The Macroeconomic Effects of AI Uncertainty*

Juan Reyes†‡

King's Business School | King's College London
Economic Statistics Centre of Excellence

December 2025

Abstract

This paper examines the macroeconomic implications of uncertainty shocks related to artificial intelligence (AI). I construct a novel text-based AI Uncertainty (AIU) Index from newspaper coverage. The index displays sharp increases around notable AI developments and regulatory announcements, while demonstrating limited correlation with established measures of economic uncertainty. At the aggregate level, AI uncertainty shocks generate significant contractionary effects on financial markets, labour market outcomes, and real activity. Industry-level estimates, in turn, highlight heterogeneous adjustments along both the labour quantity and price margins. These findings indicate that AI uncertainty represents a distinct and increasingly important driver of business cycle fluctuations.

Keywords: Artificial Intelligence, Uncertainty, Text Analysis, Business Cycle.
JEL classification: D81, E32, J22, J24.

*I am indebted to my supervisors, Michele Piffer and Martin Weale, for their detailed feedback and invaluable guidance, which have been central to the development of this paper. I am equally grateful to David Aikman, Jack Fosten, Francesca Loria, Aditya Mori, Simone Maso, Stuart McIntyre, Rebecca Riley, Filipa Sá, Gregory Thwaites, and Stylianos Zlatanos for their valuable comments and suggestions, which have substantially shaped the study's direction. I would also like to thank Zining Yuan and Shunshun Zhang for providing the initial data used in the early stages of this research. This paper has also benefited from comments and suggestions received from the 2025 King's Business School (KBS) Doctoral Symposium, 2025 Warwick Business School (WBS) PhD Workshop on Fintech, 2025 ESCoE Researcher Workshop, ESCoE Economic Measurement Webinar, 4th UK Digital Economics Workshop, and 2025 UNSW-ESCoE Conference on Economic Measurement. First version: October 2025

†Address all correspondence to Juan Reyes. King's Business School, King's College London, Melbourne House, Aldwych, WC2B 4PJ. E-mail Address: juan.reyes@kcl.ac.uk. Website: https://juanrufinoreyes.owlstown.net

‡Any views, errors, or omissions are solely my responsibility and should not be attributed to the institution I represent. Proper citation of this work is requested.

"To understand AI's implications, we need to measure its impact across multiple dimensions[...]." Kroese (2024), International Monetary Fund

1. Introduction

The rapid adoption of Artificial Intelligence (AI) and its potential to reshape productivity, labour markets, and economic structures have intensified efforts to assess its broader economic implications. For instance, the International Monetary Fund (IMF) estimates that, under a high total factor productivity (TFP) growth scenario, AI could increase global gross domestic product (GDP) by approximately 4.0% over the next decade (Cerutti et al., 2025).[1] Private sector assessments suggest even larger gains. Goldman Sachs (2023) estimates that the widespread AI adoption could boost global GDP by around 7.0% and raise the United States (US) labour productivity growth by 1.5 percentage points annually. Similarly, the McKinsey Global Institute projects that Generative AI (GenAI) models could contribute up to 25.6 trillion USD in additional global output by automating tasks, improving decision-making, and raising productivity in knowledge-intensive industries (Chui et al., 2023).

In contrast to these optimistic projections, Acemoglu (2025) offers a more conservative assessment. Using a task-based macroeconomic framework that considers both automation and task complementarities, the study shows that productivity gains depend on the fraction of tasks impacted by AI, the average cost savings at the task level, and the elasticity of substitution between tasks. The analysis distinguishes between easy-to-learn tasks with observable outcomes and hard-to-learn tasks with context-dependent factors. The results indicate that AI adoption could increase US GDP by only 0.9 to 1.6% over the next decade, alongside a 0.7% improvement in TFP. Acemoglu (2025) attributes this modest aggregate effect to the limited set of tasks that AI can automate at economically meaningful cost savings and the slow creation of new and productivity-enhancing tasks, which together constrain the economic gains from AI-driven automation.

The divergence between these assessments reflects deep uncertainty about the economic impact of AI. On the one hand, optimistic projections envision sizeable productivity and output gains, while on the other, more cautious analyses anticipate modest improvements accompanied by labour displacement and adjustment costs.

[1] A high TFP growth scenario refers to a counterfactual framework in which the adoption of AI leads to a sustained increase in global TFP growth relative to a baseline scenario that assumes historical average productivity growth in the absence of AI-driven gains. See Cerutti et al. (2025).

These opposing views arise from unresolved questions about the speed and scale of AI adoption, the types of tasks most susceptible to automation, the adaptability of firms and institutions, the persistence of productivity gains, and the capacity of labour markets to reallocate displaced workers, among others.

Despite growing attention to the broader economic impact of AI, most research focuses on the first-moment effects, particularly its influence on output, productivity, and labour market outcomes (Acemoglu et al., 2022; Chui et al., 2023; Bonney et al., 2024; Acemoglu, 2025; Cerutti et al., 2025). Far less is known about the second-moment effects surrounding the development and adoption of AI. As an evolving technology with ongoing breakthroughs and regulatory developments, AI generates uncertainty that can affect firm and household decisions through standard channels documented in the uncertainty literature (Bloom, 2009, 2014; Jurado et al., 2015; Baker et al., 2016).

In response, this paper examines the macroeconomic effects of AI-related uncertainty. To capture this, I develop a novel text-based measure, the AI Uncertainty (AIU) Index, based on newspaper coverage from leading news outlets in the US, the United Kingdom (UK), and selected European countries, accessed through Dow Jones Factiva. The index captures perceptions of uncertainty surrounding the emergence, adoption, and economic impact of AI. It combines expectations about future developments with assessments of ongoing or recent events, including labour market adjustments, regulatory changes, sectoral disruptions, and anticipated productivity gains. Following Baker et al. (2016), the index systematically records the frequency of news articles that jointly discuss AI, economic conditions, and uncertainty.

The AIU Index reveals two notable features. First, it responds to technological breakthroughs, regulatory actions, and institutional shifts within the AI landscape. These events include the release of large-scale foundation models, regulatory initiatives, and announcements or governance changes involving prominent individuals, companies, or organisations in AI development. Notable examples include: (1) the introduction of GPT-4 by OpenAI in March 2023, (2) the issuance of US Executive Order (EO) 14110 titled *"Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence"* in October 2023, and (3) the release of DeepSeek-R1 in January 2025. Second, it exhibits low correlations with broader measures of uncertainty, such as the CBOE Volatility Index (VIX), the Real Economic Uncertainty (REU) Index (Jurado et al., 2015), and the Economic Policy Uncertainty (EPU) Index (Baker et al., 2016).

Using the AIU Index, I evaluate the macroeconomic effects of AI-related

uncertainty within a structural vector autoregressive (SVAR) model. The baseline specification identifies the AI uncertainty shocks using two complementary approaches. The first relies on recursive identification, while the second uses an external instrument (SVAR-IV) based on the tail realisations of the AIU Index (Stock and Watson, 2012; Mertens and Ravn, 2013; Carriero et al., 2015). Across both strategies, the results display a consistent pattern. AI-induced uncertainty shocks lead to a short-term increase in equity prices and a reduction in hours worked. In contrast, the effects on wages, employment, and output remain statistically modest. This pattern suggests that, in the baseline specification, adjustment operates primarily through financial market responses and temporary reductions in labour input rather than sustained changes in compensation or hiring.

The muted aggregate response in the baseline specification may partly reflect an identification challenge, since AI-related news coverage could convey both uncertainty and shifts in expectations about future productivity. The same developments that raise concerns can also generate optimistic narratives about technological progress, a configuration that mirrors the identification issues highlighted in the studies of Piffer and Podstawski (2017) and Cascaldi-Garcia and Galvao (2021). This raises the possibility that the estimated effect conflates two opposing forces. Uncertainty depresses economic activity through standard "wait-and-see" behaviour, whereas positive expectations about AI-driven productivity gains support optimism and risk-taking. To evaluate whether these mechanisms overlap in the baseline estimates, I orthogonalise AI-related coverage into an uncertainty component and a growth-expectations component. This decomposition isolates the pure uncertainty dimension of AI and highlights that it induces sharp contractions in equity prices, wages, hours worked, employment, and output, indicating that AI uncertainty is a distinct and quantitatively important driver of the business cycle.

To analyse the industry-level transmission of AI uncertainty shocks, I follow the approach of Känzig (2023) and estimate local projections using the AI uncertainty shock identified from the SVAR-IV as a regressor (Jordà, 2005). The results reveal notable heterogeneity across US industries, with most sectors closely following the aggregate contraction pattern. Hours worked and employment generally decline, with the magnitude of adjustment varying across industries. Wage responses exhibit greater dispersion, varying in both magnitude and sign. While most industries experience wage declines, the financial industry diverges from the aggregate and displays a positive wage response. This heterogeneity in wage adjustments across industries

correlates with differences in industry AI exposure, suggesting that industry-specific vulnerability to AI-driven disruption shapes wage responses to AI uncertainty shocks.

Based on these results, the study contributes to four strands of the literature. First, it extends the existing research on the macroeconomic consequences of uncertainty by introducing a novel and emerging driver rooted in the development of AI. Second, it adds to the literature on measuring specific source uncertainty (Jurado et al., 2015; Baker et al., 2016; Caldara et al., 2020; Husted et al., 2020). Third, it augments the literature that employs text-based methods in empirical economics (Hansen et al., 2018; Gentzkow et al., 2019; Caldara and Iacoviello, 2022; Abiad and Qureshi, 2023; Ash and Hansen, 2023). Finally, it contributes to the growing body of literature on the economics of AI (Felten et al., 2021; Acemoglu et al., 2022; Acemoglu, 2025). To the best of my knowledge, this is the first study to develop a dedicated measure of AI-related uncertainty and to assess its macroeconomic implications.

The remainder of the paper is structured as follows. Section 2 describes the construction of the AIU Index. Section 3 compares the index with different measurements of uncertainty. Section 4 outlines the data, model, and identification of the AI uncertainty shock. Section 5 presents the empirical results. Section 6 concludes and outlines avenues for future research.

2. Measuring AI Uncertainty

This section outlines the methodology for constructing the AIU Index. The process involves four stages. First, I retrieve news articles that reference AI in the context of economic uncertainty. Second, I process the retrieved articles and transform them into structured textual data through feature extraction methods. Third, I compute a standardised index from the frequency of qualifying articles. Lastly, I identify salient AI-related developments and assess their correspondence with observed shifts in the index. The following subsections provide a detailed discussion of each stage.

2.1. Search Procedure

Article Selection. The monthly AIU Index is constructed based on news articles retrieved from the Factiva database, a comprehensive media research platform owned by Dow Jones. The initial selection comprises all articles tagged under the "artificial intelligence" subject code, which reflects the internal taxonomy of the platform based

on the substantive content of each article.[2] These subject-based classifications offer a systematic and consistent means of identifying AI-focused media coverage, forming an initial corpus of 34,187 daily news articles spanning from M1:1979 to M4:2025.

The sample is then restricted to publications classified under "Top Newspaper" section (e.g., *The Wall Street Journal, The New York Times, Financial Times, The Guardian*) as well as those commonly used in prior studies that construct a set of text-based indices from news articles (Baker et al., 2016; Caldara et al., 2020; Abiad and Qureshi, 2023) (Table 2.1). This step ensures consistency in editorial focus and economic reporting over time. By focusing on widely circulated and internationally recognised sources, the index remains anchored in a stable and comparable set of publications.

Table 2.1: Factiva News Outlet

Country:	**News Outlet:**
United States	The Boston Globe, The Baltimore Sun, Chicago Tribune, Investor's Business Daily, The New York Times, New York Post, Pittsburgh Post-Gazette, USA Today, The Wall Street Journal, The Washington Post
United Kingdom	Daily Mail, The Daily Telegraph, Financial Times, The Guardian, The Independent, Reuters News, The Times
Euro Area	Agence France Presse, DW News, Euronews

Note: Table 2.1 lists the news articles used to construct the AIU Index. All outlets are sourced from the Factiva database and classified under the "Top Newspaper" category, alongside those frequently used in constructing text-based indices in the literature. Only articles from these outlets are included in the index after applying keyword filters and excluding non-relevant content types.

Following Abiad and Qureshi (2023), which develops a news-based Oil Price Uncertainty (OPU) Index using Factiva, I apply a filtering strategy and exclude content types that are unlikely to reflect detailed economic reporting. In particular, I rule out articles that fall under categories such as sports, editorials, abstracts, advertorials or sponsored content, advice, analyses, audio-visual links, blogs, event calendars, chronologies, columns, commentaries or opinions, corporate digests, country profiles, transcripts, tables, surveys or polls, statistics, reviews, rankings, prospectuses, press releases, personal announcements, people profiles, front-page headlines, obituaries, letters, interviews, images, and headline-only listings. This approach is intended

[2]This includes categories such as *machine learning*, *risk topics - AI*, *automation*, and *generative AI*. A detailed description of the classification and filtering procedure is provided in Appendix A.

to enhance the signal-to-noise ratio by removing articles that are either stylistically peripheral or lack substantive economic content.

Keyword Filtering. After implementing the article restrictions, I apply a Boolean keyword filter to identify articles that explicitly connect AI developments to economic issues under conditions of uncertainty.[3] Specifically, articles must contain at least one term from each of the following three categories: (1) `artificial intelligence`, (2) `economy`, and (3) `uncertainty`. AI-related terms include keywords commonly associated with technological developments in the field. Economic terms capture macroeconomic concepts such as employment, productivity, or output. Uncertainty terms, on the other hand, reflect the volatility, unpredictability, risk, doubt, and related concepts. The comprehensive list of keyword stems for each category is provided in Table 2.2.

Table 2.2: Keywords used per Category

Category:	**Keywords:**
Artificial Intelligence	`"artificial intelligence" OR "artificial general intelligence" OR "deep learning" OR "generative ai" OR "large language model" OR "machine learning" OR "neural network" OR "anthropic" OR "amazon" OR "amd" OR "apple" OR "chatgpt" OR "claude" OR "deepseek" OR "gemini" OR "google" OR "grok" OR "llama" OR "meta" OR "microsoft" OR "nvidia" OR "openai" OR "perplexity" OR "sora"`
Economy	`"econom*" OR "employ*" OR "growth" OR "job*" OR "layoff" OR "macroeconom*" OR "microeconom*" OR "output" OR "productivit*" OR "recession" OR "unemploy*" OR "wage*"`
Uncertainty	`"uncert*" OR "ambigu*" OR "fluctu*" OR "risk*" OR "unknown*" OR "unpredict*" OR "volat*"`

Note: Table 2.2 presents the keyword-based filtering criteria used to construct the AIU Index. Articles are required to contain at least one term from each of the three categories. Keyword stems (e.g., "uncert*") capture linguistic variants and enhance recall across different writing styles. Boolean operators ensure that selected articles pertain to both AI and economic uncertainty. The list of alternative keywords used in constructing the index is provided in Appendix A.

While some of the uncertainty-related terms capture measurable forms of risk, economic theory distinguishes between these and *Knightian uncertainty*. The former

[3] A Boolean keyword filter combines terms using logical operators, i.e. `"AND"` and `"OR"`. Articles are retained only if they include at least one keyword from each category. This ensures the final sample reflects joint coverage of AI developments, economic relevance, and uncertainty.

refers to situations in which the probability distribution of outcomes is known or can be reasonably estimated. In contrast, the latter characterises environments where such probabilities are indeterminate due to incomplete information (Knight, 1921). Although this conceptual distinction is foundational in theory, it is often conflated in applied settings and public discourse, where observable indicators such as volatility serve as proxies for uncertainty. The AIU Index is designed to account for both quantifiable variation and broader forms of uncertainty, insofar as they are jointly reflected in media narratives surrounding AI and its economic relevance. Accordingly, the AIU Index counts uncertainty-related terms only when they appear in the same context as keywords related to both AI and the economy. This approach ensures that the index captures a broad spectrum of narratives, encompassing both measurable risk and less tractable forms of ambiguity linked explicitly to AI-related economic concerns.

Further, to ensure consistency and keyword matching, the filtering procedure includes Unicode normalisation, standardisation of character encoding, conversion to lowercase, and removal of punctuation marks and non-informative symbols.[4]

In summary, applying editorial source restrictions, content-type exclusions, and Boolean keyword filters reduces the initial pool of over 34,187 AI-tagged articles to 2,056 daily news articles. This refined sample forms the basis for the subsequent text processing and feature extraction.

2.2. Text Analysis and Feature Extraction

The corpus of news articles is processed in Portable Document Format (PDF), with each file containing up to 100 full-length articles, which is the maximum allowed per batch download from the Factiva database. Each article follows a consistent structure and includes metadata such as article type, title, author, word count, publication data, newspaper name, copyright statement, and a unique document identifier. Figure 2.1 provides an example of this format.

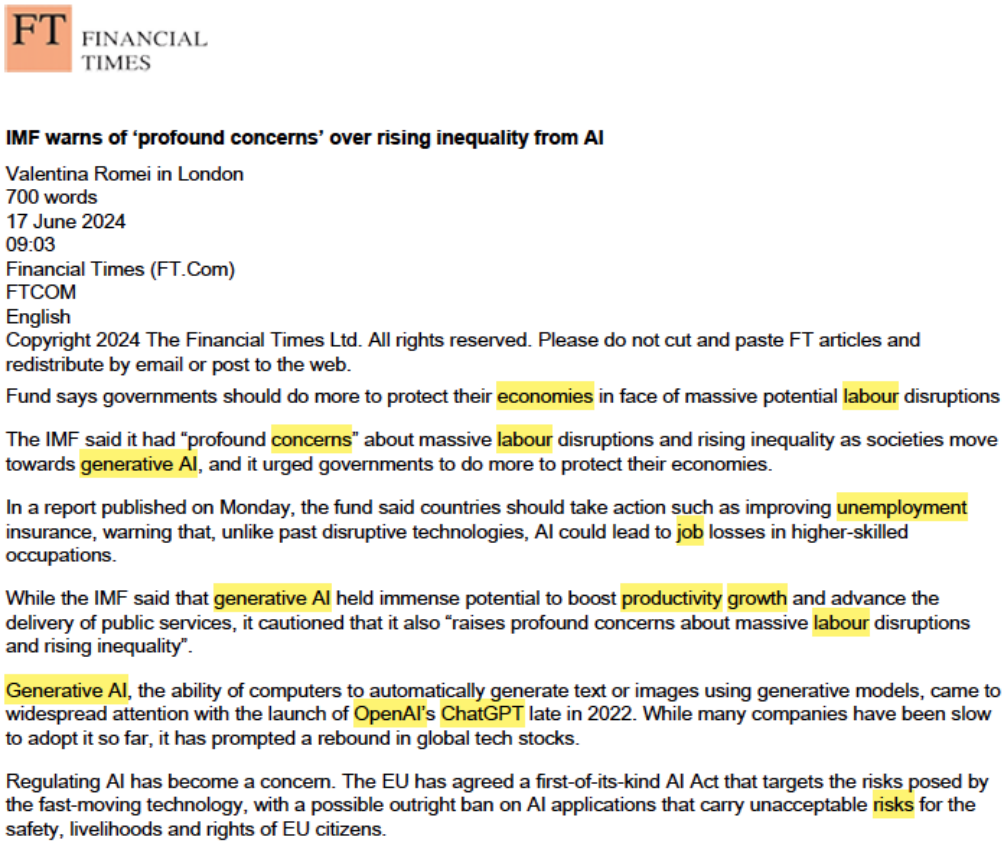
The uniformity of the article layout facilitates efficient preprocessing. The fixed placements of structural elements, such as the copyright statement near the top and the document code at the end, enable reliable segmentation of individual articles within each file. This structure also allows for the precise separation of metadata from the main body of the text. While the metadata is retained for filtering, classification, and

[4]Unicode normalisation standardises characters that may have multiple valid digital representations but appear visually identical. For example, accented letters or quotation marks may be encoded differently across sources. Normalisation converts such characters to a consistent form, improving keyword matching and ensuring encoding consistency.

documentation purposes, only the main article text serves as the input for constructing the index.

As part of the construction process, I manually review 120 randomly selected articles spanning the full sample period. This assessment confirms that the automated script correctly identifies article boundaries and records metadata fields. Conducting this review demonstrates that the extraction procedure performs as intended and that the dataset used to build the AIU Index is internally consistent.

