tasks: How Technology Displaces and Reinstates Labor”. *Journal of Economic Perspectives*, 33(2):3–30, 2019b.
- Acemoglu, D., Akcigit, U., and Kerr, W. “Networks and the Macroeconomy: An Empirical Exploration”. *NBER Macroeconomics Annual*, 30(1):273–335, 2016.
- Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. “Artificial Intelligence and Jobs: Evidence from Online Vacancies”. *Journal of Labor Economics*, 40(S1):S293–S340, 2022.
- Aghion, P., Jones, B. F., and Jones, C. I. “Artificial Intelligence and Economic Growth”. NBER Working Paper 23928, National Bureau of Economic Research, 2017.
- Aldasoro, I., Doerr, S., Gambacorta, L., and Rees, D. “The Impact of Artificial Intelligence on Output and Inflation”. BIS Working Papers 1179, Bank for International Settlements, April 2024.
- Aldasoro, I., Gambacorta, L., Pál, R., Revoltella, D., Weiss, C., and Wolski, M. “AI Adoption, Productivity and Employment: Evidence from European Firms”. CEPR Discussion Paper 21082, Centre for Economic Policy Research, January 2026.
- Alexopoulos, M. “Read All about It!! What Happens Following a Technology Shock?”. *American Economic Review*, 101(4):1144–1179, 2011.
- Andrews, D. W. K. “Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation”. *Econometrica*, 59(3):817–858, 1991.
- Antolín-Díaz, J. and Rubio-Ramírez, J. F. “Narrative Sign Restrictions for SVARs”. *American Economic Review*, 108(10):2802–2829, 2018.
- Arias, J. E., Rubio-Ramírez, J. F., and Waggoner, D. F. “Inference Based on Structural Vector Autoregressions Identified with Sign and Zero Restrictions: Theory and Applications”. *Econometrica*, 86(2):685–720, 2018a.
- Arias, J. E., Rubio-Ramírez, J. F., and Waggoner, D. F. “Inference Based on Structural Vector Autoregressions Identified With Sign and Zero Restrictions: Theory and Applications”. *Econometrica*, 86(2):685–720, 2018b.
- Auer, R. A., Levchenko, A. A., and Sauré, P. “International Inflation Spillovers through Input Linkages”. *The Review of Economics and Statistics*, 101(3):507–521, 2019.
- Babina, T., Fedyk, A., He, A., and Hodson, J. “Artificial Intelligence, Firm Growth, and Product Innovation”. *Journal of Financial Economics*, 151:103745, 2024.

- Baqae, D. R. and Farhi, E. “Macroeconomics with Heterogeneous Agents and Input-Output Networks”. Working Paper 24684, National Bureau of Economic Research, 2018.
- Baqae, D. R. and Farhi, E. “The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten’s Theorem”. *Econometrica*, 2019.
- Barsky, R. B. and Sims, E. R. “News Shocks and Business Cycles”. *Journal of Monetary Economics*, 58(3):273–289, 2011.
- Beaudry, P. and Portier, F. “Stock Prices, News, and Economic Fluctuations”. *American Economic Review*, 96(4):1293–1307, 2006.
- Bekaert, G., Hoerova, M., and Lo Duca, M. “Risk, Uncertainty and Monetary Policy”. *Journal of Monetary Economics*, 60(7):771–788, 2013.
- Blanchard, O. “The State of Macro”. *Annual Review of Economics*, 2:605–629, 2010.
- Bloom, N. “The impact of uncertainty shocks”. *Econometrica*, 77(3):623–685, 2009.
- Brynjolfsson, E., Rock, D., and Syverson, C. “Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics”. NBER Working Paper 24001, National Bureau of Economic Research, 2019.
- Brynjolfsson, E., Li, D., and Raymond, L. “Generative AI at Work”. *The Quarterly Journal of Economics*, 140(2):889–942, 2025.
- Caballero, G. “Supply Chains and Sectoral Adjustment”. *American Economic Review*, 109(2):365–403, 2019.
- Carvalho, V. M. “From Micro to Macro via Production Networks”. In Acemoglu, D., Arellano, M., and Dekel, E., editors, *Advances in Economics and Econometrics: Tenth World Congress*, volume 2, pages 73–116. Cambridge University Press, Cambridge, 2014.
- Carvalho, V. M. and Gomes, J. F. “Supply Chain Disruptions and Inflation Dynamics”. *American Economic Review*, 107(4):1057–1094, 2017.
- Carvalho, V. M. and Voigtländer, N. “Input Diffusion and the Evolution of Production Networks”. Working Paper 20025, National Bureau of Economic Research, 2014.
- Cerra, V. and Swanson, E. “The Role of Supply Chain Bottlenecks in the Transmission of Inflation”. *Journal of Monetary Economics*, 93:1–16, 2017.
- Cerutti, E. M., Garcia Pascual, A. I., Kido, Y., Li, L., Melina, G., Mendes Tavares, M., and Wingender, P. “The Global Impact of AI: Mind the Gap”. IMF Working Paper 25/76, International Monetary Fund, April 2025.
- Christensen, J. H. E. and Rudebusch, G. D. “Can Fiscal, AI, or Monetary News Explain the Rise in r^* ?”. Working Paper 2026-19, Federal Reserve Bank of San Francisco, August 2026.

- Comin, D. and Hobijn, B. “An Exploration of Technology Diffusion”. *American Economic Review*, 100(5):2031–2059, 2010.
- Crafts, N. “Artificial Intelligence as a General-Purpose Technology: An Historical Perspective”. *Oxford Review of Economic Policy*, 37(3):521–536, 2021.
- David, J., de Souza, G., and Schommer, A. “Firm Growth from the Artificial Intelligence Boom”. Chicago Fed Letter 518, Federal Reserve Bank of Chicago, Chicago, February 2026.
- De Gasperi, R. and Ricco, G. “Global Transmission of the U.S. Monetary Policy”. *Working Paper*, 2023. Please verify author names and full reference details.
- Eaton, J. and Kortum, S. “International Technology Diffusion: Theory and Measurement”. *International Economic Review*, 40(3):537–570, 1999.
- Eisfeldt, A. L., Schubert, G., and Zhang, M. B. “Generative AI and Firm Values”. NBER Working Paper 31222, National Bureau of Economic Research, 2023.
- Engel, C. “Exchange Rates, Interest Rates, and the Risk Premium”. *American Economic Review*, 106(2):436–474, 2016.
- Fang, H., Gu, X., Yan, H., and Zhu, W. “AI Patents in the United States and China: Measurement, Organization, and Knowledge Flows”. NBER Working Paper 35022, National Bureau of Economic Research, April 2026.
- Forbes, K. J. “Has Globalization Changed the Inflation Process?”. BIS Working Paper 791, Bank for International Settlements, Basel, 2019.
- Fujita, M. “The Role of Supply Chains in Business Cycles”. *Review of Economic Dynamics*, 17:1–20, 2014.
- Galí, J. “Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework”. *Princeton University Press*, 2008. Book chapter or working paper; you can cite the book if appropriate.
- Gazzani, A. and Natoli, F. “The Macroeconomic Effects of AI-based Innovation”. Temi di Discussione (Working Papers) 1542, Bank of Italy, July 2026.
- Gordon, R. J. “The Bottleneck and the Transmission of Inflation”. *Brookings Papers on Economic Activity*, 2010(1):1–59, 2010.
- Gormsen, N. J. and Lazarus, E. “Duration-Driven Returns”. *Journal of Finance*, 78(3): 1393–1447, 2023.
- Gourinchas, P.-O. and Rey, H. “International Financial Adjustment”. *Journal of Political Economy*, 115(4):665–703, 2007.

- Gürkaynak, R. S., Sack, B., and Swanson, E. T. “Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements”. *International Journal of Central Banking*, 1(1):55–93, 2005.
- Hartmann, P. and Maver, V. “Implications of Artificial Intelligence for Monetary Policy: A First Conceptual Assessment”. SUEF Policy Brief 1080, SUEF – The European Money and Finance Forum, January 2025.
- IEA. “Energy and AI”. Technical report, International Energy Agency (IEA), Paris, April 2025.
- Jaimovich, N. and Rebelo, S. “Can News about the Future Drive the Business Cycle?”. *American Economic Review*, 99(4):1097–1118, 2009.
- Jaimovich, N. and Rebelo, S. “Can Technology Explain the Recent Evolution of the Labor Share?”. *American Economic Review*, 102(2):606–611, 2012.
- Jarociński, M. and Karadi, P. “Deconstructing Monetary Policy Surprises—The Role of Information Shocks”. *American Economic Journal: Macroeconomics*, 12(2):1–43, 2020.
- Jarociński, M. and Karadi, P. “Deconstructing Monetary Policy Surprises—The Role of Information Shocks”. *American Economic Journal: Macroeconomics*, 12(2):1–43, April 2020.
- Jiang, Z., Krishnamurthy, A., and Lustig, H. “Foreign Safe Asset Demand and the Dollar Exchange Rate”. *Journal of Finance*, 76(3):1049–1089, 2021.
- Jordà, Ò. “Estimation and Inference of Impulse Responses by Local Projections”. *American Economic Review*, 95(1):161–182, 2005.
- Jordà, Ò., Schularick, M., Taylor, A. M., and Ward, F. “Global Financial Cycles and Risk Premiums”. *IMF Economic Review*, 67(1):109–150, 2019.
- Känzig, D. R. “The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements”. *American Economic Review*, 111(4):1092–1125, 2021.
- Keller, W. “International Technology Diffusion”. *Journal of Economic Literature*, 42(3):752–782, 2004.
- Kuttner, K. N. “Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market”. *Journal of Monetary Economics*, 47(3):523–544, 2001.
- Lenza, M. and Primiceri, G. E. “How to estimate a vector autoregression after March 2020”. *Journal of Applied Econometrics*, 37(4):688–699, 2022.
- MacKinlay, A. C. “Event Studies in Economics and Finance”. *Journal of Economic Literature*, 35(1):13–39, 1997.
- Miranda-Agrippino, S. and Rey, H. “U.S. Monetary Policy and the Global Financial Cycle”. *Review of Economic Studies*, 87(6):2754–2776, 2020.

- Mokyr, J. “The Past and Future of Innovation Policy”. *Journal of Economic Perspectives*, 29(3):147–172, 2015.
- Pagan, A. “Econometric Issues in the Analysis of Regressions with Generated Regressors”. *International Economic Review*, 25(1):221–247, 1984.
- Ready, R., Roussanov, N., and Ward, C. “Commodity Trade and the Carry Trade: A Tale of Two Countries”. *Journal of Finance*, 72(6):2629–2684, 2017.
- Rey, H. “Dilemma not Trilemma: The Global Financial Cycle and Monetary Policy Independence”. *NBER Working Paper No. 21162*, 2015.
- Seftian, A. A., Constantinides, M., and Quercia, D. “The Potential Impact of AI Innovations on U.S. Occupations”. *PNAS Nexus*, 3(9), 2024.
- Uhlig, H. “What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure”. *Journal of Monetary Economics*, 52(2):381–419, 2005.
- Webb, M. “The Impact of Artificial Intelligence on the Labor Market”. Working paper, Stanford University, 2020. SSRN 3482150.
- Yang, C.-H. “How Artificial Intelligence Technology Affects Productivity and Employment: Firm-Level Evidence from Taiwan”. *Research Policy*, 51(6):104536, 2022.

A Data

A.1 Construction of the AI equity factor and robustness

This section describes the construction of the AI equity factor used in the main analysis and presents a battery of robustness checks. We build a primary long–short factor from a hand-selected universe of AI-exposed equities and then compare it with i) an equal-weighted variant; ii) orthogonalisation against semiconductor returns; iii) a fully independent, market-priced factor derived from AI-themed exchange-traded funds (ETFs). Throughout, the sample runs at the daily frequency and all returns are continuously compounded log returns of adjusted closing prices obtained from Yahoo Finance.

A.1.1 Benchmark factor: market-cap-weighted long–short portfolio

The primary factor, AI_t^{MC} , is the return of a long–short portfolio that is long firms whose business models are exposed to AI and short firms in adjacent technology segments with more limited AI exposure. The long leg spans three tiers: AI compute and semiconductors (NVDA, AMD, AVGO, SMC, TSM, MU, ASML, QCOM), large-cap AI platforms (MSFT, GOOGL, AMZN, META, ORCL), and AI-oriented enterprise software (PLTR, AI, CRM, NOW, ADBE, ANET). The short leg comprises non-AI semiconductor peers (INTC, TXN, STX, WDC), legacy enterprise hardware (IBM, HPE, CSC), and traditional IT services (ACN, GLW). The short leg therefore serves as a technology-sector hedge, reducing broad technology beta and isolating relative AI-related valuation news.

The daily factor is available from 4 January 2012 to 23 July 2026, for a total of 3,658 return observations. The constituent universe is allowed to evolve with listings: firms that are not yet publicly traded receive zero portfolio weight, and weights within each leg are normalised across securities available at each rebalancing date. This is particularly relevant for later entrants such as ANET, HPE, PLTR, and AI.

Within each leg, portfolio weights are market-capitalisation weighted and rebalanced on the first trading day of each calendar quarter. Let $\mathcal{L}_t$ and $\mathcal{S}_t$ denote the sets of available long- and short-leg securities at date t , and let $w_{i,t}^{\mathcal{L}}$ and $w_{j,t}^{\mathcal{S}}$ be the corresponding weights. Market capitalisations are constructed from daily prices and time-varying shares outstanding; annual

shares outstanding are linearly interpolated between reporting dates. The factor return is

$$\text{AI}_t^{\text{MC}} = \underbrace{\sum_{i \in \mathcal{L}_t} w_{i,t}^{\mathcal{L}} r_{i,t}}_{R_t^{\mathcal{L}}} - \underbrace{\sum_{j \in \mathcal{S}_t} w_{j,t}^{\mathcal{S}} r_{j,t}}_{R_t^{\mathcal{S}}}, \quad (\text{A.1})$$

where $r_{i,t}$ is the log return of security i on day t . As a robustness variant we also construct an equal-weighted factor AI_t^{EW} . In the extended sample, equal weights are computed across the constituents actually trading on each date, so that pre-IPO observations are not treated as zero returns.

A.1.2 An independent, market-priced factor from AI ETFs

A natural concern is that the primary factor reflects discretionary stock selection rather than a priced AI-related return component. To address this we construct a fully independent measure from a basket of five AI- and robotics-themed ETFs that share no securities-selection methodology with our stock-level factor: Global X Artificial Intelligence & Technology (**AIQ**), Global X Robotics & Artificial Intelligence (**BOTZ**), ROBO Global Robotics & Automation (**ROBO**), First Trust Nasdaq AI & Robotics (**ROBT**), and Defiance Quantum (**QTUM**). The common sample begins on 6 September 2018, when all five funds are available, and contains 1,979 daily observations through 23 July 2026.

