IFS

Institute for Fiscal Studies

Working paper

25/48

Anastasios Evgenidis
Apostolos Fasianos

# AI news shocks and the macroeconomy: evidence from UK patent data

Economic and Social Research Council

# AI News Shocks and the Macroeconomy: Evidence from UK Patent Data

Anastasios Evgenidis* and Apostolos Fasianos†

October 15, 2025

**Abstract**

We examine the macroeconomic effects of artificial intelligence (AI) innovations in the United Kingdom over the period 1982–2022. Using supervised machine learning to classify around 550,000 patent documents, we construct a novel quarterly measure of AI innovation and employ an instrumental-variable approach to identify exogenous AI news shocks. Embedding this measure in a Bayesian VAR, we show that positive AI news shocks generate persistent increases in output, investment, wages, and hours worked, along with a fall in the unemployment rate. Unlike the S-shaped responses of TFP typically associated with the gradual diffusion of traditional technologies, we find that TFP rises on impact, reflecting the rapid spread of software- and data-driven AI technologies, and remains elevated. The accompanying rise in wages and fall in unemployment suggests that the productivity-enhancing effects of AI outweigh its potential displacement effects.

**JEL Classification:** E32, C32, O33

**Keywords:** Artificial Intelligence, Technological Shocks, Patents, Bayesian VAR

Acknowledgements: This research received financial support from the British Academy/Leverhulme Small Research Grant (Grant No: SRG2223\231632) and Brunel University London (BRIEF Small Research Grant). We also thank seminar and conference participants at the 10th International Conference on Computational Social Science, the 100th Conference of the Western Economic Association, and the Joint Research Centre at the European Commission in Seville, for useful comments and suggestions. We are also grateful to Michalis Bletsas (MIT), who served as an external domain expert for the classification and validation of the machine-learning algorithm employed in this study.

*Royal Holloway University of London, IFS

†Brunel University of London, University of Macedonia

# 1 Introduction

Artificial Intelligence (AI) is advancing at a remarkable pace, with breakthroughs in machine learning (ML) and large language models (LLMs) fueling expectations of a new era of economic progress. These developments have sparked debate among economists and policymakers: will AI serve as the steam engine of our time, driving growth and productivity, or will it instead reinforce the longstanding Solow paradox, that new technologies are visible everywhere but in the productivity statistics (Brynjolfsson et al., 2017; Syverson, 2017)?

At the same time, concerns persist that AI could put large segments of the workforce at risk of displacement. Theoretical work highlights three forces through which adoption shapes productivity and employment (Acemoglu and Restrepo, 2018*a*; Acemoglu et al., 2022*a*; Acemoglu, 2025). The displacement effect captures automation of existing tasks and jobs. The productivity effect increases labour demand in both AI and non-AI sectors while raising capital productivity. Finally, the reinstatement effect reflects the emergence of new tasks and activities. The productivity and reinstatement effects that generate employment, acts as a countervailing force against the displacement effect, which leads to job losses, and which of these forces ultimately prevails remains an empirical question.[1] Empirical research on AI's impact remains at an early stage, constrained by the still-limited deployment of many AI technologies and by data scarcity. As a result, most studies have focused on partial equilibrium effects, typically using measures of job exposure to potential AI adoption.[2,3]

This paper provides a first attempt to assess AI's impact at the macroeconomic level

[1]Acemoglu and Restrepo (2018*a*) stress that it is primarily the reinstatement effect, through task creation, that offsets automation's labour-reducing impact and sustains employment growth.

[2]Brynjolfsson et al. (2025) find that generative-AI tools boost call-center productivity, particularly for lower-skilled workers. Acemoglu et al. (2022*b*), using establishment-level job vacancy data, document rapid growth in AI-related postings, accompanied by reductions in hiring for non-AI positions. Bonfiglioli et al. (2025) show that AI exposure reduces employment in U.S. commuting zones for low-skill and production workers, while benefiting high-wage and STEM occupations. Babina et al. (2024) measure firm-level AI investments using employee resumes and find that AI adoption drives growth in sales, employment, and market value, primarily through product innovation.

[3]Evidence on the macroeconomic impact of AI diffusion is also scarce. One relevant contribution is Aldasoro et al. (2024), who find that AI adoption increases long-run output in a U.S. multi-sector DSGE model.

by integrating a novel measure of AI innovations into the framework of technology news shocks and expectation-driven business cycles.[4] The rationale is that technology news shocks reflect newly disclosed information about future productivity growth that can drive economic fluctuations (Beaudry and Portier, 2006*a*, 2014). While the role of such shocks is well established in the technology news shocks literature (Jaimovich and Rebelo, 2009; Barsky and Sims, 2011; Baron and Schmidt, 2014), we lack evidence on how the economy responds to AI-specific shocks. Such shocks are particularly relevant, as many AI breakthroughs remain in early adoption but often diffuse faster than traditional IT technologies (Bick et al., 2024; Dahlke et al., 2024) and carry the potential to shape multiple aspects of the economy, including productivity, economic growth, inflation, investment, and the stock markets (Aldasoro et al., 2024; Filippucci et al., 2024). In our context, we define public information about new AI discoveries as *AI news*, which conveys information about expected progress in the field, and the unexpected element of this news as *AI news shocks*.

We aim to empirically answer the following question: how does the UK economy respond to AI news shocks? The UK provides a compelling test case, given its prolonged productivity slowdown alongside major efforts to position itself as an AI forerunner (Williams et al., 2025).[5] To this aim, we construct a novel measure of AI innovations, based on patent text data for the UK, over the past four decades. The intuition is that AI-related patent applications inherently convey information about potential future technological developments in the field. This empirical framework requires overcoming two major challenges: first, the measurement of AI-specific innovation; and second, potential endogeneity, since patenting activity may be influenced by current business cycle conditions.

To tackle the first challenge, we apply supervised machine learning techniques, specifically a Naïve Bayes classifier trained on a manually classified sample of patent texts. This

[4]This perspective also allows us to overcome practical limitations about low availability of data on deployment of AI.