Figure 2.1: Sample Article and Document Structure



FT FINANCIAL TIMES

IMF warns of 'profound concerns' over rising inequality from AI

Valentina Romei in London
700 words
17 June 2024
09:03
Financial Times (FT.Com)
FTCOM
English
Copyright 2024 The Financial Times Ltd. All rights reserved. Please do not cut and paste FT articles and redistribute by email or post to the web.

Fund says governments should do more to protect their economies in face of massive potential labour disruptions

The IMF said it had "profound concerns" about massive labour disruptions and rising inequality as societies move towards generative AI, and it urged governments to do more to protect their economies.

In a report published on Monday, the fund said countries should take action such as improving unemployment insurance, warning that, unlike past disruptive technologies, AI could lead to job losses in higher-skilled occupations.

While the IMF said that generative AI held immense potential to boost productivity growth and advance the delivery of public services, it cautioned that it also "raises profound concerns about massive labour disruptions and rising inequality".

Generative AI, the ability of computers to automatically generate text or images using generative models, came to widespread attention with the launch of OpenAI's ChatGPT late in 2022. While many companies have been slow to adopt it so far, it has prompted a rebound in global tech stocks.

Regulating AI has become a concern. The EU has agreed a first-of-its-kind AI Act that targets the risks posed by the fast-moving technology, with a possible outright ban on AI applications that carry unacceptable risks for the safety, livelihoods and rights of EU citizens.

Note: Figure 2.1 shows the structure of an article downloaded from Factiva, which includes metadata such as title, author, publication data, etc. This formatting facilitates automated article segmentation and improves the efficiency and accuracy of text processing for AIU Index construction.

2.3. Index Construction

As highlighted in earlier studies on news-based uncertainty indices, one of the primary methodological challenges is controlling for variation in total article volume, which can distort comparisons across time or between news outlets. Raw counts of articles containing relevant keywords are highly sensitive to changes in publication frequency, editorial scope, and archival completeness (Baker et al., 2016; Abiad and Qureshi,

2023). To ensure comparability over time and across outlets, I adopt the standard procedure developed by Baker et al. (2016) for the EPU Index. Following this approach, I construct the AIU Index through four steps: (1) normalising article counts by outlet volume, (2) standardising variance relative to a baseline period, (3) aggregating across outlets, and (4) renormalising to index levels.

Volume Normalisation. For each news outlet, I compute the monthly share of qualifying articles, particularly those that mention at least one term from each of the three categories discussed in Section 2.1, relative to the total number of articles published by that outlet in the same month:

$$\theta_{it} = \frac{\alpha_{it}}{\tau_{it}}, \tag{2.1}$$

where θ_{it} denotes the share of articles that meet the inclusion criteria for outlet i in month t, α_{it} is the number of articles containing at least one term from each of the three keyword categories, and τ_{it} is the total number of articles published. This step ensures comparability across outlets with different publication volumes.

Variance Standardisation. Each outlet series is then scaled by its own standard deviation, σ_i, computed over a pre-specified baseline window T_{base}:

$$Y_{it} = \frac{\theta_{it}}{\sigma_i}, \quad \sigma_i = \text{stdev}(\theta_{it}, t \in T_{base}). \tag{2.2}$$

This procedure standardises the variance of each outlet series while preserving its mean. The objective is not to re-centre the distribution but to place all outlets on a comparable variance scale. This prevents outlets with more volatile coverage from exerting disproportionate influence on the aggregate index.[5]

Aggregation. The variance-standardised series, Y_{it}, are averaged across outlets available in month t:

$$Z_t = \frac{1}{N_t} \sum_{i \in S_t} Y_{it}, \tag{2.3}$$

where S_t is the set of outlets with observations in month t and $N_t = |S_t|$.[6]

[5]The baseline period is set to M1:2016 to M12:2022, representing a relatively stable phase prior to the significant advances in GenAI models. See Appendix A for details.

[6]The notation $|S_t|$ denotes the cardinality (number of elements) of the set S_t, not an absolute value.

Renormalisation to Index Levels. Finally, the aggregated series, Z_t, is renormalised based on the baseline period:

$$AIU_t = 100 \times \frac{Z_t}{\bar{Z}_{base}}, \tag{2.4}$$

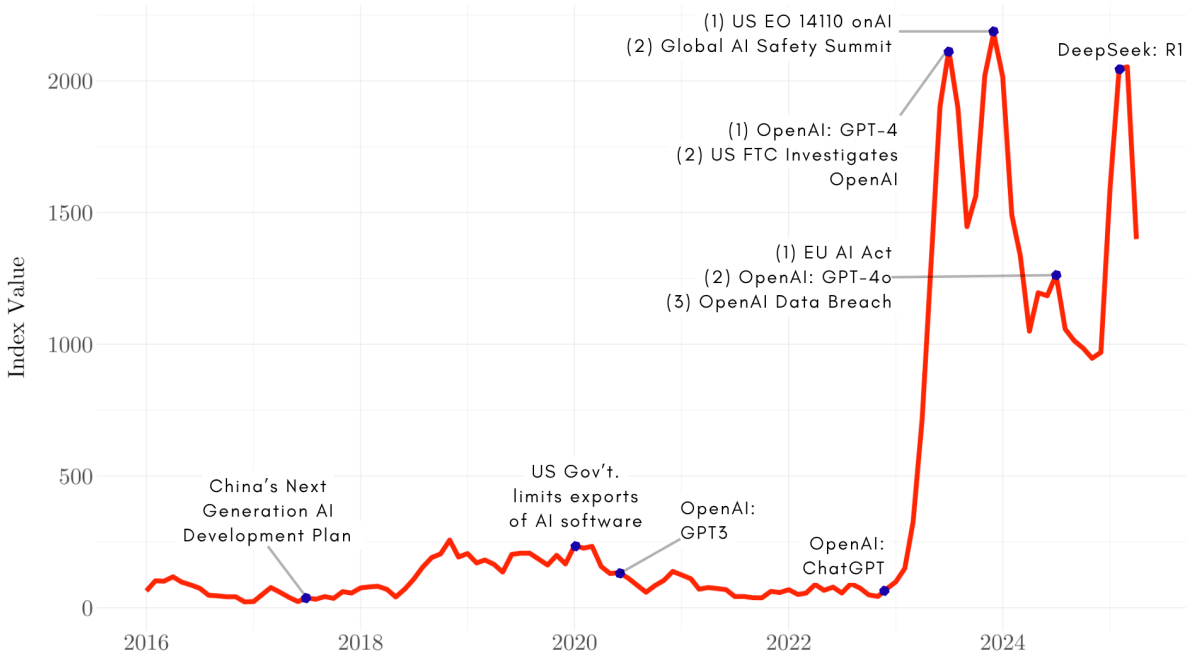
where AIU_t is the value of the AIU Index in month t, and $\bar{Z}_t$ is the mean of Z_t across the months of the baseline period.

This four-step procedure ensures that no single outlet disproportionately influences the index and mitigates potential biases arising from differences in publication frequency or archival depth. Appendix B reports the contribution of each news outlet to the AIU Index.

2.4. Event Identification

To illustrate the way the AIU Index reflects periods when AI becomes a focal macroeconomic concern, I follow the narrative approach of Romer and Romer (2010). I review the news articles underlying the most significant index movements and connect them to major AI developments linked to heightened economic uncertainty.

Figure 2.2: AIU Index (3-Month Moving Average)



Note: Figure 2.2 plots the AIU Index from M1:2016 to M4:2025. For presentation purposes, a three-month moving average is applied. Spikes in the index were investigated by manually reviewing the underlying news articles to identify the events driving the largest movements.

Evidence of such correspondence supports the interpretation of the index as reflecting shifts in public discourse on AI and its perceived economic implications. Figure 2.2 presents the AIU Index for the period M1:2016 to M4:2025, together with selected episodes where sharp movements in the series coincide with notable AI-related developments. This demonstrates that the index does not merely reflect fluctuations in media intensity but instead traces turning points in public discourse in which AI is portrayed simultaneously as a driver of growth and a source of uncertainty.

Among these episodes, the largest increases in the index align with advances in AI technology and episodes of regulatory action that were closely associated with concerns about macroeconomic uncertainty. A notable example occurred in March 2023, when OpenAI released GPT-4. Media coverage highlighted the potential productivity gains from GenAI and its contribution to economic growth, while also raising concerns about labour displacement and regulatory preparedness. These concerns were built on the earlier release of ChatGPT in late 2022, which demonstrated human-level performance across tasks such as programming, legal reasoning, writing, and mathematical problem-solving, among others.

Another significant spike followed the issuance of EO 14110 in October 2023, titled *"Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence"*. Contemporary media reports stressed that the order marked a pivotal moment in AI governance, triggering debate over its regulatory implications and broader economic stakes. Coverage pointed to risks of regulatory fragmentation, enforcement challenges, and uncertainty about the pace and scope of AI adoption. The attention surrounding this order illustrates that regulatory action itself can become a focal point of macroeconomic concern.

A further increase in the index appeared in January 2025 with the release of DeepSeek-R1. Marketed as a lower-cost alternative to leading foundation models such as ChatGPT, the launch intensified public discussion about how more affordable GenAI models could broaden access and spur innovation. At the same time, media coverage raised concerns about workforce displacement, the competitive pressures on semiconductor suppliers, and the challenges of regulating the rapid adoption of AI. The release was widely described as a turning point in global AI competition, prompting a reassessment of the US policy and export controls. Financial markets also reacted, with major technology firms, including Nvidia, experiencing sharp declines in stock values. Analysts noted that DeepSeek relied on domestically produced processors, a move aimed at strengthening China's technological self-reliance.

These episodes support the interpretation that the AIU Index captures evolving perceptions of economic uncertainty arising from AI developments. The timing and content of the underlying articles indicate that the index reflects shifts in public discourse in which AI is framed as both an engine of productivity and growth and a source of uncertainty related to labour markets, institutional capacity, and policy responses. Appendix C provides additional evidence by linking other notable spikes in the index to AI-related events.

3. AIU Index Benchmarking

To place the AIU Index in context, I compare its dynamics with established benchmark measures of uncertainty. These include the EPU Index (Baker et al., 2016), the REU Index (Jurado et al., 2015), the S&P 500 Volatility Index (VIX), and the NASDAQ 100 Volatility Index (VXN), which collectively capture broader sources of uncertainty related to policy changes, macroeconomic volatility, and financial market conditions, respectively. I then conduct statistical comparisons to examine the extent of overlap and distinction in their movements relative to the AIU Index.

Pearson Correlation. First, I obtain the pairwise Pearson correlation coefficients between the AIU Index and each benchmark measure:

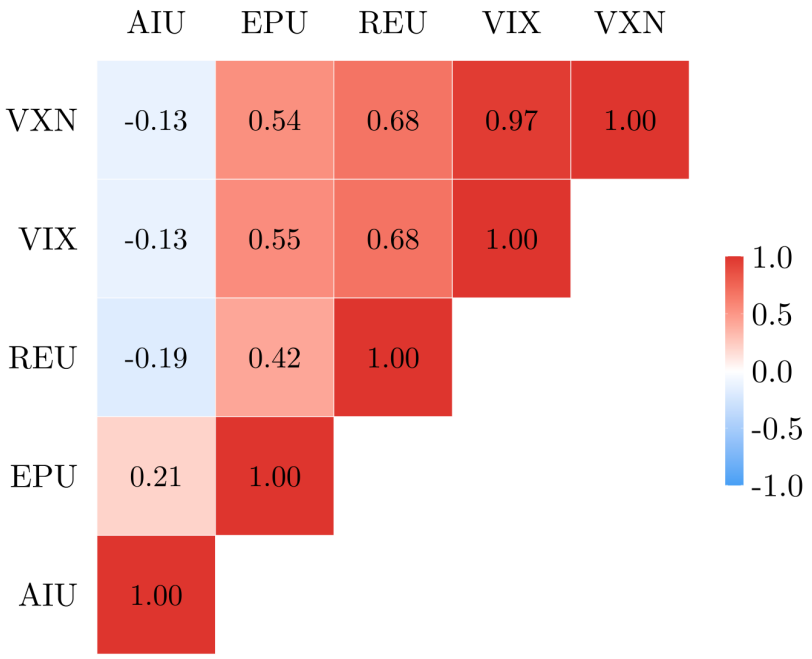
$$\rho = \frac{\text{Cov}(\Gamma_i, AIU)}{\sigma_{\Gamma_i}\sigma_{AIU}}, \tag{3.1}$$

where ρ is the Pearson correlation coefficient, $\text{Cov}(\Gamma_i, AIU)$ is the covariance between each benchmark measure, Γ_i, and the AIU Index, computed over the sample period $t = 1, \ldots, T$. While σ_{Γ_i} and σ_{AIU} are their respective standard deviations.

Figure 3.1 presents the correlation matrix. The AIU Index displays low correlations with each benchmark: (1) -0.1 with VXN, (2) -0.1 with VIX, (3) -0.2 with the REU Index, and (4) 0.2 with the EPU Index. The negative correlations with the VXN, VIX, and REU Index suggest that movements in the AIU Index are not systematically aligned with fluctuations in financial market volatility or real activity uncertainty. In contrast, the modest positive correlation with the EPU Index indicates some degree of co-movement with policy-related uncertainty, although the relationship is weaker than that observed among the benchmark indices themselves. Overall, the evidence suggests that the AIU Index captures a dimension of uncertainty that is only partially related to these broader measures. By comparison, the benchmark indices are more strongly correlated with

one another, with coefficients ranging from 0.4 to 1.0, reinforcing that the AIU Index provides distinct informational content.

Figure 3.1: Correlation Matrix of Uncertainty Measures



Note: Figure 3.1 reports the Pearson correlation coefficients between AIU Index and benchmark uncertainty measures. The benchmarks are the EPU Index, REU Index, VIX, and VXN. The sample spans from M1:2016 to M4:2025, reflecting the availability of news articles used to construct the AIU Index and the benchmark uncertainty measures.

Linear Regression. In addition to the Pearson correlations, I estimate three separate ordinary least squares (OLS) regressions. In each regression, one of the benchmark uncertainty measures serves as the dependent variable, while the AIU Index is the explanatory variable:

$$\Gamma_{it} = \alpha_i + \lambda_i \text{AIU}_t + \varepsilon_{it}, \tag{3.2}$$

where Γ_{it} denotes the EPU, REU, VIX, and VXN at time t, and AIU_t is the AIU Index. The coefficient λ_i captures the association between the AIU Index and each benchmark measure.

The regression results reported in Table 3.1 complement the correlation analysis and indicate a weak relationship between the AIU Index and existing benchmarks. Across all specifications, the estimated coefficients are close to zero. Although some coefficients are statistically significant, their magnitudes are small and economically unimportant. The explanatory power is also limited, with R^2 values consistently below 0.1. Further, the results suggest that the AIU Index reflects movements in uncertainty

that are not captured by standard measures of policy, macroeconomic, or financial uncertainty.

Table 3.1: OLS Regressions of Uncertainty Measures on AIU Index

	Dependent Variable:			
	EPU	**REU**	**VIX**	**VXN**
	(1)	(2)	(3)	(4)
AIU	0.03^{**}	-0.00^{**}	-0.00	-0.00
	(0.01)	(0.00)	(0.00)	(0.00)
Constant	219.63^{***}	0.75^{***}	19.15^{***}	23.10^{***}
	(8.34)	(0.02)	(0.79)	(0.81)
Observations	112	112	112	112
R^2	0.04	0.04	0.02	0.02

Note: Table 3.1 reports OLS regression results of benchmark uncertainty measures on the AIU Index. Robust standard errors in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$. The sample covers M1:2016 to M4:2025, reflecting the availability of news articles used in constructing the AIU Index and benchmark uncertainty measures.

4. Data, Model Specification, and Inference

This section outlines the data, model specification, and identification strategies used to evaluate the macroeconomic effects of AI uncertainty. The analysis is structured around three approaches. I first estimate a baseline SVAR with recursive ordering. I then relax the exogeneity assumption using an SVAR-IV. Finally, I apply local projections to examine industry-level responses. The following subsections provide further detail on these approaches.

4.1. Data

The SVAR is estimated using six monthly US variables spanning the period from M1:2016 to M4:2025. Table 4.1 reports the variables, their definitions, and the transformations applied before estimation. Appendix D plots the variables.

The selection of variables is grounded in the broader literature on economic uncertainty (Bloom, 2009, 2014; Jurado et al., 2015; Baker et al., 2016; Piffer and Podstawski, 2017; Caldara et al., 2020). However, two key modifications are introduced

to align the specification with the objectives of this study. First, I estimate the SVAR in levels, following Sims et al. (1990). This is to retain long-run information that could be lost through differencing or filtering methods such as the Hodrick–Prescott (HP) filter (Hodrick and Prescott, 1997). Second, inflation and monetary policy variables are omitted to maintain focus on broader macroeconomic aggregates.

Table 4.1: Data Description

Variable:	Description:	Transform:
AIU Index	AI Uncertainty Index	Level
S&P 500	Index Value at Market Close	Log Level
Wage	Average Hourly Earnings of All Employees, Total Private	Log Level
Hours	Average Weekly Hours of All Employees, Total Private	Log Level
Employment	All Employees, Total Private	Log Level
Industrial Production	Industrial Production: Total Index	Log Level

Note: Table 4.1 presents the monthly variables used in the SVAR estimation, together with their descriptions and applied transformations. All macroeconomic variables, except for the AIU Index, are sourced from the US Bureau of Labour Statistics (BLS) and Federal Reserve Bank of St. Louis (FRED).

4.2. Structural Vector Autoregression

The macroeconomic effects of AI-related uncertainty are estimated within an SVAR:

$$\mathbf{Y}_t = \alpha + \sum_{i=1}^{p} \mathbf{A}_i \mathbf{Y}_{t-i} + \mathbf{u}_t, \qquad \mathbf{u}_t = \mathbf{B}\,\varepsilon_t. \tag{4.1}$$

where $\mathbf{Y}_t$ is the $n \times 1$ vector of endogenous variables, α is an $n \times 1$ vector of intercepts, and $\mathbf{A}_i$ are $n \times n$ coefficient matrices for the p lags. The reduced-form residuals $\mathbf{u}_t$ have covariance matrix $\mathbf{\Sigma}_u \equiv E[\mathbf{u}_t \mathbf{u}_t'] = \mathbf{B}\mathbf{B}'$ with $\mathbf{B}$ denoting the contemporaneous impact matrix. Structural shocks ε_t are assumed to be mutually orthogonal and normalised to unit variance:

$$E[\varepsilon_t \varepsilon_t'] = \mathbf{I}, \tag{4.2}$$

where equations (4.1) and (4.2) define the reduced-form VAR and its structural representation. Since $\mathbf{B}$ cannot be uniquely identified from $\mathbf{\Sigma}_u$, recovering the column associated with the AI uncertainty shock requires additional restrictions. I consider two alternative identification strategies. The first is a recursive approach, which

imposes timing restrictions through the ordering of variables. The second is an instrumental-variable approach (SVAR-IV), which exploits an external instrument to isolate innovations in AI-related uncertainty without relying on recursive assumptions.

SVAR with Recursive Identification. I initially examine the macroeconomic effects of AI uncertainty using an SVAR with recursive identification. This approach provides a natural starting point because it is tractable and transparent, particularly in the absence of strong theoretical priors on the structural role of AI-related uncertainty in the economy. It also follows established practice in the literature, which models uncertainty as an exogenous source of economic fluctuations (Bloom, 2009; Jurado et al., 2015; Baker et al., 2016; Caldara et al., 2020).

Under recursive identification, the contemporaneous impact matrix $\mathbf{B}$ is restricted to be lower triangular. Variables that appear earlier in the ordering may affect those ordered after them contemporaneously, but not the reverse. Moreover, the AI uncertainty shock is identified by placing the AIU Index first, consistent with the assumption that innovations in AI-related uncertainty can contemporaneously influence all macroeconomic variables while remaining insulated from same-period movements in fundamentals. This ordering also reflects the informational nature of the index, which captures expectations about future economic conditions.

Formally, the vector of endogenous variables is ordered as:

$$\begin{pmatrix} \textit{AIU Index} \\ \log(\textit{S\&P 500}) \\ \log(\textit{Wage}) \\ \log(\textit{Hours}) \\ \log(\textit{Employment}) \\ \log(\textit{Industrial Production}) \end{pmatrix}, \tag{4.3}$$

where the AIU Index is followed by indicators for the equity market, labour market, and real economic activity.

