ECONOMIC STATISTICS CENTRE OF EXCELLENCE

A COLLABORATION WITH Office for National Statistics

The Macroeconomic Effects of AI Uncertainty

Juan Reyes

ESCoE Discussion Paper 2026-02

February 2026

ISSN 2515-4664

DISCUSSION PAPER

The Macroeconomic Effects of AI Uncertainty
Juan Reyes
ESCoE Discussion Paper No. 2026-02
February 2026

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 that displays sharp increases around notable AI developments. The index also demonstrates limited correlation with established measures of economic uncertainty. Using SVAR-IV with an instrument that orthogonalises first- and second-moment coverage of AI, I find that positive AI uncertainty shocks generate significant contractionary effects on equity prices, hours worked, and wages, with smaller and less persistent effects on employment and output. Industry-level estimates, in turn, highlight heterogeneous adjustments along both the labour quantity and price margins. These findings indicate that AI uncertainty is a distinct source of economic fluctuations, with a response pattern that differs from that of conventional uncertainty shocks.

Keywords: Artificial Intelligence, Uncertainty, Text Analysis, Business Cycle

JEL classification: D81, E32, J22, J24.

Juan Reyes, King’s College London and ESCoE
juan.reyes@kcl.ac.uk

Published by:
Economic Statistics Centre of Excellence
King’s College London
Strand
London
WC2R 2LS
United Kingdom
www.escoe.ac.uk

The Macroeconomic Effects of AI Uncertainty[*]

Juan Reyes[†‡]

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

This Version: February 2026
(*Please Check Here for the Latest Version*)

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 that displays sharp increases around notable AI developments. The index also demonstrates limited correlation with established measures of economic uncertainty. Using SVAR-IV with an instrument that orthogonalises first- and second-moment coverage of AI, I find that positive AI uncertainty shocks generate significant contractionary effects on equity prices, hours worked, and wages, with smaller and less persistent effects on employment and output. Industry-level estimates, in turn, highlight heterogeneous adjustments along both the labour quantity and price margins. These findings indicate that AI uncertainty is a distinct source of economic fluctuations, with a response pattern that differs from that of conventional uncertainty shocks.

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

[*] I am indebted to my PhD 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, Francesco Zanetti, and Stylianos Zlatanos, and the two anonymous reviewers of the ESCoE Discussion Papers series for their insightful comments and suggestions, which have substantially shaped the direction of the study. I would further like to thank Zining Yuan and Shunshun Zhang for providing the initial data used in the early stages of this research. This paper has benefited from comments and discussions with conference and seminar participants of the 2025 King's Business School (KBS) Doctoral Symposium, the 2025 Warwick Business School (WBS) PhD Workshop on Fintech, the 2025 ESCoE Researcher Workshop, the ESCoE Economic Measurement Webinar, the 4th UK Digital Economics Workshop, and the 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 or juanrufinomreyes@gmail.com. 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. Estimates on the macroeconomic impact of AI vary substantially. These range approximately 4.0% growth in global gross domestic product (GDP) over the next decade under a high total factor productivity (TFP) scenario (Cerutti et al., 2025), to projections of 7.0% growth in global GDP and increases of 1.5 percentage points in annual United States (US) labour productivity (Goldman Sachs, 2023).[1] Similarly, other estimates suggest that adoption of generative AI (GenAI) models could contribute up to 25.6 trillion USD in global output (Chui et al., 2023). A more conservative evaluation, however, points to substantially smaller aggregate gains. Using a task-based macroeconomic framework, Acemoglu (2025) estimates that AI adoption will raise US GDP by only 0.9 to 1.6% over the next decade.

These wide-ranging estimates reflect more than differences in modelling assumptions. They point to fundamental uncertainty about how AI will evolve, be regulated, and be integrated into the economy. AI development is characterised by accelerating, ambiguous technical change, fragmented and contested regulatory responses, opaque competitive dynamics, and substantial organisational adjustment costs that emerge only through deployment, which can vary across sectors. In this environment, economic agents make decisions without reliably mapping AI adoption to future productivity, task allocation, or labour demand. When such uncertainty is pervasive, standard economic theory implies contractionary responses. A large body of literature shows that elevated uncertainty depresses investment, hiring, and consumption through real-option and precautionary saving channels (Bloom, 2009, 2014; Jurado et al., 2015; Baker et al., 2016). Against this backdrop, it remains unclear whether uncertainty surrounding AI operates through similar mechanisms and whether it constitutes a distinct source of macroeconomic fluctuations rather than a reflection of broader economic or policy uncertainty.

Despite the growing interest in the economic effects of AI, its role as an independent

[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)

source of macroeconomic uncertainty has received limited systematic analysis. This can be understood for two main reasons. First, much of the existing literature focuses on first-moment effects of AI, particularly 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). These studies proceed under the implicit assumption of known AI capabilities and deployment paths, and therefore focus on how realised AI adoption affects economic outcomes. Second, commonly used uncertainty measures, such as the Real Economic Uncertainty (REU) Index (Jurado et al., 2015), the Economic Policy Uncertainty (EPU) Index (Baker et al., 2016), and financial volatility indices (i.e., S&P 500 Volatility Index, NASDAQ 100 Volatility Index), capture broad macroeconomic or policy-related volatility. These indicators do not isolate the distinct dimension of uncertainty that is specific to technological change, complicating efforts to assess whether AI-related uncertainty has distinct macroeconomic effects.

In this paper, I address these limitations by developing a novel measure of AI uncertainty and examining its macroeconomic effects. I construct the AI Uncertainty (AIU) Index using text-based methods applied to newspaper coverage from leading news outlets in the US, the United Kingdom (UK), and selected European countries. Following Baker et al. (2016), the index is based on systematic identification of news coverage that simultaneously references AI, economic conditions, and uncertainty. The resulting series displays pronounced movements around major AI-related developments, including the introduction of GPT-4 by OpenAI in March 2023, the issuance of US Executive Order 14110 titled *"Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence"* in October 2023, and the release of DeepSeek-R1 in January 2025. The AIU Index also shows low correlations with broader measures of uncertainty, suggesting it captures a distinct dimension of economic uncertainty.

Using the AIU Index, I then estimate the macroeconomic effects of uncertainty surrounding AI within a structural vector autoregression identified with an external instrument (SVAR-IV). The identification strategy exploits variation in AI-related news coverage that reflects uncertainty about the economic implications of AI, while remaining orthogonal to broader macroeconomic and policy developments. To disentangle uncertainty from positive news about AI-driven productivity gains, I construct an instrument that isolates the first- and second-moment components embedded in AI-related coverage. This approach builds on evidence that news and uncertainty shocks are often confounded in empirical applications (Piffer and Podstawski, 2018; Cascaldi-Garcia and Galvao, 2021).

The analysis yields three main findings. First, AI uncertainty shocks generate contractionary effects on the economy. Equity prices decline sharply and persistently, while labour market adjustment occurs primarily along the intensive margin. Hours worked fall temporarily, whereas wages contract more persistently and continue to weaken over time. Employment displays transitory responses, while industrial production experiences modest and short-lived contractions. This response, however, differs from conventional uncertainty shocks, which tend to generate broader and more uniform declines across macroeconomic aggregates.

Second, the adjustment dynamics following an AI uncertainty shock are consistent with labour behaving as a quasi-fixed factor (Oi, 1962), while departing from its standard implications. The combination of stable employment response alongside declining hours worked indicates that labour adjustment occurs primarily along the intensive margin rather than through changes in employment levels. This pattern is consistent with employment retention in the presence of hiring, training, and organisational adjustment costs. At the same time, wages display sustained downward adjustment rather than remaining rigid, which contrasts with standard quasi-fixed labour models. One interpretation is that uncertainty about AI weakens worker bargaining conditions by complicating the assessment of outside employment opportunities (Leduc and Liu, 2024). Moreover, unlike conventional uncertainty episodes that tend to subside as conditions stabilise, AI-related uncertainty evolves alongside ongoing technological change, reducing the scope for recovery driven by the resolution of uncertainty. Employment may therefore be maintained not in anticipation of cyclical improvement, but to preserve organisational capital and firm-specific knowledge that is valuable during periods of technological transition.

Third, responses to AI uncertainty display substantial heterogeneity across industries. Using local projections (Jordà, 2005), I find considerable variation in both the magnitude and direction of labour market responses. Hours worked and employment decline in most industries, although magnitudes differ substantially. Wage responses are particularly heterogeneous, with some sectors experiencing sustained declines while others display muted or even positive adjustments. This cross-industry variation in wage is systematically related to AI exposure, measured by the share of tasks exposed to automation (Felten et al., 2021). Industries with greater exposure to AI display larger adjustments following an AI uncertainty shock. This heterogeneity, together with composition effects, reconciles the differences between aggregate and industry-level estimates. Industries with larger employment shares exhibit smaller responses, while high-exposure industries with smaller shares display more pronounced adjustments that attenuate in employment-weighted aggregation.

This paper contributes to numerous strands of the literature. By examining uncertainty associated with AI, it extends research on uncertainty shocks and business cycle fluctuations beyond broad macroeconomic or policy related measures (Bloom, 2009; Jurado et al., 2015; Baker et al., 2016). The development of a technology focused uncertainty index constructed from newspaper coverage also advances work on source-specific uncertainty and text-based measurement (Jurado et al., 2015; Baker et al., 2016; Caldara et al., 2020; Husted et al., 2020; Caldara and Iacoviello, 2022; Abiad and Qureshi, 2023), while complementing the growing use of textual data in empirical macroeconomics (Hansen et al., 2018; Gentzkow et al., 2019; Ash and Hansen, 2023). In addition, the analysis contributes to the growing literature on the economics of AI (Felten et al., 2021; Acemoglu et al., 2022; Acemoglu, 2025) by shifting attention from realised effects to the uncertainty surrounding adoption. Evidence on heterogeneous responses across industries further informs research on labour market adjustment to technological change. To the best of my knowledge, this is the first study to construct a dedicated measure of AI related uncertainty and to assess its macroeconomic effects using identified structural shocks.

The remainder of the paper is structured as follows. Section 2 discusses the definition and 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

The AIU Index measures AI-related economic uncertainty as reflected in news coverage. It aims to capture perceived uncertainty surrounding the economic implications of AI, covering a broad set of considerations, including effects on tasks and occupational exposure, productivity and wage dynamics, labour market adjustment and worker reallocation, changes in the sectoral organisation of production, and the economic effects of regulatory responses to AI. The index reflects both near-term concerns, such as disruptions associated with model releases or policy actions, and longer-run questions about aggregate productivity gains and distributional outcomes. By tracking the frequency of relevant newspaper coverage, the index provides a timely measure of how perceptions of AI-related economic uncertainty evolve in response to new developments.

I construct the AIU Index by applying dictionary-based text analysis to news articles from leading national and international news outlets. The approach follows established methods for constructing news-based uncertainty indicators, including the EPU Index, the Trade Policy Uncertainty (TPU) Index (Caldara et al., 2020), and the Oil Price Uncertainty (OPU) Index (Abiad and Qureshi, 2023).

The development of the AIU Index proceeds in four stages. First, I retrieve news articles that reference AI in an economic context under conditions of uncertainty. Second, I process the articles into structured text using 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 consists of 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 provide 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 restriction ensures consistency in editorial focus and economic reporting over time. By focusing on widely circulated and internationally recognised sources, the index remains anchored to a stable and comparable set of publications.

Following Abiad and Qureshi (2023), I further exclude content types unlikely to reflect detailed economic reporting. In particular, I remove articles that are classified under the categories of sports, editorials, abstracts, advertorials or sponsored

[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.

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

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.

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 are retained only if they 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. Table 2.2 provides the comprehensive list of keywords for each category.

[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.

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 uncertainty-related terms correspond to measurable forms of risk, economic theory distinguishes them from *Knightian uncertainty*. The former 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 capture both quantifiable variation and broader forms of uncertainty, insofar as they are jointly reflected in media narratives surrounding AI and its economic relevance. Accordingly, uncertainty-related terms are counted only when they appear alongside 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.

The filtering procedure also includes Unicode normalisation, character encoding standardisation, conversion to lowercase, and removal of punctuation and non-informative symbols.[4]

[4]Unicode normalisation standardises characters that may have multiple valid digital representations

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 standardised metadata fields, such as article type, title, author, word count, publication data, newspaper name, copyright statement, and a unique document identifier assigned by Factiva.

The uniform layout of the documents facilitates efficient and reliable preprocessing. The fixed placements of structural elements, such as the copyright statement near the beginning of each article and the document identified at the end, enable accurate segmentation of individual articles within each file. This standardised structure also allows for clean separation of metadata from the main body of the text. While metadata are retained for filtering, classification, and documentation purposes, only the main article text is used in constructing the index.

As a validation step, I manually review 120 randomly selected articles spanning the full sample period. This review confirms that the automated extraction procedure correctly identifies article boundaries and accurately records metadata fields. The manual review provides assurance that the preprocessing pipeline performs as intended and that the resulting dataset used to construct the AIU Index is internally consistent and free from systematic extraction errors.

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 articles counts 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

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.

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]

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

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

[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.

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. Narrative Evidence

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 verify whether they align with events that generated heightened economic uncertainty.

Figure 2.1 presents the AIU Index for the period M1:2016 to M4:2025, together with selected episodes in which pronounced movements in the series coincide with major AI-related developments. These episodes are highlighted to guide the subsequent narrative analysis and illustrate the types of events examined when assessing the correspondence between index movements and contemporaneous AI-related news coverage.

Figure 2.1: AIU Index (3-Month Moving Average)

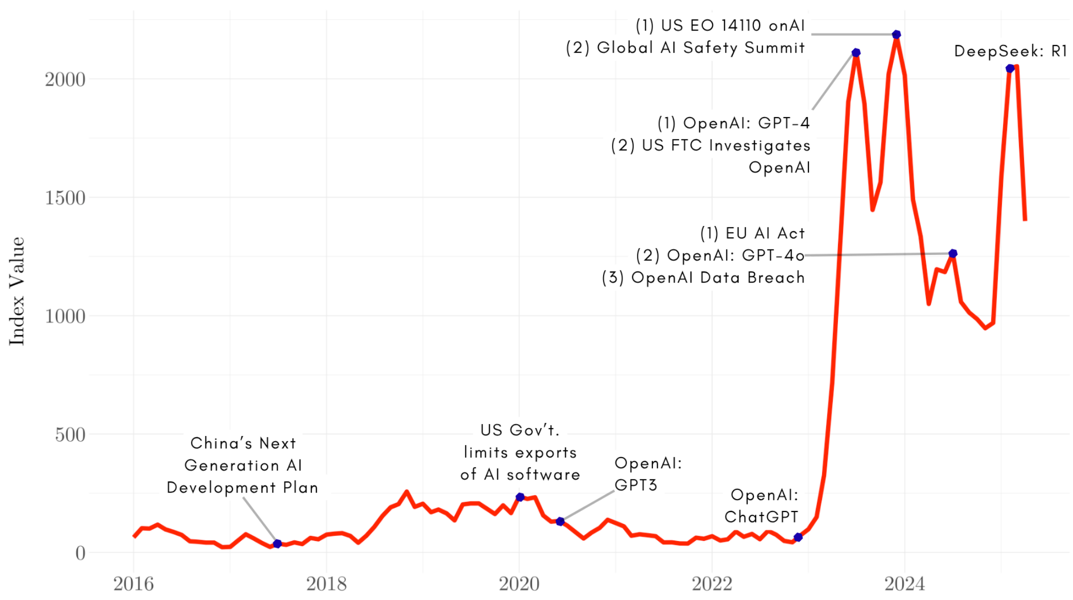

Note: Figure 2.1 plots the AIU Index from M1:2016 to M4:2025. For presentation purposes, a three-month moving average is applied. Spikes in the index were investigated by manually reviewing the underlying news articles to identify the events driving the largest movements.

Among these episodes, the largest increases in the index coincide with advances in AI technology and episodes of regulatory action that were accompanied by heightened uncertainty regarding their economic implications. A notable example occurs in March 2023, following the release of GPT-4 by OpenAI. Contemporary news articles discussed the potential productivity gains from generative AI and its contribution to economic growth, while also highlighting concerns about labour displacement, skill obsolescence, and the adequacy of existing regulatory frameworks. These discussions built on the earlier release of ChatGPT in late 2022, which had demonstrated human-level performance across tasks such as programming, legal reasoning, writing, and mathematical problem-solving, thereby intensifying debate about the economic consequences of rapid AI adoption.

Another significant increase in the index follows the issuance of Executive Order 14110 in October 2023, titled *"Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence"*. Articles published at the time discussed the order as a turning point in AI governance and raised uncertainty regarding its regulatory scope, enforcement, and potential implications for innovation and economic activity. Commentary pointed to risks of regulatory fragmentation, compliance costs, and uncertainty about how policy interventions might shape the pace and direction of AI adoption. This episode illustrates how regulatory action itself can become a focal source of macroeconomic uncertainty.

A further increase in the AIU Index appears in January 2025, coinciding with the release of DeepSeek-R1. News articles discussed the model as a lower-cost alternative to leading foundation models such as ChatGPT and examined how greater affordability could broaden access and accelerate diffusion. At the same time, coverage raised concerns about workforce displacement, competitive pressures on incumbent technology firms, and the challenges of regulating rapid adoption at scale. The release was frequently described as a turning point in global AI competition and prompted discussion about the implications for US industrial policy and export controls. Financial market reactions were also noted, with major technology firms, including Nvidia, experiencing sharp declines in equity prices. Commentators further emphasised that DeepSeek relied on domestically produced processors, reinforcing concerns about technological rivalry and supply-chain resilience.

These episodes provide narrative evidence that movements in the AIU Index coincide with periods in which AI-related developments are discussed alongside heightened economic uncertainty. The timing and content of the underlying articles indicate that fluctuations in the index reflect shifts in public discourse in which AI is framed not only as a source of potential productivity gains, but also as a driver of uncertainty related to labour markets, regulatory 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

As discussed in Section 2, the AIU Index shares methodological features akin to other perception-based uncertainty measures that rely on text analysis of news coverage, most notably the EPU Index. Both indices use news articles to capture how uncertainty is reflected in public discourse, albeit with a different thematic focus. Given this similarity in construction, it is not *a priori* clear whether the AIU Index captures information distinct from that contained in the EPU Index, or whether it largely reflects the same news-based sources of economic uncertainty.