[5]For example, the government has pledged over £1 billion through the *AI Sector Deal* and the *National AI Strategy*. More recently, it allocated £7 million from the UKRI Technology Missions Fund, via the Innovate UK BridgeAI programme, to support R&D in key AI areas.

allows us to identify AI-related patents and construct a quarterly time series that serves as our proxy for innovative activity in AI.[6] Our approach leverages two complementary sources of patent data: the *European Patent Office (EPO) Full Text Database* for the United Kingdom and the *EPO PATSTAT* database. By merging these two sources, we construct a unique dataset comprising around 550,000 patent documents filed in the UK since 1982, with very broad coverage across all major text fields and classification codes. A novel feature of our approach is that the EPO Full Text Database for the UK enables us to classify patents as AI or non-AI, providing access to claims and descriptions. This is important for capturing the technological content of each patent more comprehensively; descriptions provide a very detailed account of the invention and its technical novelty, while claims, written in precise legal language and with greater linguistic and syntactic complexity, more accurately reflect the underlying technology (Qiu and Wang, 2023; Bekkers et al., 2011; Tong and Frame, 1994; Berger et al., 2012). Taken together, descriptions and claims offer rich textual input for our analysis over and above the limited breadth of titles and abstracts, which are most commonly used in the literature. In line with recent best practices in textual analysis (Kumar and Sharma, 2022), both our manual classification and the model's predictions are assessed and confirmed by an external domain expert in AI.

The second challenge concerns the endogeneity of patent applications, specifically, the possibility that the decision to file a patent may be influenced by expectations about the UK macroeconomic outlook. To mitigate this concern, we construct an instrumental variable (IV) that isolates the component of patent filings orthogonal to prior macroeconomic expectations, as reflected in the Bank of England's (BoE) forecasts, and other potential confounding factors. This approach makes it more likely that our IV correlates exclusively with contemporaneous AI news shocks, which is essential for correct

[6]Note that not all AI innovations are patented. Some firms rely on trade secrets to protect intellectual property, particularly when disclosure through patents could undermine competitive advantage (Bhattacharya and Guriev, 2006). By definition, however, trade secrets are not publicly disclosed and thus fall outside the scope of our analysis. Since the concept of news shocks in macroeconomics is based on publicly available information, patent filings, especially in research-intensive fields such as AI, are well suited for capturing the relevant signals of inventive activity. For our purposes, this is sufficient, as the objective is to identify aggregate innovation shocks at the macroeconomic level rather than to provide a full account of all AI innovations.

identification. Similar IV strategies have been used in the technology news literature (Miranda-Agrippino, 2023), but our focus is on AI-specific patent activity.

Equipped with a plausibly exogenous instrument, we embed it into a rich-information Bayesian Vector Autoregression (BVAR) model to recover an *AI patent-based news shock* (AIPNS) series that captures news raising expectations about future progress in AI technology. The BVAR includes a comprehensive set of macroeconomic, labor market, and forward-looking variables allowing us to capture the dynamic effects of the shock. The selection of the variables is informed by the literature on the macroeconomic impact of technological news shocks (Baron and Schmidt, 2014; Bluwstein et al., 2020) and includes total factor productivity (TFP), GDP, inflation, the unemployment rate, investment, real wages, hours worked, and an aggregate stock market index. Our results are presented in the form of impulse response functions (IRFs).

Our first result highlights the distinctive nature of AI as a productivity-enhancing technology with rapid deployment potential. The response of TFP rises immediately following the AI news shock, peaking at around 0.15% before gradually converging to a new long-term equilibrium. This immediate jump contrasts with the delayed, S-shaped diffusion pattern typically documented for broad technological news shocks (Barsky et al., 2015; Kurmann and Sims, 2021; Miranda-Agrippino, 2023). The difference reflects the nature of AI patents, which are concentrated in software, programming, and data-processing innovations that can be deployed more quickly than non-AI, broader IT technologies (Bick et al., 2024). Consistent with this hypothesis and the broader technology news shocks literature, when we construct the same indicator for non-AI patents, the response of TFP is initially flat and only builds gradually over several quarters. This contrast indicates that AI represents not merely another wave of technological progress, but one characterized by rapid transmission to productivity.[7]

Second, in addition to TFP, AI news shocks trigger a broad-based and persistent macroeconomic expansion, as evidenced by increases in output, investment, and consumption. This expansion operates primarily through the labor market, which emerges

[7] In this context, rapid refers to effects observed within three months following the shock, given that our data are at quarterly frequency.

as the central channel of aggregate transmission. Specifically, positive AI news shocks are employment-enhancing, as they reduce the unemployment rate while increasing wages and hours worked. These findings speak to theoretical debate on how AI affects labor markets and the broader economy (Acemoglu and Restrepo, 2018*a*; Acemoglu et al., 2022*a*; Acemoglu, 2025). In particular, our results point, albeit indirectly, to the importance of the reinstatement channel over and above the productivity channel. If the creation of new tasks were not significant, the unemployment response would have been flat (or positive), implying that the displacement channel dominates or merely counterbalances productivity gains. Instead, the observed strong decline in the unemployment rate, coupled with rising labor input as reflected in higher wages and hours worked, is consistent with the role of the reinstatement channel. This interpretation aligns with Acemoglu and Restrepo (2018*a*) who argue that, while productivity gains from AI can increase labor demand, these effects alone are generally insufficient to offset displacement. It is the emergence of new tasks, reshaping labor composition and creating additional employment opportunities, that directly counterbalances displacement, generating net employment gains.

Third, variance decompositions indicate that AI news shocks account for a non-negligible share of macroeconomic and labor market fluctuations, highlighting their role in shaping business-cycle dynamics even before widespread adoption. Our main results remain robust when using a complementary identification strategy that does not rely on an IV. In this approach, the AI news shock is identified as the innovation that best explains future unpredictable movements in our patent-based AI innovation measure, which serves as an alternative proxy for expected technological progress in AI.

We ensure that our results are not confounded by other policy or technological forces through a series of robustness checks. Specifically, we augment the first-stage regression with alternative controls, including fiscal and monetary policy shocks, broader trends in non-AI technological innovation, and public R&D spending factors that may influence firms' patenting behavior. Across all these specifications, our main findings remain robust. Another potential issue is whether our IV could be affected by weighting patents differently, since raw patent patent counts may misrepresent the scientific significance of AI innovations (Kelly et al., 2021; Verendel, 2023). To address this, we weight patents by their forward citations using PATSTAT, which are widely used as a proxy for the scientific significance of an invention (Criscuolo and Verspagen, 2008).