SVAR with Instrumental Variable. To relax the timing restriction assumption, I also employ SVAR-IV to identify AI uncertainty shocks (Stock and Watson, 2012; Mertens and Ravn, 2013). This approach uses an external instrument that is correlated with innovations in AI-related uncertainty but orthogonal to all other shocks in the system. In contrast to the recursive specification, it does not depend on ordering

assumptions, since contemporaneous relationships among the endogenous variables remain unrestricted.

Identification requires an external instrument, z_t, that satisfies two conditions:

$$E[z_t, \varepsilon_t^{AI}] \neq 0, \tag{4.4}$$

$$E[z_t, \varepsilon_t^{j}] = 0 \quad \text{for} \quad j \neq AI, \tag{4.5}$$

where ε_t^{AI} is the AI uncertainty shock and ε_t^{j} are the remaining structural shocks. Condition (4.4) ensures relevance, while condition (4.5) imposes exogeneity.

Under these conditions, the impulse vector $\mathbf{s}_1$, the column of $\mathbf{B}$ associated with the AI uncertainty shock, can be recovered from the covariance between the instrument and the reduced-form residuals:

$$\mathbf{s}_1 \propto \text{Cov}(z_t, \mathbf{u}_t), \tag{4.6}$$

where $\text{Cov}(z_t, \mathbf{u}_t)$ denotes the vector of covariances between the instrument and the reduced-form residuals.

To fix the scale, I impose the normalisation that the structural AI uncertainty shock has unit variance. With this restriction, the identified shock series is obtained as:

$$\hat{\varepsilon}_t^{AI} = \frac{\hat{\mathbf{s}}_1' \hat{\mathbf{\Sigma}}_u^{-1} \mathbf{u}_t}{\sqrt{\hat{\mathbf{s}}_1' \hat{\mathbf{\Sigma}}_u^{-1} \hat{\mathbf{s}}_1}}, \tag{4.7}$$

where $\hat{\mathbf{\Sigma}}_u$ is the estimated covariance matrix of reduced-form residuals. This normalisation ensures that $\hat{\varepsilon}_t^{AI}$ has unit variance. Therefore, the impulse responses can be interpreted as the effects of a one standard deviation AI uncertainty shock.

Tail Realisation Instrument. To implement the SVAR-IV, one of the instrument construction method that follow is the approach of Carriero et al. (2015). In their study, the instrument is derived from extreme realisations of the VXO. In my setting, the same logic is applied to the AIU Index, where I construct a binary instrument based on its tail realisations.

The instrument takes the value 1 when the AIU Index exceeds a high quantile of its historical distribution and 0 otherwise:

$$z_t^{(\tau)} = \begin{cases} 1 & \text{if } AIU_t \geq q_\tau(AIU), \\ 0 & \text{otherwise}, \end{cases} \tag{4.8}$$

where q_τ denotes the τ-th quantile of the AIU Index. In particular, instruments are constructed using $\tau \in \{0.95, 0.90, 0.75\}$. This strategy captures periods of AI-related uncertainty that are more likely to represent exogenous shocks rather than systematic responses to macroeconomic conditions. These discrete realisations therefore provide a plausible source of external variation for identification. The tail realisations of the index are shown in Appendix E.

Residual-Based Instrument. The second approach in constructing the instrument for the SVAR-IV applies a residual-based method to orthogonalise the two distinct dimensions of AI-related coverage. One dimension captures narratives that present AI as a source of economic uncertainty. The other reflects discussions that emphasise the broader economic role of AI, such as its contribution to productivity and economic growth. Distinguishing between these components is important because they correspond to different channels through which AI could potentially affect the economy. Prior studies have found that uncertainty shocks often co-move with news or TFP shocks, which are unanticipated revisions in expectations about future productivity (Piffer and Podstawski, 2017; Cascaldi-Garcia and Galvao, 2021). The residual-based procedure addresses this limitation by constructing an instrument that isolates the uncertainty component.

This approach relies on two text-based indices drawn from the same corpus of news articles: (1) the AIU Index and (2) the AI Economic (AIE) News Index. The AIU Index, introduced earlier, identifies periods when AI is described as a source of economic uncertainty through references to AI, economic terms, and direct mentions of uncertainty. The AIE News Index follows the same construction methodology but does not require references to uncertainty (Appendix F). It therefore captures the overall discussion on the economic role of AI, including pieces both with and without uncertainty-related terms. By construction, the AIU Index is thus nested within the AIE News Index. This nested structure and common corpus ensure comparability across the two indices and permit meaningful decomposition of the SVAR-IV.

The decomposition proceeds through auxiliary regression:

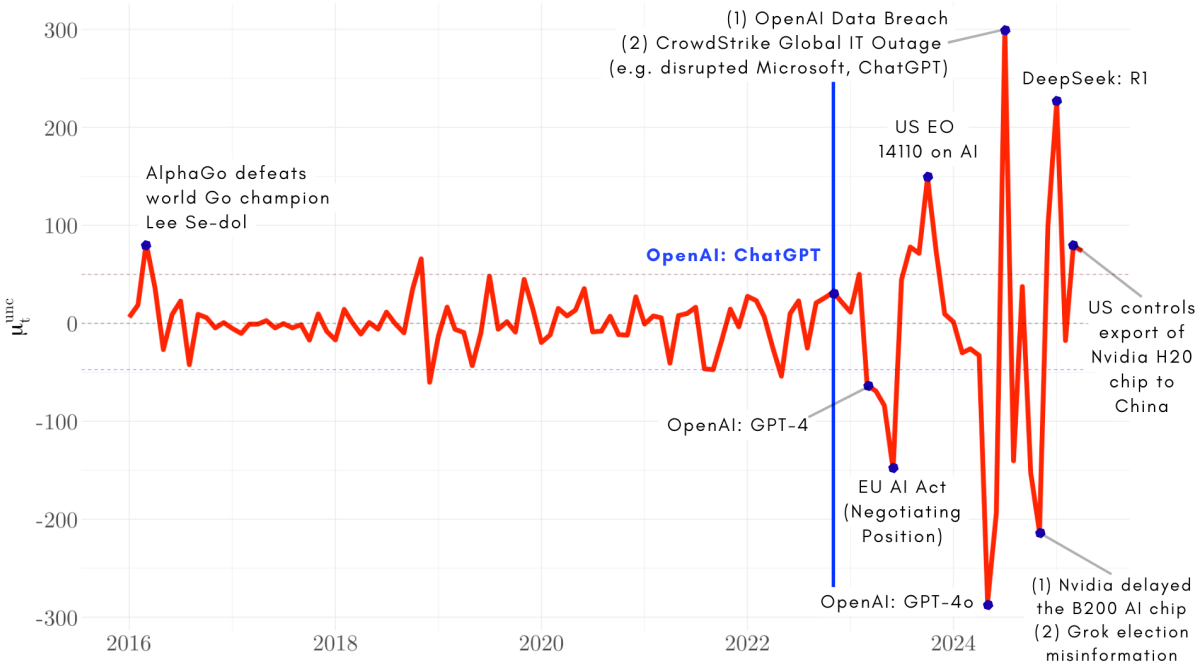
$$AIU_t = \alpha_1 + \beta_1 AIE_t + \mu_t^{unc}, \tag{4.9}$$

where μ_t^{unc} represents the residual variation in AI-related uncertainty after controlling for the AIE News Index. The estimated coefficient of $\hat{\beta}_1 = 0.9$ indicates that the two indices move nearly one-to-one, consistent with the nested structure in which

uncertainty articles represent a subset of total AI-economy coverage. Further, the regression achieves an $R^2 = 0.99$, reflecting this strong proportional relationship (Appendix G). The residual, therefore, isolates deviations from this baseline, that is, periods when AI uncertainty coverage appears at unusually high or low rates relative to the overall AI coverage volume.

Figure 4.1 presents the resulting series. The residual remains stable through late 2022, indicating that the AIU and AIE News indices moved in proportion to each other during this period. From late 2022 onward, particularly when ChatGPT was released in November 2022, the series exhibits considerably greater volatility with large positive and negative movements.

Figure 4.1: Residual-Based AI Uncertainty



Note: Figure 4.1 displays the residual-based measure of AI uncertainty μ_t^{unc}, constructed by removing the component of the AIU Index explained by the AIE News Index (Equation 4.9). The annotated points correspond to selected AI-related developments that received extensive news coverage. The sample spans M1:2016 to M4:2025, consistent with the availability of the underlying news corpus.

Positive values arise when uncertainty-related terms constitute a disproportionately large share of coverage. The DeepSeek episode in January 2025 exemplifies this pattern. As discussed in Section 3, coverage during this period highlighted themes involving AI-related geopolitical competition, regulatory uncertainty, and labour market disruption. These themes increased the relative weight of uncertainty-related language within AI-economy coverage, which explains the positive residual even in a period characterised by wide-ranging coverage.

Negative residuals reflect the opposite pattern. These episodes arise when coverage focuses more on concrete technological capabilities, demonstrated performance, and commercial integration, which lowers the relative share of uncertainty-focused content. The periods associated with GPT-4 in March 2023 and GPT-4o in May 2024 illustrate this behaviour. While these developments generated substantial coverage, the articles during these periods emphasised technical benchmarks, product features, and clearly defined applications. As a result, the proportion of uncertainty-related content was lower than predicted by the baseline relationship. The residual captures variation in narrative framing rather than mere changes in reporting quantity.

COVID-19 Pandemic Adjustment. To address the sharp but temporary increase in volatility during the COVID-19 pandemic, the estimation of both the baseline SVAR and the SVAR-IV incorporates time-varying volatility of structural shocks. Following Lenza and Primiceri (2022), I account for elevated volatility in March, April, and May 2020 by estimating a sequence of scaling factors derived from the behaviour of reduced-form residuals. This adjustment allows the model to absorb the pandemic-related surge in volatility without discarding observations, thereby preserving the full sample while limiting its influence on inference.[7].

Estimation and Inference. Both SVAR with recursive identification and SVAR-IV are estimated using OLS, with the optimal lag length selected based on standard information criteria. To construct confidence bands around the impulse response functions (IRFs), I employ a wild bootstrap procedure, as described in Gonçalves and Kilian (2004). This method addresses both estimation and identification uncertainty by generating pseudo-samples in which the signs of reduced-form residuals are randomly resampled across time. The SVAR is re-estimated for each pseudo-sample using the same identification scheme, and the procedure is repeated until 1,000 valid replications are obtained.[8]

[7]The adjustment permits a temporary increase in the variance of structural shocks during the COVID-19 pandemic. I estimate the scaling parameters $\bar{s}_0$, $\bar{s}_1$, and $\bar{s}_2$, corresponding to March, April, and May 2020, along with the decay rate ρ, by maximum likelihood. The volatility path s_t is defined as $s_t = \bar{s}_j$ for the first three months, and for $j \geq 1$, $s_{t+j} = 1 + (\bar{s}_2 - 1)\rho^j$, which allows variance to gradually return to pre-pandemic levels. All endogenous variables and their lags are then rescaled by s_t. Further details are provided in Lenza and Primiceri (2022)

[8]To ensure that the impulse responses and confidence intervals are not driven by extreme observations in the AIU Index, particularly those primarily concentrated from late 2022 onwards, I re-estimate the baseline SVAR and SVAR-IV after excluding periods with large standardised residuals from the AI uncertainty equation. This robustness check evaluates whether a small number of unusual observations materially influence the estimates. The outlier exclusion procedure is comprehensively discussed in Appendix H.

4.3. Local Projections

To provide a more comprehensive assessment of the macroeconomic effects of AI uncertainty, I extend the analysis to the industry level using local projections, following the approach of Känzig (2023). Within this framework, the structural AI uncertainty shock, $\hat{\varepsilon}_t^{AI}$, identified through the SVAR-IV, serves as the exogenous driver of the industry-level local response. This extension enables the tracing of heterogeneity in adjustment paths across industries and the identification of distributional patterns that could not be captured in aggregate estimates.

The estimation uses monthly industry-level indicators of wages, hours worked, and employment, which capture adjustments along both the intensive and extensive margins of the US labour market. Industries included in the analysis are defined according to the North American Industry Classification System (NAICS). Information on the variables used in the local projections is presented in Appendix I.

The dynamic responses of these outcomes to an AI-related uncertainty shock are estimated using the following regression:

$$y_{i,t+h} = \alpha_{i,h} + \beta_{i,h}\,\hat{\varepsilon}_t^{AI} + \gamma_{1,i,h}\,y_{i,t-1} + \cdots + \gamma_{p,i,h}\,y_{i,t-p} + e_{i,t+h}, \tag{4.10}$$

where $y_{i,t+h}$ denotes the outcome variable for industry i at horizon h, $\alpha_{i,h}$ is the intercept, $\beta_{i,h}$ captures the response to a one standard deviation AI uncertainty shock, $\gamma_{i,h}$ accounts for the autoregressive persistence, and the error term is denoted by $e_{i,t+h}$. Local projections are estimated separately for each horizon $h = 0, 1, \ldots, H$, producing a sequence of $\beta_{i,h}$ coefficients that trace the dynamic response of each industry to the AI uncertainty shock. Confidence intervals for the local projections are constructed using Newey-West standard errors, following Jordà (2005).

5. Results

This section reports the empirical results. I first assess instrument relevance and exogeneity, then present the findings from the recursive SVAR, the SVAR-IV, and the local projections. The subsections that follow discuss each set of results.

5.1. Strength of Instrument

Following Gertler and Karadi (2015) and Piffer and Podstawski (2017), I assess the validity of the instruments by examining both relevance (Equation 4.4) and exogeneity

(Equation 4.5). First, I evaluate relevance by regressing each reduced-form residual on the proposed instruments:

$$u_{it} = \alpha + \beta_i \hat{w}_t^{AI} + \delta_{it}, \quad i = 1, 2, \ldots, n, \tag{5.1}$$

where u_{it} denotes the reduced-form residual from equation i at time t, and β_i measures the relationship between the instrument and that residual, δ_{it} is the error term, and n is the number of equations in the SVAR-IV.

Then, to evaluate exogeneity, I examine whether the proposed instruments capture variation that differs from structural uncertainty driven by macroeconomic or financial conditions. The analysis compares the tail realisation and residual-based instruments with structural uncertainty shocks derived from the EPU Index, the REU Index, the VIX, and the VXN. For each benchmark measure, I estimate an SVAR that includes the uncertainty indicator together with the variables described in subsection 4.1. The benchmark uncertainty shock is identified using recursive identification with the uncertainty measure ordered first. I then regress the instrument on the corresponding benchmark structural uncertainty shock:

$$\theta_t^{IV} = \alpha_j + \beta_j \varepsilon_t^j + \eta_{jt}, \tag{5.2}$$

where θ_t^{IV} refers to the proposed instruments (i.e., tail realisation or residual-based), ε_t^j is the structural uncertainty shock derived from benchmark j, and η_{jt} is the error term.

Tail Realisation Instrument. Using the 0.95 quantile of the AIU Index yields an F-statistic of 63.7 in the targeted equation, as reported in Table 5.1. This indicates strong relevance relative to the conventional threshold of 10.0 (Stock and Yogo, 2002). Instruments constructed from the 0.90 and 0.75 quantiles are also considered for robustness, with estimates reported in Appendix J. The 0.90 quantile remains strong with an F-statistic of 34.3, whereas the 0.75 quantile generates an F-statistic of 3.1. The latter indicates that lower tail thresholds do not provide sufficient variation to serve as reliable instruments.

Across all considered quantiles, the tail realisation instruments display negligible and statistically insignificant correlations with the residuals from the non-targeted equations. The estimated coefficients are small in magnitude, statistically insignificant, and associated with low R^2 values and weak F-statistics. Consequently, the upper tail of the AIU Index is strongly correlated only with the residual in the target equation and remains uncorrelated with the remaining reduced-form innovations. This pattern

is consistent with the interpretation that the upper tail isolates the structural AI uncertainty shock while remaining orthogonal to other structural shocks in the system.

Table 5.1: Instrument Relevance (0.95 Quantile)

	AI Unc.	S&P 500	Wage	Hours	Emp.	Ind. Prod.
β	753.04***	0.01	0.00	-0.00	-0.00	-0.00
Std. Errors	(94.38)	(0.02)	(0.00)	(0.00)	(0.01)	(0.01)
t-Statistics	7.98	0.34	0.17	-0.98	-0.01	-0.13
F-Statistics	63.66	0.11	0.03	0.96	0.00	0.02
R^2	0.37	0.00	0.00	0.01	0.00	0.00

Note: Table 5.1 reports the regression of the reduced-form residuals u_{it} from each VAR equation on the proposed instrument $\hat{w}_t^{AI}$, constructed from the 0.95 quantile of the AIU Index. The specification is $u_{it} = \alpha + \beta_i \hat{w}_t^{AI} + \eta_{it}$. Robust standard errors are reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table 5.2: Instrument Exogeneity (0.95 Quantile)

	Dependent Variable: $\tau = 0.95$			
	(1)	(2)	(3)	(4)
EPU	−0.02			
	(0.02)			
REU		0.01		
		(0.02)		
VIX			−0.02	
			(0.02)	
VXN				−0.02
				(0.02)
Constant	0.05**	0.05**	0.05**	0.05**
	(0.02)	(0.02)	(0.02)	(0.02)
R^2	0.01	0.00	0.01	0.01

Note: Table 5.2 reports regressions of the tail realisation instrument (0.95 quantile of the AIU Index) on structural uncertainty shocks from benchmark measures. The specification is $\theta_t^{IV} = \alpha_j + \beta_j \varepsilon_t^j + \eta_{jt}$. Structural shocks ε_t^j are extracted from separate SVARs for each uncertainty measure (EPU, REU, VIX, VXN) using recursive identification with the uncertainty measure ordered first. Robust standard errors in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table 5.2 reports the exogeneity assessment. The coefficients on benchmark uncertainty shocks are uniformly small in magnitude and statistically insignificant, and the corresponding R^2 values are below 1.0%. Benchmark uncertainty shocks,

therefore, account for only a small fraction of the variation in the tail-based instrument. The absence of systematic correlation with shocks identified from the EPU Index, the REU Index, the VIX, and the VXN indicates that the instrument is not simply reproducing general macroeconomic or financial uncertainty. Overall, these results provide a strong indication that the tail realisation instrument captures variation specific to AI-related uncertainty rather than general economic uncertainty conditions. On this basis, the 0.95 quantile of the AIU Index is retained as the baseline tail realisation instrument, and the 0.90 quantile is included as a robustness check.

Residual-Based Instrument. The residual-based instrument μ_t^{unc} yields a coefficient of 1.1 with an F-statistic of 10.3 in the targeted equation, as shown in (Table 5.3). This also satisfies the conventional criterion for instrument relevance (Stock and Yogo, 2002).

Table 5.3: Instrument Relevance (Residual-Based)

	AI Unc.	**S&P 500**	**Wage**	**Hours**	**Emp.**	**Ind. Prod.**
β	1.10***	-0.00	-0.00	-0.00**	-0.00	-0.00
Std. Errors	(0.34)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
t-Statistics	3.22	-1.38	-0.29	-2.23	-0.26	-0.49
F-Statistics	10.34	1.91	0.08	4.97	0.07	0.24
R^2	0.09	0.02	0.00	0.04	0.00	0.00

Note: Table 5.3 reports the regression of the reduced-form residuals u_{it} from each VAR equation on the residual-based instruments derived from Equation (4.9). Specifically, the instrument corresponds to μ_t^{unc} to identify "pure" AI uncertainty shock. The estimation follows the specification $u_{it} = \alpha + \beta_i \hat{v}_t^{AI} + \eta_{it}$. Robust standard errors are reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

In all non-targeted equations, the estimated coefficients associated with the residual-based instrument remain small, statistically insignificant, and accompanied by low R^2 values and weak F-statistics. The residual-based instrument is therefore strongly related to the residual in the targeted equation and effectively unrelated to the remaining reduced-form errors. This pattern indicates that the predictive content of the instrument is concentrated in the targeted equation, in this case, the AI uncertainty equation, rather than spread across multiple equations, consistent with the identifying assumption that the instrument loads on a single structural shock.

To assess exogeneity, Table 5.4 reports regressions of the residual-based instrument on benchmark structural uncertainty shocks. Benchmark shocks derived from the EPU Index, the REU Index, the VIX, and the VXN yield coefficients that are close to zero

and statistically insignificant across all specifications. In addition, the associated R^2 values remain at or below 2.0%, and the standard errors are of similar magnitude to the point estimates, which signals limited explanatory power. These results indicate that benchmark uncertainty shocks do not account for meaningful variation in the residual-based instrument.