Because the individual ETFs are highly collinear, we extract their common component by principal component analysis (PCA). Let $z_{k,t}$ denote the standardised (zero-mean, unit-variance) daily return of ETF k , and let $\mathbf{z}_t = (z_{1,t}, \dots, z_{5,t})'$. The first principal component,

$$\text{PC1}_t = \mathbf{v}_1' \mathbf{z}_t, \quad (\text{A.2})$$

where $\mathbf{v}_1$ is the leading eigenvector of the return correlation matrix, explains 92.14% of the basket's variance and loads almost identically on all five funds (loadings between 0.441 and 0.454). We rescale PC1_t into daily return units by projecting it onto the equal-weighted basket return, and denote the result $\widetilde{\text{PC1}}_t$.

The preferred ETF-based AI factor is defined as the excess of this common AI-ETF

Table A.1: Benchmark AI factor constituents: composition and sample-end market data

Ticker	Tier	Weight (%)	Shares (bn)	Mkt cap (\$bn)
<i>Long leg</i>				
NVDA	Semis	19.93	24.270	5,066.6
GOOGL	Big tech	18.27	12.170	3,866.3
MSFT	Big tech	11.88	7.440	2,839.0
AMZN	Big tech	10.61	10.560	2,467.4
TSM	Semis	9.58	5.190	2,156.9
AVGO	Semis	7.11	4.630	1,817.1
META	Big tech	6.50	2.550	1,545.6
MU	Semis	4.76	1.110	1,099.1
AMD	Semis	3.64	1.620	874.3
ASML	Semis	2.99	0.391	704.1
ORCL	Big tech	1.61	2.730	327.7
PLTR	Software	1.13	2.170	267.7
CRM	Software	0.65	0.960	150.7
QCOM	Semis	0.64	0.850	145.4
ADBE	Software	0.39	0.440	93.4
ANET	Software	0.21	0.310	54.7
NOW	Software	0.09	0.200	18.4
SMCI	Semis	0.01	0.059	1.8
AI	Software	0.01	0.165	1.3
<i>Short leg</i>				
INTC	Semis (non-AI)	23.93	4.260	427.0
CSCO	Legacy tech	20.98	4.070	458.9
TXN	Semis (non-AI)	11.95	0.910	258.0
IBM	Legacy tech	11.44	0.910	186.7
WDC	Semis (non-AI)	9.00	0.340	189.8
STX	Semis (non-AI)	8.50	0.210	191.8
GLW	IT services	8.10	0.830	129.5
ACN	IT services	3.61	0.630	87.4
HPE	Legacy tech	2.49	1.280	61.0

Notes: The table reports the constituents of the AI long-short factor. “Weight” is the market-capitalisation weight at the end of the sample, 23 July 2026, inherited from the latest quarterly rebalance. “Shares” (billions) and “Mkt cap” (\$bn) are measured on the same date. Within each leg, constituents are ordered by weight. The long leg is concentrated in the largest AI-related firms: NVIDIA and Alphabet together account for approximately 38.2% of the long portfolio at the end of the sample.

component over the broad market,

$$\text{AI}_t^{\text{ETF}} = \widetilde{\text{PC}}1_t - r_t^{\text{SPY}}, \quad (\text{A.3})$$

where r_t^{SPY} is the return on the S&P 500 ETF. The choice of market benchmark is conse-

Table A.2: Benchmark AI factor summary statistics by specification and sub-period

Specification	Period	N	Ann. mean (%)	Ann. std (%)	Sharpe	Corr(AI,EQ)	Skew.	Kurt.
Market-cap (TV)	Full	3,658	8.87	16.94	0.52	0.147	0.073	3.204
	Pre-Nov 2022	2,745	7.42	14.80	0.50	0.108	0.160	4.035
	Post-Nov 2022	913	13.23	22.17	0.60	0.246	-0.033	1.293
Equal-weight	Full	3,658	4.67	16.65	0.28	0.117	-0.138	2.730
	Pre-Nov 2022	2,745	6.93	14.77	0.47	0.057	-0.066	2.876
	Post-Nov 2022	913	-2.13	21.33	-0.10	0.268	-0.167	1.467
Long-only (TV)	Full	3,658	20.30	24.14	0.84	0.859	-0.296	6.450
	Pre-Nov 2022	2,745	15.84	23.34	0.68	0.865	-0.496	7.419
	Post-Nov 2022	913	33.68	26.39	1.28	0.858	0.104	4.317

Notes: The table reports summary statistics for three specifications of the AI factor: market-cap time-varying (TV) long-short, equal-weight long-short, and long-only (market-cap TV). The full sample runs from 4 January 2012 to 23 July 2026. Sub-periods split the sample at the launch of ChatGPT on 30 November 2022. Annualised means and standard deviations use 252 trading days. “Corr(AI,EQ)” is the contemporaneous correlation with the S&P 500. Kurtosis is reported in excess of the normal. Over the extended sample, the benchmark market-cap long-short factor has an annualised mean of 8.87%, volatility of 16.94%, and a correlation of only 0.147 with the S&P 500. The corresponding long-only portfolio has a correlation of 0.859 with the market, highlighting the role of the short leg in removing broad equity exposure.

Table A.3: Evolution of long-leg market-cap weights (selected constituents)

Date	NVDA	MSFT	GOOGL	AMZN	AMD
2012	0.3	26.2	32.3	0.6	0.6
2013	0.2	21.0	29.6	0.7	0.2
2014	0.2	20.6	32.8	0.7	0.2
2015	0.2	22.7	25.5	0.5	0.1
2016	0.3	23.1	31.4	0.9	0.1
2017	0.7	20.9	26.2	8.7	0.5
2018	1.2	23.1	14.0	11.1	0.4
2019	0.8	27.1	13.4	14.5	0.7
2020	0.9	29.5	11.8	12.3	1.5
2021	1.5	27.4	10.1	14.3	2.0
2022	2.0	25.1	19.2	18.0	2.5
2023	5.8	28.9	18.5	15.0	1.7
2024	11.1	25.5	15.9	14.9	2.1
2025	20.3	18.6	13.9	14.1	1.2

Notes: The table reports the market-capitalisation weight (in %) of selected long-leg constituents on the first trading day of January of each year. Weights are rebalanced on the first trading day of January, April, July, and October. The extended series makes explicit that the composition of the long leg changes substantially over time, rather than imposing the end-of-sample AI portfolio retrospectively with fixed weights.

quential. A regression of $\widetilde{\text{PC1}}_t$ on the Nasdaq-100 ETF return yields a slope of 0.969 and an R^2 of 0.842, indicating that most daily variation in the diversified AI-ETF basket is shared with the broader technology sector. Consequently, subtracting QQQ or residualising the ETF factor with respect to QQQ removes much of the component that responds on AI event days. We therefore use the broader S&P 500 as the benchmark and treat the Nasdaq-based variants as diagnostics.

A.1.3 Narrative validation

We validate the sign of each factor using a hand-collected chronology of AI-related events, classified as technological innovations, adoption and commercialisation events, or policy actions, and assigned an ex ante expected sign (+, −, or ambiguous). For each dated event we take the factor return on the nearest trading day and record whether its sign matches the prior. The chronology contains 31 events, seven of which are classified as ambiguous and excluded from the scoring, leaving 24 signed events. The event chronology is concentrated in the generative-AI era, while the daily market factor itself is available from 2012.

Table A.4 reports the resulting event-day hit rates. The primary market-cap factor recovers the correct sign on all 24 signed events, while the equal-weighted variant does so on 22 (92 %). The independent ETF factor in Equation (A.3) recovers the expected sign on 18 of 24 events (75 %). Under a fair-coin null of no directional content, 18/24 correct corresponds to a one-sided exact binomial p -value of approximately 0.011. By contrast, the Nasdaq-benchmarked ETF variants obtain only 12–13 correct signs out of 24. The semiconductor- orthogonalised versions of the primary stock-level factor remain much closer to the benchmark, with 22/24 correct signs after removing SMH and 24/24 after removing SOXX.

A.1.4 Robustness checks

Composition and weighting. The event-day performance of the equal- and market-cap-weighted factors (92 % versus 100 %) indicates that the directional event signal is not specific to the market-cap weighting scheme. The extended 2012–2026 factor also uses a dynamically available constituent set, so later-listed firms enter only once their returns become observable.

Table A.4: Event-day hit rates by factor specification

Factor	Hits	Events	Hit rate
Market-cap long-short (primary)	24	24	100 %
Equal-weighted long-short	22	24	92 %
ETF: PC1 – SPY (independent)	18	24	75 %
MC factor orthogonal to SMH	22	24	92 %
MC factor orthogonal to SOXX	24	24	100 %
ETF: PC1 – QQQ (diagnostic)	12	24	50 %
ETF: residual of PC1 on QQQ	13	24	54 %
ETF: residual of PC1 on QQQ and SPY	13	24	54 %

Notes. Hit rate is the fraction of the 24 signed AI events on which the factor return has the ex ante expected sign; ambiguous events are excluded. The semiconductor-orthogonalised rows remove, respectively, the daily return on SMH or SOXX from the benchmark market-cap factor. The final three ETF rows are diagnostics. The common AI-ETF component has a slope of 0.969 and $R^2 = 0.842$ on QQQ, so Nasdaq-based hedging removes a large share of the common ETF variation.

Semiconductor cycle. Because semiconductors are a key transmission channel for AI, a natural concern is that the factor merely proxies the chip cycle. We regress the primary factor on semiconductor-sector ETFs over the full 2012–2026 sample. The exposure is small: the regression on SMH gives $\beta = 0.086$ and $R^2 = 2.3\%$, while the regression on SOXX gives $\beta = 0.057$ and $R^2 = 1.1\%$. Orthogonalising the benchmark factor with respect to each semiconductor ETF and repeating the narrative validation leaves the event-sign recovery at 22/24 (92 %) and 24/24 (100 %), respectively. Thus, the identifying event signal is not accounted for by generic semiconductor-sector variation.

Independent corroboration and its limitations. The independent ETF factor recovers the expected sign on three quarters of the signed events despite using a completely different portfolio construction. Its six misses occur when the ETF-factor return is economically small ($|r| < 0.004$) and cluster in model-release or issuer-specific announcements, including three NVIDIA-specific events. By contrast, the stock-level and ETF measures agree on several broad AI repricing episodes: both are negative on the DeepSeek R1 release and positive on the ChatGPT launch, the Stargate announcement, and the US AI Action Plan. At the same time, the relatively low full-sample correlation between the ETF and benchmark stock-level factors (0.157) shows that the ETF measure should be viewed as independent corroboration rather than as an interchangeable proxy.

Stability across the ChatGPT structural break. The benchmark market-cap factor remains only weakly correlated with the S&P 500 both before and after the launch of ChatGPT, with correlations of 0.109 and 0.246, respectively. Its annualised Sharpe ratio is also similar across the two periods (0.50 before and 0.60 after). The correlation between the independent ETF factor and the benchmark factor is 0.282 before 30 November 2022 but falls to 0.047 afterwards. This weaker post-ChatGPT comovement is consistent with the fact that the diversified thematic ETFs retain substantial broad-technology exposure even as AI exposure diffuses across their holdings, and motivates treating the ETF measure as a robustness check rather than the baseline factor.

Taken together, these checks address four principal concerns about the factor: i) sensitivity to the portfolio weighting scheme, ii) changes in the investable constituent set over the extended sample, iii) latent semiconductor-cycle exposure, and iv) discretionary stock selection, supporting its use as a market-based measure of AI-related news distinct from broad equity-market and semiconductor sector movements.

A.2 Macroeconomic data

A.2.1 US aggregates

Aggregate data for the US are taken from the Federal Reserve Economic Data (FRED) database. All series are available at monthly frequency and are seasonally adjusted at source. The real GDP series is interpolated at the monthly frequency using the Chow-Lin procedure with monthly industrial production as control. The sample used in the local projection exercise starts in January 2015 and ends in June 2026.

A.2.2 Euro area aggregates

Data for the euro area are taken from the following sources:

- **Employment, hours worked, industrial production, industrial producer price index (PPI), import prices:** Eurostat.
- **HICP core and headline, EUR/USD exchange rate:** ECB Statistical Data Warehouse.
- **Energy price index:** OECD.
- **Euro area equity prices (EUROSTOXX 50) and ILS 2-year/10-year (midpoint**

of bid/ask prices): Refinitiv.

- **Natural gas price (NL), two- and ten-year euro area OIS rates and differentials with the US:** Bloomberg.

All series are seasonally adjusted at source, except PPI and import prices, which we adjust using an STL decomposition. Quarterly employment and hours-worked series are converted to monthly frequency by linear interpolation. The sample starts in January 2015 and ends in July 2026, except for employment and hours worked, which are available through March 2026.

A.2.3 US sectoral aggregates

This section describes the sectoral aggregates used to study the heterogeneous transmission of AI news shocks across US industries. We compare sectors with strong and weak production-network connections to AI-related industries by constructing separate high- and low-exposure aggregates for producer prices, industrial production, and employment, distinguishing between upstream and downstream exposure.

Sectoral data. Sectoral price, production, and employment series are collected at monthly frequency from US statistical sources. Producer price indices are obtained from the Bureau of Labor Statistics (BLS). Industrial production indices are drawn from the Federal Reserve’s *G.17 Industrial Production and Capacity Utilization* release, mostly through FRED. Employment series are obtained from the BLS Current Employment Statistics. The sample runs from January 2015 to June 2026.

Detailed BLS PPI series are not always available on a seasonally adjusted basis. We therefore seasonally adjust the relevant PPI series using an STL decomposition. For the construction of the sectoral aggregates, all series are expressed in log levels and rebased around January 2015. Since the rebasing amounts to subtracting a series-specific constant in logs, it does not affect the estimated local projection responses.

Input-output exposure measures. We construct production-network exposure measures using the Bureau of Economic Analysis (BEA) Input-Output Accounts. Specifically, we use BEA Table 259, “The Use of Commodities by Industries – Summary,” which reports the

value of commodities supplied by each row sector and used by each column industry. We use the 2023 table.

Let Z_{ij} denote the value of inputs supplied by sector i and used by sector j . We define a set of AI-related BEA industries, $\mathcal{A}^{BEA}$, consisting of computer and electronic products, computer systems design and related services, publishing industries except internet (including software), data processing and internet-related services, and broadcasting and telecommunications.