Figure 3.1: AIU Index and Selected Uncertainty Measures

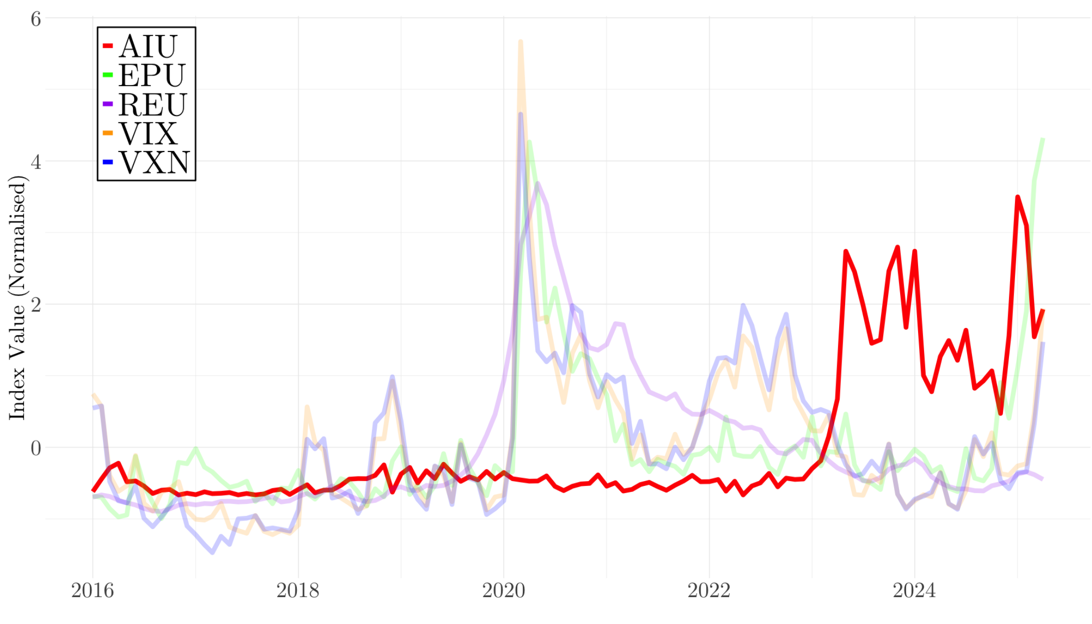

Note: Figure 3.1 plots the AIU Index together with the REU Index (Jurado et al., 2015), the EPU Index (Baker et al., 2016), the VIX, and the VXN. All series are normalised to facilitate comparison.

AI-related uncertainty may also be reflected in measures that do not rely on news-based methods, through distinct economic channels. The REU Index reflects the volatility in macroeconomic forecasts, providing a broad measure of aggregate economic uncertainty (Jurado et al., 2015). Developments in AI that raise uncertainty about productivity growth, labour market adjustment, or structural change could affect expectations about future economic outcomes, thereby increasing forecast dispersion and generating co-movement between the AIU Index and the REU Index. In addition, developments about AI may influence financial market expectations by affecting anticipated firm profitability, sectoral valuations, or growth prospects, particularly in technology-intensive industries. In such cases, AI uncertainty could be reflected in higher implied equity market volatility, leading to co-movement with market-based measures such as the VIX and VXN.

Figure 3.1 plots the AIU Index alongside measures of economic uncertainty. The figure shows that movements in the AIU Index differ systematically from those of benchmark measures. This is particularly evident during the COVID-19 pandemic in 2020, when the REU Index, EPU Index, VIX, and VXN exhibit a sharp, synchronised increase, whereas the AIU Index shows only a limited response. By contrast, from late 2022 onward, following the release of ChatGPT, the AIU Index records a sequence of pronounced increases that are not mirrored by the other uncertainty measures, which remain comparatively subdued or decline. Differences in both timing and magnitude across these indicators suggest that perceived uncertainty related to AI is not captured by traditional measures of macroeconomic, policy, or financial uncertainty.

To formally assess this distinction, I examine the extent of co-movement between the AIU Index and these established measures of uncertainty using correlation analysis and linear regressions.

3.1. Statistical Tests

Pearson Correlation. I begin by computing the pairwise Pearson correlation coefficients between the AIU Index and each benchmark measure:

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

where ρ is the Pearson correlation coefficient, $\mathrm{Cov}(\Gamma_i, AIU)$ 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 3.2 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 comovement with policy-related uncertainty, although the relationship is weaker than that observed among the benchmark indices. 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, confirming that the AIU Index provides distinct informational content.

Figure 3.2: Correlation Matrix of Uncertainty Measures

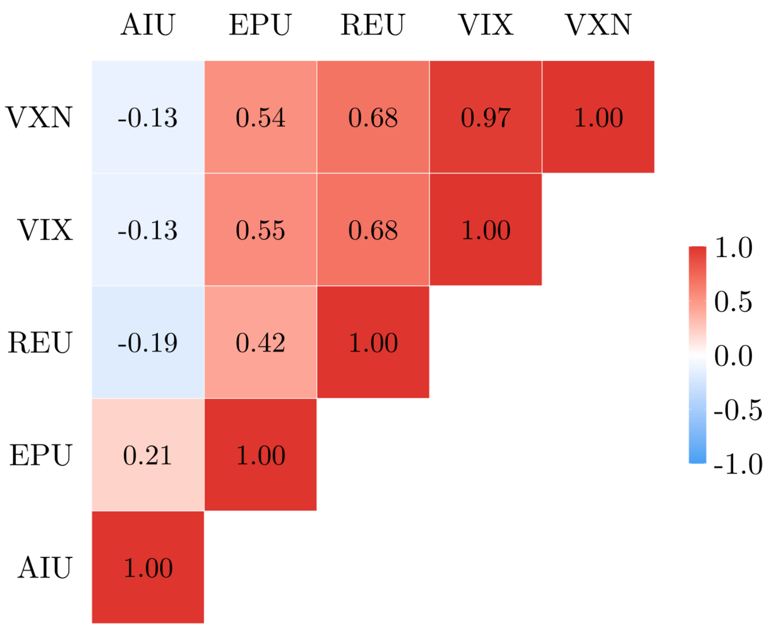

Note: Figure 3.2 reports the Pearson correlation coefficients between AIU Index and benchmark uncertainty measures. The benchmarks are the EPU Index, REU Index, VIX, and VXN. The sample spans from M1:2016 to M4:2025, reflecting the availability of news articles used to construct the AIU Index and the benchmark uncertainty measures.

Linear Regression. To complement the correlation analysis, I estimate a set of 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.

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.

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.

3.2. Interpretation

The AIU Index captures a dimension of uncertainty that is statistically independent from established measures. This distinction stems from the structural nature of AI-related uncertainty, its transmission of mixed economic signals, and its distinctive policy dimension.

Structural Nature. One explanation lies in the structural nature of AI-related uncertainty. Unlike conventional measures that reflect volatility in outcomes conditional on established economic relationships, AI-related uncertainty concerns questions such as which occupations will remain viable, how production technologies will evolve, and what forms of complementarity between human labour and artificial intelligence will emerge. This constitutes structural uncertainty in the sense that it reflects disagreement about the underlying economic relationships themselves. Such structural uncertainty need not correlate with indices capturing cyclical or policy-driven volatility, which measure movements within a known distribution of outcomes. The weak correlation between the AIU Index and benchmark uncertainty indices is consistent with this conceptual distinction. As shown in Figure 3.1, increases in AI-related uncertainty coincide with technological developments such as the release of ChatGPT in late 2022 onward, whereas the REU Index, EPU Index, VIX, and VXN display sharp spikes during episodes of broad economic and financial stress such as the COVID-19 pandemic.

Mixed Economic Signals. A second explanation is that AI-related uncertainty comprises signals of opposing sign. T Technological breakthroughs simultaneously raise expectations of productivity improvements and efficiency gains while intensifying concerns about potential adverse effects such as labour displacement, adjustment costs, and distributional consequences. This informational structure differs from macroeconomic, policy, or financial uncertainty shocks, which often transmit predominantly contractionary signals (Bloom, 2009, 2014; Jurado et al., 2015; Baker et al., 2016).

The ambiguous nature of AI-related signals manifests in the substantial differences of macroeconomic impact estimates, which range from modest aggregate gains (Acemoglu, 2025) to substantial increases in output and productivity (Chui et al., 2023; Goldman Sachs, 2023; Cerutti et al., 2025). The coexistence of optimistic and adverse assessments accounts for the weak correlation between AI uncertainty and conventional measures that predominantly capture downside risks. During episodes of financial stress or heightened policy uncertainty, news coverage tends to emphasise negative economic developments and contractionary pressures. Reporting on AI, by contrast, frequently combines optimistic projections of technological progress with concern about distributional consequences and labour market disruption. This dual nature of AI-related uncertainty is further examined in Section 4.

Disctinct Policy Dimension. The modest positive correlation between the AIU Index and EPU Index reflects a partial overlap in their informational content. This overlap arises when AI-related developments enter policy discussions, generating media coverage of legislative initiatives, executive orders, and regulatory frameworks that register in both indices (as discussed in Section 2). Hence, policy uncertainty may also encompass AI-related concerns when technological disruption is a prominent feature of policy debates.

The limited magnitude of this correlation, substantially below the correlations observed among benchmark uncertainty measures, indicates that policy considerations constitute only one dimension of AI-related uncertainty. The AIU Index predominantly reflects uncertainty regarding technological developments, sectoral transformation, and labour market effects, which operate largely independently of conventional policy channels. This independence suggests that AI-related uncertainty encompasses economic and technological dimensions that extend beyond the scope of traditional policy uncertainty.

4. Data, Model Specification, and Inference

This section describes the data, model specification, and identification strategies used to assess the macroeconomic effects of AI uncertainty. I estimate a baseline recursive SVAR, extend it using an SVAR-IV to relax exogeneity, and apply local projections to analyse industry-level responses.

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.

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).

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, 2018; 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.

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

Specifically, $\mathbf{Y}_t$ is specified as:

$$\begin{pmatrix} \textbf{\textit{AIU Index}} \\ \log(\textit{S\&P 500}) \\ \log(\textit{Wage}) \\ \log(\textit{Hours}) \\ \log(\textit{Employment}) \\ \log(\textit{Industrial Production}) \end{pmatrix}, \tag{4.2}$$

where the AIU Index is followed by indicators for the equity market, labour market, and real economic activity. This ordering is maintained across both identification strategies discussed below.

Equation (4.1) defines 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, 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. The AI uncertainty shock is identified by placing the AIU Index first in the ordering given in equation (4.2). This reflects the assumption that innovations in AI-related uncertainty can contemporaneously influence all macroeconomic variables while remaining insulated from the same period movements in fundamentals, consistent with the informational nature of the index.

SVAR with Instrumental Variable. To relax the timing restriction assumption, I also employ an instrumental variable approach to identify AI uncertainty (SVAR-IV) (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:

$$E[z_t, \varepsilon_t^{AI}] \neq 0, \tag{4.3}$$

$$E[z_t, \varepsilon_t^{j}] = 0 \quad \text{for} \quad j \neq AI, \tag{4.4}$$

where ε_t^{AI} is the AI uncertainty shock and ε_t^{j} are the remaining structural shocks. Condition (4.3) ensures relevance, while condition (4.4) 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:

$$\mathbf{s}_1 \propto \text{Cov}(z_t, \mathbf{u}_t), \tag{4.5}$$

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:

$$\hat{\varepsilon}_t^{AI} = \frac{\hat{\mathbf{s}}_1' \hat{\mathbf{\Sigma}}_u^{-1} \mathbf{u}_t}{\sqrt{\hat{\mathbf{s}}_1' \hat{\mathbf{\Sigma}}_u^{-1} \hat{\mathbf{s}}_1}}, \tag{4.6}$$

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.[7]

Tail Realisation Instrument. To implement the SVAR-IV, one of the instrument construction methods follows the approach of Carriero et al. (2015) where the instrument is derived from extreme observations of the VXO. In this paper, 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 certain quantile of its historical distribution and 0 otherwise:

$$z_t^{(\tau)} = \begin{cases} 1 & \text{if } AIU_t \geq q_\tau(AIU), \\ 0 & \text{otherwise,} \end{cases} \tag{4.7}$$

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 F.

Residual-Based Instrument. The second approach distinguishes between two dimensions of AI-related news coverage. One dimension captures narratives in which

[7] As a robustness check, I also implement an internal instrument approach (Plagborg-Møller and Wolf, 2021; Känzig, 2023), which directly incorporates the instrument into the VAR and remains valid even when the structural shock is non-invertible. This alternative specification requires that the instrument be orthogonal to leads and lags of all structural shocks, but does not impose invertibility. Details on the internal instrument methodology and results are provided in Appendix E.

AI is discussed as a source of economic uncertainty. The other reflects broader discussions of the economic role of AI, including its implications for productivity, growth, and technological progress. Distinguishing between these components is important because they correspond to different channels through which AI could potentially affect the economy. Prior studies show that uncertainty often co-moves with unanticipated revisions in expectations about future economic conditions (Piffer and Podstawski, 2018; Cascaldi-Garcia and Galvao, 2021).

The approach draws on two text-based indices constructed from the same corpus of newspaper articles. The first is the AIU Index, which measures economic uncertainty related to AI using references to AI, economic terms, and explicit mentions of uncertainty. The second is the AI Economic (AIE) News Index, which follows the same construction methodology but does not require uncertainty-related keywords and therefore captures the overall discussion of AI in an economic context.

By construction, all articles contributing to the AIU Index form a subset of those contributing to the AIE News Index. This nested structure allows the uncertainty index to be decomposed into a component that scales with the overall level of reporting on AI and the economy and a residual component that captures deviations in uncertainty-related content, conditional on the volume of such reporting. Both indices are expressed as normalised frequencies to ensure comparability over time.

The decomposition is implemented through the following regression:

$$AIU_t = \alpha_1 + \beta_1 AIE_t + \mu_t^{unc}, \tag{4.8}$$

where μ_t^{unc} represents the variation in AI-related uncertainty coverage that is orthogonal to total AI-economy coverage. Appendix H reports the corresponding regression results. The estimated coefficient $\hat{\beta}_1 = 0.93$ indicates a close proportional relationship between the two indices. On average, a one-unit increase in general coverage about AI and the economy is associated with a 0.93 unit increase in uncertainty-related coverage. An R^2 of 0.99 indicates that this relationship is highly stable over the sample period, leaving only 1.0% of the variation in AIU Index unexplained.

This residual variation corresponds to periods in which uncertainty-related content is unusually elevated or subdued relative to the overall level of reporting on AI and economic conditions. As such, it reflects changes in the framing of AI-related discourse rather than shifts in the volume of coverage alone. Examining the time-series behaviour of this residual helps clarify the nature of the variation isolated by the instrument.

Figure 4.1: Pure AI Uncertainty Component (μ_t^{unc})

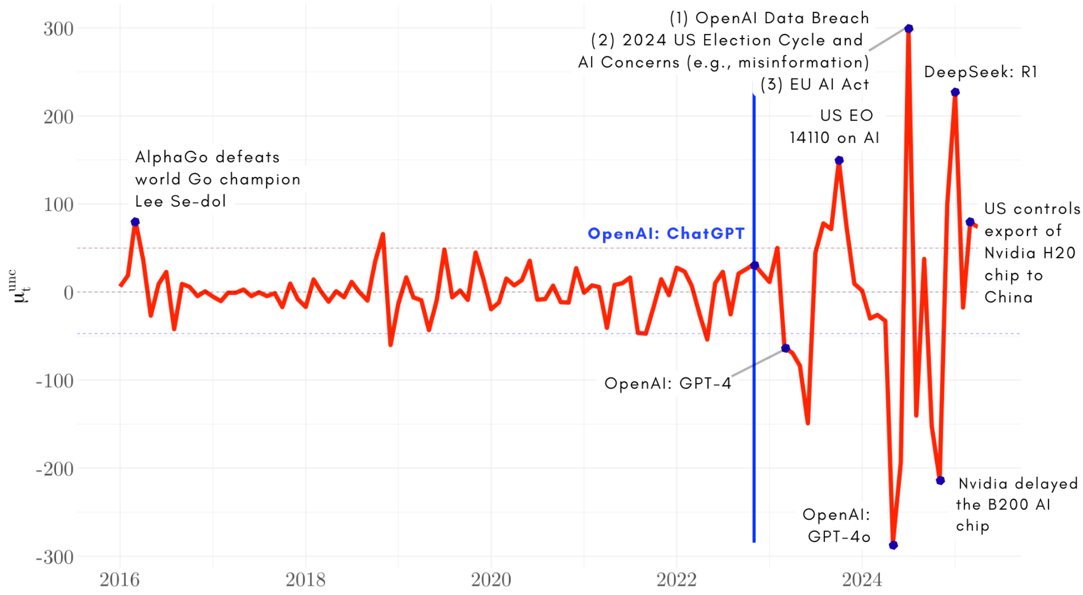

Note: Figure 4.1 displays the residual-based measure of AI uncertainty μ_t^{unc}, constructed by removing the component of the AIU Index explained by the AIE News Index (Equation 4.8). The annotated points correspond to selected AI-related developments that received extensive news coverage. The sample spans M1:2016 to M4:2025, consistent with the availability of the underlying news corpus.

Figure 4.1 presents the resulting residual series. The residual remains relatively stable through late 2022, indicating that uncertainty-related and general economic reporting on AI moved largely in proportion during this period. From late 2022 onward, beginning with the release of ChatGPT in November 2022, the series displays substantially greater volatility. Positive residuals arise when uncertainty-related themes become more prominent within coverage of AI and the economy, as observed during episodes involving regulatory debate, geopolitical competition, or labour market concerns. Negative residuals arise when coverage focuses primarily on concrete technological capabilities, demonstrated performance, or commercial applications, reducing the relative emphasis on uncertainty.

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.[8]

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.[9]

4.3. Local Projections

Aggregate estimates summarise economy-wide responses but may obscure meaningful variation across sectors. Industries differ along numerous dimensions that are relevant to the transmission of uncertainty related to AI, including exposure to automation enabled by AI, the composition of tasks performed, the structure of employment relationships, and the costs of labour adjustment. Examining whether responses to AI uncertainty vary systematically with these observable characteristics can shed light on the channels through which uncertainty affects labour market outcomes and clarify the sources of aggregate responses. To investigate this potential heterogeneity, I extend the analysis to the industry level using local projections, following Jordà (2005).

Data and Industry Classification. I use monthly industry-level data on average

[8]Following Lenza and Primiceri (2022), I adjust for the temporary increase in the variance of macroeconomic shocks during the COVID-19 pandemic. I estimate scaling parameters $\bar{s}_0$, $\bar{s}_1$, and $\bar{s}_2$ for March, April, and May 2020, respectively, along with a decay rate ρ, via maximum likelihood. The volatility scaling factor evolves as $s_t = \bar{s}_j$ for $j \in \{0, 1, 2\}$ (the first three pandemic months), and $s_{t^*+j} = 1 + (\bar{s}_2 - 1)\rho^{j-2}$ for $j \geq 3$, where $t^* =$ March 2020, allowing variance to decay exponentially toward pre-pandemic levels. All endogenous variables except the AIU Index are rescaled by s_t prior to estimation, reflecting the observation that AIU exhibited relatively stable variance during this period compared to other macroeconomic indicators. Further details are provided in Lenza and Primiceri (2022).

[9]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 I.

hourly earnings, hours worked, and employment from the US Bureau of Labour Statistics (BLS). Industries included in the analysis are classified according to the North American Industry Classification System (NAICS) at the two-digit level, yielding eight broad industries: (1) Manufacturing, (2) Trade, Transportation, and Utilities, (3) Information, (4) Financial Activities, (5) Professional and Business Services, (6) Education and Health Services, (7) Leisure and Hospitality, and (8) Other Services. Details on the variables used in the local projections are provided in Appendix J.

Estimation. Following (Känzig, 2023), 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 is treated as a common source of uncertainty affecting all industries, while allowing the magnitude and persistence of responses to differ across sectors. This approach avoids the small-sample limitations that would arise from estimating separate structural vector autoregressions for each industry.