Taken together with the results for the tail realisation instrument, the evidence indicates that both instruments isolate movements specific to AI-related uncertainty rather than broader macroeconomic or financial uncertainty. This joint evidence supports the interpretation of the identified shocks in the subsequent structural analysis and mitigates concerns about weak or contaminated instruments.

Table 5.4: Instrument Exogeneity (Residual-Based)

	Dependent Variable: μ_t^{unc}			
	(1)	(2)	(3)	(4)
EPU	10.37			
	(6.56)			
REU		8.85		
		(6.58)		
VIX			7.71	
			(6.59)	
VXN				8.70
				(6.58)
Constant	−0.06	−0.06	−0.06	−0.06
	(6.53)	(6.55)	(6.56)	(6.55)
R^2	0.02	0.02	0.01	0.02

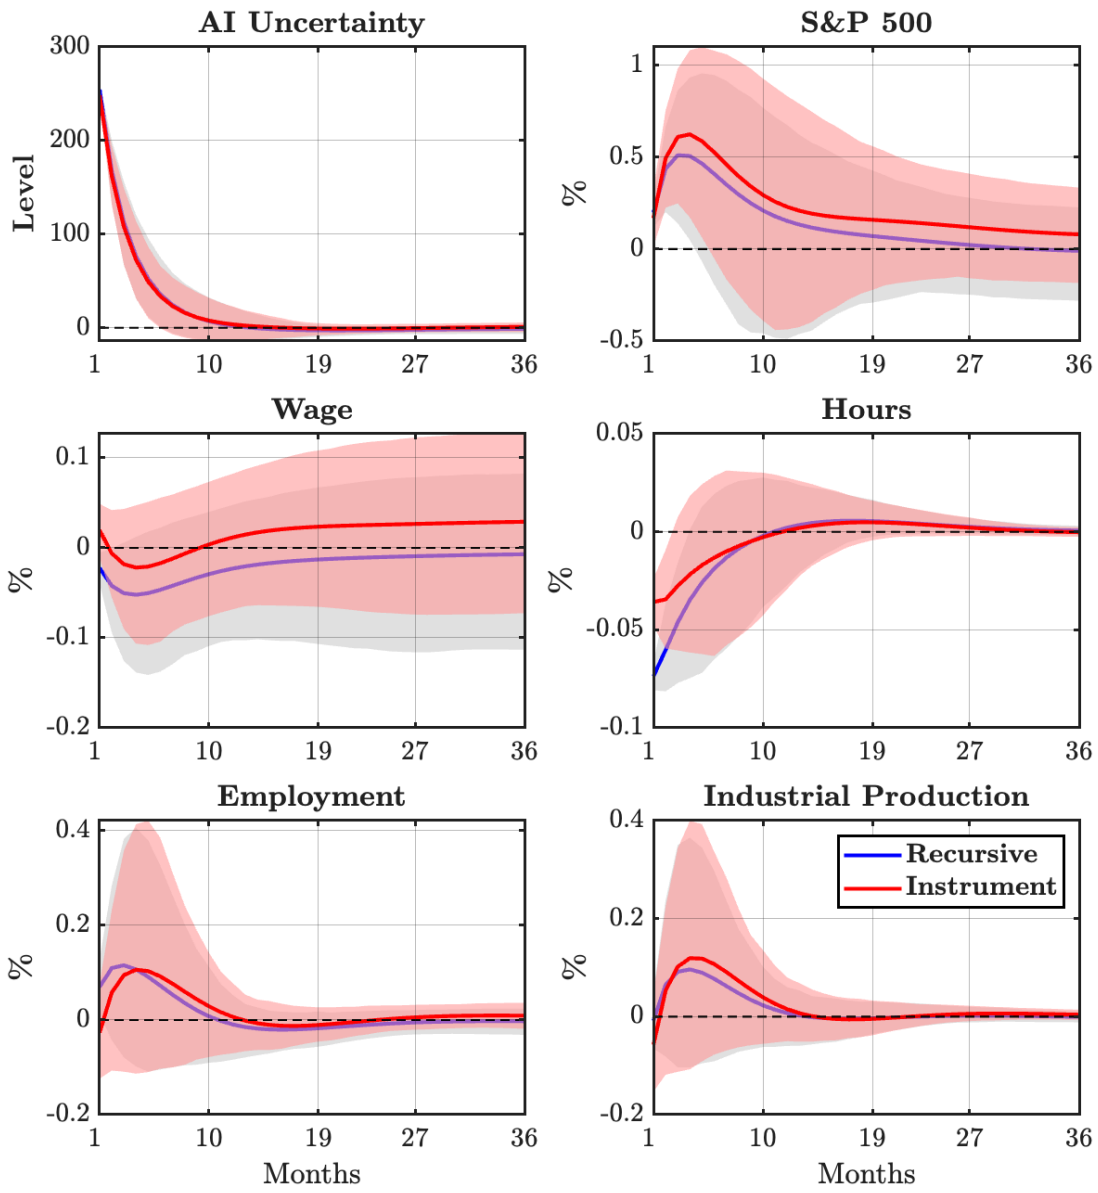
Note: Table 5.4 reports regressions of the residual-based instrument (μ_t^{unc}) on structural uncertainty shocks from benchmark measures. The specification is $\theta_t^{IV} = \alpha_j + \beta_j \varepsilon_t^j + \eta_{jt}$ (Equation 5.2). Structural shocks ε_t^j are extracted from separate SVARs for each benchmark uncertainty measure (EPU Index, REU Index, VIX, and VXN) using recursive identification with the uncertainty measure ordered first. Robust standard errors in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

5.2. Aggregate-Level Impulse Response

Composite AI Uncertainty Shock. The impulse responses in Figure 5.1 correspond to a one standard deviation AI uncertainty shock identified using a recursive SVAR and the SVAR-IV based on the tail realisation of the AIU Index. Both approaches yield similar qualitative patterns, which supports the robustness of the results. However,

to maintain a consistent benchmark for interpretation in this subsection, I focus the discussion on the SVAR-IV results obtained with the tail realisation instrument.

Figure 5.1: Impulse Responses to AI Uncertainty Shock



Note: Figure 5.1 displays the impulse responses to a one standard deviation shock in AI uncertainty estimated with an SVAR under recursive identification (blue line) and SVAR-IV (red line). The instrument used is a binary IV equal to 1 when the AIU Index is greater than or equal to its 0.95 quantile and 0 otherwise. Shaded areas denote 68% confidence bands based on 1,000 wild bootstrap replications.

Financial markets display the most pronounced adjustment. The S&P 500 expands by approximately 0.6% on impact and remains moderately above baseline throughout the 36-month horizon, stabilising near 0.2%. Despite wide confidence bands, the response is consistently positive throughout the horizons. This response differs from the contractionary equity responses typically associated with broader economic uncertainty shocks (Bloom, 2009; Jurado et al., 2015; Baker et al., 2016).

The labour market displays more moderate and asymmetric responses. Hours worked decline slightly on impact, by less than 0.1%, and return to baseline within a few months. This behaviour is robust across identification strategies. Wages remain close to zero throughout the horizon, while employment registers a brief increase in the early months before returning to baseline within the first year. These results suggest that AI-related uncertainty is primarily absorbed through short-lived adjustments in hours worked rather than sustained changes in wages or employment.

Real activity displays a similar degree of stability. Industrial production increases slightly on impact, but the effect is transitory, and the confidence bands include zero for most horizons. Relative to the more pronounced adjustments in equity prices and hours worked, the limited response of industrial production suggests that real activity plays only a minor role in the near-term transmission of AI-related uncertainty shocks.

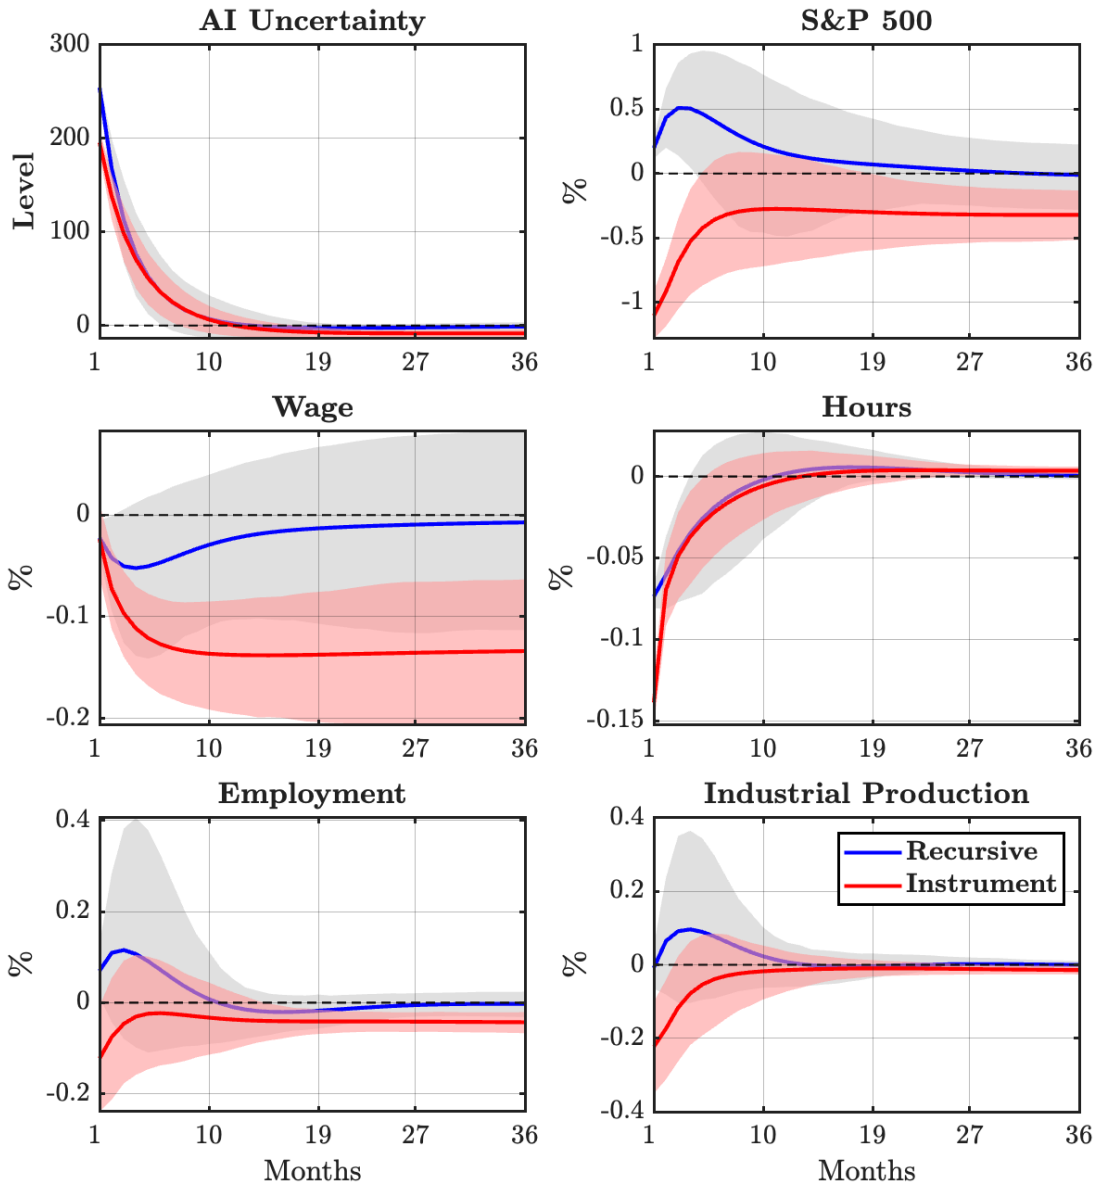
Overall, the results highlight a transmission mechanism that differs from the well-established impulse responses in the uncertainty literature. The combination of a positive and persistent equity response, a temporary decline in hours worked, and negligible effects on wages and industrial production is consistent with the view that the AI uncertainty shock identified using a dummy variable is unlike to represent a pure uncertainty shock. Since the instrument employed is equal to 1 when the AIU Index exceeds the 0.95 quantile, it could also select a period of unusually intense AI-related news coverage. These episodes likely combine uncertainty with expectations about the potential of AI to boost productivity or economic growth. In this setting, the identified shock reflects contractionary elements linked to uncertainty, together with expansionary forces associated with anticipated efficiency gains (Piffer and Podstawski, 2017; Cascaldi-Garcia and Galvao, 2021).

The positive and persistent response of equity prices, together with modest adjustments in real activity, indicates that expectations of future productivity improvements play a central role in the transmission of the identified AI uncertainty shock. This evidence underscores the importance of distinguishing between shocks that primarily capture pure AI uncertainty and shocks that reflect broader narratives of technological progress.

Pure AI Uncertainty Shock. Isolating the uncertainty component within AI-related coverage reveals effects that differ markedly from those generated by the composite AI uncertainty shock. Using the residual-based instrument introduced in Section 4, I extract the portion of the AIU Index that is orthogonal to innovation-oriented narratives and identify a shock that captures pure uncertainty about AI rather than

expectations of technological progress. This procedure removes the confounding present in the baseline specification and delivers a cleaner measure of AI-specific uncertainty. The resulting impulse responses are presented in Figure 5.2.

Figure 5.2: Impulse Responses to AI Uncertainty Shock



Note: Figure 5.2 displays the impulse responses to a one standard deviation pure AI uncertainty shock (red line) identified using SVAR-IV. The residual-based instrument isolates distinct components of AI-related coverage, separating discussions that frame AI as a source of uncertainty from broader narratives linked to productivity and innovation. The identified shock, therefore, captures adverse assessments related to labour displacement, regulatory challenges, and sectoral disruption, among others. Shaded areas denote 68% confidence bands based on 1,000 wild bootstrap replications.

Financial markets display a marked and persistent decline. The S&P 500 falls by approximately 1.1% on impact and remains below the baseline throughout the 36-month horizon. Relative to the baseline case, where equity prices increase, this reversal indicates that the earlier positive response in valuations is driven by the

innovation-related component embedded in the AIU Index. Once this component is removed, the remaining uncertainty shock exerts sustained downward pressure on financial markets.

Labour market responses also become more pronounced. Hours worked decline by nearly 0.2% on impact and remain below baseline before reverting around the 10-month horizon. Wage, in contrast, displays a persistent decline. The initial contraction of about 0.1% deepens to nearly 0.2% and remains below the baseline throughout the horizon. At the extensive margin, employment decreases by around 0.1% on impact, but this adjustment proves transitory. These results indicate that AI-related uncertainty generates larger and more persistent adjustments in labour utilisation, with the strongest amplification in hours worked and wages.

Output also contracts in response to pure AI uncertainty. Industrial production falls by roughly 0.2% on impact before gradually reverting toward baseline. Although the decline is less persistent than the responses of S&P 500 and labour market outcomes, it nonetheless signals that uncertainty about AI disrupts real activity in the short run.

Taken together, these results indicate that the movement in the composite AI uncertainty shock reflects the combined influence of two distinct elements embedded in AI-related coverage. The uncertainty component generates contractionary adjustments in equity valuations, wages, labour utilisation, and output. Expectations of productivity improvements create an expansionary influence that is most visible in equity markets. When these elements appear together in the AIU Index, their effects partially offset each other, leading to the mixed responses observed under the baseline specification. The decomposition shows that AI uncertainty on its own has a measurable and economically relevant impact on macroeconomic activity, while references to potential productivity gains shape the overall response captured by the AIU Index.

5.3. Forecast Error Variance Decomposition

The forecast error variance decomposition (FEVD) highlights the contrasting roles of composite and pure AI uncertainty shocks and clarifies transmission channels that appear muted in the baseline specification. The composite AI uncertainty shock accounts for relatively modest and persistent shares of forecast variance across variables. By contrast, the pure AI uncertainty shock explains a larger share of fluctuations in the short run and concentrates its effects in specific margins of adjustment, as reported in Table 5.5.

Table 5.5: Forecast Error Variance Explained by AI Uncertainty Shock

	AI Unc.	S&P 500	Wage	Hours	Emp.	Ind. Prod.
Comp. AI Uncert.						
h=1	95.56	1.34	0.33	2.90	0.55	0.40
h=6	87.32	2.81	3.31	5.26	4.84	3.92
h=12	79.49	2.06	3.57	6.36	6.27	4.91
h=24	76.18	2.17	2.50	6.61	6.21	4.87
h=36	74.99	2.16	2.17	6.61	6.17	4.87
Pure AI Uncert.						
h=1	62.25	10.90	2.60	24.00	0.26	1.22
h=6	60.67	6.40	14.89	16.15	0.21	0.69
h=12	55.06	5.59	23.97	14.02	0.21	0.65
h=24	52.81	7.02	27.80	13.61	0.45	0.68
h=36	52.40	8.51	28.40	13.61	0.77	0.71

Note: Table 5.5 reports the FEVD at horizons 1 to 36. Each entry indicates the percentage share of forecast error variance in the macroeconomic variables explained by the baseline AI uncertainty shock, identified using the tail realisation of the AIU Index within the SVAR-IV framework, and by the pure AI uncertainty shock, identified using the residual-based instrument within the same strategy. All values are expressed as percentages.

In financial markets, the composite AI uncertainty shock explains about 2.0% of the forecast variance in the S&P 500 from the one-year horizon onward, with contributions that remain stable at medium and longer horizons. The pure AI uncertainty shock, on the other hand, explains nearly 11.0% of the variance on impact, and its contribution gradually declines to just under 9.0% by the 36-month horizon. This contrast shows that the composite measure understates the influence of an AI uncertainty shock on equity markets.

Labour market outcomes display an even sharper divergence. At the 12-month horizon, the composite AI uncertainty shock accounts for around 3.0% of the variance in wages, 6.0% in employment, and 7.0% in hours worked, with all contributions declining over time. The pure AI uncertainty shock, on the other hand, shifts explanatory power decisively towards hours worked and wages. On impact, it explains 24.0% of the variance in hours worked, before declining to approximately 14.0% by the end of the horizon. Meanwhile, its contribution to wages increases over time and dominates at longer horizons, reaching 28.0% after three years, while its effect on employment remains negligible throughout. However, its effect on employment remains negligible throughout. These variance shares indicate that pure AI uncertainty

plays a substantial role in fluctuations in hours worked and wage dynamics, but not in movements in employment.

For output, the distinction in variance shares is equally clear. The composite AI uncertainty shock explains almost 5.0% of the forecast variance of industrial production from the 12-month horizon onward, whereas the pure AI uncertainty shock accounts for roughly 1.0% at all horizons. This pattern indicates that the explanatory power of the baseline measure for output largely reflects its overlap with narratives of technological progress and productivity gains. Once these narratives are excluded, the contractionary element of pure AI uncertainty exerts only a limited influence on output fluctuations. In other words, while AI uncertainty generates measurable short-run disruptions in real activity, its contribution to the overall variance of output is small relative to its stronger effects on equity valuations, wages, and hours worked.

Evidence from the impulse responses (Figure 5.2) and the FEVD establishes that pure AI uncertainty shock primarily affects flexible margins and operates over shorter horizons. Relative to the broad uncertainty shocks identified in earlier studies (Bloom, 2009; Jurado et al., 2015; Baker et al., 2016; Piffer and Podstawski, 2017; Cascaldi-Garcia and Galvao, 2021), the pattern documented here features front loaded and limited persistence rather than protracted contractions characteristic of traditional uncertainty episodes. Earlier contributions typically find that broad economic or financial uncertainty shocks induce substantial and persistent declines in employment and industrial production, with effects either materialising immediately or reaching maximum depth several months after the shock. In contrast, the AI uncertainty shock considered here induces an abrupt adjustment concentrated in flexible margins. Equity valuations and hours worked react on impact, wages follow a slower yet persistent adjustment path, industrial production exhibits a brief contraction, and employment remains close to baseline throughout the horizon. This profile highlights the importance of distinguishing across sources of uncertainty when assessing their macroeconomic implications.