We consider two distinct measures of production-network exposure. First, upstream exposure captures the extent to which a sector supplies AI-related industries:

$$\text{Exp}_i^{up} = \frac{\sum_{j \in \mathcal{A}^{BEA}} Z_{ij}}{\sum_j Z_{ij}}. \quad (\text{A.4})$$

Hence, Exp_i^{up} measures the share of sector i 's sales that is directed towards AI-related industries. A high value therefore identifies sectors that are important suppliers to the AI production ecosystem and may be particularly exposed to AI-induced demand pressures, capacity constraints, and input bottlenecks.

Second, downstream exposure captures the extent to which a sector relies on inputs supplied by AI-related industries:

$$\text{Exp}_j^{down} = \frac{\sum_{i \in \mathcal{A}^{BEA}} Z_{ij}}{\sum_i Z_{ij}}. \quad (\text{A.5})$$

A high value of Exp_j^{down} therefore identifies industries that intensively use AI-related inputs and is more closely related to the diffusion and adoption of AI-related technologies through the production network.

For ease of comparison, both measures are normalised by their respective maximum across BEA sectors,

$$\widetilde{\text{Exp}}_i^{up} = \frac{\text{Exp}_i^{up}}{\max_k \text{Exp}_k^{up}}, \quad \widetilde{\text{Exp}}_i^{down} = \frac{\text{Exp}_i^{down}}{\max_k \text{Exp}_k^{down}}, \quad (\text{A.6})$$

so that each exposure measure ranges between zero and one. This normalisation has no bearing on the sectoral ranking.

Mapping statistical series into BEA industries. The BEA input-output classification is coarser than the classifications underlying the BLS PPI, BLS CES, and Federal Reserve industrial-production series. We therefore map each statistical series to the closest BEA summary industry. For example, semiconductor and computer-related series are mapped to “Computer and electronic products,” software publishers to “Publishing industries, except internet (includes software),” data processing to “Data processing, internet publishing, and other information services,” computer systems design to “Computer systems design and related services,” and electric power to “Utilities.”

When several statistical series map into the same BEA summary industry, we first collapse them into a single BEA-sector series by taking their equal-weighted average in rebased log levels. This prevents an industry from receiving a larger weight merely because more detailed statistical series happen to be available for that sector. The production-network classification is therefore performed over unique BEA industries rather than over individual statistical series.

Construction of high- and low-exposure aggregates. We construct the exposure groups separately for each outcome domain—producer prices, industrial production, and employment—and for upstream and downstream exposure. We first retain BEA industries for which at least 90% of the monthly observations are available over the local projection sample. Within each resulting cross-section, industries are ranked according to either $\widetilde{\text{Exp}}_i^{up}$ or $\widetilde{\text{Exp}}_i^{down}$.

Our baseline specification defines the *high AI-network-exposure* group as the set of industries whose network exposure lies at or above the 75th percentile of the relevant cross-sectional exposure distribution. Conversely, the *low AI-network-exposure* group contains industries whose exposure lies at or below the 25th percentile.²¹ Industries belonging to the middle 50% of the distribution are omitted in order to sharpen the contrast between sectors with particularly strong and particularly weak production-network linkages to AI-related industries. The classification is performed separately for upstream and downstream exposure and for each outcome domain. Group membership is determined once from the time-invariant

²¹In the robustness exercise reported in [Section D.2.4](#) below, we set these thresholds to 90th and 10th percentiles respectively.

BEA input-output exposure measures and is held fixed throughout the sample.

For each domain d and exposure direction $g \in \{up, down\}$, we construct the high- and low-exposure indices as equally weighted geometric means of the constituent sectoral series. Let $\mathcal{H}_{d,g}$ and $\mathcal{L}_{d,g}$ denote the corresponding sets of BEA industries, and let $y_{i,t}$ denote the rebased log level of sector i . The high-exposure index is

$$Y_{d,g,t}^H = 100 \times \exp \left(\frac{1}{|\mathcal{H}_{d,g}|} \sum_{i \in \mathcal{H}_{d,g}} y_{i,t} \right), \quad (\text{A.7})$$

with the low-exposure index $Y_{d,g,t}^L$ defined analogously. The equal weighting ensures that the results are not mechanically dominated by the largest industries. To maintain a stable composition over time, the sets of constituent industries are fixed; an aggregate is treated as missing in a given month if fewer than 80% of its constituent sectors are observed.

Finally, we also construct the relative high-to-low exposure index

$$Y_{d,g,t}^{H/L} = 100 \times \frac{Y_{d,g,t}^H}{Y_{d,g,t}^L}. \quad (\text{A.8})$$

Because the local projections are estimated in logs, the response of this index is exactly the difference between the log responses of the high- and low-exposure aggregates,

$$\Delta_h \log Y_{d,g,t}^{H/L} = \Delta_h \log Y_{d,g,t}^H - \Delta_h \log Y_{d,g,t}^L. \quad (\text{A.9})$$

It therefore provides a direct measure of the differential response of industries with high AI-network exposure relative to industries with low exposure (see [Figure 7](#)).

B The AI news factor

B.1 High-frequency validation around the GPT-4o announcement

We complement the daily event validation with a high-frequency event study around the release of OpenAI’s GPT-4o on 13 May 2024. The OpenAI Spring Update started at 13:00 ET, while US equity markets were open, providing a setting in which the relative valuation response of AI-exposed firms can be examined over a narrow intraday window. The exercise

is intended as a validation of the information content of the daily AI factor rather than as an alternative high-frequency identification strategy.

Intraday data. One-minute bid and ask quotes are obtained from the Dukascopy historical datafeed for 13 May 2024. For each stock i , we construct the one-minute midquote as

$$P_{i,\tau}^{\text{mid}} = \frac{P_{i,\tau}^{\text{bid}} + P_{i,\tau}^{\text{ask}}}{2}, \quad (\text{B.10})$$

and compute intraday log returns as

$$r_{i,\tau} = \Delta \log P_{i,\tau}^{\text{mid}}. \quad (\text{B.11})$$

The analysis uses the regular US trading session and focuses on the interval from 12:00 to 14:00 ET, corresponding to a $[-60, +60]$ minute window around the beginning of the OpenAI event. When no new quote is observed within a given minute, the most recent midquote is carried forward, implying a zero return until the next price update. Dukascopy prices correspond to stock-CFD quotes rather than consolidated US transaction prices; the exercise should therefore be interpreted as a high-frequency robustness check rather than as a microstructure analysis based on the consolidated tape.

Portfolio construction and coverage. The intraday portfolios reproduce as closely as possible the construction of the daily AI factor. We use the same long and short stock universes and the same market-capitalisation weights that are in force in the second quarter of 2024. The daily benchmark contains 19 stocks in the AI-exposed long portfolio and 9 stocks in the non-AI short portfolio. Dukascopy provides usable one-minute data for 16 of the 19 long-leg stocks and 8 of the 9 short-leg stocks. The unavailable long-leg constituents are SMCI, ASML and C3.ai (ticker AI), while STX is unavailable in the short leg.

Despite the missing securities, coverage in market-capitalisation terms is very high. The available stocks account for 97.11% of the original long-leg weight and 98.10% of the original short-leg weight. We therefore retain the original daily-factor weights and renormalise them within each leg over the subset of securities for which intraday data are available. Denoting

by $\mathcal{L}_{HF}$ and $\mathcal{S}_{HF}$ the available long and short constituents, respectively, the intraday portfolio returns are

$$r_{\tau}^L = \sum_{i \in \mathcal{L}_{HF}} \tilde{w}_i^L r_{i,\tau}, \quad (\text{B.12})$$

$$r_{\tau}^S = \sum_{j \in \mathcal{S}_{HF}} \tilde{w}_j^S r_{j,\tau}, \quad (\text{B.13})$$

where $\tilde{w}_i^L$ and $\tilde{w}_j^S$ are the corresponding daily-factor weights rescaled to sum to one within each available portfolio. The high-frequency differential is then

$$r_{\tau}^{AI,HF} = r_{\tau}^L - r_{\tau}^S. \quad (\text{B.14})$$

Thus, the intraday measure preserves the economic definition of the benchmark daily factor and differs only because a small fraction of the original portfolio is unavailable in the high-frequency data.

Event-study measure. For presentation, one-minute returns are cumulated and normalised to zero in the minute immediately preceding the event:

$$CAR_{\tau}^k = 100 \sum_{s=-59}^{\tau} r_s^k - 100 \sum_{s=-59}^{-1} r_s^k, \quad k \in \{L, S, AI, HF\}. \quad (\text{B.15})$$

The resulting series therefore measures the cumulative percentage log return relative to 12:59 ET. We additionally report the cumulative return on the S&P 500 over the same window as a broad-market benchmark. The S&P 500 does not enter either portfolio and is used only to distinguish relative AI repricing from aggregate equity-market movements.

Results. As shown in Panel b of [Figure 1](#), the high-frequency exercise reveals a sizeable relative repricing of AI-exposed equities following the announcement, with a non-monotonic intraday response. The initial decline coincides with a fall in the broader market: the S&P 500 and both portfolios decline, but the AI portfolio falls more strongly than the non-AI portfolio. NVIDIA and Microsoft account for much of the negative contribution from the long leg

to the differential. This relative underperformance subsequently reverses, initially through an Alphabet-led recovery during the voice-demo window and subsequently through more widespread gains among AI-exposed firms during the vision and mathematics demonstrations. Positive contributions moderate during the final phase. The differential reaches approximately 20 basis points around 30 minutes after the beginning of the event. The S&P 500 remains below its pre-announcement level over the same window. An equal-weighted version of the intraday differential produces a qualitatively similar initial response but a smaller subsequent increase, indicating that the subsequent positive repricing is concentrated among the larger firms in the AI portfolio. Since the benchmark daily factor is market-capitalisation weighted, the corresponding intraday specification is the closest high-frequency analogue of the measure employed in the main analysis.

Stock-level contributions. Figure B.1 decomposes the differential into the contributions of individual constituents over four non-overlapping windows corresponding approximately to successive presentation phases: announcement, access and API economics ($[0, 9)$ minutes); voice demonstrations ($[9, 14)$); vision and mathematics demonstrations ($[14, 18)$); and coding, translation and closing ($[18, 26)$). Time is measured from 13:00 ET, and the windows last 9, 5, 4 and 8 minutes, respectively.²² The timing markers in Panel b show the approximate starts of the demonstrations; translation at minute 22 falls within the final contribution window. Each contribution is the stock’s cumulative log return over the indicated window multiplied by its intraday portfolio weight, with the sign reversed for short-leg stocks.²³ Blue bars identify AI-exposed long-leg stocks and yellow bars identify non-AI short-leg stocks. A positive yellow bar therefore indicates that a decline in a short-leg stock supports the differential, rather than that the stock itself appreciates.

Panel a shows negative contributions from most AI-exposed constituents during the initial


²²The phase boundaries are approximate, at one-minute precision. These windows describe returns during successive presentation phases rather than separately identifying the causal effects of individual demonstrations. Contributions are measured over each window, not per minute; window lengths therefore matter when comparing their magnitudes.

²³The weights are the renormalised second-quarter 2024 weights described above and are held fixed across the windows. Contributions are expressed in basis points and sum to the long–short differential over the corresponding window. They reflect both portfolio weights and stock returns; a larger contribution does not necessarily imply a larger individual-stock return.

announcement phase. NVIDIA and Microsoft together contribute approximately -14.5 basis points, with Meta, Broadcom and Amazon adding roughly another -6 basis points. Alphabet is the main positive contributor within the AI portfolio. Falling non-AI stocks partly offset these negative long-leg contributions, consistent with the broader market decline visible in Panel [b](#). During the voice-demo window, Panel [b](#) shows an initially concentrated reversal: Alphabet contributes around 7 basis points, while NVIDIA remains negative and Microsoft's contribution is close to zero. Negative short-leg contributions indicate that non-AI stocks also recover, dampening the increase in the differential.

During the vision and mathematics window, Panel [c](#) shows more widespread positive contributions among AI-exposed firms. Alphabet, Microsoft, Meta, Amazon and NVIDIA contribute approximately 5.2, 2.7, 2.3, 1.7 and 1.3 basis points, respectively. A simultaneous recovery in non-AI stocks again offsets part of these gains in the differential. Finally, Panel [d](#) shows smaller and more concentrated gains during coding, translation and closing: Alphabet and NVIDIA remain the leading positive contributors, while Broadcom contributes negatively. The decomposition thus documents an initial decline, an Alphabet-led reversal, a broader AI-portfolio recovery and then more modest, uneven gains.

Figure B.1: Stock-level contributions to the AI differential by presentation phase
(a) Announcement, access and API: [0, 9)
(b) Voice demonstrations: [9, 14)



Sources: Dukascopy, author's computations.

Notes: Stock-level contributions to the AI-portfolio minus non-AI-portfolio cumulative log return, in basis points. Time is measured in minutes from 13:00 ET on 13 May 2024. Blue (AI) and yellow (N) bars denote long- and short-leg constituents, respectively; short-leg contributions enter with a negative portfolio sign. Each panel reports contributions over its own non-overlapping phase window, not cumulative contributions from the event start. The windows last 9, 5, 4 and 8 minutes, respectively, and their boundaries approximate the presentation's chronology. Positive short-leg contributions indicate falling non-AI stock prices; negative short-leg contributions indicate rising prices. Stocks are ranked separately in each panel; horizontal-axis scales differ across panels.