In addition to the aggregate analysis, we examine sectoral heterogeneity in the propagation of AI news shocks. Using concordance tables from PATSTAT and Dorner and Harhoff (2018), we allocate patents to their industries of origin and construct AI innovation measures across sub-sectors of manufacturing and services. We find that AI news shocks propagate unevenly across sectors. In services, AI news shocks account for a slightly larger share of fluctuations in aggregate TFP (around 8%) compared with manufacturing, reflecting the intangible, software-driven nature of patents in this sector which can be more readily integrated into existing processes, amplifying their productivity-enhancing effects. In contrast, AI news shocks in manufacturing, especially in computers, communication equipment, and precision instruments, play a more prominent role in driving GDP, investment, and wages, with some sub-industries accounting for up to 22 percent of output variation over a 24-quarter horizon.

In our baseline analysis, we estimate IRFs using a VAR approach. Recent studies, such as Jordà et al. (2024), show that lag truncation in VAR models can bias medium- and long-term IRFs, whereas local projections (LPs) are less susceptible to this issue. In our setting, the risk of lag truncation bias is limited, since we include a relatively large number of lags in the VAR. Nevertheless, we also use LPs to estimate IRFs to AI news shocks, providing a complementary perspective. While LP estimates exhibit wider confidence bands, their qualitative patterns and magnitudes are largely consistent with the VAR results.

Our paper contributes to several strands of the literature. We construct a quarterly measure of AI patents in the UK using machine learning techniques that can be readily integrated into time series models. This aligns the UK evidence with recent US studies that rely on textual analysis to classify AI-related innovations (Giczy et al., 2022). Prior work on Europe and the UK has mainly used International Patent Classification (IPC) codes or keyword searches (Buarque et al., 2020; Xiao et al., 2023) approaches that cannot fully capture the fast-evolving and heterogeneous terminology of AI (Baruffaldi et al., 2020; Miric et al., 2023). By contrast, our method applies a broader and more flexible definition of AI, identifying technologies such as neural networks, machine learning, and speech recognition that are often missed by query-based approaches. Importantly, our dataset incorporates both patent descriptions and claims, providing a richer training sample and enabling a more precise classification of AI versus non-AI patents.

Our paper also speaks to the literature that uses patent data to account for expectations of future technological developments. In contrast to our approach, these studies typically employ small-scale VARs, often including no more than two or three variables, with standard Cholesky identification. Such specifications pair a patent-related variable with a measure of productivity, allowing authors to isolate the effects of technology shocks (Shea, 1999; Christiansen, 2008).[8]

Cascaldi-Garcia and Vukotić (2022) and Miranda-Agrippino (2023) are more recent and closely related contributions that exploit richer data in medium-scale Bayesian VARs at a quarterly frequency. The authors construct patent-based technology news shocks and trace their effects on output, prices, labor, and financial markets.[9] We build on this strand but depart in several important ways. First, our analysis provides the first evidence on the macroeconomic impact of technological news shocks for the UK using a rich quarterly dataset that complements prior studies focused only on the US.[10] Second, rather than focusing on general technology patents, we use supervised machine learning to isolate AI-specific patents from the broader technological pool. This enables us to

[8] Shea (1999) use annual patent applications and R&D expenditures to assess the impact of technology shocks on industry aggregates, identifying shocks by placing each measure last in the VAR. Similarly, Christiansen (2008) employs patent application data in a bivariate VAR, ordering the patent variable first. Alexopoulos (2011) takes a different approach, using the number of technology-related book publications to capture the timing of commercialization, with responses estimated in bivariate VARs where the publication index is placed last.

[9] The main differences between these studies are as follows. First, Cascaldi-Garcia and Vukotić (2022) identify technology news shocks using a measure proposed by Kogan et al. (2017), which combines micro-level patent data with stock market valuations, whereas Miranda-Agrippino (2023) construct an IV for technology news shocks as the component of patent applications orthogonal to expectations based on prior patent activity. Second, the former study relies on granted patents from the USPTO, while the latter uses all patent applications, capturing news content earlier in the innovation process.

[10] This contribution owes to the construction of our quarterly AI patent-based measure, whereas existing UK patent data are aggregated annually and cannot be used for time-series analysis

estimate the productivity and macroeconomic effects of a rapidly diffusing technology that is expected to bring unprecedented shifts in the economy compared to traditional innovations typically aggregated into broad technology news shocks (Frey and Osborne, 2017; Bick et al., 2024; Acemoglu and Restrepo, 2018*b*). Third, we adjust the IV approach of Miranda-Agrippino (2023) to the context of AI innovations, constructing an instrument that purges endogenous fluctuations in AI patenting activity driven by expectations regarding macroeconomic conditions, monetary and fiscal policies, broader trends in non-AI technological innovations, and public R&D spending.

Our work also connects to the broader literature on identifying technology news shocks in VARs that abstract from patent-based measures. Studies such as Uhlig (2005), Beaudry and Portier (2006*b*, 2014), Beaudry and Lucke (2010), and Forni et al. (2014) typically identify news shocks as innovations orthogonal to contemporaneous TFP movements, supplemented by additional long-run assumptions about the determinants of productivity. Other contributions depart from this strategy and instead define news shocks as those that maximize the forecast error variance of productivity at one or multiple horizons (Neville Francis and DiCecio, 2014; Barsky and Sims, 2011; Kurmann and Sims, 2021).

The remainder of the paper is structured as follows. In Section 2, we detail the construction of our patent-based AI innovation measure. We also describe our identification strategy for AI news shocks. Section 3 presents our main results, derived from a Bayesian VAR framework, and analyzes the dynamic responses of TFP, macroeconomic, and labour market outcomes to AI news shocks. In Section 4, we validate our findings through a series of robustness checks and alternative IV proxies. Section 5 provides industry-level evidence of the impact of AI news shocks. Finally, Section 6 concludes the paper.

# 2 Data and methodology

This section describes the patent data and sets out the methodological steps used to construct the series of AI-related patents and, subsequently, the AI patent-based instrument.

## 2.1 AI Innovations in the UK

We detect AI innovations via patent texts, which combine detailed technical descriptions, legal claims, and structured metadata, enabling consistent tracking of technological developments across sectors, jurisdictions, and over time. For this reason, patent texts are among the most widely used proxies for technological innovation (e.g., Acs and Varga, 1994; Kelly et al., 2021). Although not all innovations are patented, patent filings, especially in research-intensive fields such as AI, capture early signals of inventive activity and provide a rich foundation for long-run analysis (Buesa et al., 2010; Cockburn et al., 2018; Xiao et al., 2023).