For each industry i and outcome variable $y_{i,t}$, I estimate the following local projection at horizon h:

$$y_{i,t+h} = \alpha_{i,h} + \beta_{i,h}\,\hat{\varepsilon}_t^{AI} + \gamma_{1,i,h}\,y_{i,t-1} + \cdots + \gamma_{p,i,h}\,y_{i,t-p} + e_{i,t+h}, \tag{4.9}$$

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.

5. Results

This section reports the empirical results. I first assess the relevance and exogeneity of the instruments, 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 (2018), 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 exceeds the conventional relevance threshold of 10.0 (Stock and Yogo, 2002). Although such a strong first-stage relationship may appear mechanical, two considerations suggest otherwise. First, the tail realisation must supply identifying variation beyond what is already accounted for by the lag structure of the AIU Index and the behaviour of other variables in the system. Second, relevance is highly sensitive to the choice of quantile threshold. Instruments constructed from the 0.90 and 0.75 quantiles yield F-statistics of 34.3 and 3.1, respectively (Appendix K). This indicates that only sufficiently extreme tail events deliver strong identifying variation.

A complementary assessment of instrument validity comes from the exclusion restriction. Across all quantiles, the tail realisation instruments display negligible and statistically insignificant correlations with residuals from the non-targeted equations. Estimated coefficients are small in magnitude, statistically indistinguishable from zero, and associated with low R^2 values and weak F-statistics. This pattern suggests that the instrument does not load systematically on shocks to other variables in the system. If extreme AIU realisations were systematically associated with financial market or broad labour market shocks, such orthogonality would not be observed. The absence of these correlations supports the interpretation that upper-tail realisations isolate variation specific to AI-related uncertainty.

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 ϖ_t^{AI}, constructed from the 0.95 quantile of the AIU Index. The specification is $u_{it} = \alpha + \beta_i \varpi_t^{AI} + \eta_{it}$. 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 provides additional support by testing the instrument against benchmark uncertainty measures. Coefficients on uncertainty shocks measured by the EPU Index, the REU Index, the VIX, and the VXN are uniformly small in magnitude and statistically insignificant, with corresponding R^2 values below 0.1. This evidence is consistent with the view that the instrument captures uncertainty associated with AI-related developments rather than broader sources of macroeconomic uncertainty. On this basis, the 0.95 quantile of the AIU Index is retained as the baseline tail realisation instrument, with the 0.90 quantile 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 threshold 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 the non-targeted equations, the estimated coefficients associated with the residual-based instrument are small in magnitude, statistically insignificant, and accompanied by low R^2 values and weak F-statistics. Accordingly, the instrument is strongly related to the residual in the targeted equation while remaining unrelated to the remaining reduced-form errors. This pattern implies that the identifying variation in the instrument is concentrated in the targeted equation, namely the AI uncertainty equation, rather than being distributed across multiple equations. Such behaviour is consistent with the identifying assumption that the instrument loads primarily on a single structural shock.

To assess exogeneity, Table 5.4 reports regressions of the residual-based instrument on benchmark structural uncertainty shocks. Uncertainty 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 0.2, and the standard errors are of similar magnitude to the point estimates, indicating limited explanatory power. These results indicate that benchmark uncertainty shocks do not account for meaningful variation in the residual-based instrument.

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. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Overall, the evidence suggests 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.

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

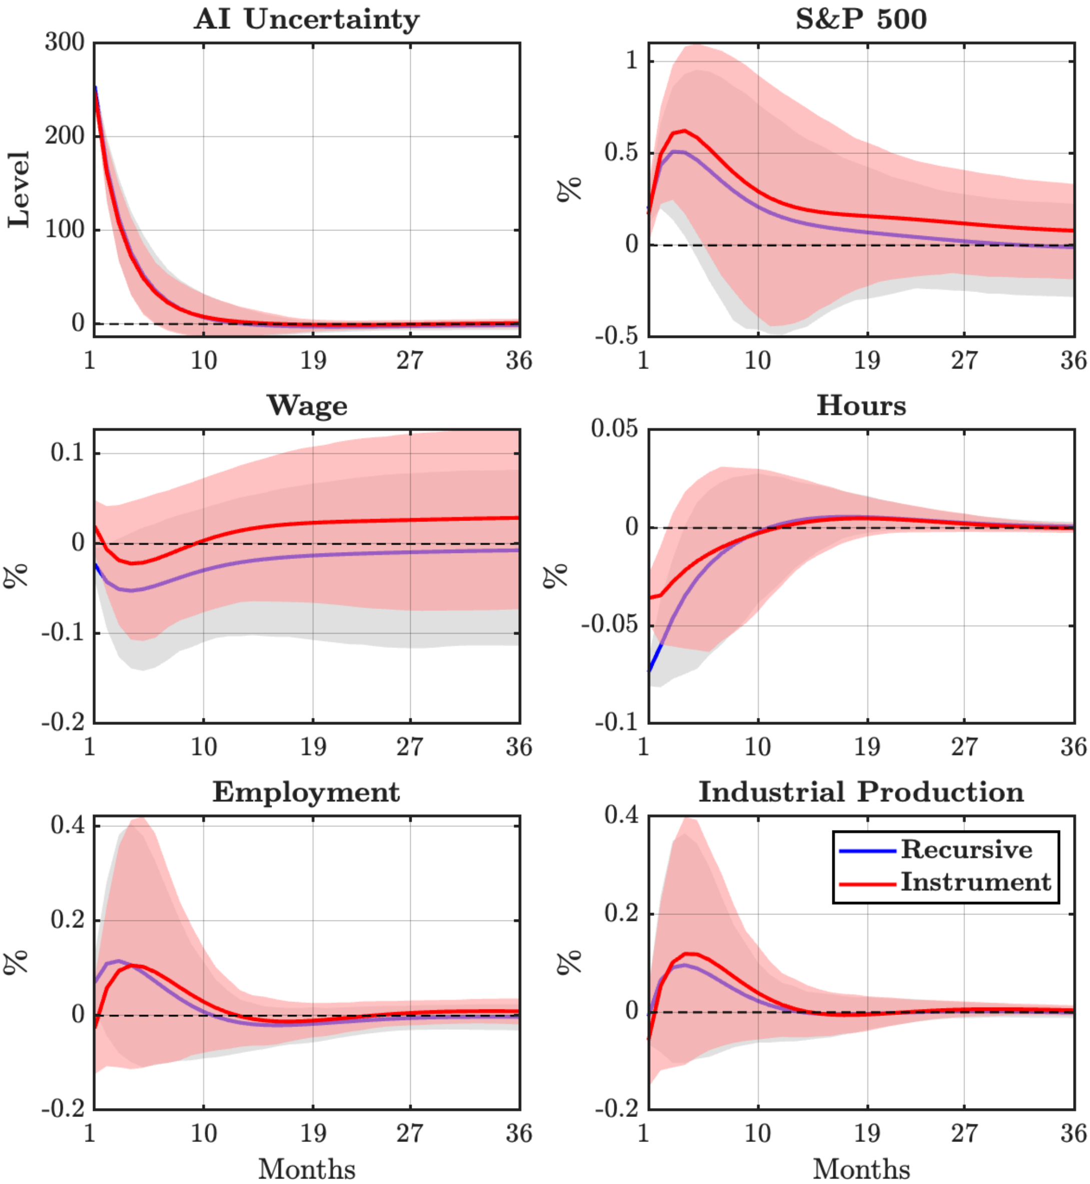

Note: Figure 5.1 displays the impulse responses to a one standard deviation shock in AI uncertainty estimated with an SVAR under recursive identification (blue line) and SVAR-IV (red line). The instrument used is a binary IV equal to 1 when the AIU Index is greater than or equal to its 0.95 quantile 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 is notable given the typically contractionary effects observed in equity markets following 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 indicate a distinctive pattern in the transmission of the composite AI uncertainty shock. The combination of a positive and persistent equity price response, a temporary contraction in hours worked, and negligible effects on wages and industrial production suggests that the AI uncertainty shock identified using the tail realisation instrument does not isolate pure uncertainty. Since the instrument conditions on the upper tail of the AIU Index distribution, it could embed both second-moment uncertainty and first-moment information about AI-related economic prospects. As a result, the identified shock reflects a mixture of contractionary uncertainty effects and expansionary productivity-related news components (Piffer and Podstawski, 2018; Cascaldi-Garcia and Galvao, 2021).

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, the resulting impulse responses are presented in Figure 5.2.

Equity markets display a pronounced and persistent contraction. The S&P 500 falls by approximately 1.1% on impact and remains below the baseline throughout the 36-month horizon. Relative to the composite AI uncertainty shock, this reversal indicates that the earlier positive response reflects the innovation-related component embedded in the AIU Index. Once this component is removed, the isolated uncertainty shock exerts sustained downward pressure on equity markets.

Figure 5.2: Impulse Responses to AI Uncertainty Shock

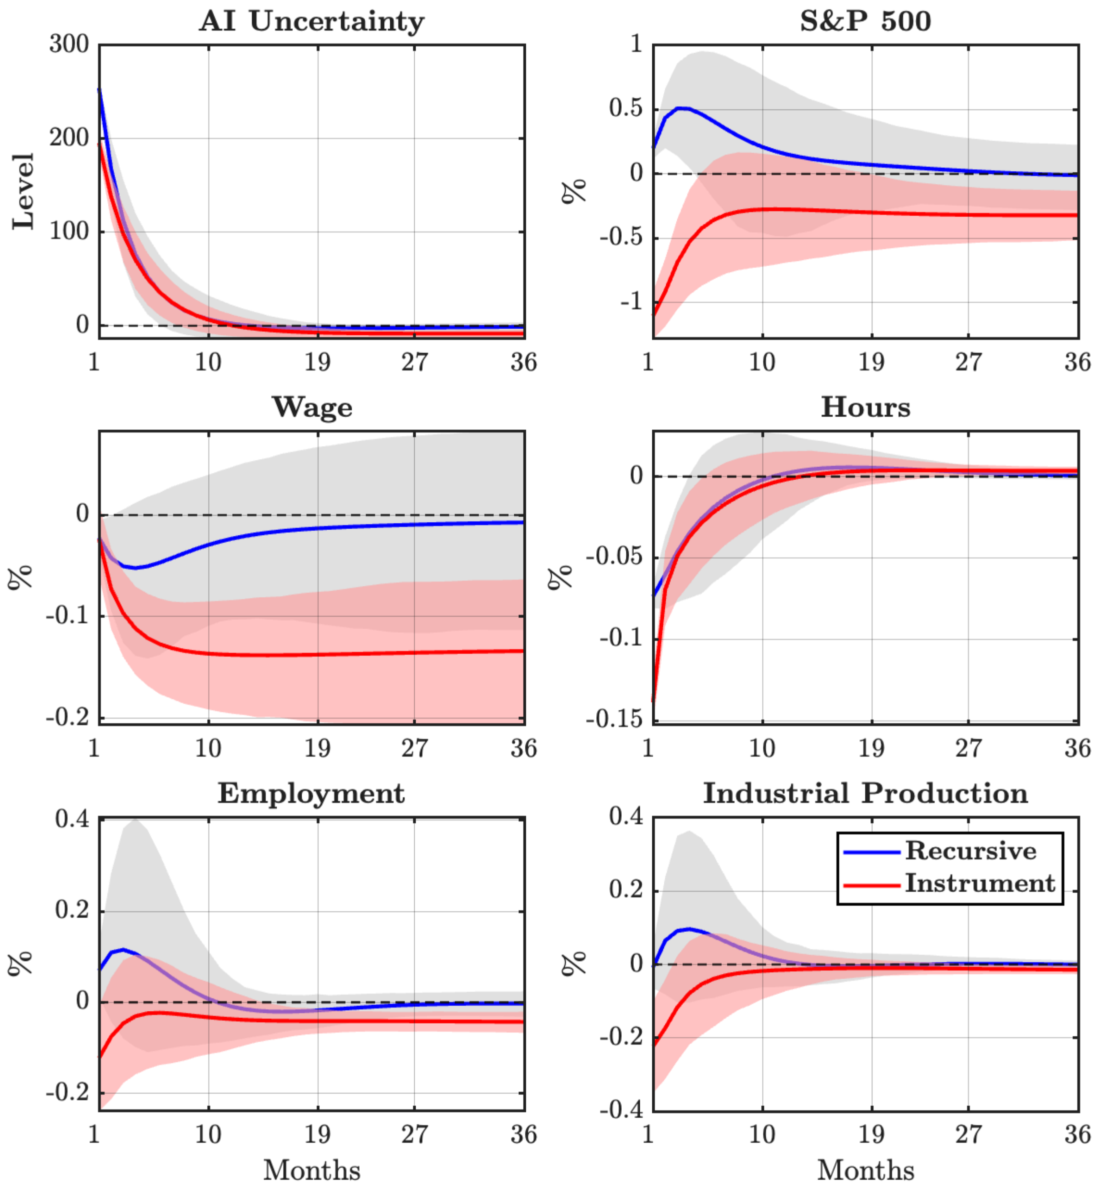

Note: Figure 5.2 displays the impulse responses to a one standard deviation pure AI uncertainty shock (red line) identified using SVAR-IV. The residual-based instrument isolates distinct components of AI-related coverage, separating discussions that frame AI as a source of uncertainty from broader narratives 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.

Labour market adjustments intensify along multiple margins. Hours worked contract by nearly 0.2% on impact and remain below baseline before reverting around the 10-month horizon. Wages display a more persistent response, with an initial decline of approximately 0.1% that deepens to nearly 0.2% and remains below baseline throughout the horizon. At the extensive margin, employment decreases by around 0.1% on impact, but this adjustment proves transitory. Wide confidence bands, however, indicate substantial uncertainty around this estimate, and the response cannot be statistically distinguished from zero. 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, while adjustment at the extensive margin remains muted.

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.

Overall, these results indicate that the movement in the composite AI uncertainty shock reflects the combined influence of two distinct components embedded in AI-related coverage. The uncertainty component induces contractionary adjustments in equity prices, wages, labour utilisation, and output. Expectations of productivity improvements create an expansionary influence that is most visible in equity markets. When both components are present in the AIU Index, their effects partially offset one another, generating the mixed responses observed under the baseline specification in Figure 5.1. The orthogonalisation of these components shows that AI uncertainty, in isolation, has a measurable and economically meaningful impact on macroeconomic activity, while productivity-related references primarily shape the aggregate response captured by the AIU Index.[10]

5.3. Forecast Error Variance Decomposition

The forecast error variance decomposition (FEVD) quantifies the importance of AI uncertainty shocks and highlights adjustment patterns not fully captured by the impulse responses. As reported in Table 5.5, the composite and pure AI uncertainty shocks differ markedly in both their overall contribution to macroeconomic fluctuations and in how these effects are distributed across variables.

In equity markets, the composite AI uncertainty shock accounts for approximately 2.0% of the forecast variance in the S&P 500 from the one-year horizon onward, with contributions that remain stable at longer horizons. The pure AI uncertainty shock, on the other hand, explains nearly 11.0% of the variance on impact, declining gradually to just under 9.0% by the 36-month horizon. This gap indicates that the composite measure substantially understates the role of an AI uncertainty shock on equity market fluctuations.

The contrast is even more pronounced in labour markets. At the 12-month horizon,

[10]The IRFs based on the internal SVAR-IV specification yield similar results. See Appendix E.

the composite AI uncertainty shock explains around 3.0% of wage variance, 6.0% of employment variance, and 7.0% of hours worked variance, with these contributions declining over time. The pure AI uncertainty shock, on the other hand, displays a sharply different variance allocation across horizons. On impact, it accounts for 24.0% of the variance in hours worked, declining to approximately 14.0% by the end of the horizon. Meanwhile, its contribution to wage variance increases steadily, reaching 28.0% at the 36-month horizon and exceeding that of all other variables at longer horizons. By contrast, employment variance remains largely unaffected by pure AI uncertainty at all horizons, which never exceeds 1.0%. This asymmetry indicates that pure AI uncertainty operates primarily through adjustments in hours worked in the short run and through wages over the medium to long run, with minimal effects on employment.

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.

For output, the distinction in variance shares is similarly informative. The composite AI uncertainty shock explains roughly 5.0% of the variance in industrial production from the 12-month horizon onward, whereas the pure AI uncertainty shock accounts for about 1.0% at all horizons. This pattern suggests that the explanatory power of the composite measure for output fluctuations largely reflects its exposure to productivity-related narratives rather than uncertainty per se. Once these narratives are orthogonalised, the contribution of pure AI uncertainty to output variance remains limited, particularly relative to its dominant role in shaping wage and hours worked dynamics.

5.4. Comparison with Other Uncertainty Shocks

Evidence from the impulse responses and the FEVD establishes that the pure AI uncertainty shock generates contractionary macroeconomic effects. The responses are front-loaded. Equity prices and hours worked respond on impact, wages adjust persistently, and industrial production exhibits a short-lived contraction, while employment remains largely unaffected. These features summarise the empirical response of the economy to AI uncertainty and provide a benchmark for comparison with other sources of uncertainty.

Empirical Comparison. To assess whether these responses differ systematically from those associated with conventional uncertainty shocks, I estimate a set of SVARs in which the REU Index, EPU Index, VIX, and VXN each replace the AIU Index as the measure of uncertainty. I employ a recursive identification scheme with the uncertainty measure ordered first, following Bloom (2009), Jurado et al. (2015), and Baker et al. (2016). The set of endogenous variables and the sample period are identical to those used in Section 4.

The results, reported in Appendix L, indicate that benchmark uncertainty shocks generate broad-based contractions across key macroeconomic indicators. Employment, industrial production, and hours worked decline markedly, equity prices fall sharply, and wage responses display mixed dynamics. These effects are largely transitory, with gradual recoveries over the horizon and most variables reverting towards baseline, consistent with the existing literature.

By contrast, the pure AI uncertainty shock, as seen in Figure 5.2, exhibits a more concentrated transmission. While equity prices respond sharply on impact, labour market adjustment occurs primarily through persistent wage compression and contraction at the intensive margin, with employment remaining close to baseline throughout the horizon. This contrast suggests that AI-related uncertainty operates through channels that differ from those associated with broader policy, financial, or macroeconomic uncertainty, despite sharing a common contractionary direction.

Economic Interpretation. The distinctive response pattern suggest different underlying process than those operating during standard uncertainty episodes. Employment stability alongside declining hours worked and persistent wage compression resemble the behaviour predicted when labour is treated as a quasi-fixed factor (Oi, 1962), but with important differences.[11] When firms face temporary downturns and have incurred substantial fixed employment costs through hiring and training, they maintain their workforce to avoid costly replacement demands when conditions improve, reducing hours instead of employment while keeping wages fixed. The response to AI uncertainty shares this similar employment-hours divergence, with hours worked falling approximately 0.2% on impact while employment holds near baseline. Two features, however, distinguish the response to an AI uncertainty shock from the quasi-fixed factor pattern.

First, wages decline persistently rather than remaining rigid. The wage response deepens over time and remains below baseline throughout the 36-month horizon (Figure 5.2), a pattern not observed with conventional uncertainty shocks. One potential explanation is that AI uncertainty weakens worker bargaining power (Leduc and Liu, 2024). Workers typically secure higher wages when they can credibly signal the availability of good employment alternatives. When AI developments create uncertainty about which skills will retain market value, workers may find it difficult to credibly evaluate their outside options. The inability to assess whether comparable positions exist or what compensation those positions would offer reduces workers' credible threat to leave. If workers face greater uncertainty than firms about future skill demands, or exhibit greater risk aversion toward skill obsolescence, this asymmetry in beliefs or preferences may lead workers to accept wage concessions in exchange for employment stability.