B.2 Event study

Table B.5: Individual event study results

Type	Event date	Trading date	Event	Exp.	AI factor	EQ returns	ΔVIX	$\Delta 2Y$	$\Delta Slope$	$\Delta ILS2Y$	$\Delta ILS10Y$	$\Delta NEER$
Innovation	30-Nov-22	30-Nov-22	OpenAI ChatGPT (GPT-3.5)	+	+1.83***	+2.94***	-0.76***	-2.08***	+0.87***	+0.27***	+0.99***	-0.36***
Adoption	23-Jan-23	23-Jan-23	Microsoft 10bn OpenAI investment	+	+0.58***	+1.09***	-0.02	+1.47***	-0.82***	-0.33***	+0.54***	-0.55***
Innovation	14-Mar-23	14-Mar-23	OpenAI GPT-4 release	+	+1.69***	+1.51***	-1.54***	+3.60***	-2.21***	+1.71***	+3.15***	+0.34***
Policy	28-Mar-23	28-Mar-23	Italy bans ChatGPT	-	-0.73***	-0.28***	-0.34***	+1.66***	-1.67***	+1.20***	+1.48***	-1.45***
Adoption	24-May-23	24-May-23	NVIDIA Q1 FY24 earnings	+	+0.83***	-0.75***	+0.83***	+1.04***	-0.55***	-0.33***	-0.71***	+0.39***
Innovation	18-Jul-23	18-Jul-23	Meta Llama 2	+	+0.34***	+0.71***	-0.18***	+0.12***	-0.21***	-0.00	-0.91***	-0.24***
Adoption	23-Aug-23	23-Aug-23	NVIDIA Q2 FY24 earnings	+	+0.13***	+1.01***	-0.55***	-1.51***	-2.21***	+0.02	-0.13***	-0.89***
Adoption	25-Sep-23	25-Sep-23	Amazon 4bn Anthropic investment	+	+0.56***	+0.35***	-0.17***	-0.22***	+3.35***	-0.91***	-0.24***	+1.49***
Policy	30-Oct-23	30-Oct-23	Biden Executive Order on AI	?	+1.55***	+1.09***	-0.85***	+0.84***	+0.01	-0.41***	-0.12***	-0.57***
Innovation	06-Dec-23	06-Dec-23	Google Gemini	-	-0.48***	-0.44***	+0.07***	+0.62***	-2.49***	-1.09***	-1.08***	-0.66***
Policy	08-Dec-23	08-Dec-23	EU AI Act political agreement	?	+0.11***	+0.28***	-0.29***	+2.55***	-1.18***	+0.68***	+1.17***	+0.31***
Adoption	22-Feb-24	22-Feb-24	NVIDIA FY24 Q4 earnings	+	+2.99***	+1.91***	-0.44***	+1.04***	-1.10***	+0.20***	-0.88***	+0.15***
Innovation	18-Mar-24	18-Mar-24	NVIDIA Blackwell architecture	+	+1.36***	+0.51***	-0.03**	+0.17***	+0.56***	+0.05***	-0.01	+0.42***
Innovation	18-Apr-24	18-Apr-24	Meta Llama 3	-	-0.26***	-0.29***	-0.06***	+0.94***	-0.03	-0.06**	+0.19***	-0.14***
Innovation	13-May-24	14-May-24	OpenAI GPT-4o release	+	+0.11***	+0.38***	-0.09***	-0.89***	+0.28***	-0.73***	-0.30***	-0.28***
Adoption	23-May-24	23-May-24	NVIDIA Q1 FY25 earnings	+	+2.52***	-0.74***	+0.25***	+1.08***	-0.26***	+0.01	-1.39***	+0.51***
Adoption	10-Jun-24	10-Jun-24	Apple Intelligence announced	+	+0.28***	+0.24***	+0.30***	-0.04***	+1.11***	+0.26***	+0.11***	+0.89***
Adoption	29-Aug-24	29-Aug-24	NVIDIA Q2 FY25 earnings	?	-2.02***	-0.04**	-0.82***	+0.85***	-0.26***	+0.27***	+0.26***	+0.81***
Adoption	21-Jan-25	21-Jan-25	Stargate Project announced	+	+0.28***	+0.82***	-0.51***	+0.42***	-1.65***	-1.17***	-0.82***	-2.69***
Innovation	27-Jan-25	27-Jan-25	DeepSeek R1 release	-	-5.57***	-1.42***	+1.70***	-2.17***	+0.00	-0.36***	-0.71***	+1.49***
Innovation	24-Feb-25	24-Feb-25	Anthropic Claude 3.7 Sonnet	?	-1.19***	-0.49***	+0.43***	-1.30***	+1.12***	+0.07***	+0.41***	+0.11***
Adoption	27-Feb-25	27-Feb-25	NVIDIA FY25 Q4 earnings	?	-3.00***	-1.59***	+1.12***	+0.41***	+0.57***	+0.84***	+0.08***	+1.72***
Innovation	05-Apr-25	04-Apr-25	Meta Llama 4	?	+1.47***	-5.82***	+8.47***	-0.65***	-0.54***	-2.86***	-2.93***	+2.77***
Innovation	22-May-25	22-May-25	Anthropic Claude 4 Opus and Sonnet	+	+1.23***	-0.01	-0.34***	-0.00	-1.09***	+0.02	+0.09***	+0.76***
Adoption	29-May-25	29-May-25	NVIDIA Q1 FY26 earnings	+	+0.73***	+0.30***	-0.04***	-0.93***	-0.01	-0.62***	-0.03*	-1.18***
Innovation	10-Jun-25	10-Jun-25	OpenAI o3-pro release	-	-0.51***	+0.49***	-0.11***	-0.02	-0.55***	-0.52***	-0.41***	+0.35***
Policy	23-Jul-25	23-Jul-25	US AI Action Plan	+	+2.65***	+0.77***	-0.64***	+1.07***	+0.02	-0.72***	-0.73***	-0.31***

Continued on next page

Type	Event date	Trading date	Event	Exp.	AI factor	EQ returns	ΔVIX	$\Delta 2Y$	$\Delta Slope$	$\Delta ILS2Y$	$\Delta ILS10Y$	$\Delta NEER$
Policy	30-Jul-25	30-Jul-25	US restrictions AI chip exports	?	+1.09***	−0.20***	−0.25***	+1.63***	−1.12***	+0.19***	−0.07***	+1.12***
Adoption	28-Aug-25	28-Aug-25	NVIDIA Q2 FY26 earnings	+	+0.10***	+0.27***	−0.21***	+0.58***	−1.39***	+0.01	−1.14***	−1.29***
Policy	13-Jun-26	12-Jun-26	US export restrictions Anthropic	−	−2.01***	+0.44***	−0.95***	+0.78***	−0.29***	−0.74***	+0.72***	−1.69***

Notes: The table reports individual event-dummy coefficients from the event-study regressions, in standardised units. Stars denote statistical significance based on Newey–West HAC standard errors: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$. The column “Exp.” reports the ex ante expected sign of the AI-factor response. For events falling on non-trading days, the trading date is the nearest available trading day used in the regression. ILS2Y and ILS10Y denote the 2-year and 10-year inflation-linked swap rates, respectively.

B.2.1 Robustness checks

The event-study evidence in [Section 2.1](#) shows that the AI factor responds systematically, and in the expected direction, around a hand-collected set of AI-related announcements. A potential concern, however, is that the factor may respond more generally to salient technology or US mega-cap news, rather than specifically to information related to AI. This concern is particularly relevant given that several large technology firms enter the long leg of the portfolio.

We address this issue through a placebo-event exercise. We collect 26 major technology-related events that are economically important but whose primary content is not related to AI. The events, reported in [Table B.6](#), span the following categories: generic big-tech product announcements, non-AI digital and corporate innovations, broad technology policy measures, semiconductor-cycle news, and major technology disruptions. The chronology therefore deliberately includes announcements affecting firms and sectors represented in the AI factor, while excluding events whose main information content concerns AI capabilities, adoption, investment, or regulation.

To remain directly comparable with the baseline event study, we focus on the *signed average abnormal response* (SAAR). For placebo events for which a directional prior can be formulated ex ante, we define

$$D_t^P = \sum_{e \in \mathcal{E}^P} s_e \mathbf{1}\{t = \tau_e\}, \quad s_e \in \{-1, +1\}, \quad (\text{B.16})$$

where $\mathcal{E}^P$ denotes the set of signed placebo events, τ_e is the corresponding trading date, and s_e is the expected direction of the technology shock. We then re-estimate the specification in [Equation \(1\)](#), replacing the AI-event dummy with D_t^P . All variables are standardised, and the control variables and Newey–West HAC inference are identical to those used in the baseline event study.

For several generic product and corporate announcements, there is no unambiguous ex ante prediction for the direction of a *relative* AI long–short portfolio. These observations are therefore retained in the placebo chronology, but are assigned no sign and do not enter the SAAR. Importantly, we do not infer their sign from the realised movement of the factor, which

would introduce ex post selection bias. The signed placebo sample consequently contains 13 events: five semiconductor-cycle events, four broad technology-policy events, and four technology disruptions.

As an additional check, we construct a “market-moving” placebo subsample. An event is classified as market-moving whenever

$$\max \left\{ \left| \widetilde{r}_t^{EQ} \right|, \left| \Delta \widetilde{VIX}_t \right| \right\} \geq 0.75,$$

where the two variables are expressed in standard-deviation units. This selection criterion is deliberately based only on broad market returns and the VIX, and therefore does not use the AI factor itself. Eleven of the 26 placebo dates satisfy this criterion, of which seven have a directional prior and therefore enter the market-moving SAAR.

Placebo responses in the baseline financial system. [Table B.7](#) first reports the SAAR generated by the signed placebo dates for the same set of financial variables considered in the baseline event study, including inflation compensation at the 2- and 10-year horizons. None of the coefficients is statistically different from zero. In particular, the AI factor displays a placebo SAAR of only -0.03 standard deviations ($p = 0.91$). Broad equity returns, the VIX, short-term yields, the slope of the yield curve, inflation compensation, and the dollar exchange rate similarly display no systematic signed response on these non-AI technology dates. The placebo SAARs of the 2- and 10-year ILS rates are 0.41 and 0.23 standard deviations, respectively, and are both statistically insignificant ($p = 0.35$ in either case); in original units these correspond to approximately 2.50 and 0.71 basis points.

The absence of a coherent signed response stands in sharp contrast with the large and directionally consistent AI-factor responses documented around the AI events in [Section 2.1](#). The complementary absolute-response specification leads to the same conclusion. For the AI factor, the placebo absolute-response AAR is 0.06 standard deviations ($p = 0.59$), compared with 0.57 on AI-event dates. A Newey–West HAC Wald test rejects equality of the two absolute responses at the 5% level ($p = 0.032$). Thus, the full set of non-AI technology announcements does not generate systematically larger AI-factor movements.

Table B.6: Non-AI technology placebo events

Date	Category	Event	Exp.	Market-moving
07-Sep-2016	Big-tech product	Apple iPhone 7 launch	?	No
04-Oct-2016	Big-tech product	Google Pixel phone launch	?	No
16-Jun-2017	Digital/corporate	Amazon Whole Foods acquisition announced	?	No
12-Sep-2017	Big-tech product	Apple iPhone X launch	?	No
14-Dec-2017	Broad tech policy	FCC net-neutrality repeal vote	?	No
15-Nov-2018	Chip cycle	NVIDIA gaming/crypto inventory shock	—	Yes
18-Dec-2018	Chip cycle	Micron memory-cycle/inventory warning	—	No
19-Mar-2019	Digital/corporate	Google Stadia cloud-gaming platform announced	?	No
18-Jun-2019	Digital/corporate	Facebook Libra/Calibra digital wallet announced	?	Yes
13-Oct-2020	Big-tech product	Apple iPhone 12 launch	?	No
20-Oct-2020	Broad tech policy	US DOJ Google antitrust lawsuit	—	No
08-Jun-2021	Tech disruption	Fastly global internet outage	—	No
24-Jun-2021	Digital/corporate	Microsoft Windows 11 announced	?	No
04-Oct-2021	Tech disruption	Facebook/Instagram/WhatsApp global outage	—	Yes
28-Oct-2021	Digital/corporate	Facebook rebrands as Meta / metaverse strategy	?	Yes
18-Jan-2022	Digital/corporate	Microsoft Activision Blizzard acquisition announced	?	Yes
24-Mar-2022	Broad tech policy	EU Digital Markets Act political agreement	—	Yes
21-Jun-2022	Tech disruption	Cloudflare global outage	—	Yes
30-Jun-2022	Chip cycle	Micron PC/smartphone demand slowdown	—	Yes
09-Aug-2022	Broad tech policy	US CHIPS and Science Act signed	+	No
07-Sep-2022	Big-tech product	Apple iPhone 14 launch	?	Yes
06-Oct-2022	Chip cycle	AMD PC-demand revenue warning	—	Yes
26-Jan-2023	Chip cycle	Intel PC-cycle earnings shock	—	Yes
05-Jun-2023	Digital/corporate	Apple Vision Pro spatial-computing launch	?	No
06-Sep-2023	Broad tech policy	EU DMA gatekeeper designations	—	No
19-Jul-2024	Tech disruption	CrowdStrike Windows outage	—	No

Notes: The table reports the non-AI technology events used in the placebo exercise. “Exp.” denotes the ex ante expected sign of the generic technology shock for events with an unambiguous directional interpretation. Events marked “?” are retained in the placebo chronology but are excluded from the signed dummy and the SAAR calculation. “Market-moving” identifies events for which the maximum absolute standardised movement in broad equity returns or the VIX is at least 0.75 standard deviations. The AI factor is not used in defining this subsample.

Table B.7: Placebo-event SAAR across financial variables

Variable	Placebo SAAR	<i>p</i> -value
AI factor	−0.033	0.914
EQ returns	−0.125	0.649
Δ VIX	−0.102	0.542
Δ 2Y	+0.059	0.859
Δ Slope	−0.178	0.375
Δ ILS2Y	+0.412	0.355
Δ ILS10Y	+0.234	0.354
Δ NEER	−0.034	0.859

Notes: Coefficients are signed average abnormal responses (SAARs) in standard-deviation units. The signed placebo dummy includes the 13 events with an ex ante directional prior. *p*-values are based on Newey–West HAC standard errors using the same bandwidth rule as in the baseline event study. The ILS variables are daily changes in 2- and 10-year inflation compensation.

Removing semiconductor-market exposure. A related concern is that the factor may capture semiconductor-market developments rather than AI-specific news. [Section A.1.4](#) shows that the unconditional exposure of the benchmark factor to semiconductor-sector ETFs is small, with R^2 values of approximately 2.3% against **SMH** and 1.1% against **SOXX**. We complement that exercise here by repeating the placebo event study using semiconductor-orthogonalised factors. For $k \in \{\text{SMH}, \text{SOXX}\}$, we estimate

$$AI_t = a_k + b_k r_t^k + u_t^k \quad (\text{B.17})$$

over the full daily sample and define

$$AI_t^{\perp k} \equiv \hat{u}_t^k. \quad (\text{B.18})$$

The residual factors therefore remove the linear contemporaneous component of the benchmark AI factor associated with broad semiconductor returns. We standardise each residual factor before re-estimating the event-study regressions.

For comparison with the placebo responses, we additionally pool the signed AI events used in the baseline analysis into a single signed AI-event dummy and estimate its SAAR using the same specification. [Table B.8](#) reports the results.