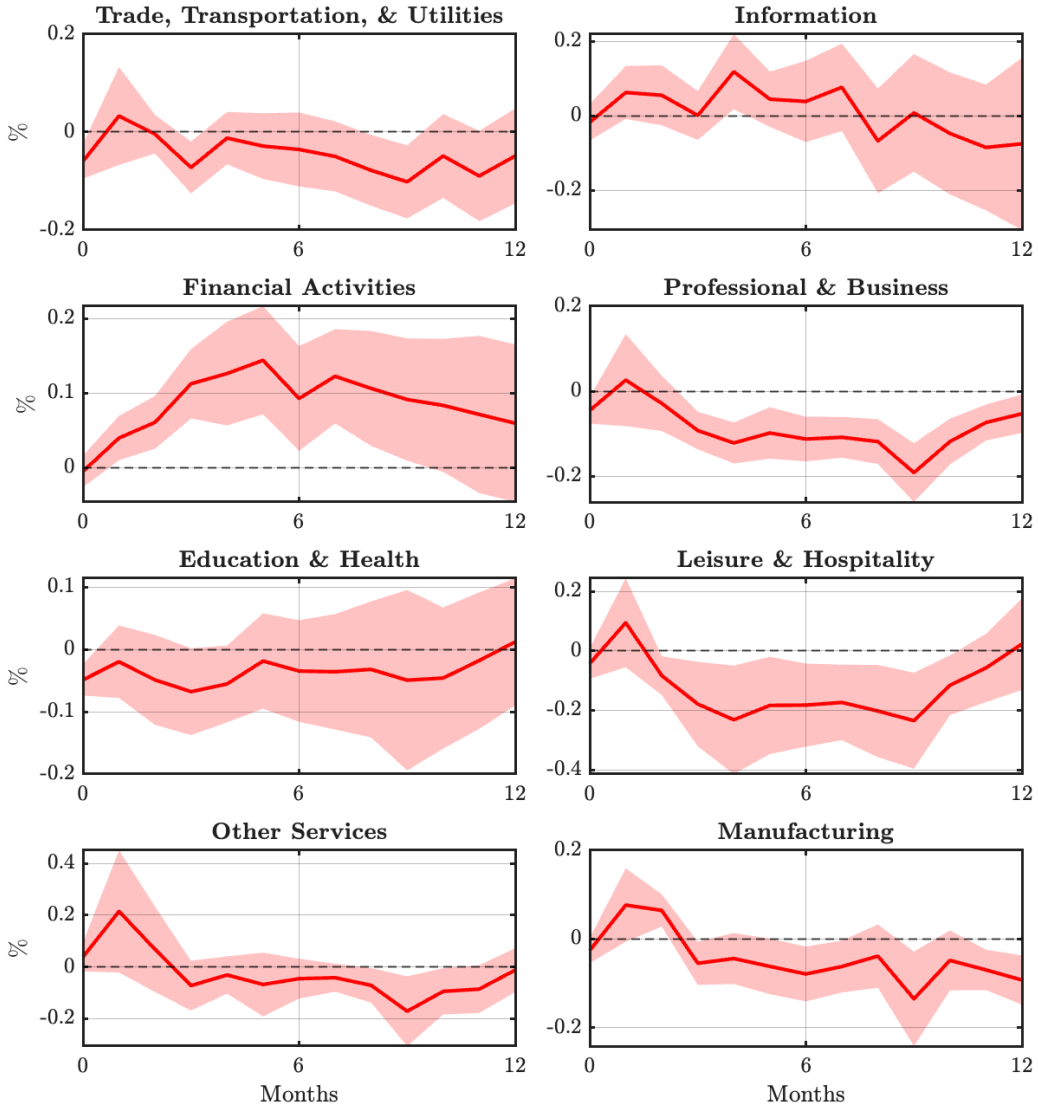
5.4. Industry-Level Impulse Response

To provide a comprehensive assessment of the macroeconomic effects of AI uncertainty shocks, I extend the analysis to the industry level. Following the approach of Känzig (2023), I use the AI uncertainty shock identified in the SVAR-IV using the residual-based instrument and use it as a regressor in a local projection framework (Jordà, 2005). This approach is used to estimate the dynamic responses of wage, hours

worked, and employment across US industries classified under the North American Industry Classification System (NAICS).

Wages. Industry-level wage responses to pure AI uncertainty shocks display substantial heterogeneity, although most industries display contractions consistent with the aggregate response (Figure 5.3).

Figure 5.3: Response of Industry-Level Wage to AI Uncertainty Shock



Note: Figure 5.3 presents the industry-level wage impulse responses to a one standard deviation pure AI uncertainty shocks, identified using an SVAR-IV and estimated using local projections. The industries are classified according to NAICS. The shaded regions represent 68% confidence intervals, computed using Newey–West standard errors.

Professional and business services, along with leisure and hospitality, register the most pronounced declines, with wages falling approximately 0.2% within six months and remaining depressed throughout the 12-month horizon. Manufacturing industry,

on the other hand, display a distinct two-phase pattern, with an initial expansion of roughly 0.1% on impact followed by a reversal to similar magnitude by the end of the horizon. These persistent contractions in multiple industries align with the aggregate finding that pure AI uncertainty exerts sustained downward pressure on compensation.

Financial activities, however, follow a persistent distinct pattern. Wages in this industry increase by approximately 0.1% that peak around the 5-month horizon before moderating toward baseline. This positive response contrasts sharply with both the aggregate pattern and the predominantly negative movements in other industries.

The remaining industries display more muted adjustment. Trade, transportation, and utilities exhibit a slight downward tendency, while information, education and health, and other services remain close to baseline throughout the horizon. Wide confidence intervals in these sectors indicate that observed movements are not statistically significant.

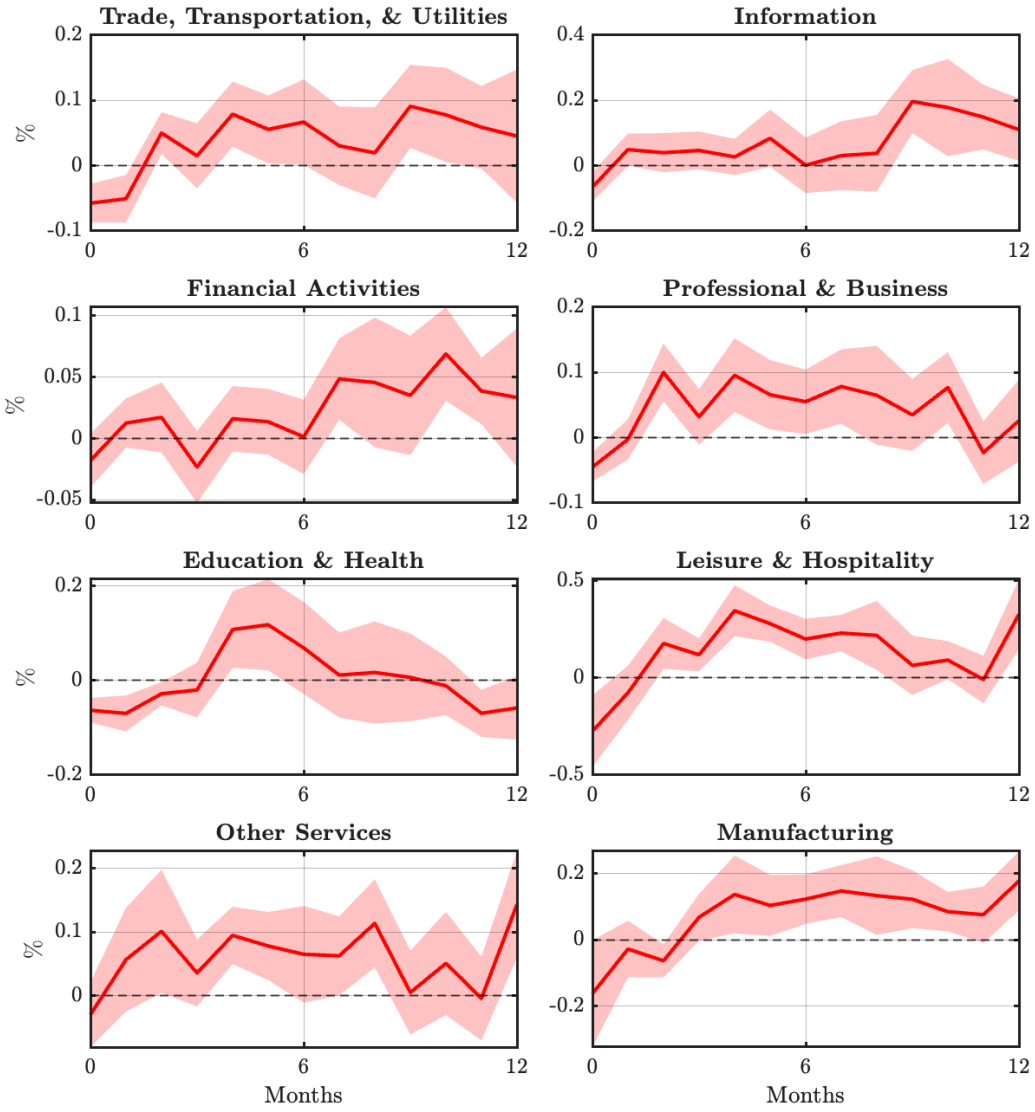
Overall, the industry-level analysis reveals substantial heterogeneity in wage responses to pure AI uncertainty shocks that the aggregate results obscure. While most sectors experience persistent declines, financial activities exhibit a countercyclical response, and several industries show no statistically significant adjustment.

Hours Worked. The response of industry-level hours worked to pure AI uncertainty shocks is heterogeneous (Figure 5.4). Most industries experience an initial contraction followed by a recovery and, in several cases, a persistent expansion above baseline.

Leisure and hospitality registers the largest initial contraction, with hours declining by approximately 0.3% on impact. Hours worked then rise sharply, reaching roughly 0.2% at the 5-month horizon, and remain above baseline throughout the horizon. Education and health also follows a similar but more moderate pattern, with an initial decline of less than 0.1%, a temporary rebound at the 5-month horizon, and a gradual return toward baseline thereafter.

Manufacturing records modest contraction on impact followed by a gradual increase, with hours worked expanding to approximately 0.2 percent and remaining persistently above baseline. The confidence interval, however, is wide and indicate substantial uncertainty around the point estimate. Trade, transportation, and utilities follow a comparable pattern, with an initial impact decline of less than 0.1%, subsequent recovery, and a small expansion at long horizons. By contrast, information industry remain close to baseline initially and then drift upward to around 0.2% toward the end of the horizon.

Figure 5.4: Response of Industry-Level Hours to AI Uncertainty Shock



Note: Figure 5.4 presents the industry-level hours worked impulse responses to a one standard deviation pure AI uncertainty shocks, identified using an SVAR-IV and estimated using local projections. The industries are classified according to NAICS. The shaded regions represent 68% confidence intervals, computed using Newey–West standard errors.

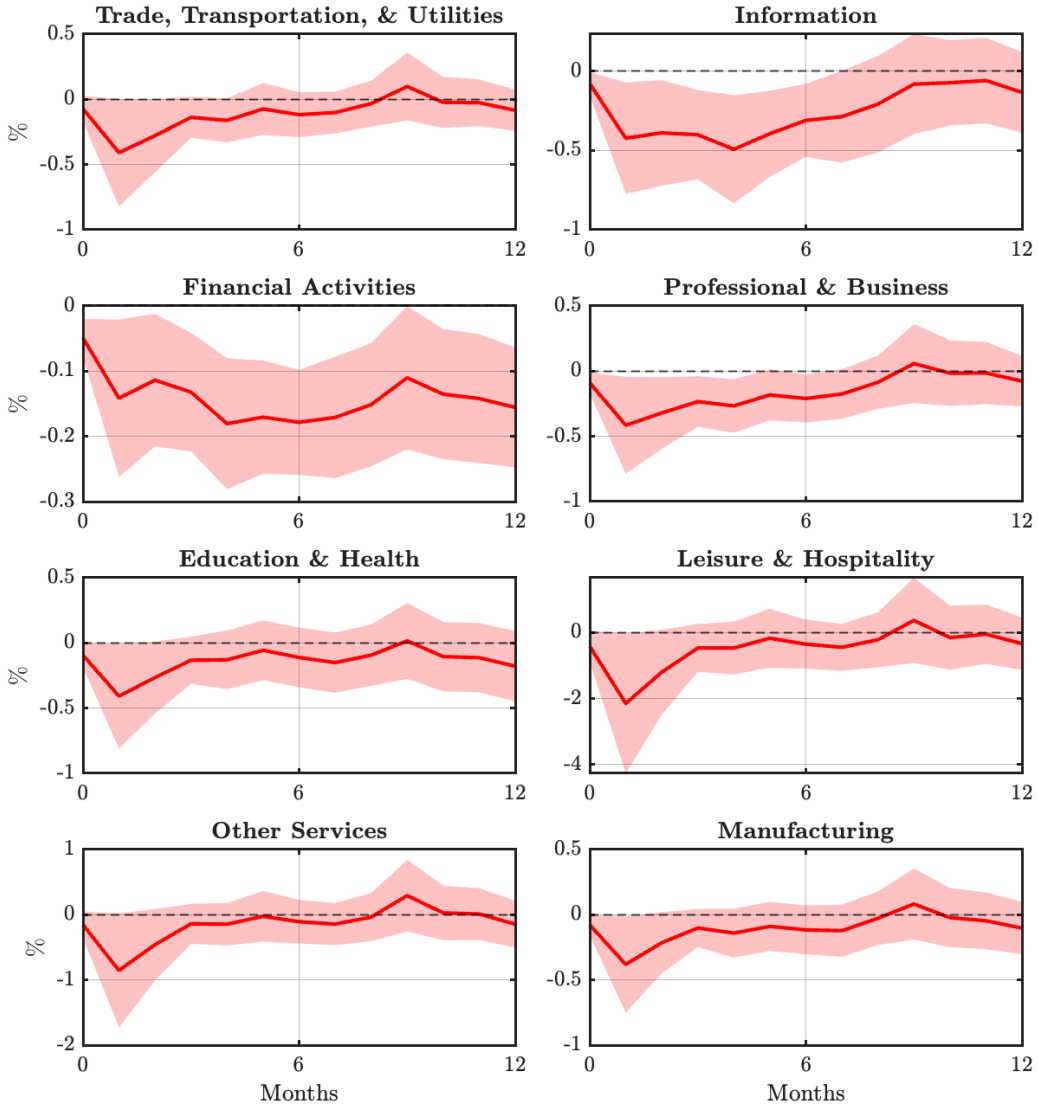
Professional and business services, financial activities, and other services record minimal movements on impact, with hours worked remaining near zero and then recording small positive responses that generally stay below 0.1%. In these industries, confidence intervals are wide relative to the point estimates. Hence, inference about the magnitude of the responses is imprecise.

The initial contractions observed across industries are consistent with the aggregate impulse response for hours worked, which also displays a decline on impact. The subsequent industry level expansions, however, differ from the aggregate pattern, where hours remain modestly below baseline throughout the horizon. This discrepancy could reflects compositional and aggregation effects that obscure the

recovery in hours worked at the industry level in the aggregate series.

Employment. The industry-level employment responses to pure AI uncertainty shocks broadly reflects the aggregate pattern of a modest and transitory declines, with most industries exhibiting small adjustments that revert toward baseline within the 12-month horizon (Figure 5.5). However, notable heterogeneity emerges in both the magnitude and persistence of contractions.

Figure 5.5: Response of Industry-Level Employment to AI Uncertainty Shock



Note: Figure 5.5 presents the industry-level employment impulse responses to a one standard deviation pure AI uncertainty shocks, identified using an SVAR-IV and estimated using local projections. The industries are classified according to NAICS. The shaded regions represent 68% confidence intervals, computed using Newey–West standard errors.

The information industry, along with professional and business services, experiences declines of approximately 0.5% before gradually returning to baseline.

Financial activities registers a more persistent contraction of roughly 0.2% that shows limited recovery throughout the horizon. The most pronounced response occurs in leisure and hospitality, where employment falls approximately 2.0% on impact, although wide confidence intervals surrounding this estimate substantially limit inferential precision. Meanwhile, the remaining industries display negligible or statistically insignificant movements.

These findings reveal that while the aggregate employment response masks substantial cross-industry variation, certain sectors face disproportionate exposure to labor market adjustments following pure AI uncertainty shocks. The heterogeneity in both magnitude and duration of employment responses suggests differential vulnerability across industries to AI-related uncertainty.

6. Conclusion

The rapid emergence of AI has intensified debate over its broader economic consequences. Optimistic assessments emphasise substantial productivity gains and stronger growth, while more cautious analyses anticipate modest improvements accompanied by labour displacement and transitional adjustment costs. The distance between these views, combined with the pace of AI diffusion, implies that uncertainty is an intrinsic feature of the current AI environment. This paper quantifies that uncertainty by developing the AIU Index, a novel text-based measure derived from newspaper coverage, and uses it to assess the macroeconomic effects of AI-related uncertainty.

The empirical analysis shows that, after separating the innovation-related component of AI coverage, pure AI uncertainty shocks generate distinctly contractionary effects, concentrated in the flexible margins of adjustment. Financial markets exhibit persistent declines, with the S&P 500 falling by around 1.1% on impact and remaining below baseline over the three-year horizon. In the labour market, hours worked contract by nearly 0.2% on impact and wages remain compressed, while employment stays close to baseline. Industrial production falls modestly in the short run and then returns toward its pre-shock level. These dynamics differ from those associated with broad uncertainty shocks documented in earlier work, which typically feature hump shaped and protracted contractions in employment and output (Bloom, 2009; Jurado et al., 2015; Baker et al., 2016). In contrast, AI-related uncertainty produces front-loaded adjustments in equity valuations, hours, and wages rather than

persistent, broad-based downturns.

Industry-level evidence reveals substantial heterogeneity around these aggregate responses. Most industries experience wage declines that mirror the aggregate pattern, while financial activities exhibit countercyclical wage increases. Hours worked contract on impact across industries, but then recover and move above baseline in several sectors, even though the aggregate series remains slightly below baseline. This pattern is consistent with compositional and aggregation effects that mute sectoral recoveries in the aggregate data. Employment responses are modest and short-lived in most sectors, although information, professional and business services, and leisure and hospitality experience more pronounced contractions. These cross-sectional results show that AI-related uncertainty does not propagate uniformly through the economy and that sectoral exposure to AI technologies shapes the incidence and form of adjustment.

While the results highlight important differences across industries, they also leave open questions about the mechanisms driving these heterogeneous responses. The present framework does not fully capture why certain industries absorb AI uncertainty through flexible margins while others show muted or even expansionary responses. To address this limitation, I am constructing industry-specific AI uncertainty indices that allow more precise identification of sector-level exposure to AI-related developments. This extension incorporates richer industry-level dynamics by accounting for differential exposure to AI technologies, variation in adjustment costs, and sector-specific production structures that mediate the transmission of uncertainty shocks.

Overall, the findings establish that AI-related uncertainty represents an emerging source of business cycle fluctuations, with transmission channels distinct from those of traditional uncertainty episodes. Even before the long-run effects of AI adoption on productivity and growth are realised, uncertainty about AI generates contractionary effects on financial markets, labour outcomes, and real activity. The combination of aggregate and industry-level evidence shows that AI uncertainty operates mainly through flexible margins of adjustment and that its impact is uneven across sectors, consistent with treating AI uncertainty as a separate component within the broader set of shocks that drive business cycle fluctuations.

References

Abiad, A. and Qureshi, I. A. (2023). The Macroeconomic Effects of Oil Price Uncertainty. *Energy Economics*, 125.

Acemoglu, D. (2025). The Simple Macroeconomics of AI. *Economic Policy*, 40(121):13–58.

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). Artificial Intelligence and Jobs: Evidence from Online Vacancies. *Journal of Labor Economics*, 40(S1):S293–S340.

Ash, E. and Hansen, S. (2023). Text Algorithms in Economics. *Annual Review of Economics*, 15:659–688.

Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring Economic Policy Uncertainty. *Quarterly Journal of Economics*, 131(4):1593–1636.

Bloom, N. (2009). The Impact of Uncertainty Shocks. *Econometrica*, 77(3):623–685.

Bloom, N. (2014). Fluctuations in Uncertainty. *Journal of Economic Perspectives*, 28(2):153–176.

Bonney, K., Breaux, C., Buffington, C., Dinlersoz, E., Foster, L., Goldschlag, N., Haltiwanger, J., Kroff, Z., and Savage, K. (2024). The Impact of AI on the Workforce: Tasks versus Jobs? *Economics Letters*, 244.