Second, the nature of employment retention differs fundamentally. The standard quasi-fixed factor response reflects firm expectations that productivity shocks are temporary, making it optimal to retain workers to avoid costly replacement demands, including hiring and training costs, when conditions normalise (Oi, 1962). The option value of retaining workers thus derives from the anticipated reversal of the shock. AI uncertainty provides no such anchor. The uncertainty persists rather than resolves, evolving alongside technological development without converging to a known state.

[11] A quasi-fixed factor has employment costs that are partially variable (wages) and partially fixed (hiring and training costs incurred per worker hired). Firms are reluctant to reduce employment of workers with substantial fixed employment costs during temporary downturns to avoid losing their investment and incurring costly replacement demands when conditions improve. See Oi (1962).

Firms may retain workers not in anticipation of economic recovery, but because organisational learning during periods of technological transition requires preserving firm-specific human capital. Integrating AI into existing production processes requires tacit knowledge embedded in current workers regarding operational procedures, task interdependencies, and organisational routines. Severing employment relationships eliminates this organisational capital before firms can identify which worker capabilities complement or substitute for AI technologies. Firms may therefore maintain employment levels to preserve organisational knowledge while adjusting along the intensive margin through hours and wages.

The FEVD provides supporting evidence that wages and employment adjust through different channels. At the 36-month horizon, pure AI uncertainty explains 28.4% of wage variance but only 0.8% of employment variance (Table 5.5). This pronounced asymmetry suggests that wages and employment respond differently across margins, in ways not fully captured by standard real-options frameworks, which often suggest that uncertainty affects both wages and employment. Moreover, the persistence of these effects, rather than the temporary contractions and recoveries observed with conventional uncertainty shocks, is consistent with the structural characteristics of AI-related uncertainty discussed in Section 3.

While several mechanisms may contribute to these patterns, AI-related uncertainty is associated with a distinct macroeconomic adjustment characterised by stable employment, declining hours, persistent wage compression, transitory output contractions, and strong equity price responses. These findings distinguish AI-related uncertainty from conventional uncertainty shocks and provide macroeconomic evidence to inform future theoretical work.

5.5. Industry-Level Impulse Response

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).

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. The Manufacturing industry, on the other hand, displays a distinct two-phase pattern, with an initial expansion of roughly 0.1% on impact, followed by a reversal of a 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.

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

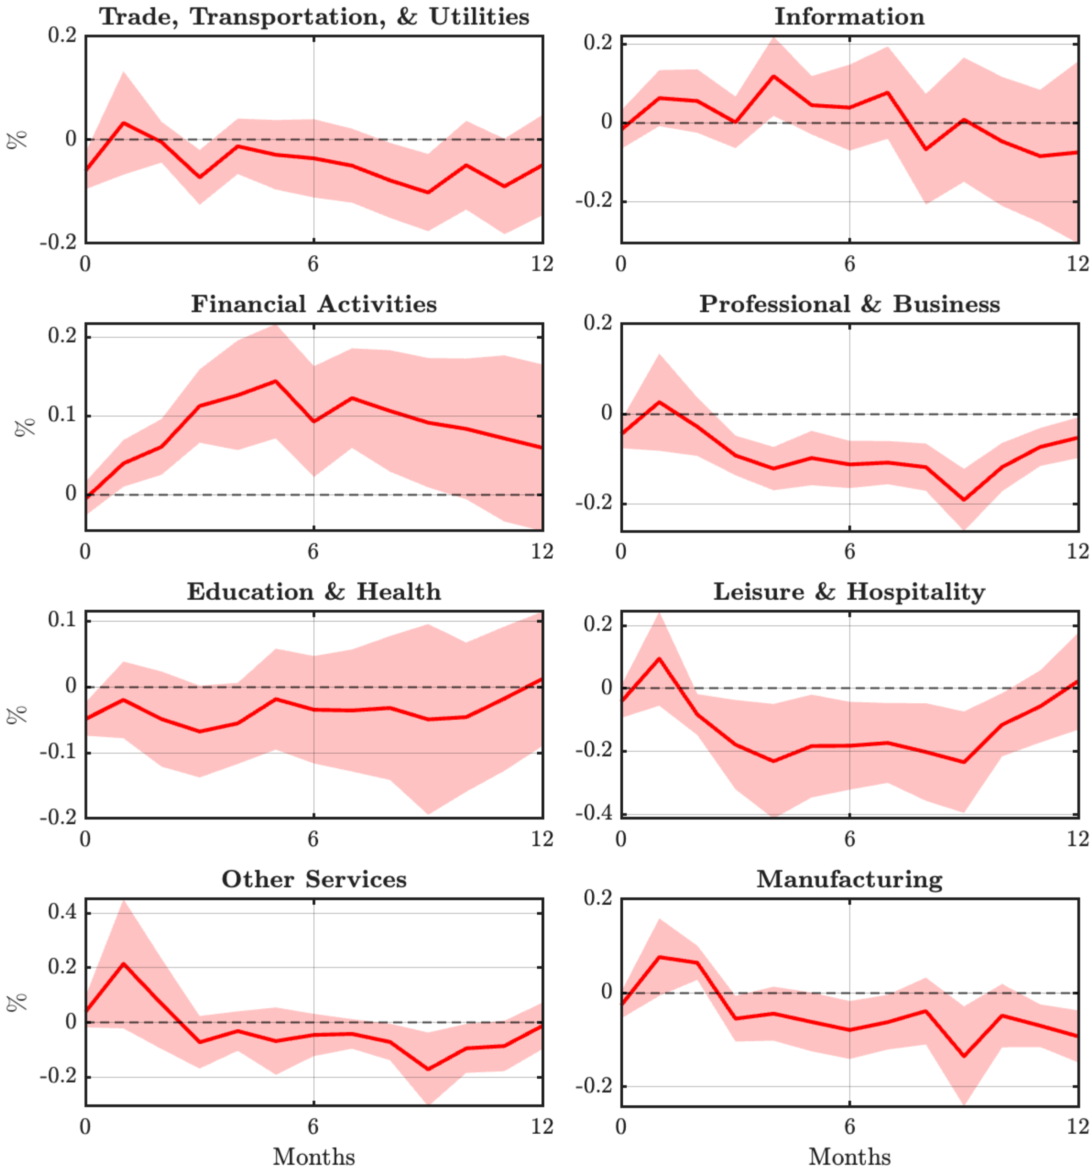

Note: Figure 5.3 presents the industry-level wage impulse responses to a one standard deviation pure AI uncertainty shocks, identified using an SVAR-IV and estimated using local projections. The industries are classified according to NAICS. The shaded regions represent 68% confidence intervals, computed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

Financial Activities, however, follow a persistent, distinct pattern. Wages in this industry increase by approximately 0.1%, peaking 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 follow 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 a modest contraction on impact, followed by a gradual increase, with hours worked expanding to approximately 0.2% and remaining persistently above baseline. The confidence interval, however, is wide and indicates 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, the information industry remains close to baseline initially and then drifts upward to around 0.2% toward the end of the horizon.

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 reflect compositional and aggregation effects that obscure the recovery in hours worked at the industry level in the aggregate series.

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

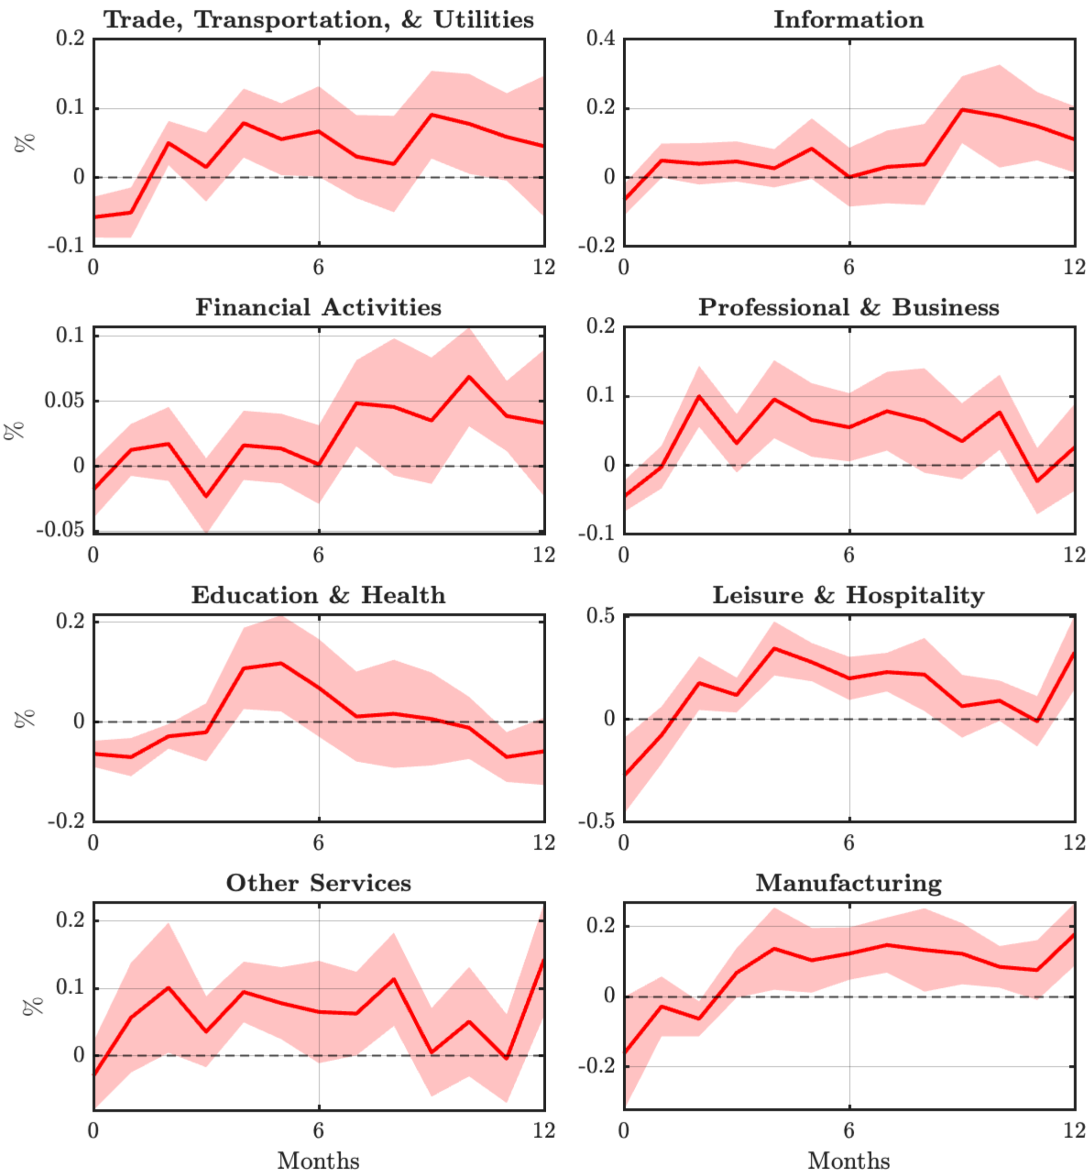

Note: Figure 5.4 presents the industry-level hours worked impulse responses to a one standard deviation pure AI uncertainty shocks, identified using an SVAR-IV and estimated using local projections. The industries are classified according to NAICS. The shaded regions represent 68% confidence intervals,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

Employment. The industry-level employment responses to pure AI uncertainty shocks broadly reflect the aggregate pattern of a modest and transitory decline, 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.

The Information industry, along with Professional and Business Services, experiences declines of approximately 0.5% before gradually returning to baseline. Financial Activities register a more persistent contraction of roughly 0.2%, with 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.

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

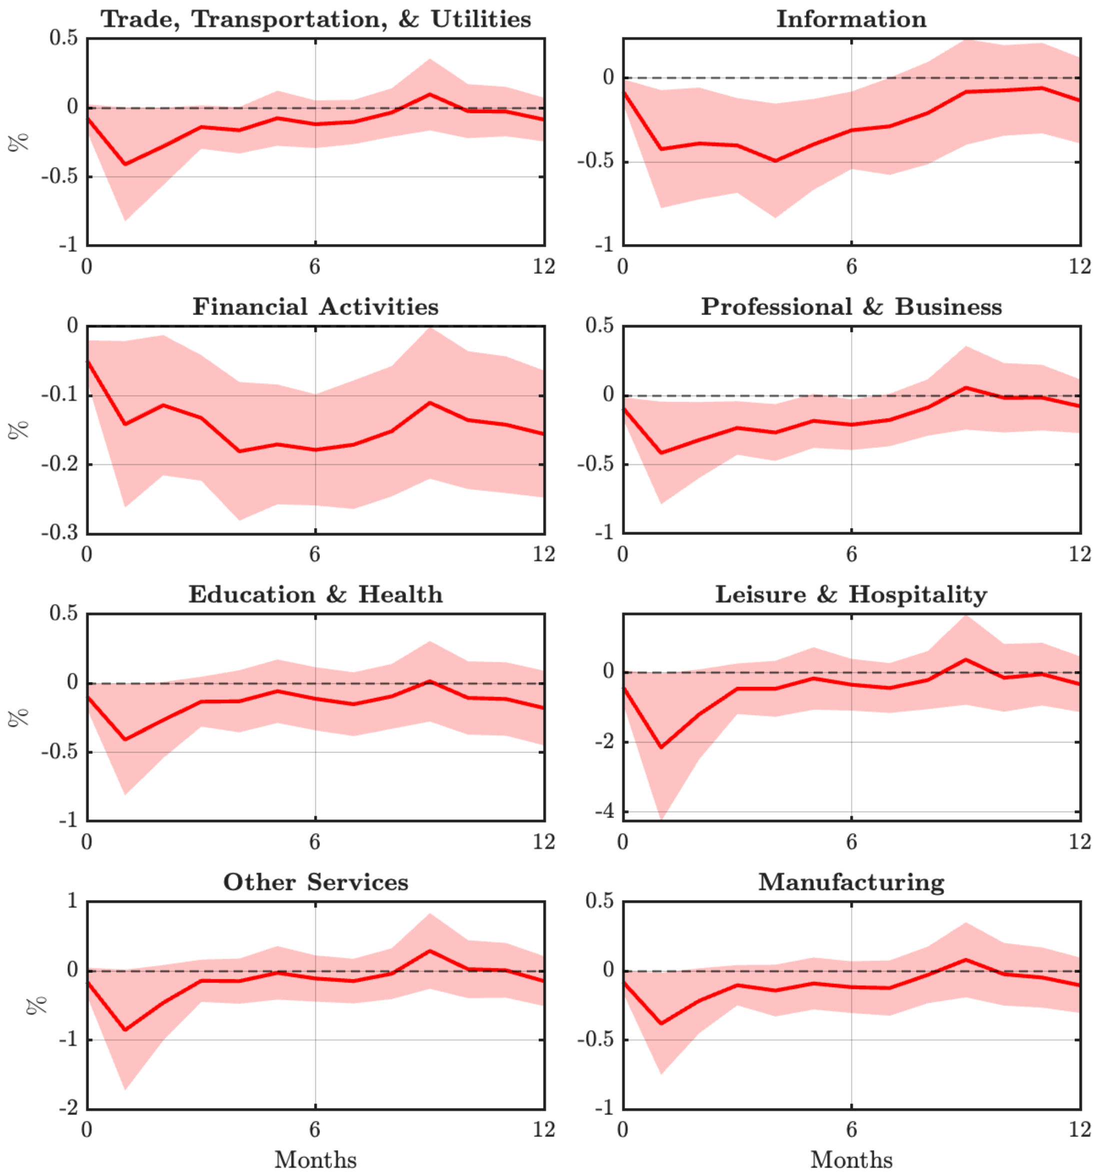

Note: Figure 5.5 presents the industry-level employment impulse responses to a one standard deviation pure AI uncertainty shocks, identified using an SVAR-IV and estimated using local projections. The industries are classified according to NAICS. The shaded regions represent 68% confidence intervals,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

These findings reveal that while the aggregate employment response masks substantial cross-industry variation, certain sectors face disproportionate exposure to labour 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.

5.6. The Role of AI Exposure and Compositional Effects

The industry-level estimates show variation that is not fully captured in the aggregate findings. While aggregate responses to AI uncertainty shocks show predominantly contractionary effects across labour market outcomes, the industry-level analysis reveals substantial heterogeneity. Responses differ sharply across industries. In particular, Financial Activities exhibits expansionary movements in wage, whereas several other sectors experience sizeable contractions. This divergence raises two natural questions: (1) what explains the differences in industry-level responses?, and (2) why do aggregate effects remain contractionary despite expansionary adjustments in some sectors?

To explore these questions, I examine the relationship between the exposure of these industries to AI and wage responses using the AI Industry Exposure (AIIE) Index from Felten et al. (2021), which measures the overlap between the required occupational abilities of an industry and the current capabilities of AI. The index is originally constructed at the 4-digit NAICS level. However, for this paper, I aggregate it to 2-digit industries to match the labour market data used in the local projections. It is standardised across industries with a mean of zero, where positive values indicate that the task composition of an industry is more susceptible to AI-related changes than the average industry, and negative values indicate less susceptibility.

Figure 5.6 plots industry-level wage responses obtained from local projection estimates at the 6-month horizon against AI exposure, with marker sizes proportional to average employment shares over the sample period, calculated as:

$$\bar{s}_i = \frac{1}{T}\sum_{t=1}^{T}\left(\frac{L_{i,t}}{\sum_{j=1}^{N} L_{j,t}}\right) \tag{5.3}$$

where $L_{i,t}$ is employment in industry i at time t and $\sum_{j=1}^{N} L_{j,t}$ is total employment across all $N = 8$ industries at time t, and T denotes the number of time periods in the sample.

The figure yields two complementary insights that together explain the aggregate-industry divergence. First, there is a strong positive correlation between AI exposure and wage responses at the 6-month horizon ($\rho = 0.79, p = 0.02$). Industries with higher AI exposure tend to display more positive wage responses. Financial

Activities records a positive response of approximately 0.06%, whereas Manufacturing exhibits a decline of roughly 0.11%. This relationship, however, varies across horizons. As shown in Appendix M, there is no systematic relationship on impact ($\rho = -0.12$, $p = 0.78$), suggesting that initial responses are driven by common uncertainty channels affecting industries broadly. A strong positive correlation emerges at the 3-month horizon ($\rho = 0.72$, $p = 0.04$), and strengthens at 6 months, and then weakens by the 12-month horizon ($\rho = 0.32$, $p = 0.44$). This pattern indicates that AI exposure plays a time-varying role, becoming most salient at intermediate horizons when industries process AI-related uncertainty, before other factors dominate at longer horizons.