Table B.8: AI versus placebo-event SAAR: semiconductor-orthogonalised factors

Factor	All signed placebos		Market-moving	AI events	
	SAAR	Hit rate	Placebo SAAR	AI SAAR	$p(\text{AI} = \text{P})$
Baseline	−0.033	0.54	+0.181	+1.208	0.001
Orthogonalised to SMH	+0.016	0.62	+0.260	+1.013	0.004
Orthogonalised to SOXX	+0.005	0.54	+0.232	+1.074	0.003

Notes: SAAR denotes the signed average abnormal response in standard-deviation units. The first two columns refer to the 13 signed placebo events. “Hit rate” is the fraction of signed placebo events for which the factor moves in the ex ante expected direction. “Market-moving” reports the SAAR for the seven signed placebo events belonging to the market-moving subset. The corresponding p -values are 0.712, 0.564, and 0.608 for the baseline, **SMH**-orthogonalised, and **SOXX**-orthogonalised factors, respectively. “AI SAAR” is obtained from a pooled signed dummy containing the AI events used in the baseline analysis. The final column reports the p -value of a Newey–West HAC Wald test of equality between the AI-event SAAR and the SAAR for the full signed placebo sample. The baseline factor uses 3,658 daily observations, while the semiconductor-orthogonalised specifications use the 2,904 observations for which the ETF-based factors are available.

The contrast is sharp. For the benchmark factor, the signed response on placebo dates is essentially zero, at −0.03 standard deviations, whereas the pooled response on AI-event dates equals 1.21 standard deviations. A Wald test strongly rejects equality of the two responses

($p = 0.001$). Moreover, selecting only placebo dates that generate sizeable movements in broad equities or volatility does not alter this conclusion: the market-moving placebo SAAR is 0.18 and remains statistically insignificant ($p = 0.71$).

The result is also robust to removing broad semiconductor exposure. The placebo SAAR equals 0.02 after orthogonalisation with respect to **SMH** and is essentially zero after orthogonalisation with respect to **SOXX**. In contrast, the corresponding AI-event SAARs remain sizeable, at 1.01 and 1.07 standard deviations. Equality between the AI and placebo responses is rejected in both cases, with p -values of 0.004 and 0.003, respectively. Thus, removing the component associated with broad semiconductor returns leaves the directional AI-event signal largely intact.

Placebo categories. [Table B.9](#) decomposes the signed placebo evidence by event category. Generic big-tech product announcements and non-AI digital/corporate innovations are shown for completeness but have no signed SAAR, since no directional prior was imposed. This is a deliberately conservative treatment: assigning a sign to such announcements after observing the relative performance of the AI portfolio would mechanically favour the placebo test. None of the three categories with an ex ante sign generates a statistically significant SAAR. The result is particularly informative for semiconductor-cycle events. Although chip-cycle announcements can generate sizeable movements in the long-short portfolio, their *direction* is not systematically captured by the AI factor. The absolute-response AAR for the baseline factor is positive and statistically significant (0.57 standard deviations, $p = 0.026$), whereas the signed semiconductor-cycle SAAR is -0.23 standard deviations and statistically insignificant ($p = 0.71$). The factor moves in the predicted direction on only two of five events. The corresponding hit rate remains 40% after orthogonalising the factor with respect to either **SMH** or **SOXX**. This distinction is useful: semiconductor news can move the portfolio substantially, but does not reproduce the directionally consistent response observed around AI events.

Overall assessment. Taken together, the placebo exercises provide a direct falsification test of the interpretation of the equity factor. First, non-AI technology events do not generate a systematic signed response of the factor, even when attention is restricted to episodes

Table B.9: Signed placebo-event responses by category

Category	N_s	Baseline		Orthog. SMH		Orthog. SOXX	
Broad tech policy	4	−0.066	(0.879)	+0.100	(0.824)	+0.055	(0.900)
Chip cycle	5	−0.229	(0.712)	−0.157	(0.789)	−0.171	(0.769)
Tech disruption	4	+0.244	(0.572)	+0.149	(0.707)	+0.176	(0.659)

Notes: Entries are signed average abnormal responses (SAARs); Newey–West HAC p -values are reported in parentheses. N_s denotes the number of events with an ex ante directional prior. Generic big-tech product and digital/corporate events are excluded from this table because no directional prior is imposed for these categories.

that produce sizeable movements in broad financial markets. The same is true for both 2- and 10-year inflation compensation, indicating that the ILS responses documented around AI events are not a generic feature of salient technology-news dates. Second, the contrast between AI and placebo events remains highly significant after removing the component of the factor associated with either SMH or SOXX. Third, the demanding semiconductor-cycle placebo shows that large chip-sector announcements can move the long–short portfolio, as is natural given its composition, but they do not reproduce the directionally consistent response observed around AI news. The evidence therefore supports the interpretation of the factor as capturing AI-specific information rather than generic mega-cap, broad technology, or semiconductor-cycle news.

C BVAR

C.1 Event-informed posterior weighting

The baseline identification combines contemporaneous sign and magnitude restrictions with narrative restrictions on the sign of the AI news shock around the event dates described in [Section 2.1](#). The latter restrictions discard structural decompositions that are inconsistent with the direction of the AI-related information revealed by these events. Conditional on satisfying these restrictions, however, all admissible structural decompositions would otherwise receive equal posterior weight. This may leave substantial identification uncertainty if the admissible set contains rotations in which the AI news shock plays only a limited role in explaining movements in the AI factor even on dates associated with major AI announcements.

To exploit more fully the information contained in the event chronology, we therefore complement the hard restrictions with a soft event-informed re-weighting of the admissible

structural decompositions. The procedure builds on the narrative-identification logic of [Antolín-Díaz and Rubio-Ramírez \(2018\)](#), but differs from a standard narrative restriction in an important respect: rather than imposing an additional zero-one restriction, we use the historical decomposition on the selected event dates to assign different posterior weights to otherwise admissible rotations.

Structural contributions on event dates. Let $m = 1, \dots, M$ index posterior draws of the reduced-form BVAR together with an admissible structural rotation, and let

$$u_t = B_0^{(m)} \varepsilon_t^{(m)} \quad (\text{C.19})$$

denote the corresponding structural representation, where $\varepsilon_t^{(m)}$ is the vector of mutually orthogonal structural shocks and $B_0^{(m)}$ is the contemporaneous impact matrix. Shock $s = 1$ is the identified AI news shock.

For a variable i , the contribution of structural shock s to the unexpected movement on date t is

$$c_{i,s,t}^{(m)} = B_{0,is}^{(m)} \varepsilon_{s,t}^{(m)}. \quad (\text{C.20})$$

We focus on the contribution to the AI factor, corresponding to $i = AI$. For each event $e \in \mathcal{E}$ occurring at date τ_e , we therefore compute

$$c_{s,e}^{(m)} \equiv c_{AI,s,\tau_e}^{(m)}. \quad (\text{C.21})$$

Importantly, these contributions are computed for all structural shocks, not only for the three economically labelled shocks. Hence,

$$\sum_{s=1}^N c_{s,e}^{(m)} = u_{AI,\tau_e}^{(m)}, \quad (\text{C.22})$$

so that the historical decomposition exactly reproduces the reduced-form innovation in the AI factor draw by draw.

AI-event dominance measure. For each event and structural draw, we measure the extent to which the AI shock dominates competing explanations of the AI-factor innovation. Define

$$D_e^{(m)} = \frac{\left| c_{1,e}^{(m)} \right| - \max_{s \neq 1} \left| c_{s,e}^{(m)} \right|}{\sum_{s=1}^N \left| c_{s,e}^{(m)} \right|}. \quad (\text{C.23})$$

The statistic is scale invariant and compares the absolute contribution of the AI news shock with that of the largest competing structural disturbance. Positive values indicate that the AI news shock is the largest contributor to the AI-factor innovation on event date e , while negative values indicate that another structural shock accounts for a larger contribution. The normalisation by total absolute structural contributions prevents unusually large event-day innovations from mechanically receiving disproportionate weight.

For descriptive purposes, we also report the AI share of the absolute innovation,

$$S_e^{(m)} = \frac{\left| c_{1,e}^{(m)} \right|}{\sum_{s=1}^N \left| c_{s,e}^{(m)} \right|}. \quad (\text{C.24})$$

This statistic is used only to characterise the historical decomposition; the posterior weighting itself is based on the relative dominance measure in [Equation \(C.23\)](#).

Soft event likelihood. Rather than requiring $D_e^{(m)} > 0$ as an additional hard narrative restriction, we map the dominance statistic into a smooth score,

$$\pi_e^{(m)} = \Lambda(\kappa D_e^{(m)}) = \frac{1}{1 + \exp\left[-\kappa D_e^{(m)}\right]}, \quad (\text{C.25})$$

where $\kappa > 0$ governs the strength with which the event information discriminates across admissible structural decompositions. When $D_e^{(m)} = 0$, the event receives a score of $1/2$. The score approaches one as the AI news shock becomes increasingly dominant and approaches zero when a competing shock becomes increasingly dominant.

This specification deliberately preserves uncertainty about the structural interpretation

of individual announcement days. In particular, the procedure does not require the AI news shock to explain the majority of the observed innovation, nor does it exclude decompositions in which other disturbances also make economically important contributions.

Combining the information across the $E = |\mathcal{E}|$ selected event dates, we construct the event score for draw m as the geometric mean

$$\mathcal{L}_{\mathcal{E}}^{(m)} = \left[\prod_{e \in \mathcal{E}} \pi_e^{(m)} \right]^{1/E}, \quad (\text{C.26})$$

or, equivalently,

$$\log \mathcal{L}_{\mathcal{E}}^{(m)} = \frac{1}{E} \sum_{e \in \mathcal{E}} \log \pi_e^{(m)}. \quad (\text{C.27})$$

The $1/E$ tempering is important because it prevents posterior concentration from increasing mechanically with the number of events included in the chronology. The event score should therefore be interpreted as a probabilistic ranking of admissible decompositions rather than as a fully specified sampling likelihood for the event-study observations.

Posterior weights. Let $\mathcal{A}$ denote the set of BVAR draws satisfying the baseline contemporaneous and narrative restrictions. Conditional on belonging to $\mathcal{A}$, the conventional sign-restricted posterior assigns equal weight to the admissible draws. We instead assign the normalised event-informed weights

$$\omega_m = \frac{\mathcal{L}_{\mathcal{E}}^{(m)}}{\sum_{j \in \mathcal{A}} \mathcal{L}_{\mathcal{E}}^{(j)}}, \quad m \in \mathcal{A}, \quad (\text{C.28})$$

and set $\omega_m = 0$ for draws outside the admissible set. Thus, the hard restrictions continue to determine the support of the structural posterior, whereas the event information affects the relative probability mass within that support.

For any posterior object $g^{(m)}$, including structural shocks, impulse responses, historical decompositions and forecast-error variance decompositions, posterior summaries are subsequently computed using the weights in [Equation \(C.28\)](#). For example, the event-informed

posterior distribution of an impulse response satisfies

$$\Pr_{\omega}(IRF_h \leq x) = \sum_{m \in \mathcal{A}} \omega_m \mathbf{1}\left\{IRF_h^{(m)} \leq x\right\}. \quad (\text{C.29})$$

Medians and posterior intervals are obtained from the corresponding weighted quantiles.

Choice of the weighting parameter. Our baseline calibration sets $\kappa = 10$. The parameter is selected using only first-stage BVAR diagnostics, prior to examining the downstream macroeconomic responses. With $M = 10,000$ admissible structural draws, this calibration yields an effective sample size of approximately 7,288 draws, or 72.9% of the original posterior support. The median width of the 68% posterior interval for the daily AI news shock falls by approximately 10%, while the correlation between the event-weighted and equal-weight posterior-median daily shock series remains 0.99. Hence, the event information reduces structural identification uncertainty without materially altering the chronology of the identified AI news shock. As sensitivity checks, we consider weaker values $\kappa = 5$ and $\kappa = 2.5$. These imply progressively less posterior reweighting and retain a larger effective share of the original admissible set. The exercise therefore allows us to verify that the results are not driven by an excessively concentrated posterior or by a small number of structural rotations.

C.2 Validation against external AI release dates

This subsection complements the validation of the raw AI factor by examining the identified structural shock around the model-release dates collected by [Christensen and Rudebusch \(2026\)](#) (C&R). We first inspect daily changes in the AI factor around each release and apply the same 1.5-standard-deviation salience criterion used for narrative selection in [Section 3](#). We then assess the identified shock around qualifying dates that were not imposed as narrative restrictions. The chronology is external, but the validation sample is conditional on the raw-factor salience screen.

[Table C.10](#) reports the full set of 26 C&R dates. For events associated with a qualifying AI-factor move within one calendar day of the release, the table reports that matched daily move and its date. For events that do not satisfy the screen, the table reports the AI-factor change on the event trading day. The final two columns indicate whether the event was

already included in the BVAR narrative set and, for events that were not included, whether it would independently satisfy the same 1.5-standard-deviation screening rule.

Table C.10: C&R events and raw AI-factor movements

Event	C&R date	Move date	ΔAI (%)	In BVAR	1.5 s.d.
OpenAI ChatGPT (GPT-3.5)	30-Nov-2022	30-Nov-2022	+2.03	Yes	–
Google Bard	06-Feb-2023	07-Feb-2023	+2.14	No	Yes
OpenAI GPT-4 / Anthropic Claude 1	14-Mar-2023	14-Mar-2023	+1.84	Yes	–
Anthropic Claude 2	11-Jul-2023	10-Jul-2023	-2.33	No	Yes
xAI Grok 1	03-Nov-2023	03-Nov-2023	+0.21	No	No
Anthropic Claude 2.1	21-Nov-2023	21-Nov-2023	+0.15	No	No
Google Gemini 1.0 Nano and Pro	06-Dec-2023	06-Dec-2023	-0.47	No	No
Google Gemini 1.5 Pro	15-Feb-2024	15-Feb-2024	-0.74	No	No
Anthropic Claude 3	04-Mar-2024	04-Mar-2024	-2.09	No	Yes
OpenAI GPT-4o	13-May-2024	13-May-2024	-1.13	No	No
xAI Grok 1.5	15-May-2024	15-May-2024	+0.81	No	No
Anthropic Claude 3.5 Sonnet	20-Jun-2024	20-Jun-2024	-2.15	No	Yes
xAI Grok 2	13-Aug-2024	13-Aug-2024	+0.67	No	No
OpenAI o1-preview	12-Sep-2024	11-Sep-2024	+2.16	No	Yes
Google Gemini 2.0 Flash	11-Dec-2024	11-Dec-2024	+2.69	No	Yes
OpenAI o3-mini	31-Jan-2025	30-Jan-2025	-3.69	No	Yes
xAI Grok 3	17-Feb-2025	18-Feb-2025	-2.08	No	Yes
Anthropic Claude 3.7 Sonnet	24-Feb-2025	24-Feb-2025	-1.24	No	No
OpenAI GPT-4.5	27-Feb-2025	26-Feb-2025	+1.95	No	Yes
Google Gemini 2.5 Pro	25-Mar-2025	26-Mar-2025	-3.38	No	Yes
OpenAI GPT-4.1	14-Apr-2025	14-Apr-2025	-1.90	No	Yes
Anthropic Claude 4	22-May-2025	22-May-2025	+1.36	No	No
xAI Grok 4	09-Jul-2025	09-Jul-2025	+1.65	No	No
Anthropic Claude Opus 4.1	05-Aug-2025	05-Aug-2025	+0.06	No	No
OpenAI GPT-5	07-Aug-2025	07-Aug-2025	+0.79	No	No
Anthropic Claude Sonnet 4.5	29-Sep-2025	29-Sep-2025	+0.12	No	No

Notes: Move date is the date of the reported daily AI-factor change, ΔAI . In BVAR indicates inclusion in the narrative restrictions; 1.5 s.d. indicates whether an external event passes the salience screen. – denotes an event already included in the narrative set.