To our knowledge, no existing study systematically identifies AI innovations in UK patent documents using machine-learning methods. The UK Intellectual Property Office (UKIPO) has published reports on AI patents (e.g., UKIPO 2019, 2021), but these are annual and cover only a limited time period, which makes them unsuitable for time-series analysis.[11] Importantly, in these reports, AI patent identification is based on query methods, which carry important caveats.[12]

### 2.1.1 Patent Data

We merge two datasets containing all patent filings registered in the UK since 1982. The first is the PATSTAT database maintained by the EPO, which includes over 100 million patents registered in about 100 patent offices worldwide, going back to the nineteenth century.[13] PATSTAT provides granular information on each patent, including bibliographic details such as the title, abstracts, inventor addresses, dates, citations, and International Patent Classification (IPC) codes.[14] This dataset has been widely used in research on

[11]See UKIPO (2019) *Artificial Intelligence: A worldwide overview of AI patents and patenting by the UK AI sector*, and UKIPO (2021) *AI patents, the UK landscape and emerging trends*, available at: https://www.gov.uk/government/publications/artificial-intelligence-a-worldwide-overview-of-ai-patents and https://www.gov.uk/government/publications/ai-patents-the-uk-landscape-and-emerging-trends.

[12]The literature highlights that query methods face significant limitations in keeping pace with the rapidly evolving and heterogeneous terminology associated with AI technologies (Baruffaldi et al., 2020; Miric et al., 2023).

[13]For more information on the dataset, see https://www.epo.org/searching-for-patents/business/patstat.

[14]IPC codes follow a hierarchical classification system with six digits, each providing progressively finer details about the technology described in a patent. In the IPC system, the first letter denotes

the economics of innovation (e.g., Bloom et al., 2020; Aghion et al., 2023; Straccamore et al., 2023; Gross and Sampat, 2023).

Despite its widespread use, PATSTAT lacks several key components essential for tracking technological developments, most notably, the full-text descriptions and claims. Claims represent the legally binding part of a patent, as they define the precise scope of the intellectual property rights granted (Abood and Feltenberger, 2018; Baruffaldi et al., 2020). Additionally, abstracts in PATSTAT are often missing for older filings, particularly those from the 1980s and early 1990s. To address these limitations, we complement PATSTAT with the *EPO Full Text Database for the UK*, which provides rich textual information, including both the description and claims sections. Merging these sources enables more accurate classification of AI-related innovations, particularly for earlier periods of our sample. This approach aligns with recent advances in patent classification, which increasingly rely on claims and descriptions as primary sources of both legal and technical information respecitively (Kelly et al., 2021).

A methodological choice concerns whether patents should be classified by the inventor's location or by the filing country. The inventor's location indicates where the knowledge was originally produced; for instance, if an AI researcher based in India files a patent in the UK, the inventor location is India. By contrast, the filing country reflects the intention of the applicant to protect, commercialize, and deploy the innovation in the UK. This is known in the literature as home bias in patent filing (Jamasb and Pollitt, 2011). Accordingly, we define a patent as UK-labeled based on the country indicated in the patent registry at the time of filing. Our final UK sample comprises 542,977 patents registered with the UK patent office between 1982 and 2022.

Another methodological choice concerns whether to rely on patent filings or granted patents. In our analysis, we focus on filed patents, as they provide an earlier signal of technological progress, while granted patents are often delayed for several years due to administrative procedures and patent examiner workloads (Christiansen, 2008). From an

the broad technology section. For example, `Fxxxxx` corresponds to *Mechanical Engineering; Lighting; Heating; Weapons; Blasting*. Subsequent digits refine the classification: for instance, `F02xxx` designates *Combustion Engines* within section F.

identification standpoint, the filing date is also preferable because it captures the moment when information about an invention becomes publicly available. In this sense, the filing date represents the first public disclosure of a patent's content, making it a natural proxy for news about future technological developments (Miranda-Agrippino, 2023).

### 2.1.2 AI-related patents definition and classification

Existing approaches for identifying AI patents typically rely on keyword queries and IPC codes to match patent texts with technology classes.[15] While transparent and straightforward, these methods face significant limitations. Keyword searches can be ambiguous, as words can carry multiple meanings, synonyms, or spelling variations, and they require lengthy boolean queries. IPC codes pose similar challenges: they may not align with emerging AI technologies and patents frequently deviate from class definitions. Both approaches also depend on pre-defined dictionaries linking keywords or codes to technologies, which are often outdated, incomplete, or entirely absent, and whose construction involves a large degree of subjectivity (Baruffaldi et al., 2020; Miric et al., 2023).

We mitigate many of the limitations noted above by applying a supervised ML algorithm for the classification process (Miric et al., 2023; Choudhury et al., 2019). To this end, we manually classify a subset of our dataset, which serves as the training sample for the algorithm.[16] A definition of AI is foundational for our classification. We would ideally rely on a single definition for AI that remains largely consistent throughout our classification period. In practice, however, no universal definition exists, and the concept has evolved over the decades. (Winfield, 2020).[17] The development of AI can be broadly divided into three waves (UNCTAD, 2025). The first wave (1950s–1980s) was dominated

[15]See, for example, Fujii and Managi (2018) and Buarque et al. (2020), among others.

[16]Given the large size of our dataset, manual classification of all patents would be prohibitively costly.