Figure 5.6: Industry-Level Wage Response to AI Uncertainty Shock by AI Exposure ($h = 6$)

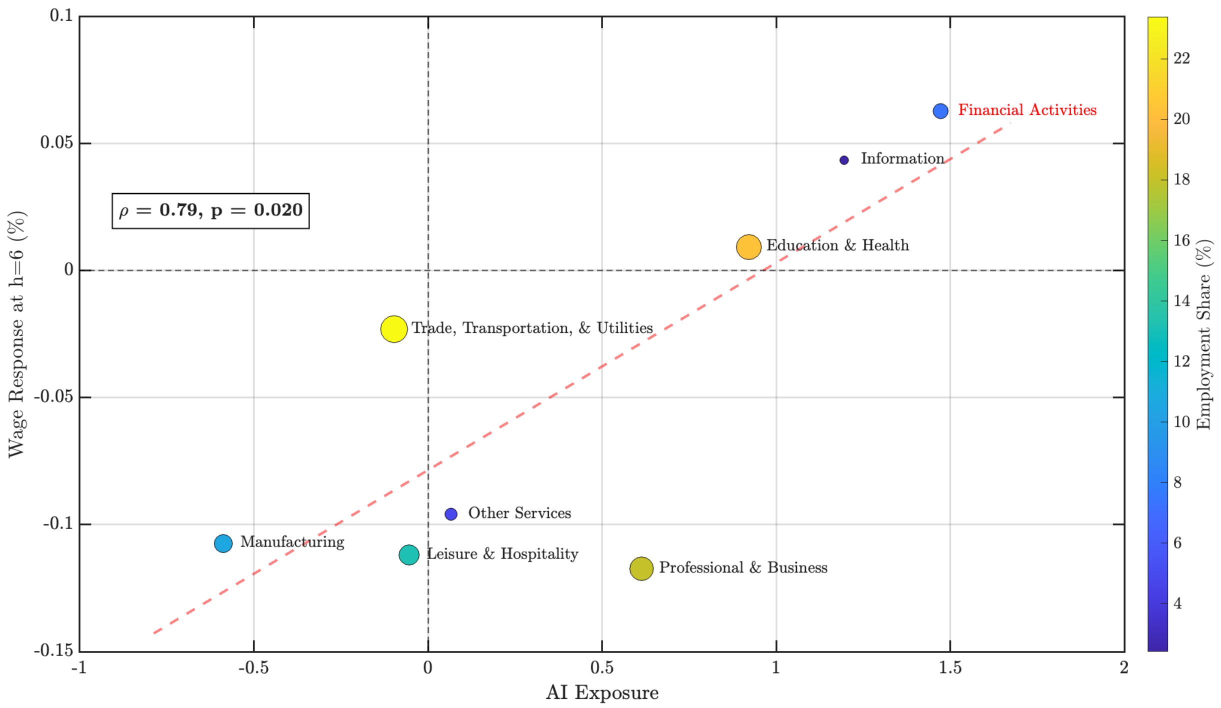

Note: Figure 5.6 plots the relationship between industry exposure to AI (x-axis) and wage responses to a one standard deviation AI uncertainty shock at the 6-month horizon (y-axis). AI exposure is measured using the AIIE Index from Felten et al. (2021), aggregated to 2-digit NAICS industries. Marker sizes represent average employment shares ($\bar{s}_i$) over the sample period. The dashed line shows the fitted linear relationship. Small employment shares in high-exposure sectors, e.g., financial activities, explain why aggregate response remains contractionary despite some industries experiencing positive responses. See Appendix M for results at other horizons.

Second, the contribution of industries with positive responses to aggregate outcomes depends critically on employment composition. Financial Activities and Information, the two industries with positive responses, together account for approximately 10.0% of total employment (i.e., 6.2% and 3.5%, respectively) as reflected in Figure 5.6. By contrast, industries with contractionary responses represent substantially larger employment shares. As a result, positive responses in small and high-exposure sectors are outweighed by modest declines in employment-intensive industries. Aggregate responses, therefore, reflect weighted averages in which large sectors with negative effects dominate, explaining why aggregate wage responses remain contractionary across horizons despite pronounced sectoral heterogeneity.

The time-varying correlation between AI exposure and wage responses suggests that industry characteristics interact with uncertainty in complex ways. High exposure sectors may face greater uncertainty about task reallocation and skill requirements, leading to distinct adjustment dynamics at intermediate horizons. However, other factors such as skill composition, capital intensity, and sector-specific cyclical patterns could also shape responses. The weakening of the exposure-response relationship at longer horizons further indicates that AI-related uncertainty may interact with industry characteristics primarily over short to medium horizons, before broader macroeconomic dynamics become more influential. Decomposing the relative importance of AI exposure versus these institutional and structural characteristics represents a valuable direction for future research. The observed association between responses and *ex ante* measures of AI exposure (Felten et al., 2021) nevertheless suggests that the identified shock captures meaningful variation in AI-related uncertainty rather than reflecting only conventional macroeconomic fluctuations.

6. Conclusion

The rapid emergence of AI has intensified debate over its macroeconomic 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 divergence between these views, together with the pace of AI development and adoption, implies that uncertainty is intrinsic to the contemporary 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 shocks.

The empirical analysis established three main findings. First, positive AI uncertainty shocks are contractionary, but their adjustment pattern displays important differences relative to conventional uncertainty shocks. Rather than producing broad and transitory contractions across real activity, AI-related uncertainty is characterised by pronounced and front-loaded adjustments in equity markets and labour market conditions, alongside comparatively muted and less persistent effects on employment and aggregate output. These dynamics share some common features with established uncertainty episodes, but display a distinct configuration that reflects the nature of uncertainty surrounding AI.

Second, the response pattern resembles behaviour predicted by models in which labour is treated as a quasi-fixed factor (Oi, 1962), but it departs from standard formulations in important ways. Wages decline persistently rather than remaining rigid, consistent with weakened worker bargaining power when AI-related uncertainty makes outside options difficult to assess (Leduc and Liu, 2024). Firms may therefore retain workers not in anticipation of cyclical recovery, but to preserve investments in firm-specific training and tacit knowledge embedded in operational procedures and organisational routines that cannot be rapidly reconstituted (Oi, 1962). A deeper distinction concerns the nature of uncertainty itself. Conventional uncertainty tends to resolve as economic conditions become clearer, whereas AI-related uncertainty persists as the technology continues to evolve and uncertainty over task automation and skill relevance remains unresolved. This ongoing uncertainty provides a natural explanation for why AI uncertainty is associated with sustained wage and labour-input adjustments, rather than the temporary contractions and subsequent recoveries that characterise many other uncertainty shocks.

Third, industry-level evidence reveals substantial heterogeneity beneath these aggregate responses. Most industries experience wage declines, whereas the financial industry displays countercyclical wage increases. Hours worked contract on impact across industries but subsequently recover and move above baseline in several sectors, even as the aggregate series remains slightly below baseline. This aggregate-industry divergence stems from two complementary factors. Compositional effects matter because industries with contractionary responses account for more than half of employment, whereas industries with expansionary responses comprise a much smaller share. Responses also vary with AI exposure (Felten et al., 2021), a relationship that is strongest at intermediate horizons. These patterns establish that AI-related uncertainty does not propagate uniformly across industries, with both employment composition and technological exposure determining the magnitude and direction of industry-level adjustments.

These findings have two closely related implications. For business cycle analysis, AI uncertainty represents an emerging source of macroeconomic fluctuations with transmission that is not well captured by the existing uncertainty measures. Even before any long-run productivity effects of AI adoption materialise, uncertainty surrounding AI is associated with contractionary movements in equity markets, labour outcomes, and real activity. From a measurement standpoint, the results support treating AI uncertainty as a separate component within the broader class of uncertainty shocks, rather than subsuming it within general economic policy or financial uncertainty. This distinction helps isolate the economic effects of AI-specific uncertainty, clarifying how it shapes macroeconomic outcomes independently of other sources of volatility.

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.

Leduc, S. and Liu, Z. (2024). Automation, bargaining power, and labor market fluctuations. *American Economic Journal: Macroeconomics*, 16:311–349.

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.

Oi, W. Y. (1962). Labor as a quasi-fixed factor. *Journal of Political Economy*, 70:538–555.

Piffer, M. and Podstawski, M. (2018). Identifying Uncertainty Shocks Using the Price of Gold. *The Economic Journal*, 128(616):3266–3284.

Plagborg-Møller, M. and Wolf, C. K. (2021). Local projections and vars estimate the same impulse responses. *Econometrica*, 89:955–980.

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.

Stock, J. H. and Watson, M. W. (2018). Identification and estimation of dynamic causal effects in macroeconomics using external instruments. In *Economic Journal*, volume 128, pages 917–948. Blackwell Publishing Ltd.

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)

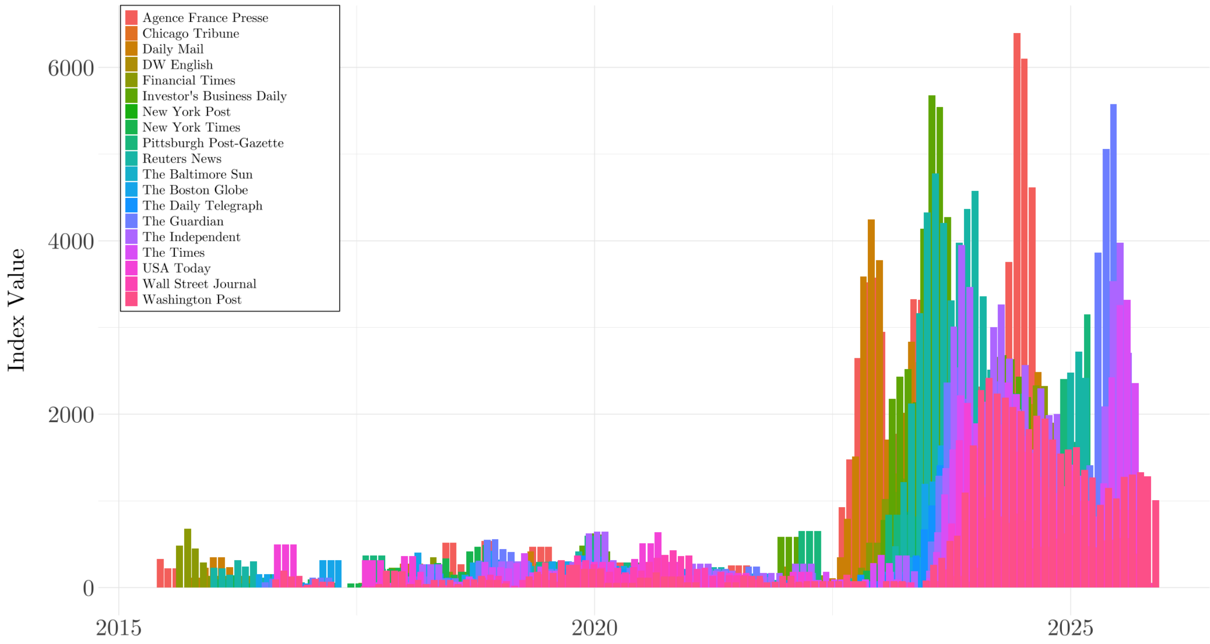

Note: Figure B plots the AIU Index per news outlet from M1:2016 to M4:2025. The index is normalised such that its mean equals 100. It reflects the standardised share of newspaper articles that simultaneously reference AI, economic conditions, and uncertainty.

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

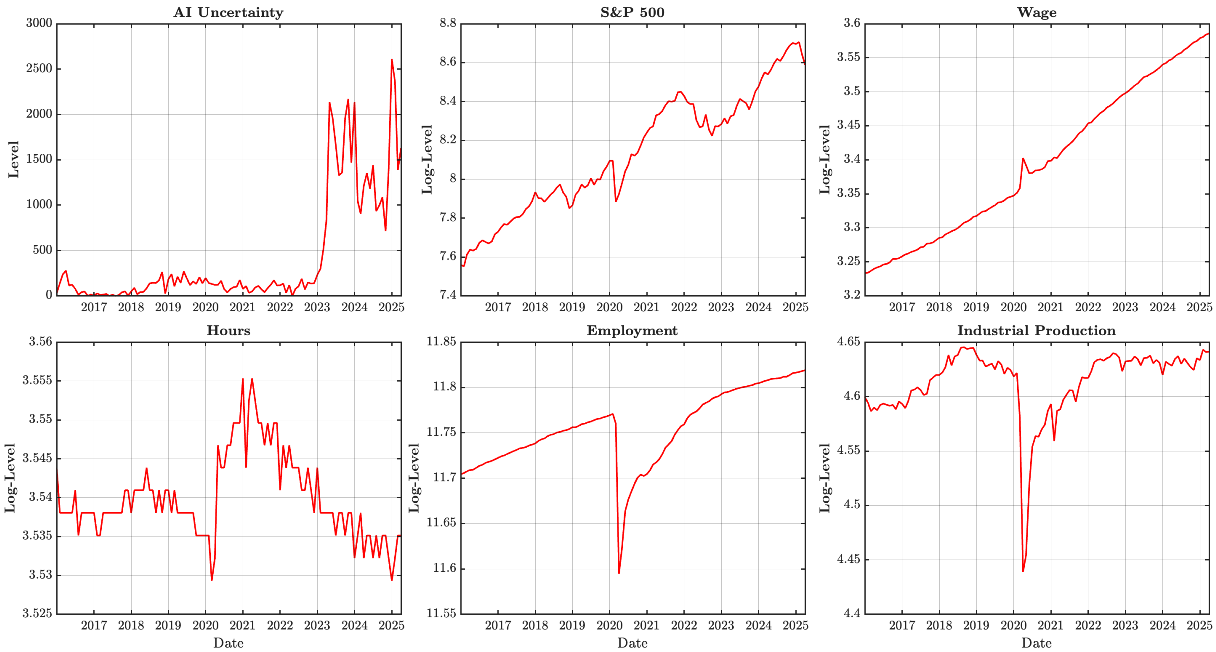

Note: Figure D plots the variables used in the analysis. The AIU Index is expressed in levels, while all other variables are expressed in log-levels. The data cover the period from M1:2016 to M4:2025. All series, except the AIU Index, are sourced from the US BLS and FRED unless otherwise noted.

Appendix E

Internal Instrument Approach

Methodology. SVAR-IV relies on an invertibility condition that requires the structural shock to be recoverable from current and past observables. Formally, invertibility implies that the structural shock can be expressed as a function of current and lagged VAR residuals. When relevant information is omitted from the VAR, this condition may fail, leading to inconsistent impulse responses (Stock and Watson, 2018).

To address potential invertibility violations, I also implement an internal instrument approach as in Plagborg-Møller and Wolf (2021) and Känzig (2023). This strategy augments the VAR with the instrument itself, thereby expanding the information set available for identification and avoiding the need to impose invertibility. The cost of this is a stronger exogeneity requirement. In particular, the instrument must be orthogonal not only to contemporaneous structural shocks, but also to all leads and lags of those shocks:

$$E[z_t, \varepsilon^j_{t+k}] = 0 \quad \text{for all} \quad k \neq 0, \tag{F.1}$$

where k indexes time leads and lags. This restriction is stronger than the contemporaneous exogeneity condition imposed in the external instrument framework, but it allows identification even when the structural shock is non-invertible.

Under this assumption, identification is achieved by augmenting the VAR with the instrument ordered first and recovering impulse responses from the first orthogonalised innovation. Specifically, I obtain the impulse response vector associated with the AI uncertainty shock as:

$$\bar{s}_1 = \frac{[\text{chol}(\bar{\Sigma})]_{\cdot,1}}{\text{chol}(\bar{\Sigma})_{1,1}}, \tag{F.2}$$

where $\bar{\Sigma}$ denotes the variance-covariance matrix of residuals from the augmented VAR. By construction, this procedure delivers consistent estimates of relative impulse responses regardless of whether the shock is invertible, or the instrument is measured with error (Plagborg-Møller and Wolf, 2021).

Results. Figure E reports the impulse responses obtained from the external SVAR-IV used in the main analysis and from the internal SVAR-IV specification. The responses are very similar, both qualitatively and quantitatively. Sign patterns are consistent across all variables, and the magnitudes and dynamics are closely aligned. Overall, however, these findings suggest that the results are robust to relaxing the assumption of invertibility.

Figure E: Impulse Responses to AI Uncertainty Shock

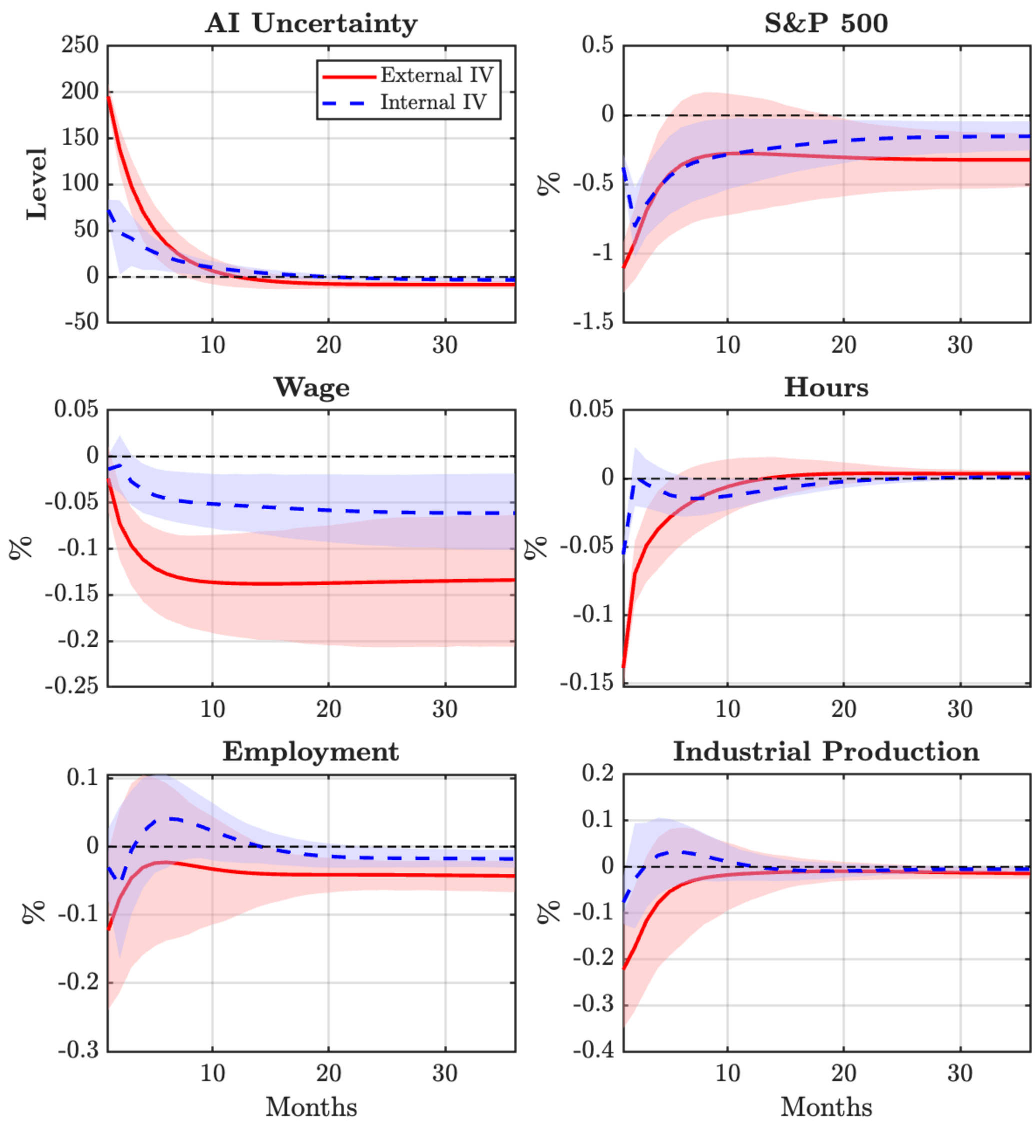

Note: Figure E presents the impulse responses to a one standard deviation AI uncertainty shocks, identified using an external (solid lines) and internal (dashed lines) SVAR-IV. Shaded areas denote 68% confidence bands based on 1,000 wild bootstrap replications.