The raw-data screen selects 13 of the 26 C&R events. Two of these—the ChatGPT launch and the GPT-4 / Claude 1 release date—were already part of the BVAR narrative set. Removing these overlaps leaves 11 salient release dates that were not imposed as narrative restrictions.

Passing the salience screen was necessary but not sufficient for inclusion in the original event study and BVAR narrative selection. We did not retain these releases as separate identifying events when their market impact was difficult to distinguish from other news or when events occurred too close together to permit a clear attribution. For example, Nasdaq-100 rebalancing news ahead of Claude 2 and Microsoft and IBM earnings announcements

ahead of o3-mini may have contributed to the qualifying preceding-day AI-factor movements. The GPT-4.5 release on 27 February 2025 also coincided with the NVIDIA earnings-related trading day already included in the event study. These selection concerns do not prevent us from examining the identified shock around the omitted releases as a supplementary validation exercise, but they qualify its interpretation.

C.2.1 Detection of shocks around the retained releases

The next step is to check whether the identified daily AI shock is also unusually large around the 11 salient C&R events that were not used as narrative restrictions. We consider the event day itself and short post-event windows extending one and two trading days after the event. [Table C.11](#) reports the posterior median AI shock over the three windows and the highest percentile reached by its absolute value relative to the corresponding full-sample non-event distribution. [Table C.12](#) then reports the same validation across the full set of admissible, event-weighted BVAR draws, summarising for each event the posterior probability of exceeding the 75th or 90th percentile in at least one of the three windows.

Table C.11: Posterior median AI shocks around external C&R events

Event	Event day	0 to +1	0 to +2	Max. abs. percentile	Above 75th	Above 90th
Google Bard	1.51	2.95	1.95	97.2	Yes	Yes
Anthropic Claude 2	-0.06	-0.22	0.60	34.2	No	No
Anthropic Claude 3	-1.32	-1.78	-2.56	90.9	Yes	Yes
Anthropic Claude 3.5 Sonnet	-1.11	-1.93	-3.60	97.4	Yes	Yes
OpenAI o1-preview	1.23	0.02	-1.39	85.4	Yes	No
Google Gemini 2.0 Flash	1.72	0.74	1.71	94.5	Yes	Yes
OpenAI o3-mini	0.25	-0.18	-0.70	39.1	No	No
xAI Grok 3	-1.02	-0.80	-2.10	83.9	Yes	No
OpenAI GPT-4.5	-1.75	-0.28	-2.74	94.8	Yes	Yes
Google Gemini 2.5 Pro	0.17	-2.17	-2.52	91.8	Yes	Yes
OpenAI GPT-4.1	-1.95	-2.02	-4.67	99.3	Yes	Yes

Notes: The first three numerical columns report the posterior median AI shock on the event trading day and cumulated through one and two subsequent trading days. “Max. abs. percentile” is the highest percentile reached by the absolute posterior median shock across the three windows relative to the corresponding full-sample non-event distribution.

The event-level evidence is strong for most of the external dates. Nine of the 11 events produce a posterior median AI shock above the 75th percentile in at least one of the three short windows, and seven of the 11 exceed the 90th percentile. The strongest cases include Bard, Claude 3.5 Sonnet, Gemini 2.0 Flash, GPT-4.5 and GPT-4.1. Two events, Claude 2 and o3-mini, are not associated with a particularly large identified AI shock; the attribution

Table C.12: Posterior probability of detecting a large AI shock

Event	Above 75th percentile in any window	Above 90th percentile in any window
Google Bard	99.3%	85.0%
Anthropic Claude 2	28.8%	6.6%
Anthropic Claude 3	87.2%	49.4%
Anthropic Claude 3.5 Sonnet	99.9%	92.0%
OpenAI o1-preview	86.2%	14.3%
Google Gemini 2.0 Flash	88.5%	61.1%
OpenAI o3-mini	3.3%	0.0%
xAI Grok 3	71.2%	21.4%
OpenAI GPT-4.5	92.4%	68.4%
Google Gemini 2.5 Pro	84.9%	45.9%
OpenAI GPT-4.1	98.0%	89.8%

Notes: Probabilities are calculated across the full set of admissible, event-weighted BVAR draws. An event is counted as detected when the absolute AI shock exceeds the indicated percentile in at least one of the event-day, one-day post-event, or two-day post-event windows.

concerns discussed above are especially relevant in these cases.

Table C.13 summarises the successive screening steps and the main cross-validation statistics. The first two rows show that half of the original C&R chronology passes the raw AI-factor salience screen. After removing the two dates already used in the BVAR narrative restrictions, 11 observations remain for the external validation. The posterior median shock catches more than four-fifths of these events above the 75th percentile and almost two-thirds above the 90th percentile.

Across the event weighted BVAR draws, the median share of retained events exceeding the 75th-percentile threshold is 81.8%, with a 68% range of 63.6–81.8%. At the more demanding 90th-percentile threshold, the median share is 45.5%, with a 68% range of 36.4–63.6%. Thus, the strong detection rate at the lower threshold also appears across BVAR draws, while detection at the higher threshold is more sensitive to identification.

Overall, the exercise provides supporting evidence that the identified shock detects salient releases beyond the imposed narrative dates, particularly at the 75th-percentile threshold. Two features qualify this interpretation. First, the screening window includes preceding-day moves, whereas the shock-validation windows begin on the event trading day; this distinction matters for releases such as Claude 2 and o3-mini. Second, the detection shares condition on large raw-factor movements and allow exceedance in any of three windows. They are descriptive validation statistics, not rejection rates from a test calibrated for this selection

Table C.13: Summary of the C&R cross-validation exercise

Statistic	Result
C&R chronology	26
Salient events	13
Narrative overlap	2
External events (N_E)	11
Posterior median detection: $\hat{s}_{0.75}$	9 of 11 (81.8%)
Posterior median detection: $\hat{s}_{0.90}$	7 of 11 (63.6%)
Posterior median of $s_{0.75}^{(d)}$	81.8%
68% posterior range of $s_{0.75}^{(d)}$	63.6–81.8%
Posterior median of $s_{0.90}^{(d)}$	45.5%
68% posterior range of $s_{0.90}^{(d)}$	36.4–63.6%

Notes: Salient events pass the 1.5-standard-deviation raw AI-factor screen; narrative overlap counts those already used as BVAR narrative restrictions. N_E is the remaining external sample. $\hat{s}_q$ is the share of these events whose absolute posterior median AI shock exceeds the corresponding non-event q -quantile in at least one window: event day, 0 to +1, or 0 to +2 trading days. $s_q^{(d)}$ is the analogous share within posterior draw d ; its median and 68% range use the event-weighted posterior draws.


and multiple-window procedure.

C.3 Additional results

C.3.1 Forecast error variance decomposition

The FEVD provides a useful validation of the structural identification. On impact, the identified AI news shock accounts for 55.8% of the forecast-error variance of the AI factor, while its contribution to broader financial variables is considerably smaller, ranging from 3.7% for the VIX to 16.3% for the NEER and reaching 10.9% for equity returns. By contrast, risk-on and monetary policy shocks play a substantially larger role in explaining movements in yields, volatility, and exchange rates, while the residual group of shocks accounts for around 40–55% of their variance. This pattern suggests that the identified AI news shock captures a relatively specific source of variation rather than a generic financial-market disturbance. The decomposition is essentially unchanged at the six-day horizon—for example, the AI contribution falls only from 55.8% to 55.6% for the AI factor and from 16.3% to 16.1% for the NEER—consistent with the high-frequency nature of the BVAR, in which the bulk of the forecast-error variance is determined on impact and subsequent daily propagation is limited.

Figure C.2: Forecast error variance decomposition



Notes: Forecast error variance decomposition of the endogenous variables at horizons 0 (Panel a) and 6 (Panel b).


C.3.2 Historical decomposition

As an additional diagnostic of the identification, we report the historical decomposition of unexpected movements in the AI factor and broad equity returns around the major events reported in Table 4. The results reinforce the visual evidence in Panel a of Figure 3, where the largest realisations of the identified AI news shock tend to coincide with salient AI-related events. In particular, the AI news shock accounts for most movements in the AI factor on event dates and is especially important around the largest positive and negative episodes, whereas broad equity innovations receive more substantial contributions from generic risk-on and monetary policy shocks. The decomposition therefore supports the interpretation that the long-short factor contains information specific to AI developments, while the response of the broader equity market reflects a mixture of AI-related and more general financial forces.

C.4 Predictability of the monthly AI news shock

As an additional validation exercise, we assess whether lagged US macro-financial conditions predict the identified monthly AI news shock. For each retained structural BVAR draw d , we

Figure C.3: Historical decompositions around major events
(a) AI factor (b) Equity returns



Notes: Historical decomposition of the changes in the AI factor and equity returns on event days reported in Table 4.

estimate

$$\varepsilon_t^{AI,(d)} = \alpha^{(d)} + X'_{t-1}\beta^{(d)} + u_t^{(d)}, \quad (C.30)$$

where X_{t-1} contains the standardised values of six US macro-financial predictors: r_{t-1}^{SP} , the previous month's S&P 500 return; $\log VIX_{t-1}$, the log VIX at the end of the previous month; i_{t-1}^{2y} and i_{t-1}^{10y} , the corresponding end-of-month 2-year and 10-year Treasury yields; and g_{t-1}^{IP} and π_{t-1}^{CPI} , industrial production growth and CPI inflation for the previous month. We also consider an augmented specification that controls for the lagged AI shock,

$$\varepsilon_t^{AI,(d)} = \alpha^{(d)} + \rho^{(d)}\varepsilon_{t-1}^{AI,(d)} + X'_{t-1}\beta^{(d)} + u_t^{(d)}. \quad (C.31)$$

We estimate the regressions separately for each posterior shock history and summarise the coefficient estimates using the event-informed weights employed in the local projections. These summaries capture variation across shock histories, not the estimation uncertainty within each regression. For the latter, we use Newey–West standard errors with three lags.

Table C.14 reports the weighted 16th, 50th (median) and 84th percentiles of the coefficient

estimates across shock histories. The regressions explain little shock variation: the weighted median R^2 is 2.7% in the baseline and 3.3% when the lagged shock is included. The latter has a median coefficient of 0.067.

Table C.14: Predictability of the identified AI shock

Regressor	16 th pct.	Median	84 th pct.
<i>Lagged macro-financial predictors</i>			
Constant	-0.0035	0.0007	0.0050
r_{t-1}^{SP}	-0.0036	-0.0024	-0.0010
$\log VIX_{t-1}$	-0.0076	-0.0058	-0.0035
i_{t-1}^{2y}	-0.0006	0.0097	0.0199
i_{t-1}^{10y}	-0.0224	-0.0122	-0.0019
g_{t-1}^{IP}	0.0004	0.0032	0.0058
π_{t-1}^{CPI}	-0.0071	-0.0022	0.0026
R^2	0.0139	0.0268	0.0528
<i>Including the lagged AI shock</i>			
Constant	-0.0034	0.0006	0.0046
ε_{t-1}^{AI}	0.0187	0.0672	0.1249
r_{t-1}^{SP}	-0.0046	-0.0029	-0.0012
$\log VIX_{t-1}$	-0.0075	-0.0058	-0.0034
i_{t-1}^{2y}	0.0002	0.0095	0.0188
i_{t-1}^{10y}	-0.0212	-0.0118	-0.0022
g_{t-1}^{IP}	0.0002	0.0030	0.0056
π_{t-1}^{CPI}	-0.0066	-0.0015	0.0034
R^2	0.0189	0.0334	0.0633

Notes: The table reports event-weighted posterior percentiles across retained structural BVAR shock histories. Macro-financial regressors are dated $t-1$. Treasury yields are measured at the end of the previous month. Industrial production growth and CPI inflation refer to the previous monthly observation. Predictors are standardized prior to estimation, while the lagged AI shock remains in its original units. Newey–West standard errors with three lags are used within each draw.

The results provide little evidence of systematic predictability. As shown in Table C.15, the weighted probability that the joint null of no predictive content from the macro-financial variables is rejected at the 5% level is only 0.9% in both specifications. Even at the 10% level, the corresponding rejection probabilities remain below 4%. Individual predictors also display very limited rejection frequencies. The lagged AI shock is significant at the 5% level in only

7.7% of the weighted posterior draws.

Table C.15: Posterior rejection probabilities

	5% level	10% level
Joint test, baseline specification	0.009	0.035
Joint test, including ε_{t-1}^{AI}	0.009	0.037
ε_{t-1}^{AI}	0.077	0.145
<i>Individual 5% rejection probabilities</i>		
r_{t-1}^{SP}	0.000	
$\log VIX_{t-1}$	0.000	
i_{t-1}^{2y}	0.039	
i_{t-1}^{10y}	0.044	
g_{t-1}^{IP}	0.002	
π_{t-1}^{CPI}	0.028	

Notes: The first three rows report the event-weighted fraction of posterior shock histories for which the corresponding null hypothesis is rejected. The lower panel reports individual 5% rejection probabilities for the baseline specification. For comparison, in the specification including the lagged AI shock, the corresponding probabilities are 0.000, 0.000, 0.036, 0.044, 0.003, and 0.030, respectively.

Overall, the exercise suggests that the identified AI shock is largely unpredictable from lagged US macro-financial conditions. The very low explanatory power, the small posterior probability of joint rejection, and the limited serial predictability of the shock itself support its interpretation as an exogenous innovation at the monthly frequency.