[17]AI has been defined in many different ways depending on context. For example, the EU AI Act (2024) offers a functional and legal description, framing AI as "a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments". By contrast, McCarthy (1997) provides a broader definition, describing it as "the science and engineering of making intelligent machines." Studies that classify patent texts also diverge in scope: some, such as Fujii and Managi (2018) and European Patent Office (2017), adopt a relatively narrow approach centered on computational models, while others, like OECD (2022) and Cockburn et al. (2018), take a wider view that includes robotics and AI-enabled devices.

by rule-based approaches, where systems relied on explicitly programmed logic and expert knowledge. The second wave (1990s–2010s) saw the rise of statistical methods, big data, machine learning, and eventually deep learning, enabling more flexible and data-driven models. The current third wave (late 2010s–present) is defined by contextual adaptation, where generative AI systems such as ChatGPT exhibit the capacity to interpret, generate, and interact with human-like adaptability across a wide range of contexts.[18]

Since the meaning and scope of AI have shifted over time, a narrow definition would be too restrictive. We therefore employ a broad, component-based framework for our manual classification, drawing on the approach of Giczy et al. (2022). Our framework identifies eight core technological components that precisely characterize AI innovations: ML, AI hardware, natural language processing, speech recognition, computer vision, knowledge processing, planning/control, and evolutionary computation (see Appendix I for a detailed description of the AI classes). A patent is considered AI-related if its central inventive contribution, as evidenced in the description, claims, and abstract, maps to at least one of these components. We read the title and abstract, examined the claim(s), and consulted the description when necessary. When multiple AI components applied, we recorded all. Borderline cases lacking explicit evidence of inference were excluded from the manually labeled set.

To build a representative training sample, we ensure coverage across time by sampling patents proportionally by year and quarter. We also balance the sample between AI and non-AI patents, yielding 1,165 patents in total (550 AI and 615 non-AI).[19] Figure A1 in Appendix B presents two representative examples of both AI and non-AI related patents, illustrating our manual classification process. Finally, to assess our manual classification

[18]Specifically, the first wave emphasized rule-based systems and symbolic AI, including early work on optimal control, the development of expert systems, and initial attempts at natural language processing such as ELIZA. This period also saw early neural network research, such as perceptrons, which faced theoretical limitations that contributed to the first AI winter (Minsky and Papert, 1969). The second was shaped by statistical learning, as researchers increasingly adopted machine learning methods, including support vector machines and early deep learning models. Although often distanced from the term "AI," these methods laid the foundation for the resurgence of neural networks in the 2010s. The third wave is characterized by contextual adaptation and generative AI, driven by advances in deep neural architectures, especially transformer-based models (Vaswani et al., 2017; Brown et al., 2020) allowing large-scale content generation in language, images and beyond.

[19]Balancing classes in text classification helps prevent bias toward the majority class and improves model generalization; see, for example, Chawla (2005).

process, an external domain expert validated a random subset of 100 patent texts (50 AI and 50 non-AI) from our training sample, confirming our classification with 98% accuracy.[20]

### 2.1.3 Tracking AI innovations using textual analysis

To classify all 542,977 UK patents between AI and non-AI related, we apply the Naïve Bayes Classifier (NBC). The NBC is a supervised ML algorithm popular for its computational efficiency and interpretability, particularly for large-scale datasets like ours. It assumes that the probability of a given word appearing in a document is independent of the presence of any other word, a simplifying assumption that enables efficient computation (Domingos and Pazzani, 1997). Its small number of parameters reduces the risk of overfitting, making it well-suited for our analysis (Murphy, 2012).[21]

We pre-process the patent texts following standard cleaning steps used in textual analysis. This involves concatenating the title and abstract, converting all letters to lowercase, removing numbers, special characters, punctuation, and stopwords, and deleting extra spaces.[22] After cleaning, we apply the Porter stemming algorithm to extract 5,462 word stems, called tokens. Stemming groups related words under the same root, for instance, "artificial," "artificially," and "artificiality" are reduced to "artific-." On average, a patent title contains about 7 tokens, an abstract about 52 tokens, and the main body (description and claims) about 550 tokens.

Next, we rank the tokens (features) using the term frequency–inverse document frequency (TF-IDF) metric, which assigns higher values to words that occur frequently within a given document but are relatively scarce across the entire corpus.[23] We then apply the mutual information criterion, retaining tokens that occur more frequently in one

[20] A letter from the external expert is enclosed with this submission; it details the validation protocol and clarifies each step of the manual classification workflow.

[21] While LLMs can admittedly achieve higher accuracy, they are substantially slower and require greater computational resources, making them impractical for large-scale, routine classification (Cunha et al., 2025).

[22] Stopwords, such as "the," "of," "to," "are," and "is", are common words that occur frequently in a language but contribute little to the text's meaning.

[23] Highly common words, e.g., "apparatus", usually lack discriminatory power, while extremely rare terms may add noise or capture idiosyncratic phrasing. Filtering by TF-IDF helps retain tokens that are more informative for distinguishing AI from non-AI patents.

class and down-weighting those with low overall frequency (Manning et al., 2009). From this process, we retain the 100 highest-ranked words along with the 5 most frequent IPC codes, producing a final search dictionary of 105 tokens used to train the NBC. Figure 1 visualizes the 105 tokens selected by the mutual information criterion. Note that the size of each token is proportional to its mutual information score. The most frequent token appears to be "image", which aligns with computer vision, the category that appears most frequently in the training sample.[24] They are followed by "data", "neural", and "network" which are general terms that could appear in all eight components of AI technology.

Figure 1: Frequently used tokens indicating an AI-related patent.