Appendix F

Figure F: AIU Index Upper Tail Realisations

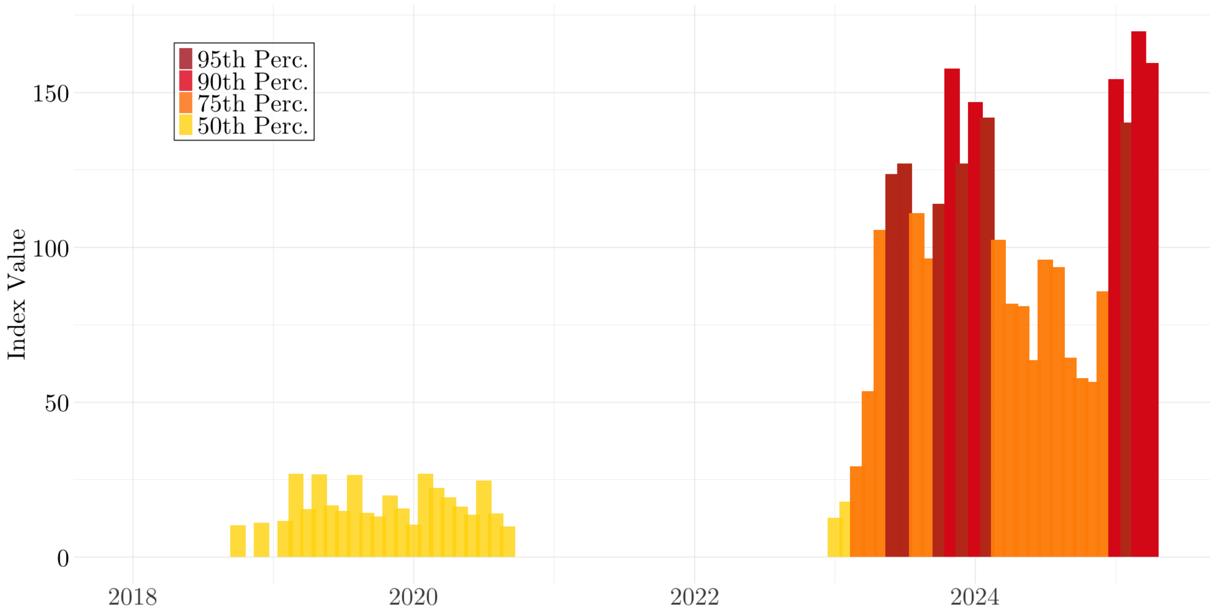

Note: Figure F plots the tail realisations (0.95, 0.90, 0.75, and 0.50 quantiles) of the AIU Index from M1:2016 to M4:2025. The AIU Index reflects the standardised share of news articles that simultaneously reference AI, economic conditions, and uncertainty.

Appendix G

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

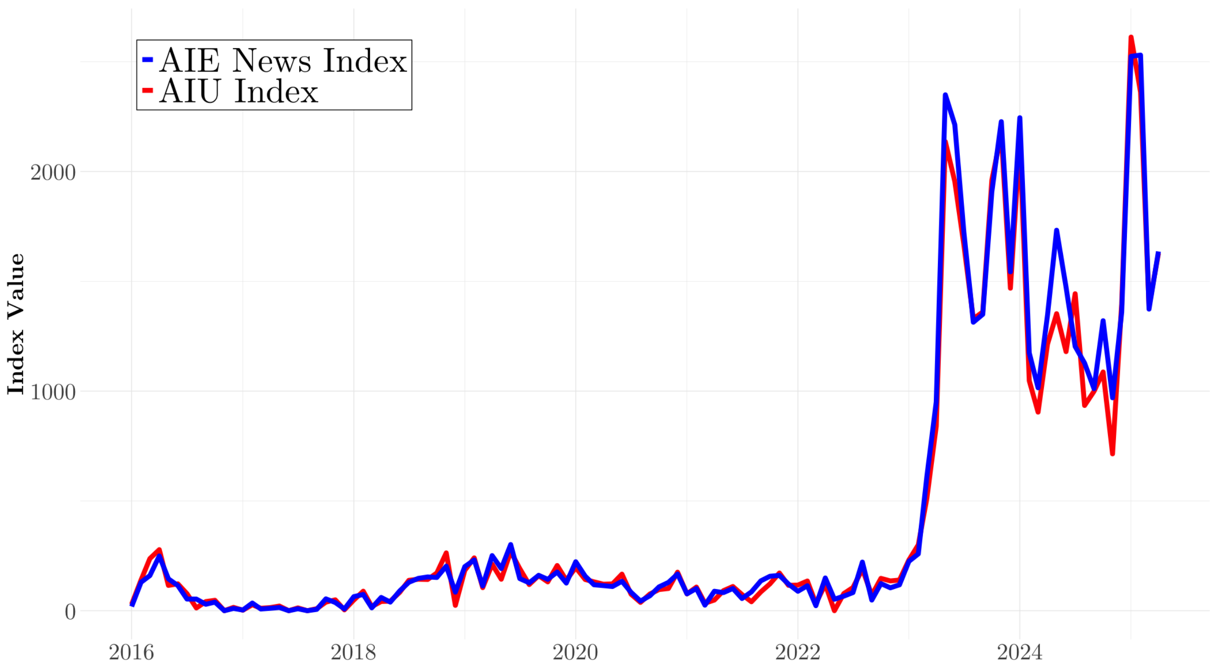

Note: Figure G plots the AIE News Index from M1:2016 to M4:2025. For presentation purposes, a three-month moving average is applied. The AIE News Index follows the same construction steps as the AIU Index but excludes uncertainty-related keywords.

Appendix H

Table H: 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 H 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 I

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 I: 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 $|z| > 2$ *observations: 6.3% of sample*

Total $|z| > 3$ *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 I reports the number of excluded observations under each rule.

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

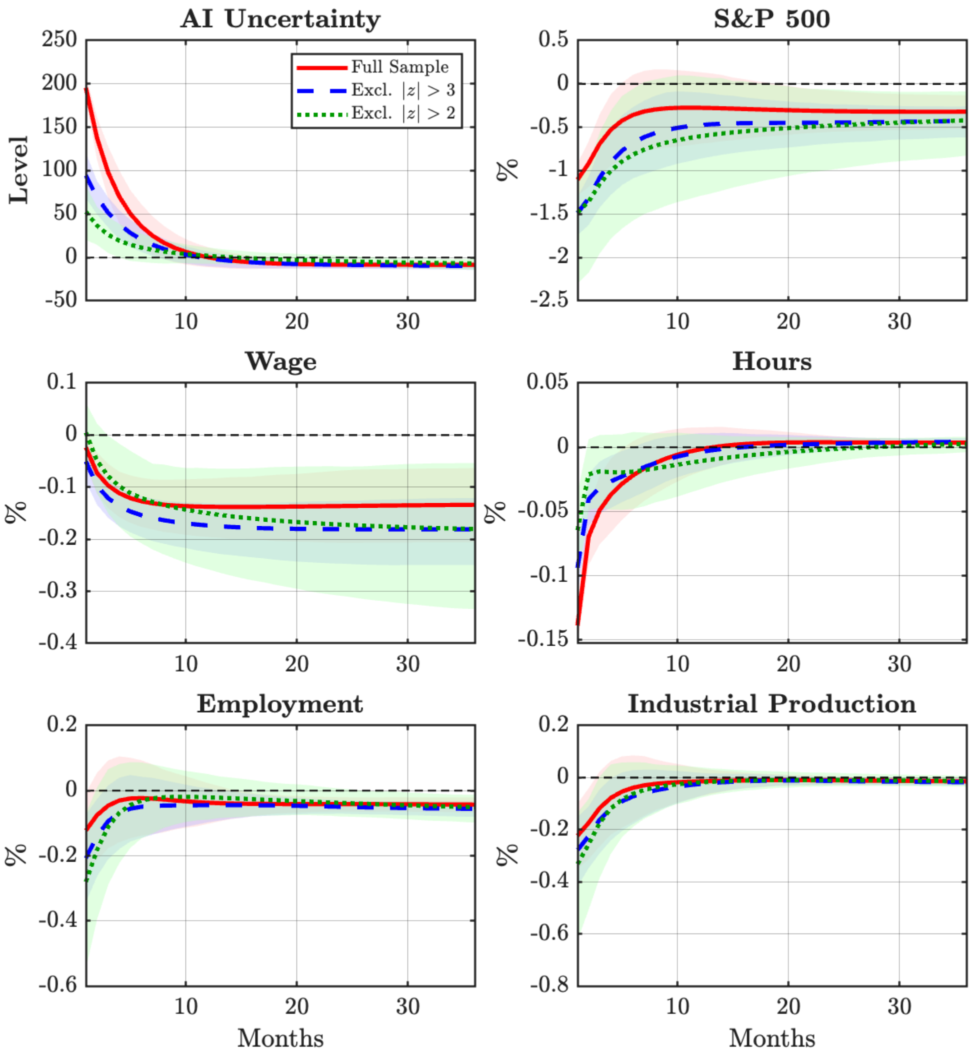

Note: Figure I displays the impulse responses to a one standard deviation shock in AI uncertainty across three specifications via SVAR-IV using the residual-based instrument discussed in Section 4: (1) full sample (red line), (2) excluding observations with standardised residuals $|z| > 3$ in the AIU equation (blue dashed line), and (3) excluding observations with $|z| > 2$ (green dotted line). 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 J

Table J: 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 J 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 K

Table K: 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 K 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 L

Figure L.1: Impulse Responses to Economic Policy Uncertainty Shock (Recursive Identification)

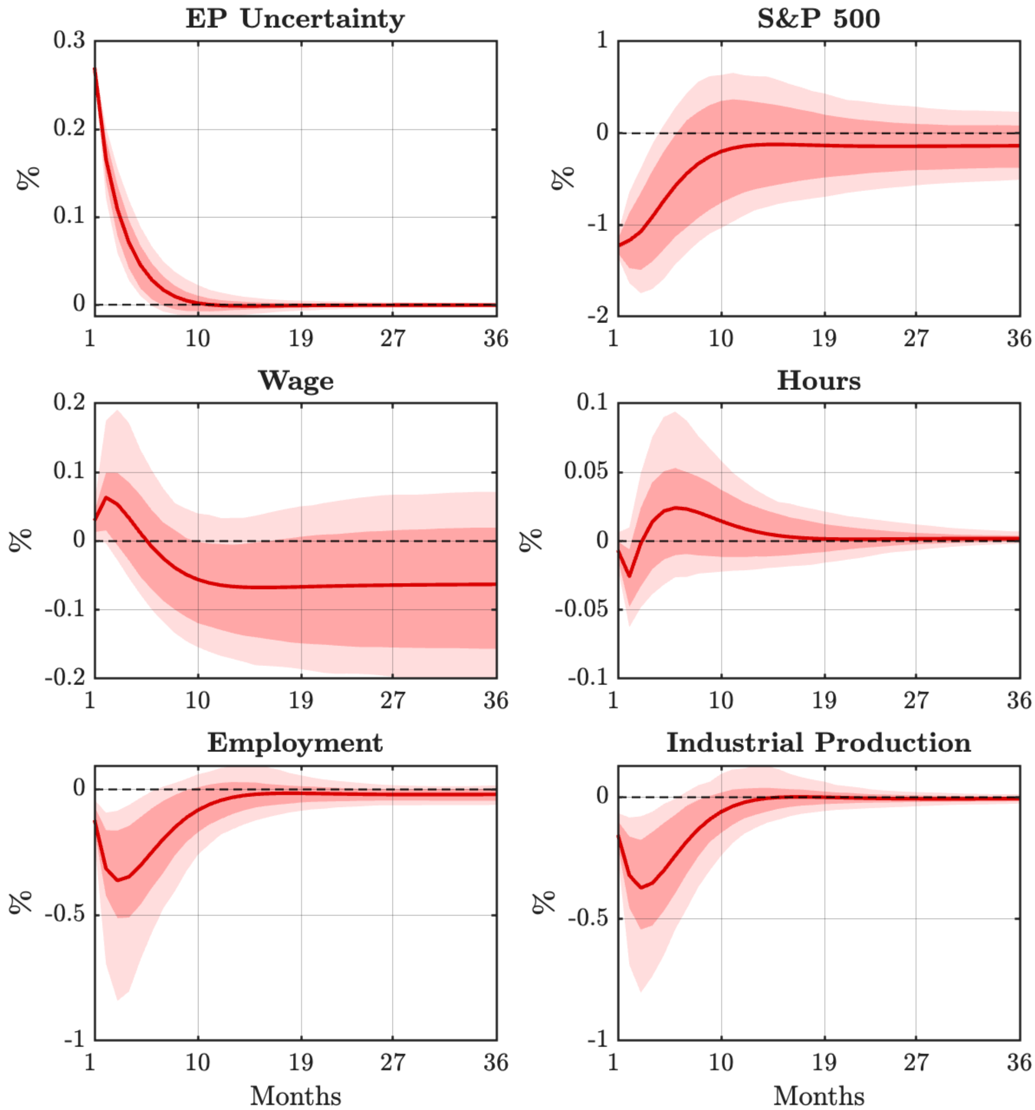

Note: Figure L.1 displays the impulse responses to a one standard deviation shock in economic policy uncertainty estimated with an SVAR under recursive identification. The light and dark shaded areas denote 68% and 90% confidence bands based on 1,000 wild bootstrap replications.

Appendix L

Figure L.2: Impulse Responses to Real Economic Uncertainty Shock (Recursive Identification)

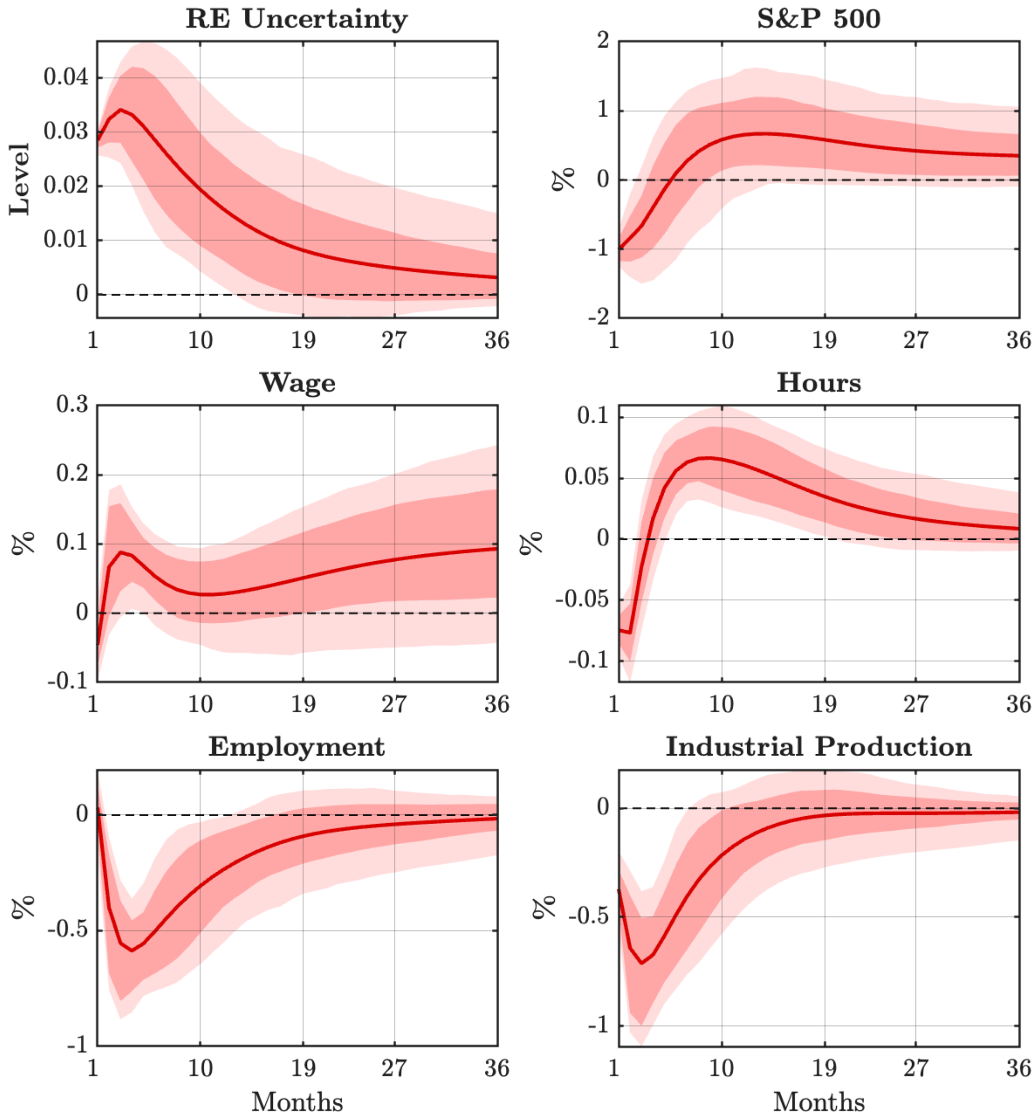

Note: Figure L.2 displays the impulse responses to a one standard deviation shock in real economic uncertainty estimated with an SVAR under recursive identification. The light and dark shaded areas denote 68% and 90% confidence bands based on 1,000 wild bootstrap replications.