D Local projections

D.1 Bayesian smooth local projections

This section describes the Bayesian smooth local projection procedure used to estimate the macroeconomic responses to the AI news shock. The estimator has two objectives. First, it regularises the impulse response across adjacent horizons, which is particularly useful at long horizons, where standard local projections are imprecisely estimated. Second, it propagates uncertainty about the structural shock estimated in the daily BVAR through the local projection stage.

D.1.1 Horizon-specific local projections

For a given realisation r of the monthly AI news shock, we estimate the sequence of horizon-specific regressions

$$y_{t+h} - y_{t-1} = \alpha_h + \beta_h^{(r)} \varepsilon_t^{\text{AI},(r)} + \mathbf{x}_t^{(r)'} \boldsymbol{\gamma}_h^{(r)} + u_{t+h,h}^{(r)}, \quad h = 0, \dots, H, \quad (\text{D.32})$$

where $\mathbf{x}_t^{(r)}$ contains lags of the AI news shock, lags of the dependent variable and the remaining controls (COVID dummy). The coefficient of interest $\beta_h^{(r)}$ is left unrestricted in sign and magnitude at every horizon. The baseline estimates cumulative responses through the dependent variable $y_{t+h} - y_{t-1}$. Let

$$\mathbf{m}^{(r)} = \left(m_0^{(r)}, \dots, m_H^{(r)} \right)' \quad (\text{D.33})$$

denote the vector of horizon-specific estimates of the impulse-response coefficient and let $\mathbf{V}^{(r)}$ denote the corresponding estimation uncertainty. We employ the Gaussian approximation

$$\mathbf{m}^{(r)} \mid \boldsymbol{\beta}^{(r)} \sim \mathcal{N} \left(\boldsymbol{\beta}^{(r)}, \mathbf{V}^{(r)} \right), \quad (\text{D.34})$$

where

$$\boldsymbol{\beta}^{(r)} = \left(\beta_0^{(r)}, \dots, \beta_H^{(r)} \right)'. \quad (\text{D.35})$$

The nuisance coefficients entering [Equation \(D.32\)](#) are regularised using Minnesota-type shrinkage, with prior tightness declining with the lag order. The contemporaneous AI-shock coefficient, which defines the impulse response, is not subject to this nuisance-parameter shrinkage.

D.1.2 Second-order random-walk (RW2) prior

Estimating each β_h independently can generate considerable horizon-to-horizon sampling variation, especially because the effective sample size declines with h . We therefore place a second-order random-walk prior on the impulse-response sequence. Let $\mathbf{D}_2$ denote the second-difference matrix, such that

$$\mathbf{D}_2 \boldsymbol{\beta}^{(r)} = \begin{pmatrix} \beta_2^{(r)} - 2\beta_1^{(r)} + \beta_0^{(r)} \\ \vdots \\ \beta_H^{(r)} - 2\beta_{H-1}^{(r)} + \beta_{H-2}^{(r)} \end{pmatrix}. \quad (\text{D.36})$$

The smoothing prior is

$$\mathbf{D}_2 \boldsymbol{\beta}^{(r)} \mid q_r^2 \sim \mathcal{N}(\mathbf{0}, q_r^2 \mathbf{I}). \quad (\text{D.37})$$

Equivalently, the smoothing precision is $\lambda_r = 1/q_r^2$. Small values of q_r^2 imply a smoother response, whereas larger values allow greater curvature across horizons. Importantly, [Equation \(D.37\)](#) does not impose the sign, level, monotonicity or location of the peak of the impulse response; it penalises only abrupt changes in its slope.

Rather than fixing the amount of smoothing, we estimate it hierarchically. We specify

$$q_r^2 \sim \mathcal{IG}(a_q, b_q), \quad (\text{D.38})$$

where the baseline prior uses $a_q = 2.5$ and is parameterised to imply $\mathbb{E}[q_r^2] = 1$. Conditional on q_r^2 , the posterior for $\boldsymbol{\beta}^{(r)}$ is Gaussian. Combining [Equation \(D.34\)](#) and [Equation \(D.37\)](#), its precision matrix is

$$\mathbf{P}_r = \mathbf{V}^{(r)-1} + \frac{1}{q_r^2} \mathbf{D}_2' \mathbf{D}_2, \quad (\text{D.39})$$

and therefore

$$\boldsymbol{\beta}^{(r)} \mid q_r^2, \mathbf{m}^{(r)} \sim \mathcal{N}(\boldsymbol{\mu}_r, \mathbf{P}_r^{-1}), \quad (\text{D.40})$$

$$\boldsymbol{\mu}_r = \mathbf{P}_r^{-1} \mathbf{V}^{(r)-1} \mathbf{m}^{(r)}. \quad (\text{D.41})$$

The conditional posterior for q_r^2 is inverse-Gamma,

$$q_r^2 \mid \boldsymbol{\beta}^{(r)} \sim \mathcal{IG}\left(a_q + \frac{H-1}{2}, b_q + \frac{1}{2} \left\| \mathbf{D}_2 \boldsymbol{\beta}^{(r)} \right\|^2\right), \quad (\text{D.42})$$

which gives a simple Gibbs sampler alternating between [Equation \(D.41\)](#) and [Equation \(D.42\)](#).

Because the first-stage BVAR itself produces a distribution of admissible shock histories, learning the smoothing parameter from a single arbitrarily selected structural draw would be undesirable. We therefore calibrate the hierarchical distribution using a representative subset of the posterior shock histories and pool the resulting draws of q^2 with equal weight. In the baseline implementation we use 500 shock histories for this calibration, run 300 Gibbs iterations for each calibration draw, and discard the first 150 iterations as burn-in. The resulting hyperposterior is then propagated through the full local projection exercise.²⁴

D.1.3 Propagation of first- and second-stage uncertainty

The AI news shock entering [Equation \(D.32\)](#) is not observed. Let

$$\mathcal{E} = \{\boldsymbol{\varepsilon}^{(r)}, \omega_r\}_{r=1}^M \quad (\text{D.43})$$

denote the event-informed posterior distribution of daily structural AI-shock histories obtained from the BVAR, where ω_r is the posterior weight attached to draw r . Each daily shock history is first aggregated to the monthly frequency,

$$\boldsymbol{\varepsilon}_m^{(r)} = \sum_{t \in m} \boldsymbol{\varepsilon}_t^{(r)}, \quad (\text{D.44})$$

before any posterior summary is taken. This is important because aggregation and posterior quantiles do not commute. The event weights are preserved after monthly aggregation.

²⁴The posterior estimates of the hyperparameters are reported in [Table D.16](#) below.

Since the downstream Bayesian local projection routine operates on equally weighted shock histories, we transform the weighted posterior into an equally weighted sample using systematic importance resampling. If R histories are required, the resulting indices

$$r_1, \dots, r_R \sim_{\text{SIR}} \{\omega_1, \dots, \omega_M\} \quad (\text{D.45})$$

are distributed according to the event-informed posterior probabilities. Higher-weight structural rotations consequently enter the LP stage more often, while low-weight rotations remain represented according to their posterior mass. Repeated indices are therefore an intentional feature of the importance resampling rather than duplicate observations to be removed.

For every resampled shock history r , we estimate Equation (D.32) and generate B draws

$$\boldsymbol{\beta}^{(r,b)} \sim p(\boldsymbol{\beta} \mid \boldsymbol{\varepsilon}^{(r)}, \mathbf{Y}), \quad b = 1, \dots, B. \quad (\text{D.46})$$

The unconditional posterior of the impulse response can therefore be written using the law of total probability as

$$\begin{aligned} p(\boldsymbol{\beta} \mid \mathbf{Y}) &= \int p(\boldsymbol{\beta} \mid \boldsymbol{\varepsilon}, \mathbf{Y}) p(\boldsymbol{\varepsilon} \mid \mathcal{I}_{\text{BVAR}}) d\boldsymbol{\varepsilon} \\ &\simeq \frac{1}{R} \sum_{r=1}^R p(\boldsymbol{\beta} \mid \boldsymbol{\varepsilon}^{(r)}, \mathbf{Y}), \end{aligned} \quad (\text{D.47})$$

where $\mathcal{I}_{\text{BVAR}}$ denotes the information and identifying restrictions used in the first-stage BVAR. The Monte Carlo approximation to Equation (D.47) is consequently

$$\left\{ \boldsymbol{\beta}^{(r,b)} : r = 1, \dots, R, \ b = 1, \dots, B \right\}. \quad (\text{D.48})$$

Variation across r captures first-stage uncertainty concerning the structural decomposition and the estimated AI-shock history. Variation across b captures second-stage uncertainty concerning the local projection coefficients conditional on that history. Posterior medians and pointwise highest-density posterior sets are computed directly from this pooled distribution.

This treatment is important for inference because the structural shock is a generated

regressor (Pagan, 1984). A conventional two-step procedure that first replaces the latent shock by a single posterior estimate and subsequently treats that estimate as observed would condition on the first-stage estimate and discard uncertainty about the structural decomposition. Our procedure instead marginalises over this first-stage uncertainty. Thus, while the estimation remains modular rather than a fully joint BVAR–LP model, the reported uncertainty bands incorporate both uncertainty about the generated shock and conditional LP estimation uncertainty. In the baseline specification we use $R = 10,000$ event-weighted monthly shock histories and retain $B = 50$ posterior local projection draws for each shock history, yielding up to 500,000 draws from the mixture posterior at each horizon. The baseline horizon is 60 months and the model contains twelve lags.

D.1.4 Posterior estimates of the hierarchical RW2 hyperparameters

As discussed above, the Bayesian local projections impose a hierarchical second-order random-walk (RW2) prior on the sequence of shock-response coefficients across horizons. The degree of smoothing is not fixed ex ante, but is estimated from the data through the hyperparameter q^2 , with the corresponding smoothing precision given by $\lambda = 1/q^2$. Higher values of λ correspond to stronger smoothing across horizons. Table D.16 reports the posterior estimates of these hyperparameters for the baseline specification ($p = 12$), separately for US aggregates, euro area aggregates, and the US sectoral-network responses.

The estimates are relatively homogeneous across variables and applications. Posterior median values of λ are generally concentrated in a fairly narrow range, implying similar but not identical degrees of smoothing across outcomes. Production and real-activity variables tend to imply somewhat lower smoothing precision than price variables, while there is no evidence that the hierarchical procedure selects extreme smoothing parameters for any particular series. This provides reassurance that the shape of the estimated impulse responses is not driven by markedly different or highly restrictive smoothing assumptions across variables.

Table D.16: Posterior estimates of hierarchical RW2 hyperparameters

Variable	λ	q^2
<i>US aggregates</i>		
PCE core	6.85 [5.11, 8.98]	0.146 [0.111, 0.196]
CPI headline	6.64 [4.94, 8.74]	0.151 [0.114, 0.202]
CPI core	6.70 [4.98, 8.79]	0.149 [0.114, 0.201]
PPI headline	6.64 [4.95, 8.73]	0.151 [0.115, 0.202]
Industrial production	5.42 [3.92, 7.33]	0.185 [0.136, 0.255]
Real GDP	5.40 [3.91, 7.31]	0.185 [0.137, 0.255]
Employment	5.82 [4.24, 7.82]	0.172 [0.128, 0.236]
Hours	6.03 [4.36, 8.06]	0.166 [0.124, 0.229]
Hourly earnings	6.37 [4.74, 8.41]	0.157 [0.119, 0.211]
CAPEX orders	5.96 [4.36, 7.93]	0.168 [0.126, 0.229]
Capacity utilisation	5.60 [4.06, 7.53]	0.179 [0.133, 0.246]
<i>Euro area aggregates</i>		
Equity prices	5.64 [4.10, 7.57]	0.177 [0.132, 0.244]
Gas prices	6.03 [4.44, 8.03]	0.166 [0.125, 0.225]
Inflation expectations (10Y)	6.33 [4.67, 8.36]	0.158 [0.120, 0.214]
Inflation expectations (2Y)	6.00 [4.40, 7.97]	0.167 [0.125, 0.227]
OIS rate (10Y)	6.44 [4.78, 8.50]	0.155 [0.118, 0.209]
OIS rate (2Y)	6.59 [4.89, 8.70]	0.152 [0.115, 0.204]
Employment	6.07 [4.46, 8.09]	0.165 [0.124, 0.224]
Energy prices	6.11 [4.50, 8.12]	0.164 [0.123, 0.222]
HICP core	6.67 [4.96, 8.77]	0.150 [0.114, 0.202]
HICP headline	6.54 [4.82, 8.64]	0.153 [0.116, 0.208]
Hours worked	5.79 [4.22, 7.75]	0.173 [0.129, 0.237]
Import prices	6.40 [4.76, 8.46]	0.156 [0.118, 0.210]
Industrial production	5.39 [3.91, 7.26]	0.185 [0.138, 0.256]
PPI	6.69 [4.97, 8.79]	0.149 [0.114, 0.201]
Exchange rate vs USD	5.66 [4.12, 7.59]	0.177 [0.132, 0.243]
EA-US spread (2Y)	5.73 [4.17, 7.69]	0.175 [0.130, 0.240]
EA-US spread (10Y)	5.75 [4.19, 7.73]	0.174 [0.129, 0.239]
<i>US sectoral aggregates</i>		
Prices – upstream, high exposure	6.28 [4.63, 8.33]	0.159 [0.120, 0.216]
Prices – upstream, low exposure	6.70 [5.00, 8.81]	0.149 [0.114, 0.200]
Prices – upstream differential	6.31 [4.66, 8.33]	0.159 [0.120, 0.215]
Prices – downstream, high exposure	6.60 [4.91, 8.68]	0.152 [0.115, 0.204]
Prices – downstream, low exposure	6.31 [4.67, 8.36]	0.159 [0.120, 0.214]
Prices – downstream differential	6.30 [4.65, 8.37]	0.159 [0.120, 0.215]
Production – upstream, high exposure	5.66 [4.12, 7.63]	0.177 [0.131, 0.243]
Production – upstream, low exposure	5.79 [4.24, 7.71]	0.173 [0.130, 0.236]
Production – upstream differential	5.63 [4.14, 7.51]	0.178 [0.133, 0.242]
Production – downstream, high exposure	5.64 [4.11, 7.55]	0.177 [0.132, 0.243]
Production – downstream, low exposure	5.91 [4.32, 7.88]	0.169 [0.127, 0.232]
Production – downstream differential	5.90 [4.31, 7.87]	0.170 [0.127, 0.232]
Employment – upstream, high exposure	6.33 [4.69, 8.37]	0.158 [0.119, 0.213]
Employment – upstream, low exposure	5.81 [4.22, 7.78]	0.172 [0.129, 0.237]
Employment – upstream differential	6.36 [4.71, 8.40]	0.157 [0.119, 0.212]
Employment – downstream, high exposure	6.30 [4.67, 8.35]	0.159 [0.120, 0.214]
Employment – downstream, low exposure	6.25 [4.59, 8.32]	0.160 [0.120, 0.218]
Employment – downstream differential	6.36 [4.69, 8.43]	0.157 [0.119, 0.213]

Notes: Posterior estimates of the hyperparameters governing the hierarchical second-order random-walk (RW2) prior used to smooth the local-projection coefficients across horizons. $\lambda = 1/q^2$ is the implied smoothing precision; q^2 is the innovation variance of the RW2 process. The 16th and 84th percentiles are in square brackets. Estimates refer to the baseline specification with 12 monthly lags.