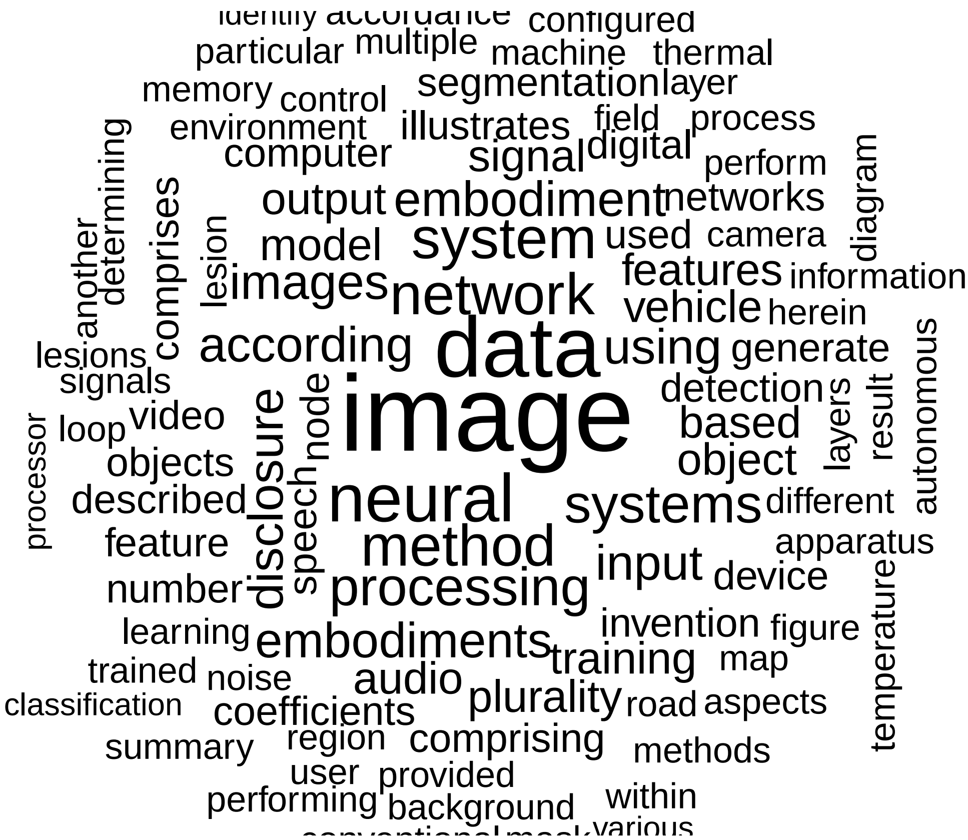


We evaluate the NBC using a 75/25 train–test split of our manually classified sample. The training set contains 874 patents, while the test set contains 291 patents.[25] Table 1 reports the confusion matrix for the test set and the corresponding performance metrics. Out of the 150 AI-related patents in the test set, 127 are correctly classified,

[24]Among the AI-classified patents, approximately 50% pertain to computer vision, 35% to speech recognition, and 16% to natural language processing.

[25]Results are robust to alternative splits such as 60/40 and 80/20; these are available upon request.

corresponding to a recall of 0.847, while 23 (15.3%) are misclassified as non-AI. Among the 141 non-AI patents, 108 are correctly classified, yielding a recall of 0.766, while 33 (23.4%) are misclassified as AI. The overall accuracy across the test set is 0.806, reflecting the percentage of total correct classifications (both AI and non-AI) out of all predictions made. The precision for the AI class is 0.794, meaning that when the model predicts a patent as AI, it is correct approximately 80% of the time. The F1-score for the AI class, which balances precision and recall, is 0.819. To strengthen the credibility of our classification, we perform an ex-post evaluation by cross-checking the model's predictions with an external domain expert in AI.[26]

Table 1: Confusion matrix and performance metrics for recognition of AI patents (75/25 split).

| Confusion Matrix (counts) | | | Performance Metrics | |
|---|---|---|---|---|
| Actual / Predicted | AI | Non-AI | Metric | Value |
| AI | 127 | 23 | Accuracy | 0.806 |
| Non-AI | 33 | 108 | Precision (AI) | 0.794 |
| | | | Precision (Non-AI) | 0.824 |
| | | | Recall (AI) | 0.847 |
| | | | Recall (Non-AI) | 0.766 |
| | | | F1-score (AI) | 0.819 |
| | | | F1-score (Non-AI) | 0.794 |

*Note*: In the confusion matrix, rows correspond to actual classes and columns to predicted classes. Metrics are computed from the test set of 291 patents.

Subsequently, we apply the trained NBC to classify the remaining unlabeled patents in our dataset. This process identifies approximately 190,000 AI-related patents spanning from 1982 to 2022. Each patent is assigned to the quarter of its earliest filing date, producing a quarterly series of AI patenting activity that forms the basis for the construction of our instrument, which is discussed in detail in the following section. Figure 2 presents our quarterly series of AI-related patents filed in the UK between 1982 and 2022, shown both as a share (right axis) and in total counts (left axis). The share of AI patents remains relatively low and stable through the 1980s and 1990s. In the early 2000s, a

[26]Involving human domain experts in the validation of ML outputs has become increasingly common in applied ML research, particularly for algorithms targeting specialized predictions. The combination of algorithmic forecasts with expert judgment enhances both interpretability and trustworthiness of the findings; see Kumar and Sharma (2022) for further discussion.

gradual uptick is visible, with AI patents accounting for slightly over 10% of all filings by 2003. This coincides with the second wave of AI, marked by developments in statistical learning and supported by rising computational power, expanding data availability, and more sophisticated algorithms. From 2010 onwards, the upward trajectory accelerates sharply, with the share of AI patents approaching 30% by 2022. This post-2010 surge reflects the third wave of contextual adaptation, marked by deep learning and, more recently, transformer-based architectures and LLMs. On the volume side, the absolute number of AI-related patents follows a similar pattern, with pronounced growth after 2017.

Figure 2: AI-related patents in the UK (1982–2022).

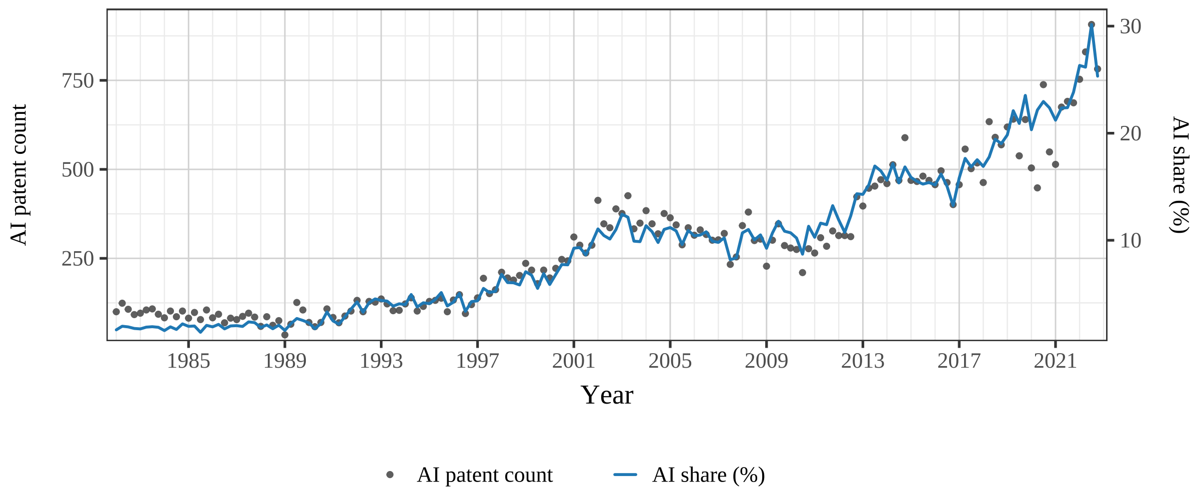


*Note:* Authors' own estimations. The dotted series indicate the predicted number of AI-related patents (levels), while the line shows their share of all patents registered in the UK patent office.