Appendix L

Figure L.3: Impulse Responses to Uncertainty Shock (VIX) (Recursive Identification)

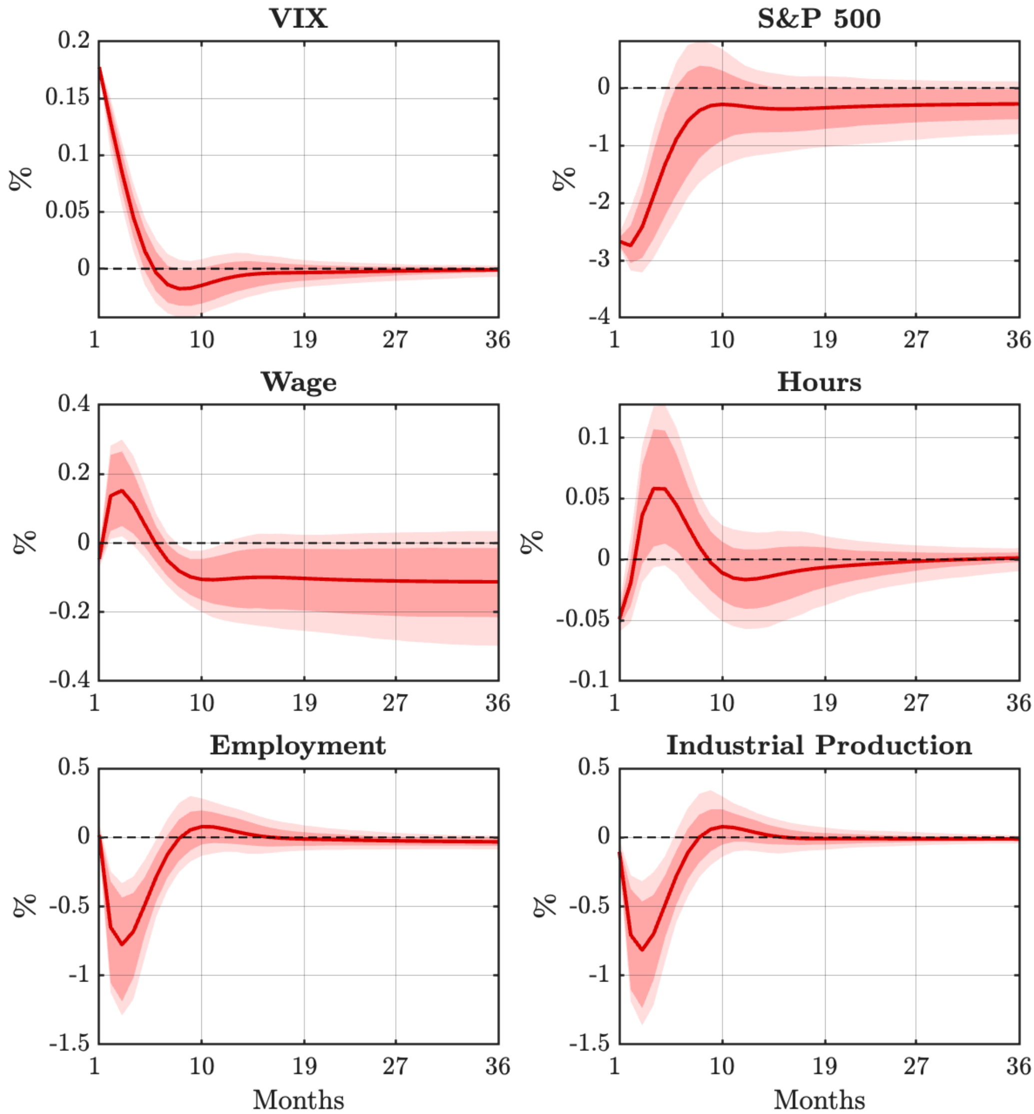

Note: Figure L.3 displays the impulse responses to a one standard deviation shock in uncertainty (measured using VIX) estimated with an SVAR under recursive identification. The light and dark shaded areas denote 68% and 90% confidence bands based on 1,000 wild bootstrap replications.

Appendix L

Figure L.4: Impulse Responses to Uncertainty Shock (VIX)
(Recursive Identification)

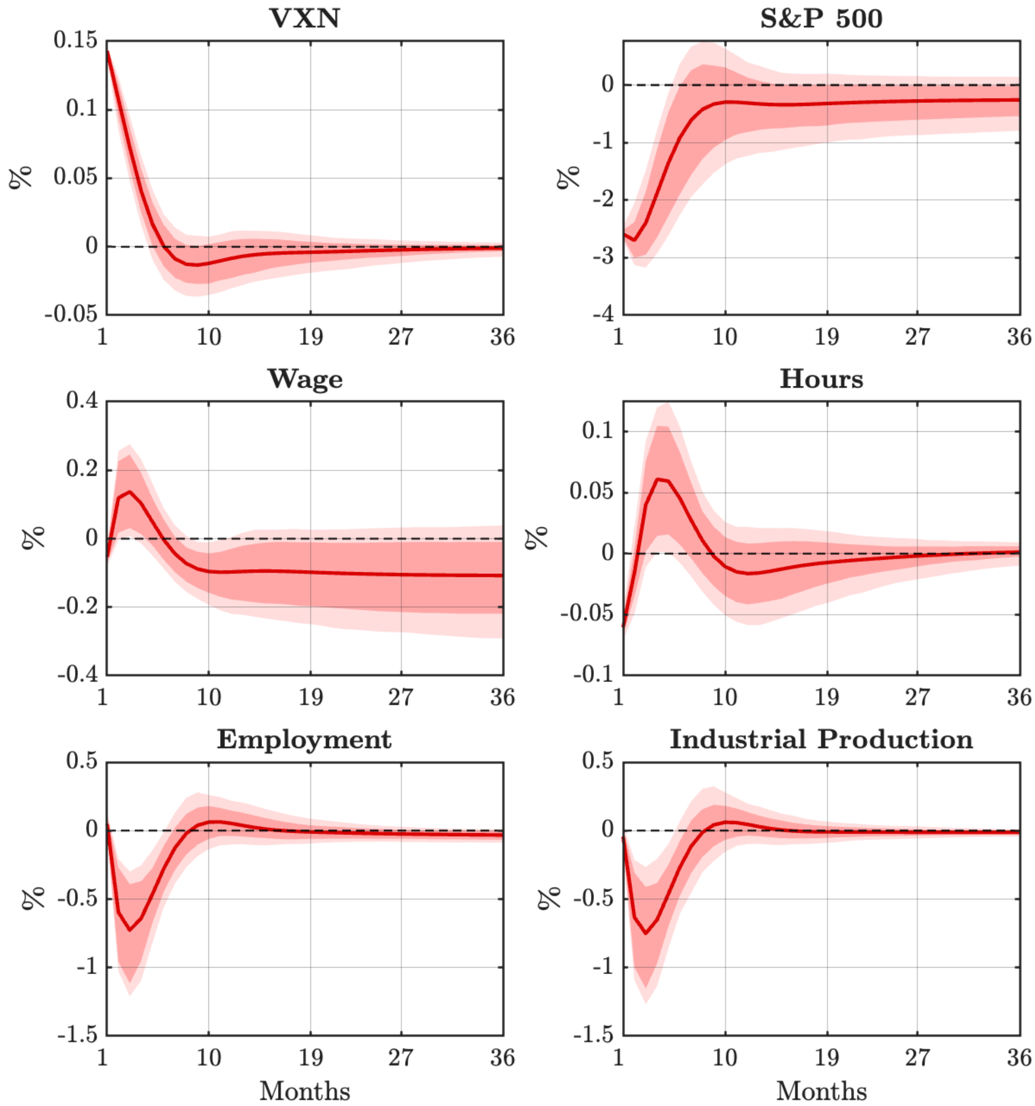

Note: Figure L.4 displays the impulse responses to a one standard deviation shock in uncertainty (measured using VXN) estimated with an SVAR under recursive identification. The light and dark shaded areas denote 68% and 90% confidence bands based on 1,000 wild bootstrap replications.

Appendix M

Figure M.1: Industry-Level Wage Response to AI Uncertainty Shock by AI Exposure ($h = 0$)

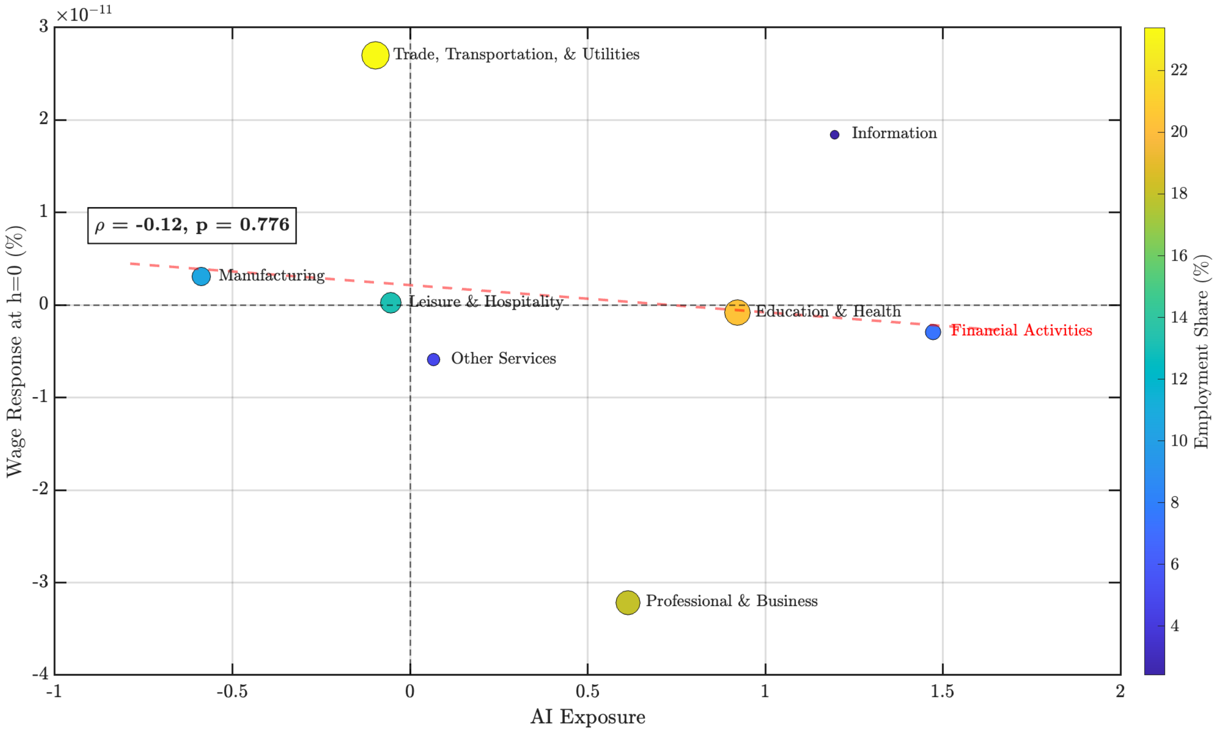

Note: Figure M.1 plots the relationship between industry exposure to AI (x-axis) and wage responses to a one standard deviation AI uncertainty shock at the 0-month horizon (y-axis). AI exposure is measured using the AIIE Index from Felten et al. (2021), aggregated to 2-digit NAICS industries. Marker sizes represent average employment shares ($\bar{s}_i$) over the sample period. The dashed line shows the fitted linear relationship.

Appendix M

Figure M.2: Industry-Level Wage Response to AI Uncertainty Shock by AI Exposure ($h = 3$)

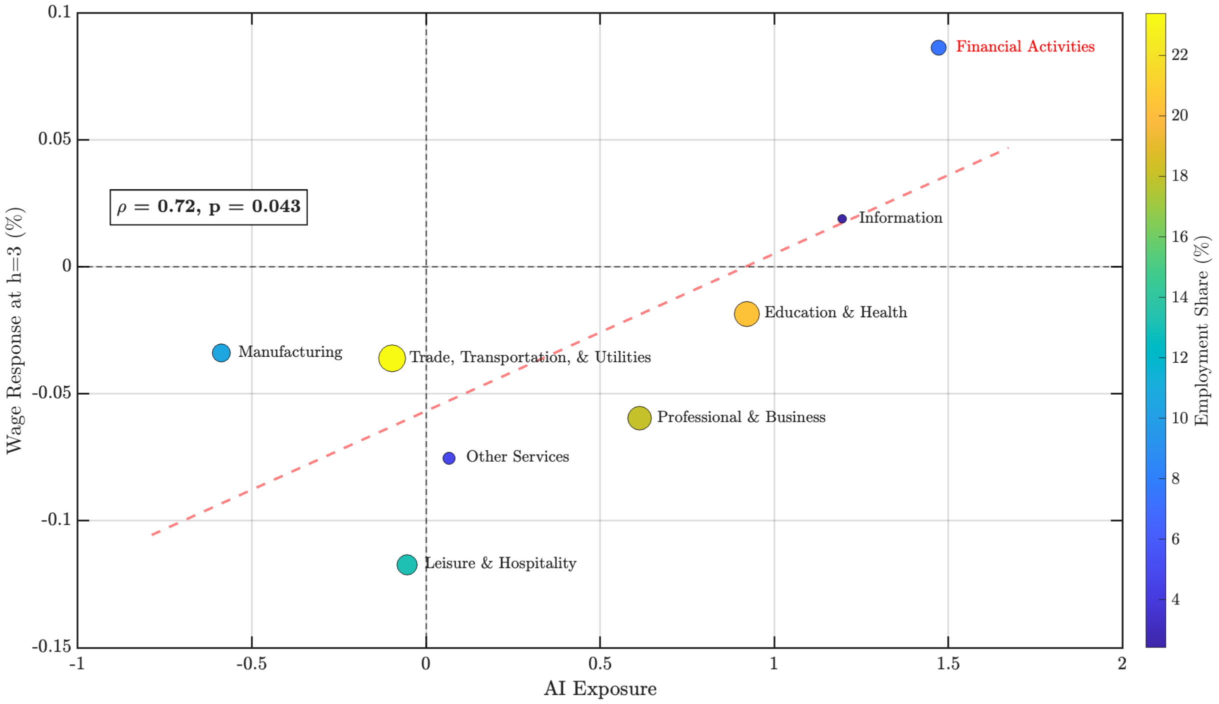

Note: Figure M.2 plots the relationship between industry exposure to AI (x-axis) and wage responses to a one standard deviation AI uncertainty shock at the 3-month horizon (y-axis). AI exposure is measured using the AIIE Index from Felten et al. (2021), aggregated to 2-digit NAICS industries. Marker sizes represent average employment shares ($\bar{s}_i$) over the sample period. The dashed line shows the fitted linear relationship.

Appendix M

Figure M.3: Industry-Level Wage Response to AI Uncertainty Shock by AI Exposure ($h = 12$)

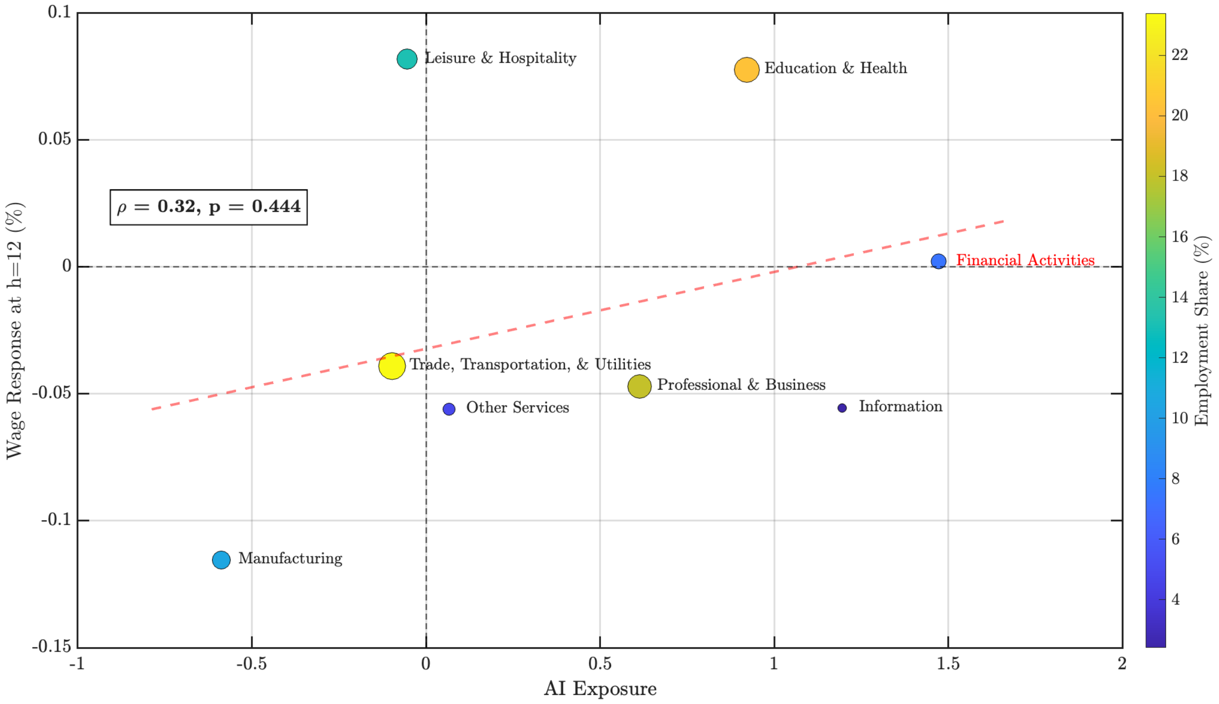

Note: Figure M.3 plots the relationship between industry exposure to AI (x-axis) and wage responses to a one standard deviation AI uncertainty shock at the 12-month horizon (y-axis). AI exposure is measured using the AIIE Index from Felten et al. (2021), aggregated to 2-digit NAICS industries. Marker sizes represent average employment shares ($\bar{s}_i$) over the sample period. The dashed line shows the fitted linear relationship.

Appendix N

Figure N.1: Impulse Responses to AI Uncertainty Shock
Recursive Identification (Different Lag Specification)

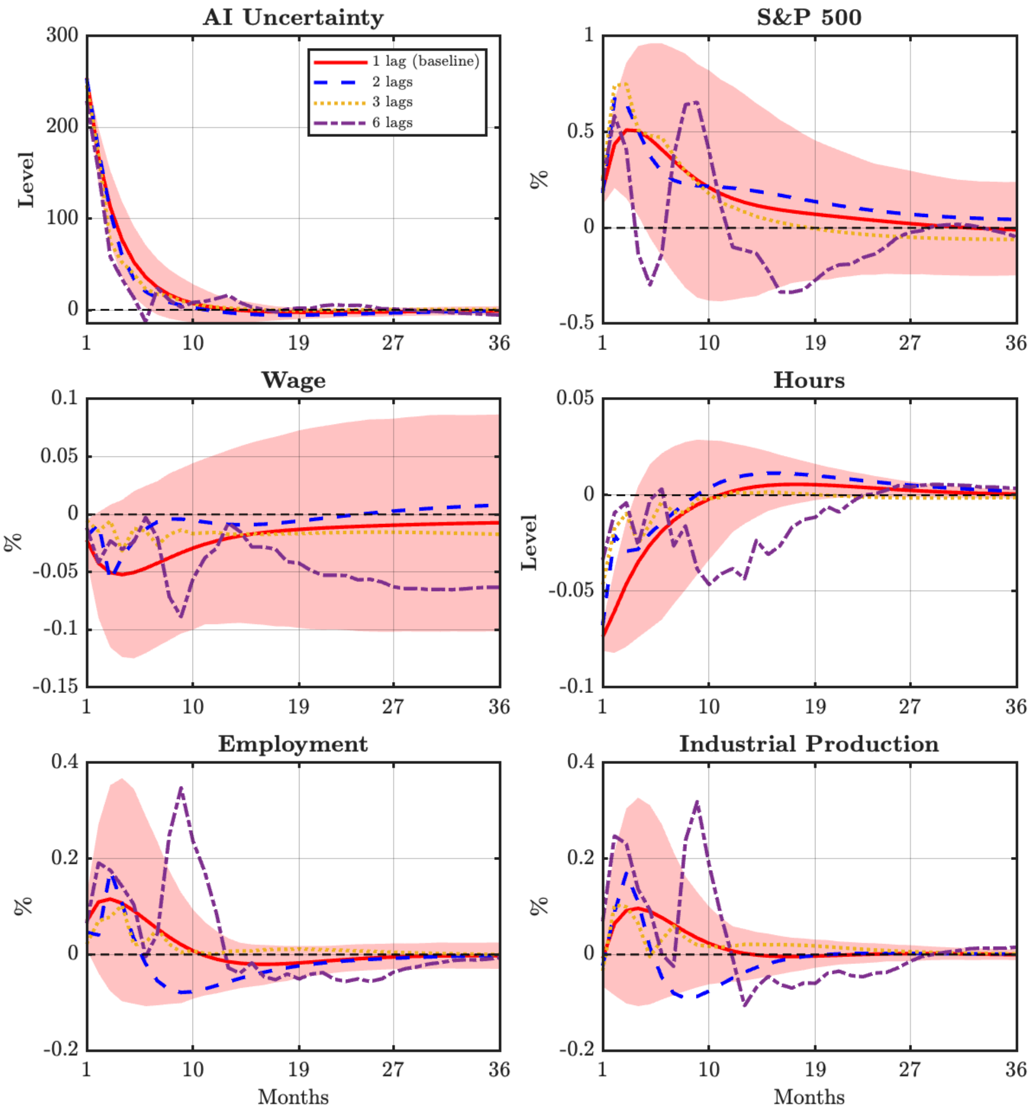

Note: Figure N.1 displays the impulse responses to a one standard deviation shock in AI uncertainty estimated with an SVAR under recursive identification with different lag specifications. Shaded areas denote 68% confidence bands based on 1,000 wild bootstrap replications.