D.2 Robustness checks

D.2.1 Shorter lag length

As a first robustness exercise, we assess the sensitivity of the local projection estimates to the choice of lag length. While the baseline specification includes $p = 12$ monthly lags, we re-estimate [Equation \(5\)](#) using $p = 6$. The shorter lag structure provides a useful check that the estimated propagation patterns are not mechanically driven by the relatively rich dynamic specification adopted in the baseline. It also reduces the number of parameters estimated at each horizon and therefore mitigates the loss of effective observations that becomes increasingly relevant at longer horizons in our relatively short monthly sample. All other elements of the estimation procedure, including the identified AI-shock draws, the shock normalisation, the cumulative specification of the dependent variable, and the two-layer resampling procedure, are kept unchanged. [Figure D.4](#), [Figure D.5](#), and [Figure D.6](#) compare the resulting impulse responses with those obtained under the baseline $p = 12$ specification.


Overall, the results are broadly stable across the two lag specifications. For the US, the shorter-lag estimates preserve the main features of the baseline response: prices rise over the medium run, capital expenditure expands relatively early, production and capacity utilisation initially display adjustment effects and subsequently recover, while the later increase in employment and earnings is accompanied by a decline in hours. The responses obtained with $p = 6$ are generally somewhat smaller and become more volatile at the longest horizons, but the timing and signs of the main adjustment phases remain broadly unchanged. A similar conclusion emerges for the euro area. The positive responses of headline prices, producer prices, import prices, and energy costs are preserved, as is the delayed improvement in real activity. The financial block is also broadly preserved: the euro depreciation, the medium-run increase in inflation expectations and nominal rates, and the phased response of the EA–US interest rate differentials are all reproduced under the shorter lag structure. Taken together, these results indicate that the domestic and cross-border propagation patterns documented in the baseline are not sensitive to the choice between six and twelve monthly lags.

Figure D.4: Impact of AI news shocks on US aggregates




Notes: Cumulative impulse response functions (IRFs) to a one-standard-deviation AI news shock. Variables are expressed as percent deviations from the baseline ($t - 1$). Darker (lighter) shaded areas represent 68% (90%) highest-density posterior sets. Blue (red) lines indicate the response with $p = 12$ (6).

Figure D.5: Impact of AI news shocks on euro area *real and price* aggregates



Notes: Cumulative impulse response functions (IRFs) to a one-standard-deviation AI news shock. Variables are expressed as percent deviation from the baseline ($t-1$). Darker (lighter) shaded areas represent 68% (90%) highest-density posterior sets. Blue (red) lines indicate the response with $p = 12$ (6).

Figure D.6: Impact of AI news shocks on euro area *financial* aggregates



Notes: Cumulative impulse response functions (IRFs) to a one-standard-deviation AI news shock. Equity prices and exchange rate are expressed as percent deviation from the baseline ($t - 1$), rates, spreads and inflation expectations in basis points. Darker (lighter) shaded areas represent 68% (90%) highest-density posterior sets. Blue (red) lines indicate the response with $p = 12$ (6).


D.2.2 No cross-horizon smoothing

As an additional robustness exercise, we assess whether the shape of the estimated impulse responses is driven by the hierarchical smoothing imposed across local-projection horizons. We therefore re-estimate the baseline specification without the second-order random-walk prior on the sequence of shock-response coefficients, while leaving the remaining elements of the Bayesian specification unchanged. We perform this exercise for the US and euro area aggregate specifications, which provide a direct assessment of the role played by cross-horizon regularisation in the estimation of the macroeconomic transmission of the shock.

For the US ([Figure D.7](#)), the resulting responses are, as expected, substantially less regular across adjacent horizons and generally less precisely estimated, particularly at longer horizons. Nevertheless, the main medium-run patterns underlying the baseline results remain visible. Core capital-goods orders continue to display a pronounced and persistent expansion, while industrial production and real GDP exhibit delayed hump-shaped increases. Employment also rises with a lag, although the responses of hours and earnings are considerably more volatile across horizons. On the price side, producer prices and the different measures of consumer inflation continue to display predominantly positive medium-horizon responses, despite considerably greater horizon-to-horizon variation. The comparison therefore suggests that cross-horizon smoothing primarily regularises the estimated response profiles rather than generating their main economic features. In particular, the delayed expansion in real activity and employment, the strong investment response, and the medium-run increase in prices are already visible in the independently estimated horizon-specific coefficients. At the same time, the substantial volatility and widening posterior intervals of the unsmoothed estimates illustrate the gains in precision and interpretability obtained from the hierarchical RW2 specification.

For the euro area ([Figures D.8](#) and [D.9](#)), the unsmoothed estimates lead to a similar conclusion. The responses become considerably more irregular across adjacent horizons and uncertainty increases, but the main features of the baseline transmission mechanism remain visible. On the real side, employment displays the same delayed expansion, becoming positive at medium horizons and peaking after roughly three years before subsequently fading. Industrial production likewise shows a sizeable hump-shaped increase at medium horizons,

Figure D.7: Impact of AI news shocks on US aggregates—unsmoothed local projections



Notes: Cumulative impulse response functions (IRFs) to a one-standard-deviation AI news shock. Variables are expressed as percent deviations from the baseline ($t - 1$). Darker (lighter) shaded areas represent 68% (90%) highest-density posterior sets.


while hours worked are considerably less precisely estimated. The price responses retain the characteristic delayed pattern found in the baseline. Both headline and core HICP are initially weak, but become positive at medium horizons, with the increase particularly visible between roughly 30 and 45 months. Producer and import prices also exhibit pronounced positive responses over this part of the horizon. Energy prices remain positive for much of the estimation window, although the horizon-by-horizon estimates are substantially more volatile and less precise.

The financial responses are also broadly consistent with the baseline evidence. Euro area equity prices increase following the shock and remain predominantly positive over much of the horizon, albeit with sizeable horizon-to-horizon fluctuations. Short- and long-term inflation-linked and OIS rates continue to display predominantly positive medium-run responses, while the corresponding euro area–US yield spreads are considerably noisier. The exchange-rate response is also less precisely estimated without cross-horizon regularisation.

Overall, as for the US, removing the RW2 prior mainly increases the volatility and uncertainty of the horizon-specific estimates rather than eliminating the principal patterns in the data. In particular, the delayed strengthening of real activity and inflation and the associated adjustment in euro area financial variables remain visible in the absence of cross-horizon smoothing.


Taken together, the US and euro area results indicate that the hierarchical RW2 prior primarily regularises otherwise noisy horizon-by-horizon estimates. The main economic patterns underlying the baseline results remain visible when each horizon is estimated independently, while the unsmoothed specification produces substantially greater short-run variation and wider posterior intervals. This supports the use of hierarchical cross-horizon smoothing as a device to improve precision and interpretability rather than as a source of the dynamic patterns documented in the baseline.

Figure D.8: Impact of AI news shocks on euro area *real and price* aggregates—unsmoothed local projections



Notes: Cumulative impulse response functions (IRFs) to a one-standard-deviation AI news shock. Variables are expressed as percent deviation from the baseline ($t - 1$). Darker (lighter) shaded areas represent 68% (90%) highest-density posterior sets.

Figure D.9: Impact of AI news shocks on euro area *financial* aggregates—unsmoothed local projections



Notes: Cumulative impulse response functions (IRFs) to a one-standard-deviation AI news shock. Equity prices and exchange rate are expressed as percent deviation from the baseline ($t - 1$), rates, spreads and inflation expectations in basis points. Darker (lighter) shaded areas represent 68% (90%) highest-density posterior sets.

D.2.3 Excluding 2020–2022

We re-estimate [Equation \(5\)](#) after removing observations from 2020–2022, a period shaped by the pandemic, reopening, geopolitical tensions, and the energy shock. This is a demanding sample restriction: it removes three years of observations and, together with the lag structure, substantially reduces the effective estimation sample. We therefore report responses only through month 24, beyond which posterior uncertainty becomes especially large.

For the US ([Figure D.10](#)), the exclusion exercise preserves evidence of a short-run increase in prices but changes its persistence. Headline and core consumer prices, core PCE prices, and producer prices rise more strongly during the first year than in the baseline. Their posterior medians subsequently decline and generally turn negative, whereas the baseline price responses build more gradually. The short-run inflationary response is therefore not confined to 2020–2022, but the persistent pass-through from producer to consumer prices is sensitive to the inclusion of those years.


The real and labour-market responses are mixed. Industrial production and average hours remain comparatively close to their baseline paths over much of the horizon. GDP, employment, and hourly earnings are more sensitive: the reduced-sample medians become weaker after the first year and in several cases turn negative. Capacity utilisation is broadly similar initially but declines more at intermediate horizons when 2020–2022 is excluded.

The clearest departure concerns investment. In the baseline, core capital-goods orders rise strongly and persistently. In the restricted sample, their median response is negative through much of the first 18 months and recovers only towards the end of the reported horizon. Thus, the investment-led sequencing that characterises the full-sample results is not preserved in this exercise.

These differences should be read alongside the substantial widening of the posterior intervals. The two specifications' posterior sets overlap extensively for many variables, limiting the precision with which their median paths can be distinguished. Overall, the exercise supports a short-run inflationary effect outside the pandemic and reopening period, but indicates meaningful sample-sensitivity in the persistence of prices and in the medium-run responses of investment, GDP, employment, earnings, and capacity utilisation.

For the euro area ([Figures D.11](#) and [D.12](#)), several early price and cost responses remain

Figure D.10: Impact of AI news shocks on US aggregates excluding 2020–2022
(a) PCE core **(b) CPI headline** **(c) CPI core**




Notes: Cumulative impulse response functions (IRFs) to a one-standard-deviation AI news shock. Variables are expressed as percent deviations from the baseline ($t-1$). Darker (lighter) shaded areas represent 68% (90%) highest-density posterior sets.

visible. Headline and core HICP, producer prices, import prices, and energy prices generally retain an inflationary short-run profile, although the restricted-sample medians are less persistent and the posterior intervals widen considerably. Industrial production and the labour-market variables remain imprecisely estimated, and their medium-run paths differ more clearly from the baseline. The evidence therefore supports the presence of cross-border price transmission outside 2020–2022, while providing weaker support for the precise timing and magnitude of the real response.


The financial results are more sample-sensitive. The responses of equity prices, the exchange rate, OIS rates, inflation expectations, and the EA–US spreads change materially in magnitude and, at some horizons, in sign. In particular, the restricted-sample rate responses do not reproduce the baseline phasing in a stable manner. The broad evidence of international financial transmission remains, but its dynamic profile depends importantly on the inclusion of the exceptional monetary and financial conditions prevailing in 2020–2022. This qualification is consistent with the wider posterior sets and is the reason the main text treats the euro area financial robustness results more cautiously than the price responses.

Figure D.11: Impact of AI news shocks on euro area *real/price* aggregates excluding 2020–2022



Notes: Cumulative impulse response functions (IRFs) to a one-standard-deviation AI news shock. Variables are expressed as percent deviations from the baseline ($t - 1$). Darker (lighter) shaded areas represent 68% (90%) highest-density posterior sets.

Figure D.12: Impact of AI news shocks on euro area *financial* aggregates excluding 2020–2022



Notes: Cumulative impulse response functions (IRFs) to a one-standard-deviation AI news shock. Equity prices and the exchange rate are expressed as percent deviations from the baseline ($t - 1$); rates, spreads, and inflation expectations are expressed in basis points. Darker (lighter) shaded areas represent 68% (90%) highest-density posterior sets.


D.2.4 Sectoral results with different exposure thresholds

We assess whether the sectoral results depend on the thresholds used to define high- and low-exposure industries. We reconstruct the upstream and downstream aggregates using sectors above the 90th percentile and below the 10th percentile of the relevant exposure distribution, rather than the 75th and 25th percentile cut-offs used in the baseline, and re-estimate the local projections without changing the remaining specification ([Figures D.13 and D.14](#)).

The narrower groups produce wider posterior sets, as fewer sectors enter each aggregate, but preserve the main qualitative patterns. Producer-price increases remain substantially stronger among less-exposed sectors both upstream and downstream, and the high-minus-low differentials are more negative than in the baseline over much of the medium and long horizon. The production reversal also remains visible: more-exposed sectors initially perform better, whereas less-exposed sectors account for more of the response at longer horizons as the shock diffuses beyond the AI ecosystem.


The employment results retain the baseline network asymmetry. Highly exposed upstream suppliers exhibit a stronger relative response over the medium run, although this advantage eventually dissipates and reverses. Downstream, employment increasingly shifts towards less-exposed sectors. Thus, the sectoral transmission mechanism is not driven by the baseline thresholds; isolating the most and least exposed industries sharpens the distinction between the shock's initial concentration around the AI production network and its subsequent diffusion to the broader economy.

Figure D.13: Impact of AI news shock on sectors



Notes: Cumulative impulse responses to a one-standard-deviation AI news shock, shown separately by downstream and upstream input-output exposure to AI-related sectors. Darker (lighter) shaded areas represent 68% (90%) highest-density posterior sets for the robustness exercise; dashed lines show the baseline median and 68% posterior set. Variables are expressed as percentage deviations from the baseline ($t - 1$).

Figure D.14: *Differential* impact of AI news shock on sectors



Notes: Differential impulse responses to a one-standard-deviation AI news shock for industries with high AI-network exposure relative to industries with low exposure. Darker (lighter) shaded areas represent 68% (90%) highest-density posterior sets for the robustness exercise; dashed lines show the baseline median and 68% posterior set. Responses are expressed in percentage points.

DeNederlandscheBank

EUROSYSTEEM

De Nederlandsche Bank N.V.
Postbus 98, 1000 AB Amsterdam
020 524 91 11
dnb.nl