Caldara, D. and Iacoviello, M. (2022). Measuring Geopolitical Risk. *American Economic Review*, 112(4):1194–1225.

Caldara, D., Iacoviello, M., Molligo, P., Prestipino, A., and Raffo, A. (2020). The Economic Effects of Trade Policy Uncertainty. *Journal of Monetary Economics*, 109:38–59.

Carriero, A., Mumtaz, H., Theodoridis, K., and Theophilopoulou, A. (2015). The Impact of Uncertainty Shocks under Measurement Error: A Proxy SVAR Approach. *Journal of Money, Credit and Banking*, 47(6):1223–1238.

Cascaldi-Garcia, D. and Galvao, A. B. (2021). News and Uncertainty Shocks. *Journal of Money, Credit and Banking*, 53(4):779–811.

Cerutti, E., Pascual, A. G., Kido, Y., Li, L., Melina, G., Tavares, M. M., and Wingender, P. (2025). The Global Impact of AI: Mind the Gap.

Chui, M., Hazan, E., Roberts, R., Singla, A., Smaje, K., Sukharevsky, A., Yee, L., and Zemmel, R. (2023). The Economic Potential of Generative AI: The Next Productivity Frontier. Technical Report, McKinsey & Company.

Felten, E., Raj, M., and Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42:2195–2217.

Gentzkow, M., Kelly, B., and Taddy, M. (2019). Text as Data. *Journal of Economic Literature*, 57(3):535–574.

Gertler, M. and Karadi, P. (2015). Monetary Policy Surprises, Credit Costs, and Economic Activity. *American Economic Journal: Macroeconomics*, 7(1):44–76.

Goldman Sachs (2023). Generative AI Could Rise Global GDP by 7 Percent. https://www.goldmansachs.com/insights/articles/generative-ai-could-raise-global-gdp-by-7-percent.html. Accessed: 05 May 2025.

Gonçalves, S. and Kilian, L. (2004). Bootstrapping Autoregressions with Conditional Heteroskedasticity of Unknown Form. *Journal of Econometrics*, 123(1):89–120.

Hansen, S., McMahon, M., and Prat, A. (2018). Transparency and Deliberation Within the FOMC: A Computational Linguistics Approach. *Quarterly Journal of Economics*, 133(2):801–870.

Hodrick, R. J. and Prescott, E. C. (1997). Postwar U.S. Business Cycles: An Empirical Investigation. *Journal of Money, Credit and Banking*, 29(1):1–16.

Husted, L., Rogers, J., and Sun, B. (2020). Monetary Policy Uncertainty. *Journal of Monetary Economics*, 115:20–36.

Jordà, O. (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, 95(1):161–182.

Jurado, K., Ludvigson, S. C., and Ng, S. (2015). Measuring Uncertainty. *American Economic Review*, 105(3):1177–1216.

Knight, F. (1921). *Risk, Uncertainty and Profit*. Houghton Mifflin Company, Boston, New York.

Kroese, B. (2024). 12th IMF Statistical Forum: Measuring the Implications of AI on the Economy. https://www.imf.org/en/News/Seminars/Conferences/2024/11/20/12th-statistical-forum. Accessed: 19 December 2024.

Känzig, D. R. (2023). The Unequal Economic Consequences of Carbon Pricing. Working Paper Series 31221, National Bureau of Economic Research.

Lenza, M. and Primiceri, G. E. (2022). How to Estimate a Vector Autoregression After March 2020. *Journal of Applied Econometrics*, 37(4):688–699.

Mertens, K. and Ravn, M. O. (2013). The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States. *American Economic Review*, 103(4):1212–1247.

Piffer, M. and Podstawski, M. (2017). Identifying Uncertainty Shocks Using the Price of Gold. *The Economic Journal*, 128(616):3266–3284.

Romer, C. D. and Romer, D. H. (2010). The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks. *American Economic Review*, 100(3):763–801.

Sims, C. A., Stock, J. H., and Watson, M. W. (1990). Inference in Linear Time Series Models with Some Unit Roots. *Econometrica*, 58(1):113–144.

Stock, J. and Yogo, M. (2002). Testing for Weak Instruments in Linear IV Regression. NBER Technical Working Paper 0284, National Bureau of Economic Research.

Stock, J. H. and Watson, M. W. (2012). Disentangling the Channels of the 2007–2009 Recession. *Brookings Papers on Economic Activity*, pages 81–135.

Appendix A

Computation of the AI Uncertainty (AIU) Index

The AIU Index is constructed according to the following steps described below:

1. **ARTICLE SELECTION AND FILTERING PROCEDURE:**

   The construction of the AIU Index begins by identifying a corpus of news articles that consistently report on the economic dimensions of AI. Articles are sourced from the Factiva database, a comprehensive repository of global news content managed by Dow Jones. The selection process is designed to ensure both topical relevance and consistency in editorial standards. Moreover, it involves three (3) stages: (1) filtering via subject classifications, (2) exclusion of non-substantive content types, and (3) restriction to English-language articles. Each of these steps is detailed below:

   Factiva-Based Thematic Filtering: The first stage involves retrieving articles tagged with Factiva subject codes related to AI. Specifically, the search is restricted to the following categories: (1) `artificial intelligence`, (2) `machine learning`, (3) `risk topics - artificial intelligence`, (4) `automation`, and (5) `generative AI`. These subject codes are part of Factiva's internal classification system, which groups articles based on their substantive relevance to the assigned topic.

   Content-Type Exclusions: To ensure that the index captures substantive reporting on AI with potential economic relevance, I exclude articles classified under content types that generally lack detailed analysis or original reporting. Following the exclusion criteria in Abiad and Qureshi (2023), who also used Factiva to construct the Oil Price Uncertainty (OPU) Index, the corpus excludes the following categories: (1) `sports`, (2) `editorials`, (3) `abstracts`, (4) `advertorials or sponsored content`, (5) `advice`, (6) `analyses`, (7) `audio-visual links`, (8) `blogs`, (9) `event calendars`, (10) `chronologies`, (11) `columns`, (12) `commentaries or opinions`, (13) `corporate digests`, (14) `country profiles`, (15) `transcripts`, (16) `tables`, (17) `surveys or polls`,

(18) `statistics`, (19) `reviews`, (20) `rankings`, (21) `prospectuses`, (22) `press releases`, (23) `personal announcements`, (24) `people profiles`, (25) `front-page headlines`, (26) `obituaries`, (27) `letters`, (28) `interviews`, (29) `images`, and (30) `headline-only listings`. These exclusions aim to enhance the signal-to-noise ratio by focusing the analysis on articles that are likely to contain original reporting or analytical discussion relevant to the AIU Index.

Language and Source Restrictions: The dataset is restricted to English-language articles published in newspapers classified under Factiva's Top Newspaper source group. Restricting coverage to English ensures consistency in keyword matching and avoids biases introduced by translation or multilingual reporting. Limiting the sample to leading newspapers helps guarantee both archival completeness and editorial reliability, while also aligning with established practice in the construction of text-based indices such as the EPU Index (Baker et al., 2016), GPR Index (Caldara and Iacoviello, 2022), and OPU Index (Abiad and Qureshi, 2023). The selected outlets are as follows:

Table A: Factiva News Outlet

Country:	**News Outlet:**
United States	The Boston Globe, The Baltimore Sun, Chicago Tribune, Investor's Business Daily, The New York Times, New York Post, Pittsburgh Post-Gazette, USA Today, The Wall Street Journal, The Washington Post
United Kingdom	Daily Mail, The Daily Telegraph, Financial Times, The Guardian, The Independent, Reuters News, The Times
Euro Area	Agence France Presse, DW News, Euronews

2. KEYWORD FILTERING:

The objective of the keyword filtering process is to isolate news articles that capture economic uncertainty associated with developments in AI. Following the methodology of Baker et al. (2016) in constructing the EPU Index, I apply a structured Boolean keyword filter designed to identify articles that meet three inclusion criteria. In particular, an article must contain at least one term from each of the following categories:

AI-related terms → This includes keywords such as: `"artificial intelligence"`, `"artificial general intelligence"`, `"deep learning"`, `"generative ai"`, `"large language model"`, `"machine learning"`, `"natural language processing"`, `"neural network"`, `"alphabet"`, `"alphago"`, `"amazon"`, `"amd"`, `"anthropic"`, `"apple"`, `"chatbot"`, `"chatgpt"`, `"copilot"`, `"claude"`, `"deepmind"`, `"deepseek"`, `"gemini"`, `"google"`, `"grok"`, `"llama"`, `"meta"`, `"microsoft"`, `"nvidia"`, `"openai"`, `"oracle"`, `"perplexity"`, `"softbank"`, `"sora"`.

Economy-related terms → This includes keywords such as: `"bank*"`, `"business*"`, `"econom*"`, `"education*"`, `"employ*"`, `"financ*"`, `"firm*"`, `"fiscal"`, `"gdp"`, `"growth*"`, `"industr*"`, `"invest*"`, `"job*"`, `"labo*"`, `"layoff"`, `"macroeconom*"`, `"manufactur*"`, `"market**"`, `"microeconom*"`, `"monetar**"`, `"output"`, `"productivit*"`, `"recession*"`, `"retail*"`, `"sector**"`, `"service*"`, `"suppl*"`, `"supply chain*"`, `"trad*"`, `"unemploy*"`, `"wage*"`, `"work*"`, `"workforc*"`.

Uncertainty-related terms → This includes keywords such as: `"uncert*"`, `"ambigu*"`, `"anxi*"`, `"concern*"`, `"dilemma"`, `"doubt*"`, `"fear*"`, `"instabil*"`, `"risk*"`, `"unclear"`, `"unknown*"`, `"unpredict*"`, `"volat*"`, `"worry"`.

The use of wildcard-based stemming, such as `"uncert*"`, allows the search to capture multiple grammatical forms and journalistic variations while preserving thematic relevance. For example, the stem `"uncert*"` includes both "uncertainty" and "uncertain".

Further, an article is included in the sample only if it contains at least one term from each of the three categories. This Boolean filtering structure applies an `"AND"` condition across the main categories and an `"OR"` condition within each category. It ensures that all included articles explicitly discuss AI within an economic context and under conditions of uncertainty. Articles that mention AI without economic or uncertainty context, or discuss economic topics unrelated to AI, are excluded to maintain the conceptual coherence of the index.

3. AGGREGATION PROCEDURE:

Similar to the keyword filtering process, the aggregation procedure used to construct the AIU Index follows the methodology of Baker et al. (2016) for the

EPU Index. It employs a systematic procedure involving article-level filtering, within-source normalisation, and cross-source aggregation. In particular, the procedure consists of three (3) steps: (1) volume normalisation, (2) variance standardisation, (3) aggregation, and (4) renormalisation to index levels.

Volume Normalisation: For each news outlet i and month t, I calculate the proportion of articles that contain at least one keyword from each of the three categories, i.e. AI-related terms, economy-related terms, and uncertainty-related terms. This is given by:

$$\theta_{it} = \frac{\alpha_{it}}{\tau_{it}} \tag{A.1}$$

where: θ_{it} is the share of qualifying articles for outlet i in month t, α_{it} is the number of articles that contain at least one term from each of the three keyword categories, and τ_{it} is the total number of articles published by outlet i in month t.

This normalisation step adjusts for differences in publication frequency by using the proportion of relevant articles, rather than absolute counts.

Variance Standardisation: To adjust for outlet-specific reporting tendencies and heterogeneity in coverage intensity, each outlet series is scaled by its own standard deviation, σ_i, computed over a pre-specified baseline window T_{base}:

$$Y_{it} = \frac{\theta_{it}}{\sigma_i}, \quad \sigma_i = \text{stdev}\big(\theta_{it},\ t \in T_{\text{base}}\big). \tag{A.2}$$

where σ_i is the standard deviation of θ_{it} over the baseline period T_{base}. This procedure standardises the variance of each outlet's time series while leaving its mean unchanged. The purpose is not to re-centre the distribution but to ensure comparability across outlets by placing them on a common variance scale. Without this adjustment, outlets with more volatile reporting behaviour could dominate the aggregate index, introducing bias unrelated to underlying trends in AI-related uncertainty. By scaling relative to outlet-specific variability, the procedure balances contributions across sources and mitigates distortions from editorial styles or uneven publication intensity.

Aggregation: The variance-standardised series Y_{it} are then averaged across the

set of outlets with observations available in month t:

$$Z_t = \frac{1}{N_t} \sum_{i \in S_t} Y_{it}. \tag{A.3}$$

where S_t is the set of outlets with an observation in month t, and $N_t = |S_t|$ is the number of such outlets. The notation $|S_t|$ denotes the cardinality (number of elements) of the set S_t, not an absolute value. Averaging across the available outlets produces a single aggregate series that reflects broad-based patterns rather than the idiosyncratic behaviour of individual sources. This approach naturally accommodates an unbalanced panel (N_t can vary over time) and smooths outlet-specific fluctuations, thereby enhancing the robustness of the AIU Index.

Renormalisation to Index Levels: Finally, the aggregated series Z_t is renormalised relative to its mean value over the baseline period:

$$AIU_t = 100 \times \frac{Z_t}{\bar{Z}_{base}}, \tag{A.4}$$

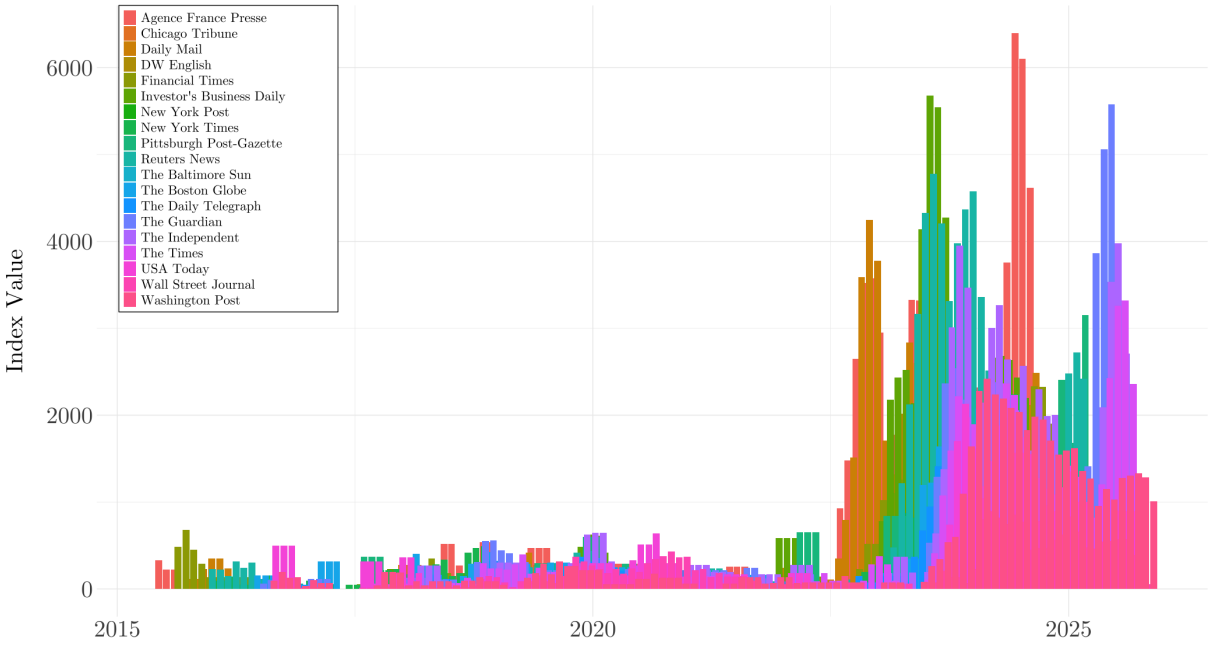
where AIU_t denotes the value of the AIU Index in month t, and $\bar{Z}_{base}$ is the mean of Z_t during the baseline period.

This final step expresses the index in percentage terms, with the baseline mean normalised to 100. The four-step procedure ensures that no single outlet disproportionately influences the index and mitigates biases arising from differences in publication frequency or archival depth.

The baseline period is set from M1:2016 to M12:2022, representing a phase of relatively stable attention preceding the sharp increase in interest in generative AI models from 2023 onwards. Standardising each outlet's time series against this benchmark improves comparability across sources with differing editorial priorities and publication volumes

Appendix B

Figure B: AIU Index per News Outlet (3-Month Moving Average)



Note: Figure B plots the AIU Index per news outlet from M1:2016 to M4:2025. The index is normalised such that its mean equals 100. It reflects the standardised share of newspaper articles that simultaneously reference AI, economic conditions, and uncertainty.

Appendix C

Table C.1.: Selected AI-Related Events (1/2)

Year:	Month:	Event:
2016	March	AlphaGo defeated world Go champion Lee Se-dol
2017	July	China Next Generation AI Development Plan
2018	June	OpenAI introduces GPT-1
2019	February	Donald Trump signs US Executive Order titled *American AI Initiative*
	February	OpenAI releases GPT-2
	April	Google shuts down its external AI Ethics Council (ATEAC)
	July	Microsoft invests USD 1 billion in OpenAI
2020	January	US Government limits exports of AI software
	May	OpenAI unveils GPT-3
2022	November	OpenAI launches ChatGPT
2023	January	Microsoft invests USD 10 billion in OpenAI
	February	Google announces and releases AI chatbot Bard
	February	Microsoft launches AI-powered Bing and Edge
	March	OpenAI releases GPT-4
	April	Italy temporarily bans ChatGPT
	April	UK invests in a supercomputer as part of AI strategy
	May	Sam Altman testifies before the US Senate
	June	European Parliament approves first draft of the EU AI Act
	July	US Federal Trade Commission (FTC) investigates OpenAI
	October	Joe Biden signs US Executive Order 14110 titled *Safe, Secure, and Trustworthy Development and Use of AI*
	November	UK AI Safety Summit held at Bletchley Park
	November	OpenAI Dev Day 2023. Launches GPT-4 Turbo, etc.
	November	OpenAI reinstates CEO Sam Altman
2024	January	OpenAI recruits team to manage AI risks during US election
	February	OpenAI unveils Sora
	March	Microsoft hires DeepMind co-founder Mustafa Suleyman
	April	Meta introduces Llama 3
	May	OpenAI releases GPT-4o
	July	CrowdStrike Global IT Outage

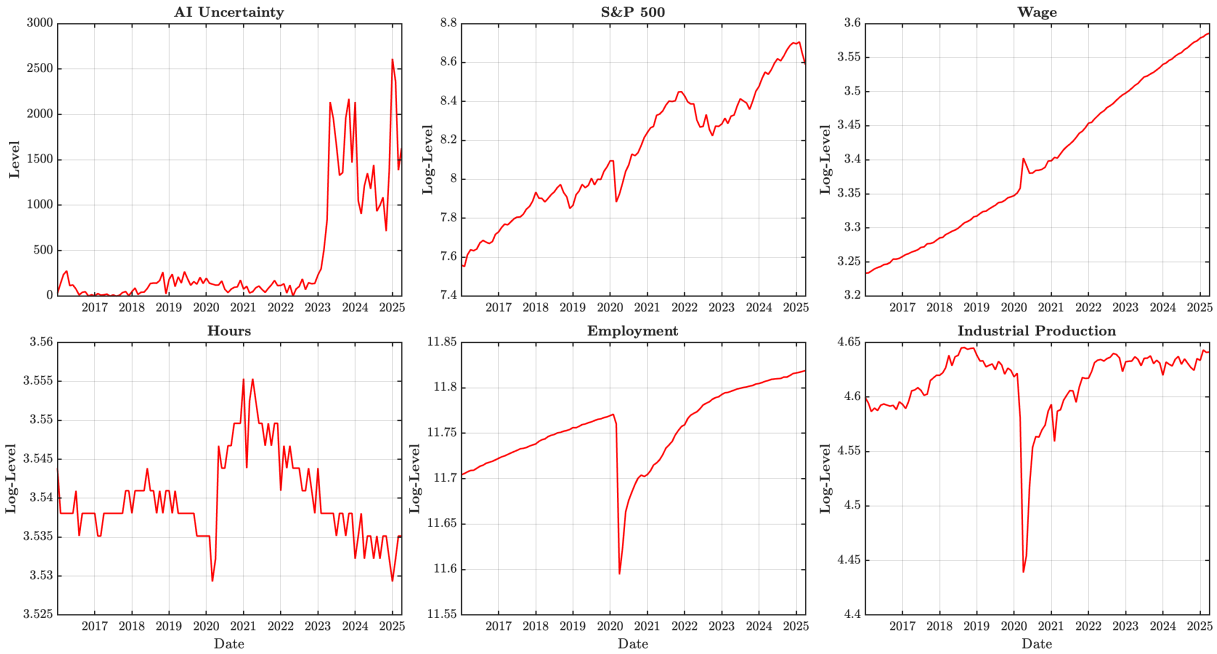
Appendix C

Table C.2.: Selected AI-Related Events (2/2)