Appendix N

Figure N.2: Impulse Responses to AI Uncertainty Shock SVAR-IV (Different Lag Specification)

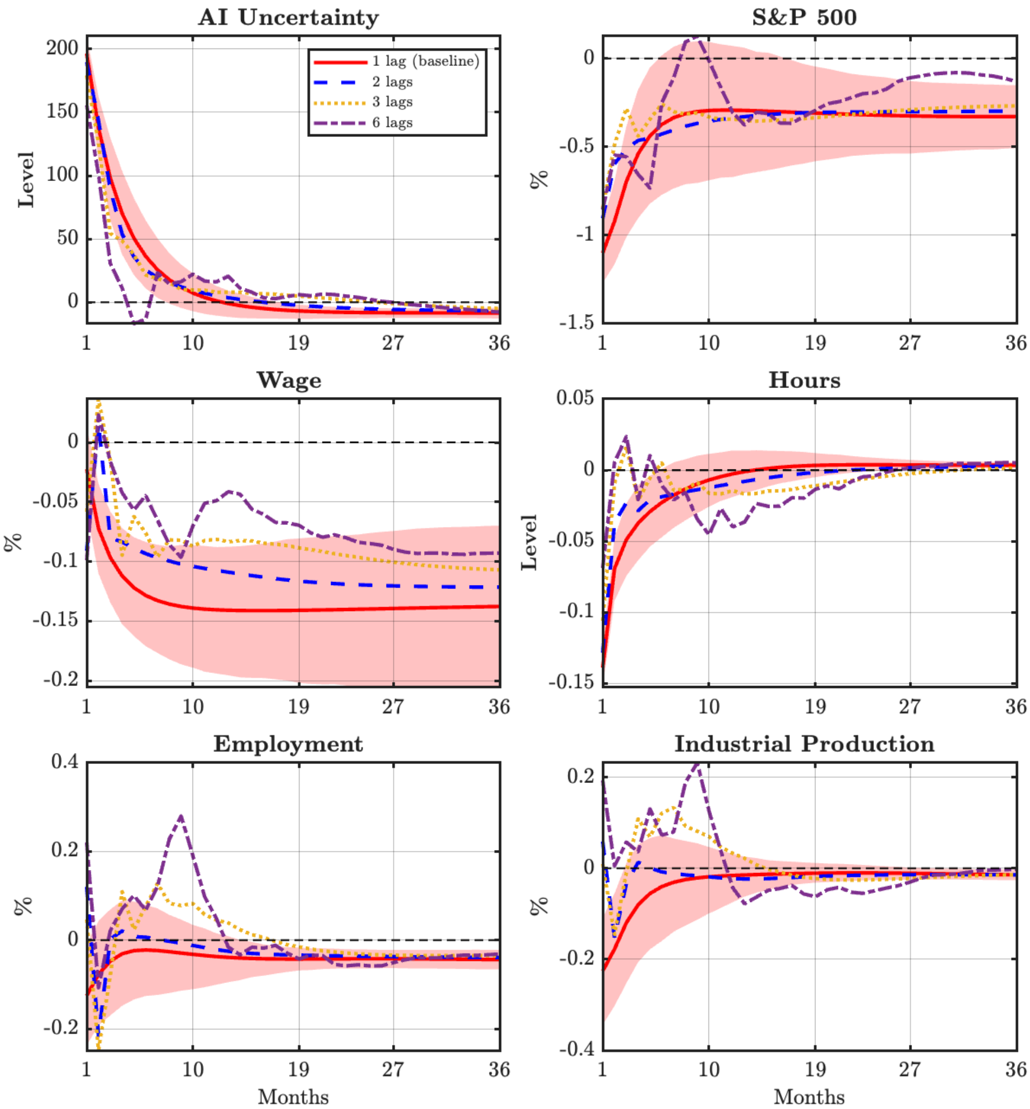

Note: Figure N.2 displays the impulse responses to a one standard deviation shock in AI uncertainty estimated using SVAR-IV with different lag specifications. The instrument is the residual-based measure of AI uncertainty from Equation (4.9). Shaded areas denote 68% confidence bands based on 1,000 wild bootstrap replications.

Appendix O

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

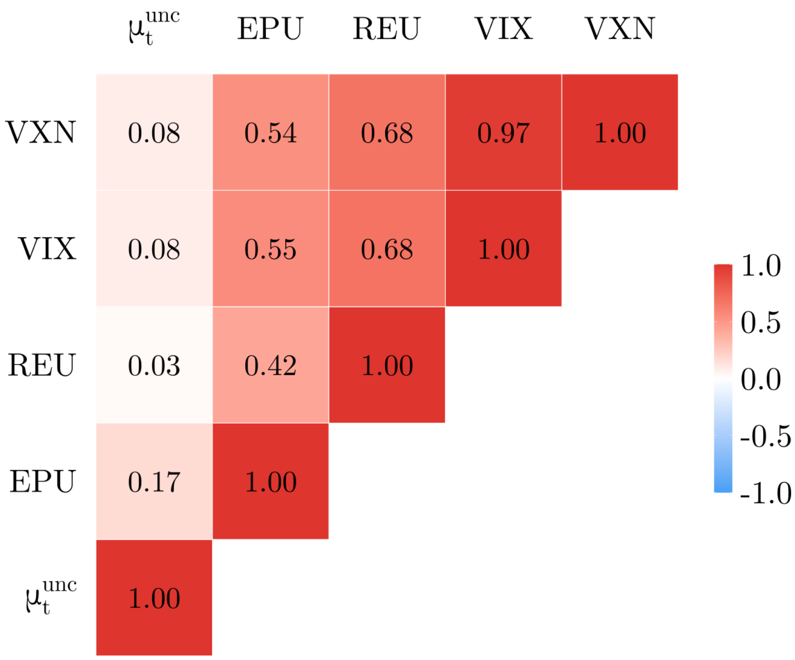

Note: Figure O reports the Pearson correlation coefficients between residual-based instrument μ_t^{unc} and benchmark uncertainty measures. The benchmarks are the EPU Index, the REU, and the VIX. The sample spans from M1:2016 to M4:2025, reflecting the availability of news articles used in constructing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

Appendix P

Figure P.1: Response of Industry-Level Wage to AI Uncertainty Shock (Different Lag Specifications)

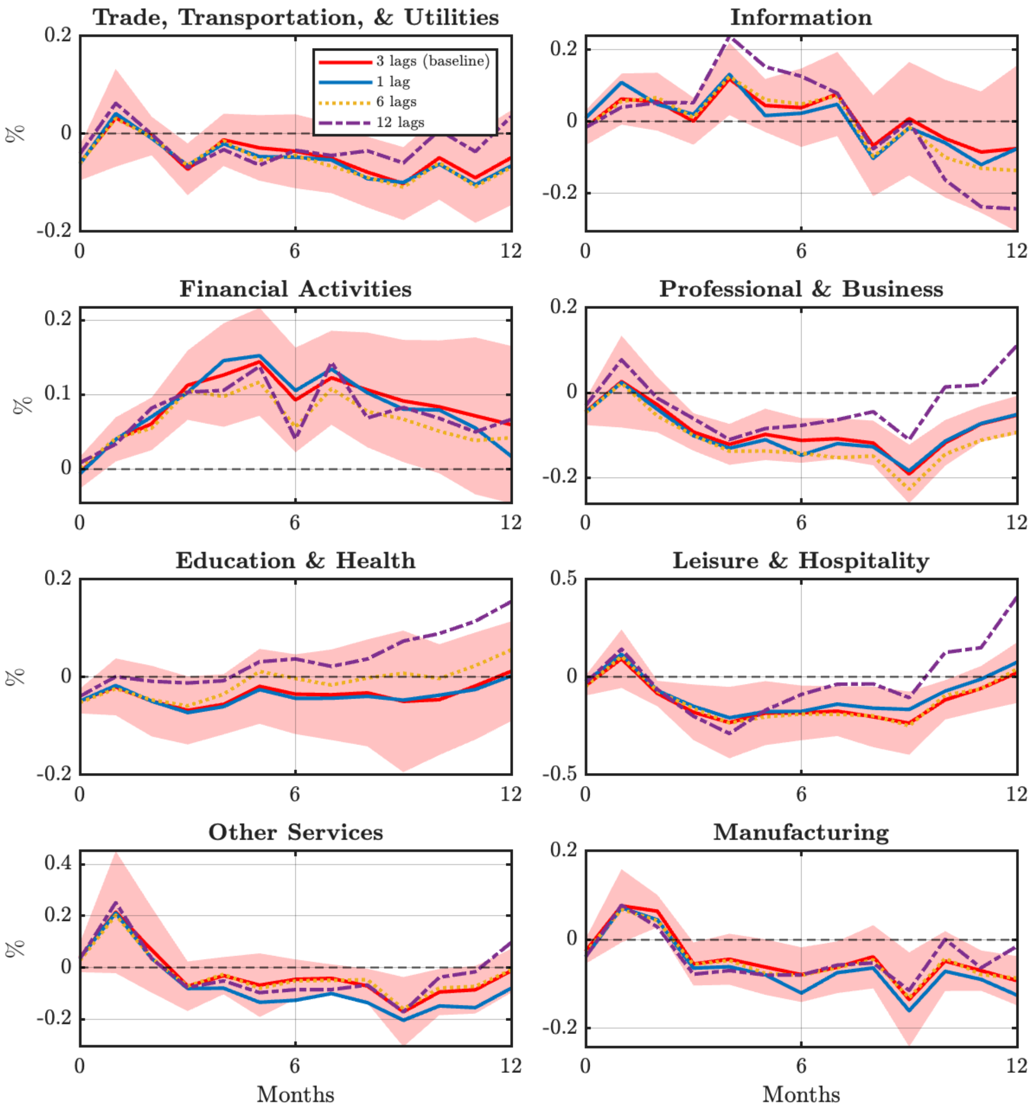

Note: Figure P.1 presents the industry-level wage impulse responses to a one standard deviation "pure" AI uncertainty shocks estimated via local projections with different lag specifications. The industries are classified according to NAICS. The shaded regions represent 68% confidence intervals, computed using Newey–West standard errors.

Appendix P

Figure P.2: Response of Industry-Level Hours Worked to AI Uncertainty Shock (Different Lag Specifications)

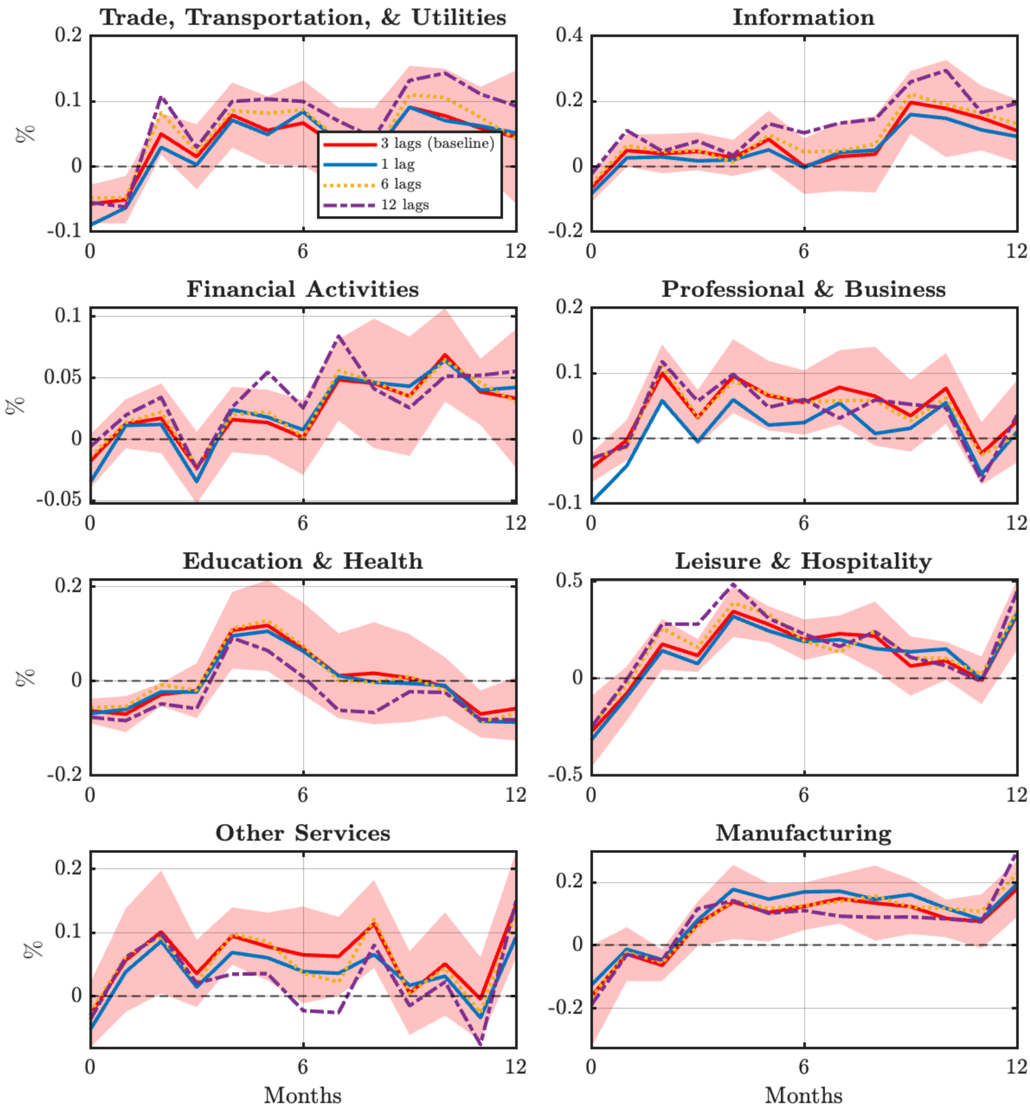

Note: Figure P.2 presents the industry-level hours worked impulse responses to a one standard deviation "pure" AI uncertainty shocks estimated via local projections with different lag specifications. The industries are classified according to NAICS. The shaded regions represent 68% confidence intervals, computed using Newey–West standard errors.

Appendix P

Figure P.3: Response of Industry-Level Employment to BAI Uncertainty Shock (Different Lag Specifications)

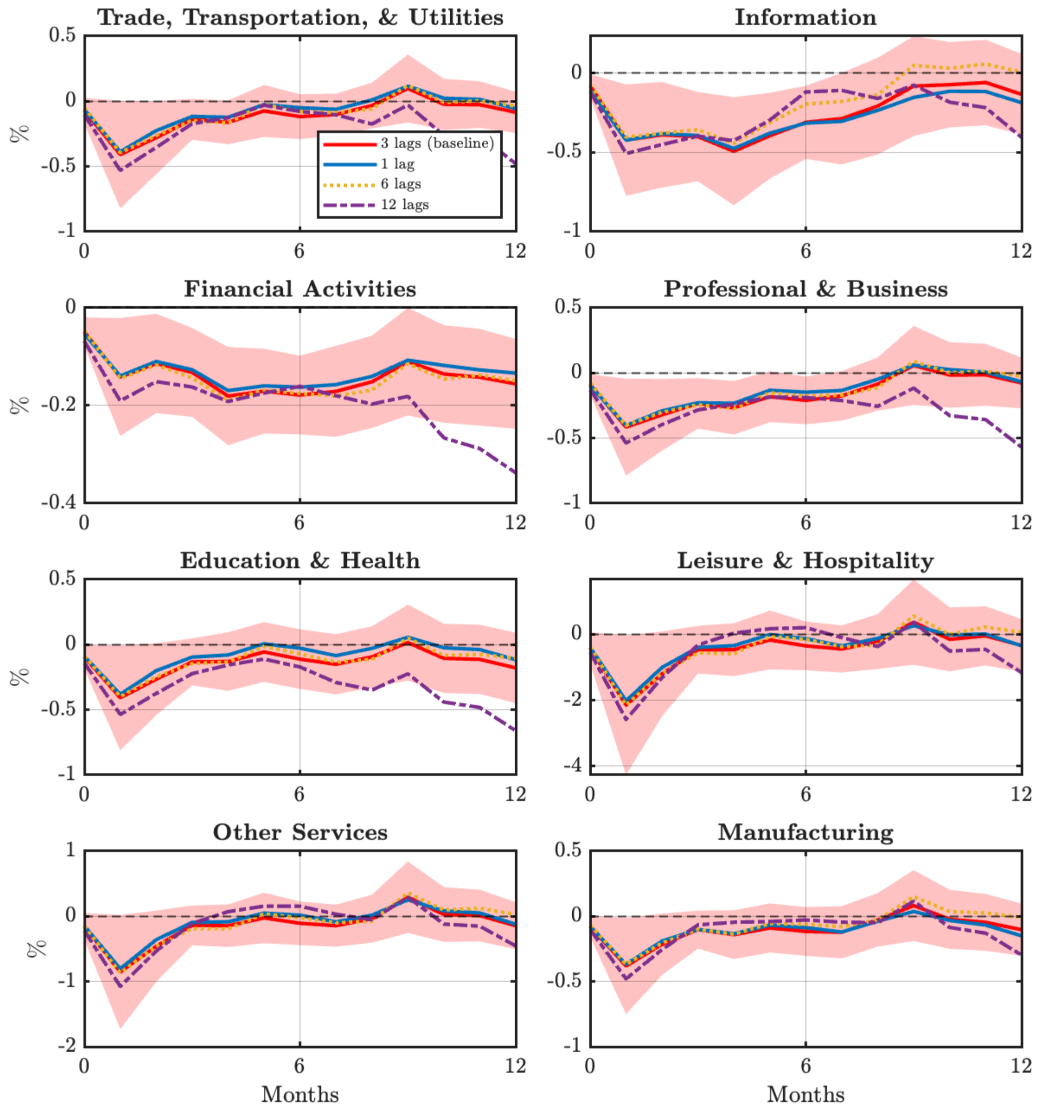

Note: Figure P.3 presents the industry-level employment impulse responses to a one standard deviation "pure" AI uncertainty shocks estimated via local projections with different lag specifications. The industries are classified according to NAICS. The shaded regions represent 68% confidence intervals, computed using Newey–West standard errors.