## 2.2 AI Patent-based instrument

Armed with our forty-year AI patent-based innovation measure, our next objective is to construct an instrument for the identification of the AI news shocks. A key challenge concerns the endogeneity of patent applications, specifically, the possibility that the decision to file a patent may be influenced by agents' expectations about the UK macroeconomy. For example, if firms anticipate an upcoming expansion in aggregate demand or expect government subsidies for digital technologies, they may accelerate the filing of AI-related patents in anticipation of higher future returns. To mitigate this concern, we construct an IV that isolates the component of AI patent applications independent of expectations about the UK macroeconomic outlook, as well as lags of AI patent applications. The proposed approach increases the likelihood that the instrument is solely correlated with current AI news shocks, which is essential for accurate identification. In Section 4, we extend this specification by incorporating additional controls to ensure that our identification strategy is robust.

We construct our instrument by taking the residuals of the following regression:

$$AIPatents_t = a + \zeta(L)AIPatents_t + \sum_{h=1,4,8} \theta_h E_t\left[x_{t+h}\right] + z_t \tag{1}$$

Equation 1 is estimated at quarterly frequency, where $AIPatents_t$ represents the annual growth rate of all AI patent applications, $\zeta(L) = \sum_{j=1}^{4} \zeta_j L^j$, $L$ is the lag operator, and $E_t\left[x_{t+h}\right]$ represents a vector of forecasts from the BoE for the two core macroeconomic variables included in $x_t$, namely GDP growth and CPI inflation.[27] We consider forecasts with horizons of one, four, and eight quarters ahead. The time index in $E_t$ refers to the publication date of the forecasts. Because the information set underlying the Monetary Policy Committee (MPC) forecasts corresponds to the previous quarter, the collection of forecasts in $E_t[x_{t+h}]$ reflects pre-existing perceptions of the UK macroeconomic outlook.[28] Specification 1 is estimated using our AI patent-based measure from 1997Q3 to 2022Q4 due to the unavailability of GDP growth forecasts before 1997Q3.

The regression results are presented in column (1) of Table A1 in the Appendix, labeled *Baseline*. Note that columns (2) to (5) incorporate additional confounding factors into the baseline specification, providing alternative instruments to ensure the robustness of the baseline specification.[29]

[27]Projections of GDP growth and CPI inflation are published in the Monetary Policy Report by the BoE. The report contains projections for GDP growth since 1997Q3. On 10 December 2003, the Chancellor of the Exchequer announced that he changed the inflation target from 2.5% for annual Retail Price Index (RPIX) inflation to annual inflation of 2.0% as measured by the CPI. The BoE provides inflation forecasts for RPIX inflation up to 2003Q4, and for CPI inflation from 2004Q1 onward. For more information, see: Bank of England Monetary Policy Report - February 2024.

[28]Although the forecasts are published at the beginning of the second month of each quarter, the information on which the MPC bases its projections, drawn from the quarterly release of GDP expenditure components and income by the ONS, becomes available six weeks prior to publication. As a result, the forecasts effectively rely on data from the previous quarter. For more information see here.

[29]These are described in detail in Section 4.

In the *Baseline* results, AI patent applications display a strong degree of autocorrelation, as indicated by the statistically significant coefficients on three of the four lagged terms. Consequently, incorporating its own lags in equation 1 mitigates the dependency of the IV on past AI patenting activity. Importantly, we find significant coefficients for CPI at $h = 8$ and for GDP growth at $h = 4$, respectively. This suggests that expectations regarding the UK's macroeconomic outlook contain information relevant for AI patent applications beyond their own lags, thereby influencing the decision to file an AI patent.[30] Figure 3 presents our constructed IV, i.e. the residuals from the baseline specification, and Table A2 in Appendix D reports the Wald test from a regression of the IV on the BoE's macroeconomic forecasts. The results indicate that the instrument is not Granger-caused by macroeconomic expectations, supporting the validity of our IV.

A potential issue is that AI patent applications may correlate with factors beyond the expected macroeconomic outlook. Specifically, monetary and fiscal policy changes may directly influence AI patent applications by affecting macroeconomic aggregates. To address this, Section 4 presents an alternative specification of Equation 1 that includes controls for monetary and fiscal policy disturbances. We also account for broader non-AI technologies and public innovation policies to ensure the instrument is not confounded by contemporaneous technological or policy developments.

[30]The strong joint significance of the control variables as evidenced by the Wald tests at the bottom of Table A1 confirms the relevance of our instrument for identifying exogenous AI technology news shocks.

Figure 3: AI patent-based instrument.