Year:	Month:	Event:
2024	July	OpenAI Data Breach
	August	US election officials warned that Grok was spreading election misinformation
	October	Nvidia delayed the initial shipments of Blackwell
2025	January	DeepSeek releases R1 model
	February	OpenAI unveils Deep Research powered by o3
	February	Paris AI Summit
	April	US Gov't. controls the exports of Nvidia H20 chips to China

Appendix D

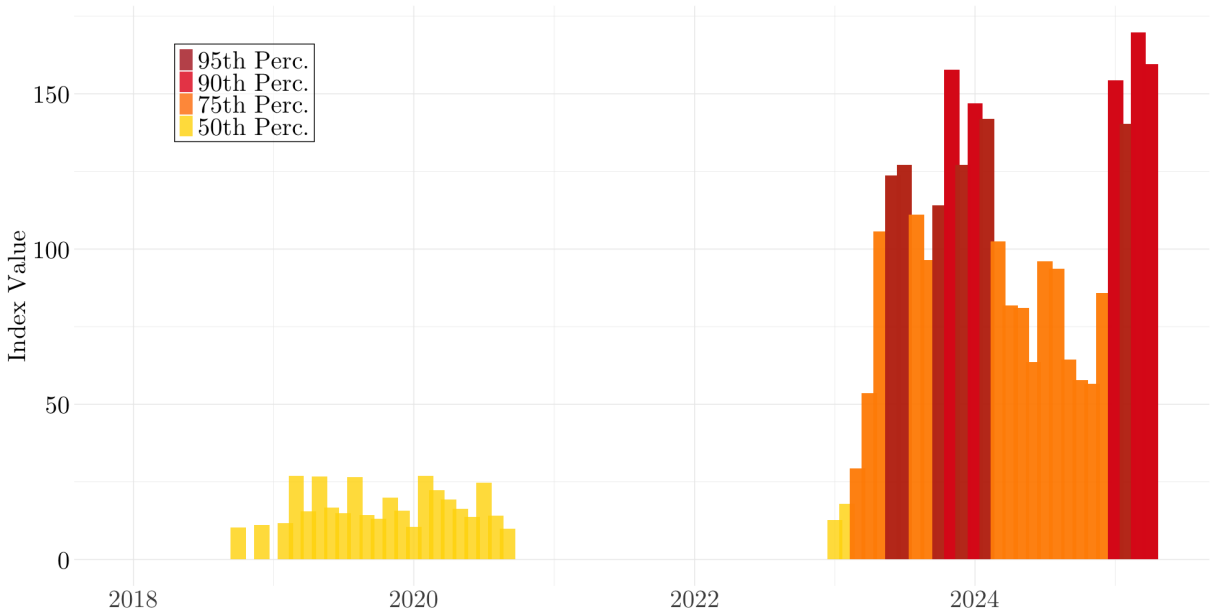
Figure D: Time Series of Variables



Note: Figure D plots the variables used in the analysis. The AIU Index is expressed in levels, while all other variables are expressed in log-levels. The data cover the period from M1:2016 to M4:2025. All series, except the AIU Index, are sourced from the US BLS and FRED unless otherwise noted.

Appendix E

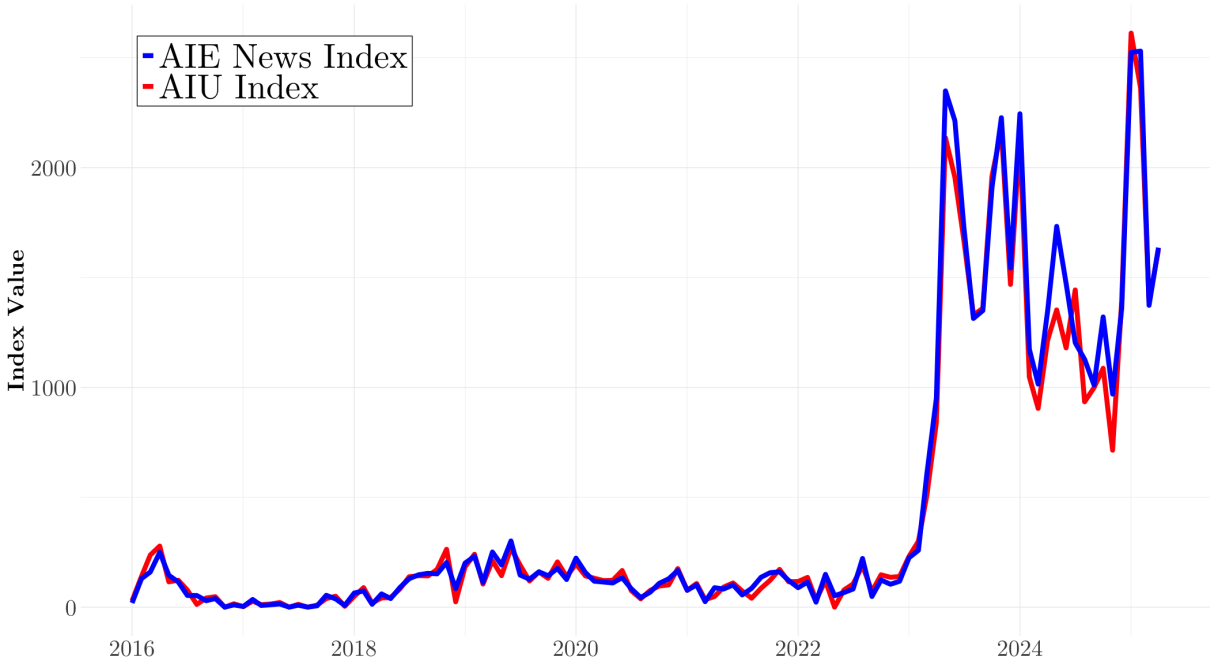
Figure E: AIU Index Upper Tail Realisations



Note: Figure E plots the tail realisations (0.95, 0.90, 0.75, and 0.50 quantiles) of the AIU Index from M1:2016 to M4:2025. The AIU Index reflects the standardised share of news articles that simultaneously reference AI, economic conditions, and uncertainty.

Appendix F

Figure F: AIE News Index (3-Month Moving Average)



Note: Figure F plots the AIE News Index from M1:2016 to M4:2025. For presentation purposes, a three-month moving average is applied. The AIE News Index follows the same construction steps as the AIU Index but excludes uncertainty-related keywords.

Appendix G

Table G: OLS Regression of AIU Index on AIE News Index

	AIU
AIE	0.93***
	(0.01)
Constant	4.68
	(7.66)
Observations	112
R^2	0.99

Note: Table G reports OLS regression results of the AIU Index on AIE News Index. Robust standard errors in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$. The sample covers M1:2016 to M4:2025, reflecting the availability of news articles used in constructing the AIU and AIE News indices.

Appendix H

Robustness to Outlier Exclusion

This appendix evaluates the sensitivity of the results to extreme observations in the AIU Index. Two econometric issues motivate the analysis:

1. **IMPULSE RESPONSE FUNCTIONS:**

   The impulse response functions (IRFs) are obtained via ordinary least squares (OLS) estimation of the vector autoregression (VAR) model, and a few extreme observations may function as influential observations with high leverage on the estimated coefficients. If this occurs, the dynamic responses reported in the main analysis could be driven disproportionately by these influential points rather than by systematic features of AI-related uncertainty.

2. **CONFIDENCE INTERVAL:**

   The wild bootstrap used to construct confidence intervals is valid under the assumption of heteroskedasticity, and its consistency requires serially uncorrelated residuals and the absence of structural breaks in the residual distribution. The most pronounced forecast errors in the AI uncertainty equation may signal changes in the variance or tail behaviour of the residuals. If the residual distribution changes during periods of extreme uncertainty, the bootstrap may not accurately approximate the sampling distribution of the impulse responses, leading to confidence intervals with incorrect coverage.

Methodology. To evaluate both concerns, I re-estimate the structural vector autoregression with instrumental variable (SVAR-IV) after excluding observations where the standardised residual in the AIU equation exceeds $|z| > 3$ or $|z| > 2$. I then compare the impulse responses and confidence intervals with those obtained from the full sample.

Following the estimation of SVAR on the full sample, I obtain the reduced-form residuals for each equation:

$$u_{i,t} = y_{i,t} - \hat{y}_{i,t}, \quad i = 1, \ldots, k, \quad t = p + 1, \ldots T, \tag{H.1}$$

where $y_{i,t}$ denotes the observed value of variable i at time t and $\hat{y}_{i,t}$ is the fitted value.

To identify observations with unusually large forecast errors, I standardise the residuals by their sample standard deviation:

$$z_{i,t} = \frac{u_{i,t}}{\hat{\sigma}_i}, \quad \hat{\sigma}_i = \sqrt{\frac{1}{T-p} \sum_{t=p+1}^{T} u_{i,t}^2} \tag{H.2}$$

where $\hat{\sigma}_i$ is the sample standard deviation of the residual from equation i, computed over $T - p$ observations after accounting for p lags. By construction in OLS estimation, VAR residuals have mean zero, which implies that $\hat{\sigma}_i$ reduces to the root mean squared error without requiring a mean adjustment. The standardised residual $z_{i,t}$ therefore expresses each forecast error in units of equation-specific standard deviations. Under approximate normality and constant variance, these standardised residuals follow a standard normal distribution.

Table H.1: AIU Equation VAR Residuals Exceeding Outlier Thresholds

Category:	**Observation No.:**	**Date:**	**Z-score:**
$\lvert z\rvert > 2$	89	2023-05	4.58
	94	2023-10	2.40
	97	2024-01	2.79
	98	2024-02	-3.39
	108	2024-12	2.07
	109	2025-01	4.82
	111	2025-03	-2.71
$\lvert z\rvert > 3$	89	2023-05	4.58
	98	2024-02	-3.39
	109	2025-01	4.82
Total $\lvert z\rvert > 2$ *observations: 6.3% of sample*			
Total $\lvert z\rvert > 3$ *observations: 2.8% of sample*			

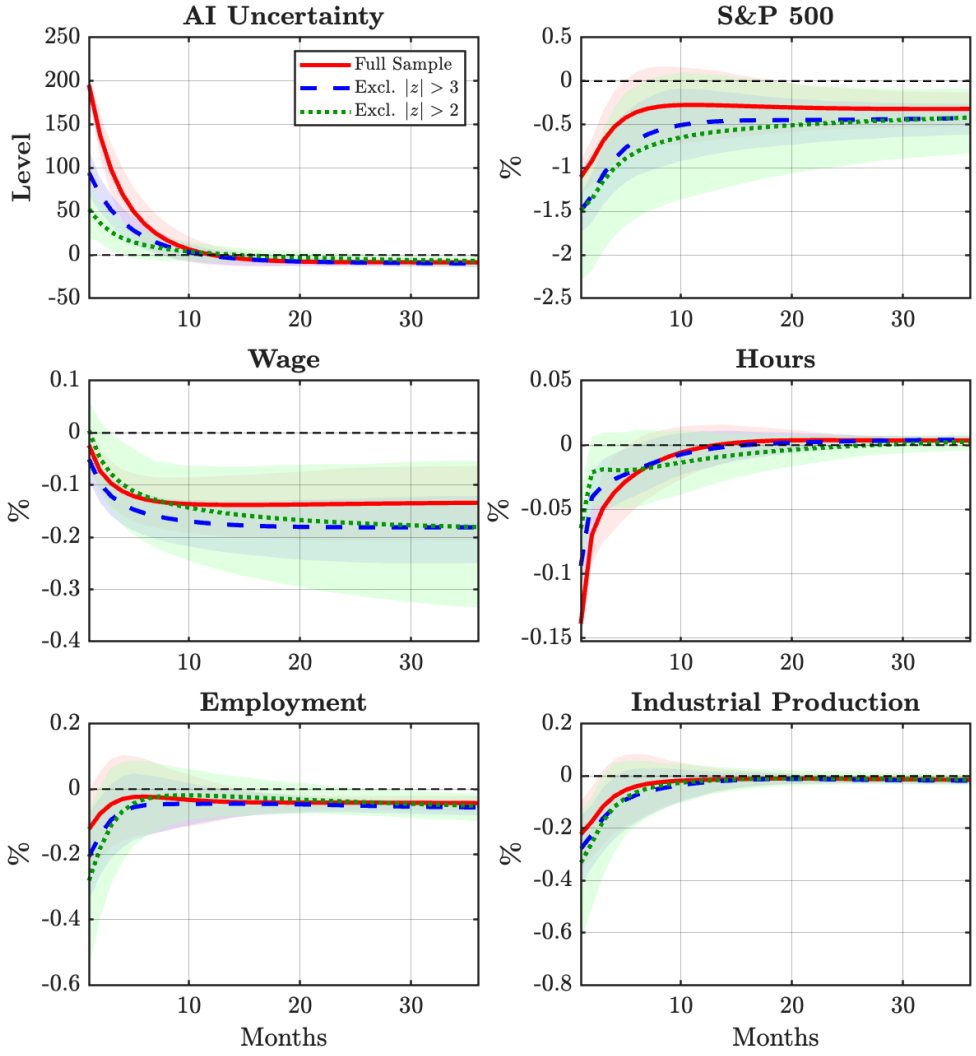
Note: Table H.1 reports the standardised residuals from the VAR(p) AI uncertainty equation exceeding normality thresholds. Standardisation: $z_t = \hat{u}_t / \hat{\sigma}$ where $\hat{\sigma}$ is the residual standard deviation.

Observations with $|z| > 3$ correspond to forecast errors more than three standard deviations from zero and are treated as extreme outliers. A second threshold excludes observations with $|z| > 2$, identifying moderately large forecast errors and serving as a

conservative robustness check. Table H.1 reports the number of excluded observations under each rule.

Results. Figure H.1 displays the impulse responses across three specifications: (1) full sample (red), (2) excluding $|z| > 3$ observations (blue), and (3) excluding $|z| > 2$ observations (green).

Figure H.1: Impulse Responses to AI Uncertainty Shock (Excl. $|z| > 3$ and $|z| > 2$)



Note: Figure 5.1 displays the impulse responses to a one standard deviation shock in AI uncertainty estimated with an SVAR under recursive identification (blue line) and SVAR-IV (red line). The instrument used is a binary IV equal to 1 when the AIU Index is greater than or equal to its 0.95 quantile and 0 otherwise. Shaded areas denote 68% confidence bands based on 1,000 wild bootstrap replications.

The impact response in the AI uncertainty equation decreases as larger residuals are removed, which reflects that excluding the upper tail realisations reduces the estimated size of the underlying innovation. However, the responses of each macroeconomic variables remains similar across specifications. The effects diminish over comparable

horizons, and the overall pattern and persistence of the responses show limited variation. This stability indicates that the propagation of AI-related uncertainty is not driven by a small number of high-leverage observations.

The analysis shows that neither the economic results nor the statistical inference is driven by a small set of extreme observations in the AIU Index. The macroeconomic responses remain consistent across specifications, and the wild bootstrap continues to produce confidence intervals that align closely with those from the full sample. These findings indicate that the SVAR-IV estimates reflect systematic macroeconomic responses to AI-related uncertainty rather than the influence of isolated high-leverage residuals.

Appendix I

Table I: Data Description – Local Projections

Variable:	**Description:**	**Transform:**
AI Uncertainty	AI Uncertainty Shock *(identified from the SVAR-IV)*	N/A
Wage	Average Hourly Earnings of All Employees	Log Level
Hours	Average Weekly Hours of All Employees	Log Level
Employment	Total Employment, All Employees	Log Level

Note: Table I reports the monthly variables used in the local projection estimation, together with their descriptions and transformations. All variables, except for AI Uncertainty, are obtained from the US Bureau of Labour Statistics (BLS). Service-providing industries include: (1) Trade, Transportation, and Utilities, (2) Information, (3) Financial Activities, (4) Professional and Business Services, (5) Education and Health Services, (6) Leisure and Hospitality, and (7) Other Services (excluding Public Administration). Goods-producing industries comprise only Manufacturing. Industry classifications follow NAICS. Natural Resources and Mining, as well as Construction, are excluded from the analysis since their dynamics are heavily influenced by external disturbances, such as oil price and monetary policy shocks, which could confound the estimated effects of AI-related uncertainty.

Appendix J

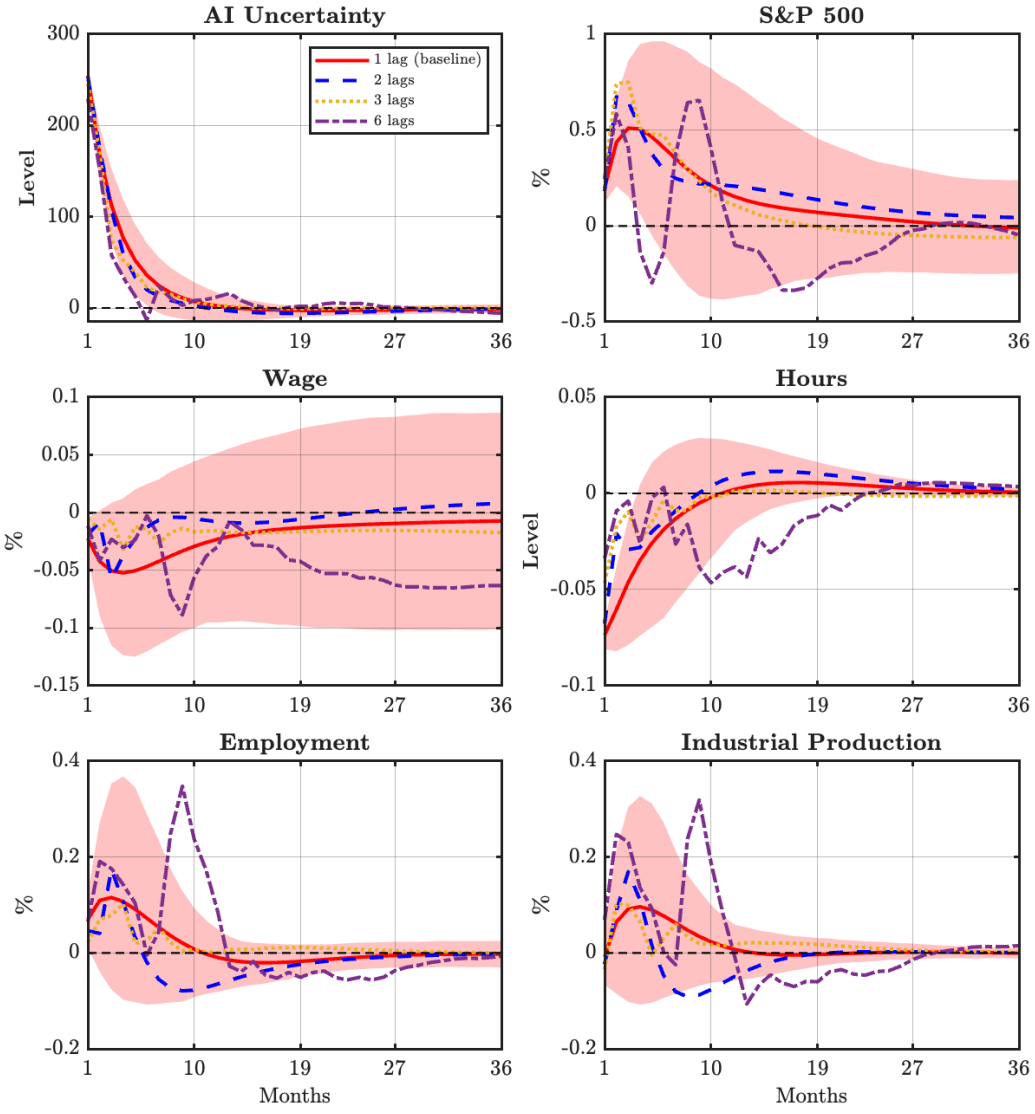
Table J: Test on the Strength of the Instrument (Tail Realisations)

	AI Unc.	S&P 500	Wage	Hours	Emp.	Ind. Prod.
0.95 Quantile						
β	753.04***	0.01	0.00	-0.00	-0.00	-0.00
Std. Errors	(94.38)	(0.02)	(0.00)	(0.00)	(0.01)	(0.01)
t-Statistics	7.98	0.34	0.17	-0.98	-0.01	-0.13
F-Statistics	63.66	0.11	0.03	0.96	0.00	0.02
R^2	0.37	0.00	0.00	0.01	0.00	0.00
0.90 Quantile						
β	421.10***	-0.00	0.00	-0.00	-0.00	-0.00
Std. Errors	(71.91)	(0.01)	(0.00)	(0.00)	(0.00)	(0.01)
t-Statistics	5.86	-0.26	0.15	-1.03	-0.25	-0.63
F-Statistics	34.29	0.07	0.02	1.07	0.06	0.40
R^2	0.24	0.00	0.00	0.01	0.00	0.00
0.75 Quantile						
β	98.88*	0.00	-0.00	-0.00	0.00	0.00
Std. Errors	(55.93)	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)
t-Statistics	1.77	0.41	-0.63	-1.00	0.71	0.50
F-Statistics	3.13	0.17	0.40	0.99	0.50	0.25
R^2	0.03	0.00	0.00	0.01	0.00	0.00

Note: Table J reports the regression results from $\hat{u}_{it} = \alpha + \beta_i w_t^{AI} + \eta_{it}$, where $\hat{u}_{it}$ denotes the reduced-form residual from the VAR equation corresponding to each variable. The purpose is to test the strength of the proposed proxy w_t^{AI}, constructed from the upper tail of the AIU Index. Robust standard errors are reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Appendix K

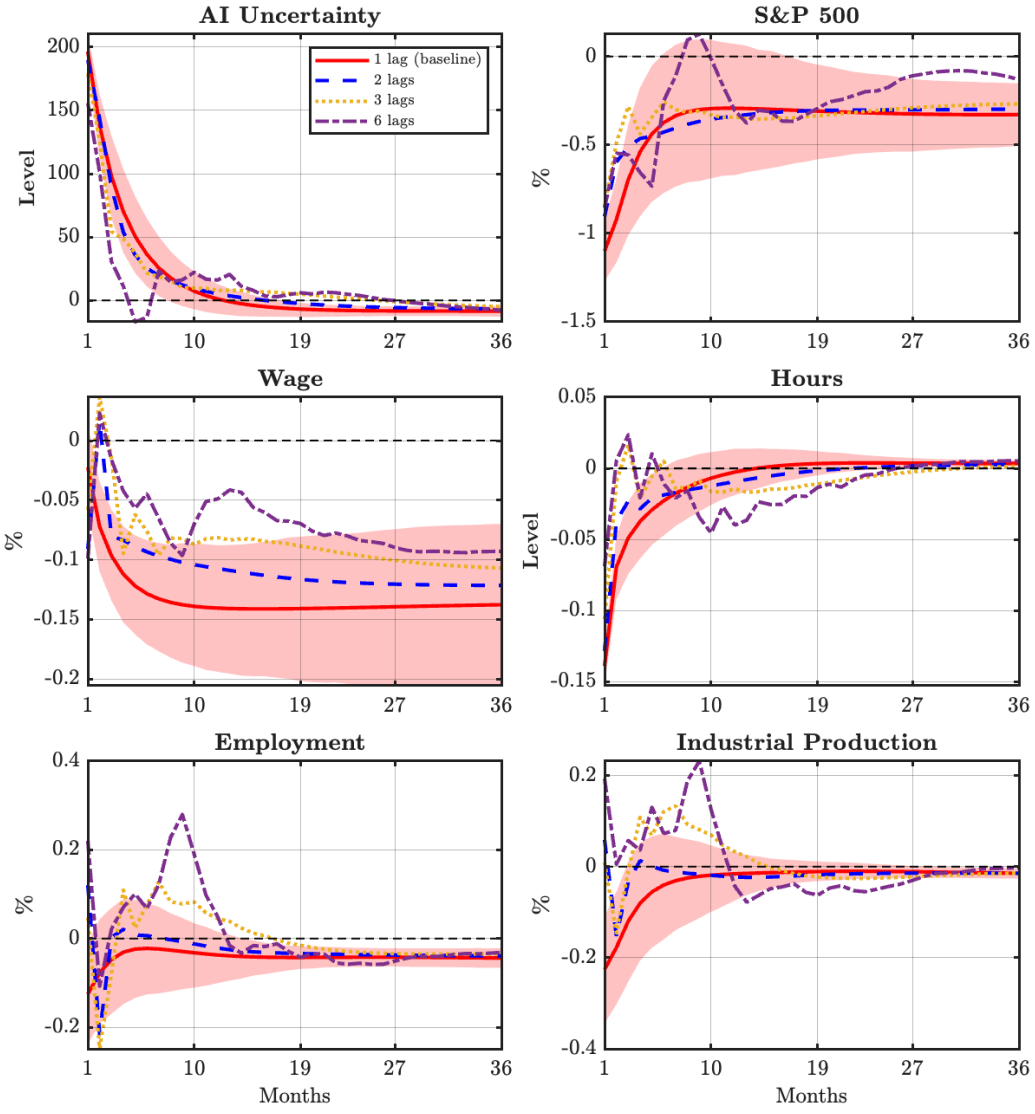
Figure K: Impulse Responses to AI Uncertainty Shock
Recursive Identification (Different Lag Specification)



Note: Figure J displays the impulse responses to a one standard deviation shock in AI uncertainty estimated with an SVAR under recursive identification with different lag specifications. Shaded areas denote 68% confidence bands based on 1,000 wild bootstrap replications.

Appendix L

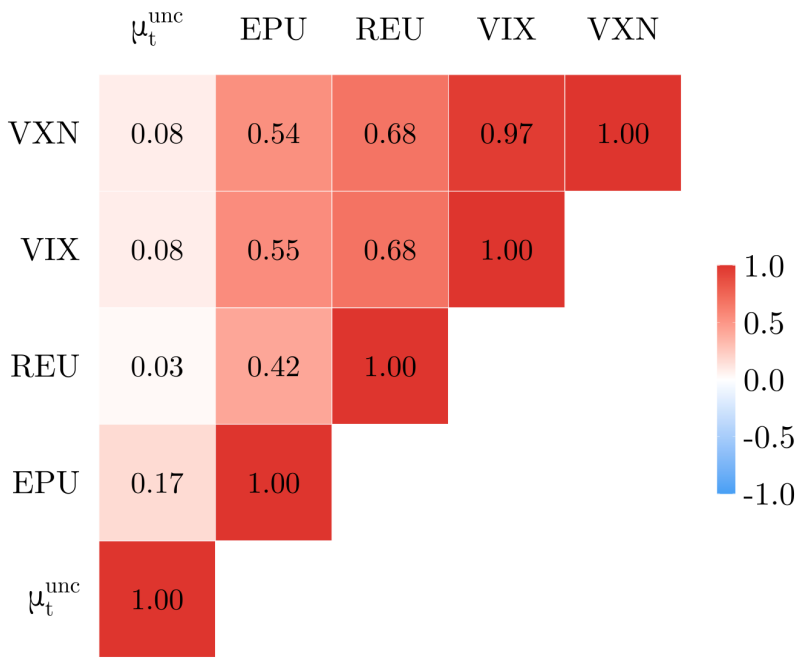
Figure L: Impulse Responses to AI Uncertainty Shock
SVAR-IV (Different Lag Specification)



Note: Figure K displays the impulse responses to a one standard deviation shock in AI uncertainty estimated using SVAR-IV with different lag specifications. The instrument is the residual-based measure of AI uncertainty from Equation (4.9). Shaded areas denote 68% confidence bands based on 1,000 wild bootstrap replications.

Appendix M

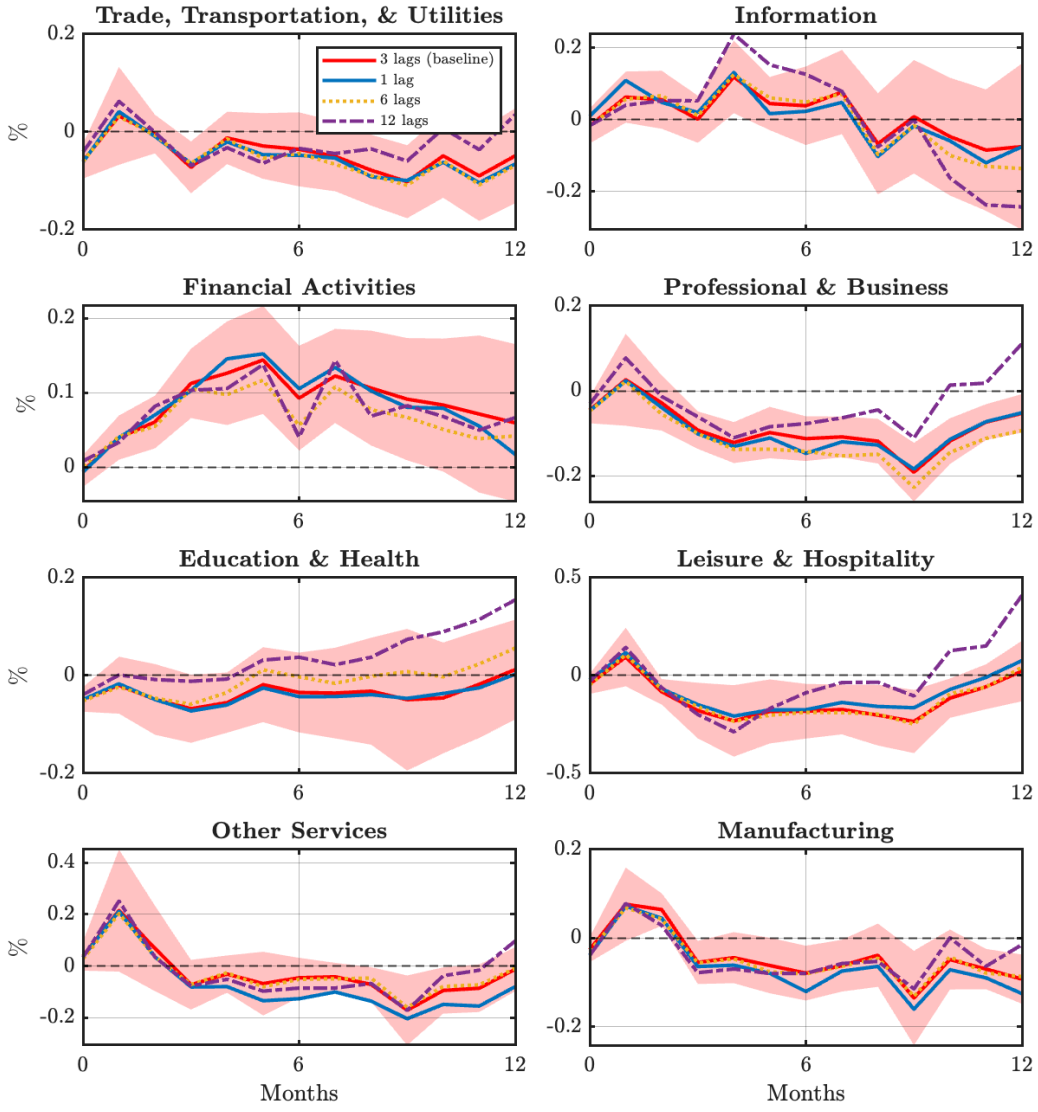
Figure M: Correlation of μ_t^{unc} with different Uncertainty Measures



Note: Figure J reports the Pearson correlation coefficients between residual-based instrument μ_t^{unc} and benchmark uncertainty measures. The benchmarks are the EPU Index, the REU, and the VIX. The sample spans from M1:2016 to M4:2025, reflecting the availability of news articles used in constructing the AIU Index and benchmark uncertainty measures.

Appendix N

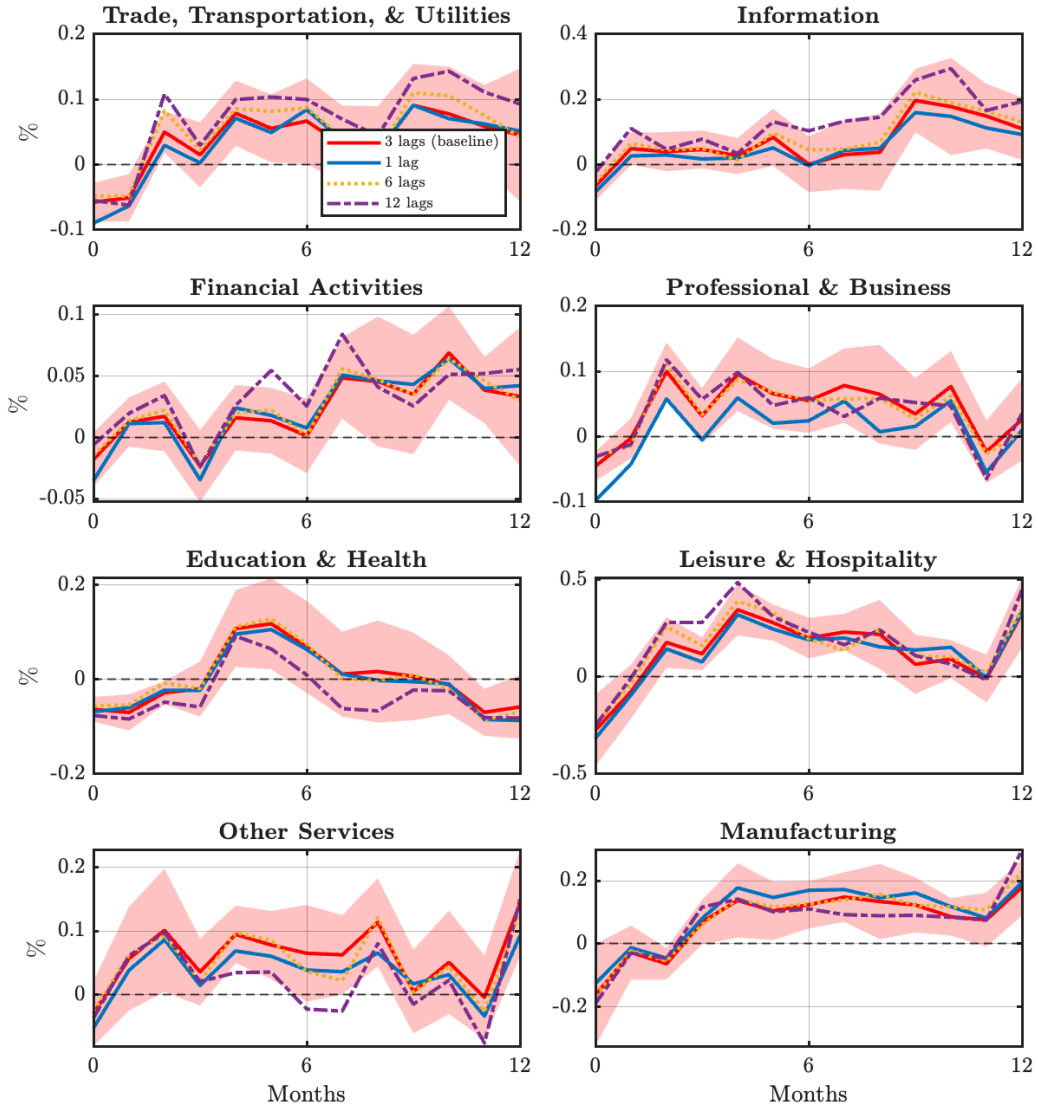
Figure N: Response of Industry-Level Wage to AI Uncertainty Shock (Different Lag Specifications)



Note: Figure M presents the industry-level wage impulse responses to a one standard deviation "pure" AI uncertainty shocks estimated via local projections with different lag specifications. The industries are classified according to NAICS. The shaded regions represent 68% confidence intervals, computed using Newey–West standard errors.

Appendix O

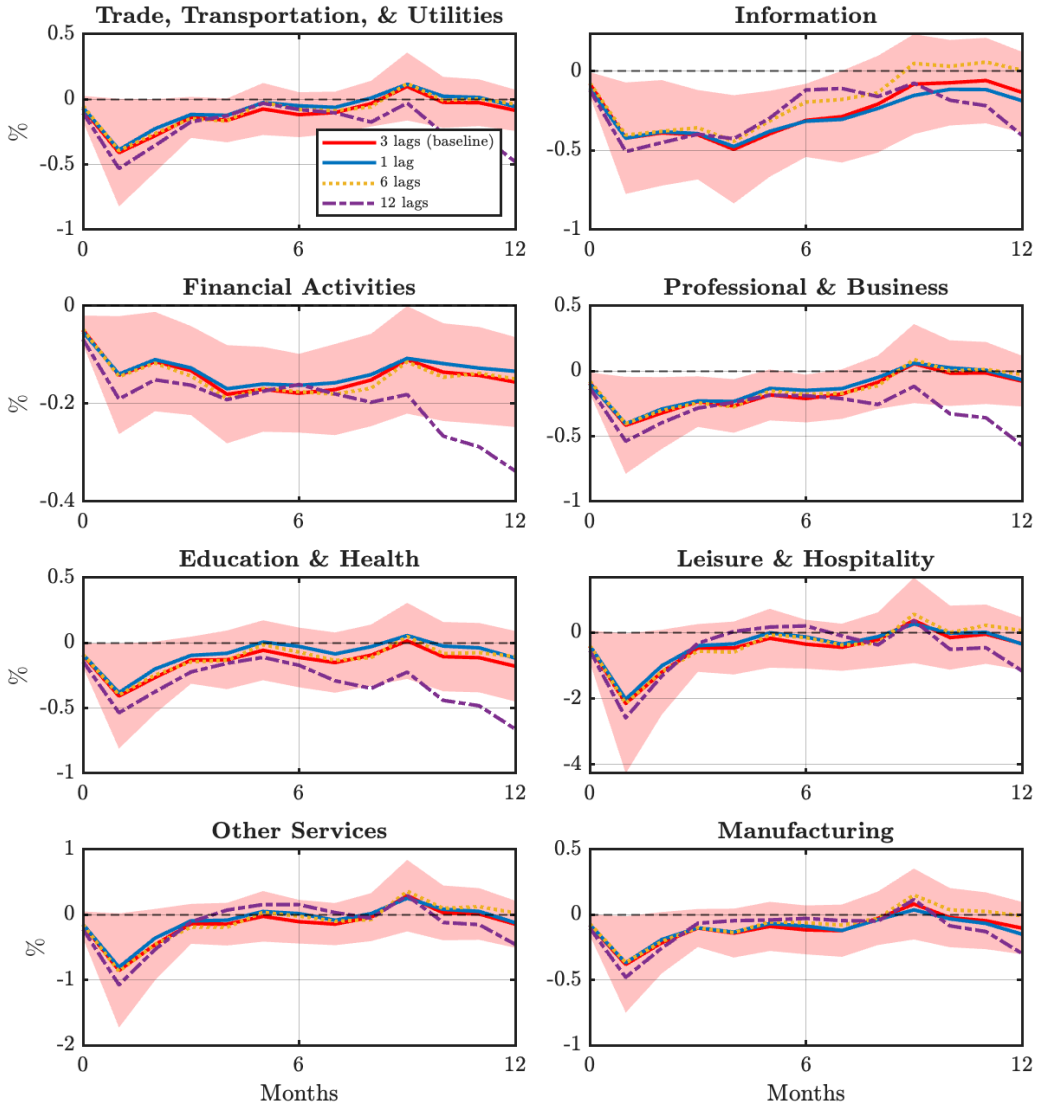
Figure O: Response of Industry-Level Hours Worked to AI Uncertainty Shock (Different Lag Specifications)



Note: Figure N presents the industry-level hours worked impulse responses to a one standard deviation "pure" AI uncertainty shocks estimated via local projections with different lag specifications. The industries are classified according to NAICS. The shaded regions represent 68% confidence intervals, computed using Newey–West standard errors.

Appendix P

Figure P: Response of Industry-Level Employment to BAI Uncertainty Shock (Different Lag Specifications)



Note: Figure O presents the industry-level employment impulse responses to a one standard deviation "pure" AI uncertainty shocks estimated via local projections with different lag specifications. The industries are classified according to NAICS. The shaded regions represent 68% confidence intervals, computed using Newey–West standard errors.