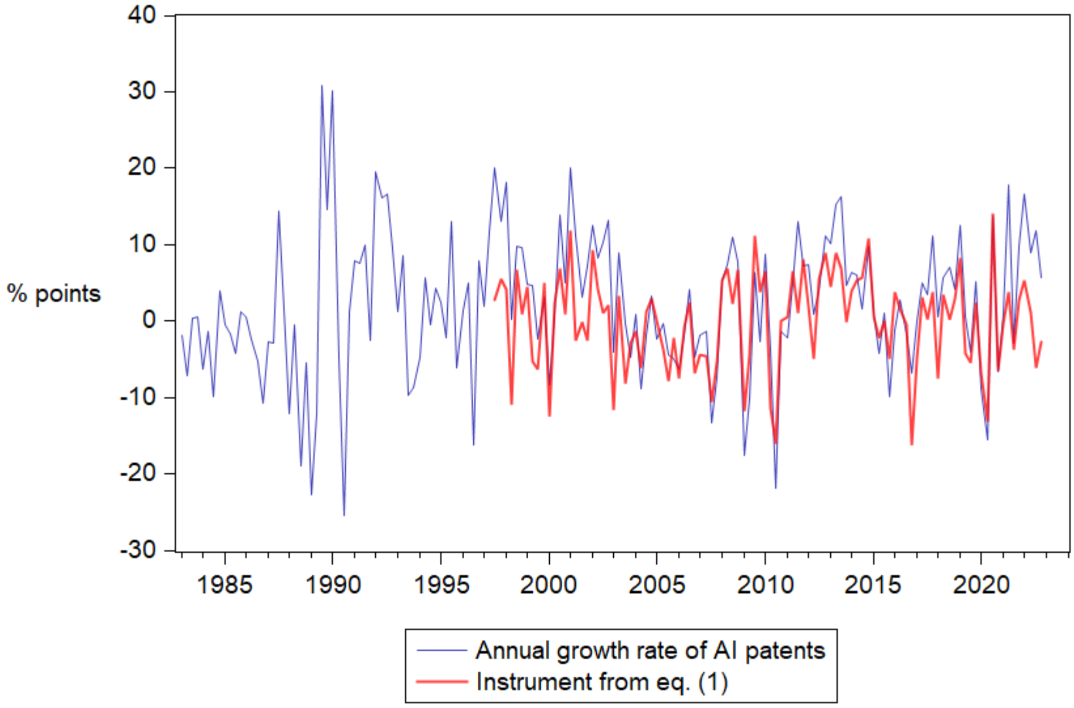


*Note:* This figure shows the AI patent-based instrument, i.e. residuals from eq. (1), plotted as a red line, alongside the growth rate of AI patent applications, shown as a blue line.

## 2.3 Proxy SVAR and identification

We next employ our AI patent-based IV to infer the dynamic effects of news about future technological progress in AI on several macroeconomic variables and labour market outcomes. To achieve this, we use a proxy Structural VAR (proxy-SVAR) (Stock and Watson, 2012, 2018; Mertens and Ravn, 2013). Consider the following reduced-form VAR of lag order $p$:

$$y_t = c_0 + \sum_{j=1}^{p} B_j y_{t-j} + u_t \tag{2}$$

where $u_t \sim N(0, \Sigma_u)$ represents a vector of reduced-form shocks, $y_t$ is an $n$-dimensional vector of endogenous economic variables of interest, $c_0$ is a constant vector, and $B_j$ contains the autoregressive reduced-form coefficients. To estimate IRFs and forecast error variance decompositions (FEVDs), we require that the information in our VAR is sufficient to recover all the structural shocks. In particular, we need to specify a matrix $A_0^{-1}$ such that pre-multiplying Equation 2 yields:

$$A_0^{-1} y_t = A_0^{-1} c_0 + A_0^{-1} \sum_{j=1}^{p} B_j y_{t-j} + A_0^{-1} u_t \tag{3}$$

The last term is the vector of $n$ structural disturbances, $\varepsilon_t = A_0^{-1} u_t$ with $\varepsilon_t \sim N(0, \Sigma_\varepsilon)$. Note that:

$$u_t = A_0 \varepsilon_t, \tag{4}$$

where $A_0$ collects the contemporaneous effects of $\varepsilon_t$ on $y_t$. If one considers a suitable identification scheme, $u_t = A_0 \varepsilon_t$ guarantees that the structural disturbances can be recovered from the observables in the VAR. In our case, consider that the only shock that is economically significant is an AI patents-based news shock (AIPNS), $\varepsilon_{AI,t}$. If this shock is the first shock of $\varepsilon_t$, it means that obtaining the first column of $A_0$ suffices to identify $\varepsilon_{AI,t}$. The identification of this column is where our instrument $z_t$ can be employed.

Conditional on equation 4, the conditions for identification in our proxy SVAR are:

$$\mathbb{E}[z_t \varepsilon_{AI,t}] = \phi, \quad \phi \neq 0 \quad \text{(relevance)} \tag{5}$$

$$\mathbb{E}[z_t \varepsilon_{i,t}] = 0, \quad i \neq AI \quad \text{(exogeneity)} \tag{6}$$

Condition (5) states that the instrument $z_t$ is correlated with the structural news shock $\varepsilon_{AI,t}$ via the following relationship: $z_t = \beta \varepsilon_{AI,t} + \sigma v_t, \quad v_t \sim N(0,1)$, where $\beta$ and $\sigma$ are parameters used to construct a measure of reliability proposed by Mertens and Ravn (2013), which captures the correlation between the instrument and the structural shock of interest. Condition (6) states that $z_t$ is not correlated with the remaining structural shocks and establishes exogeneity of the instrument.[31] Under these conditions, the impact responses of all VAR variables to $\varepsilon_{AI,t}$ are estimated from the projection of the VAR innovations $\hat{u}_t$ on $z_t$. We estimate our model using Bayesian techniques with four lags.[32]

[31] Note that these two conditions are very similar to those required of an instrument in a standard 2SLS setting.

[32] The use of four lags is standard in quarterly VARs and also in related studies on technology news shocks (Miranda-Agrippino, 2023; Cascaldi-Garcia and Vukotić, 2022).

We employ a normal-inverse Wishart prior following Caldara and Herbst (2019) and Negro and Schorfheide (2011), with 50,000 posterior draws, using the last 10,000 replications for inference.

We study the propagation of AI news shocks to the labor markets and the macroeconomy using a nine-variable VAR. The variables included are real GDP per capita, unemployment rate, TFP, total hours worked, real wages, GDP deflator, real consumption per capita, real investment per capita, and the FTSE All-Share Index.[33,34,35] The variables follow those commonly employed in VARs on the news shocks literature (Barsky and Sims, 2011; Beaudry and Portier, 2014; Miranda-Agrippino, 2023; Cascaldi-Garcia and Vukotić, 2022), with data drawn from the Office for National Statistics (ONS) and the Federal Reserve Economic Data (FRED). All series, except the unemployment rate, enter the VAR in log levels. The model is estimated over 1971Q1–2022Q4, while the IV identification sample, used to project the VAR residuals on the instrument, covers the period 1997Q3–2022Q4.

# 3 Results

## 3.1 Baseline estimates

The IRFs to a one standard deviation AI patent-based news shock, are depicted in Figure 4. These IRFs represent the mode of the posterior distribution of the parameters, with solid blue lines indicating median estimates. Shaded areas correspond to 68% posterior bands. We consider a horizon of 24 quarters (6 years).

The first result is that TFP responds upon impact, reaching up to a 0.15% increase, followed by a gradual decline towards its new long-term equilibrium level, which remains

[33] Real variables are deflated using the GDP deflator and transformed in per capita terms by dividing by the trend in population.

[34] Hours worked are measured in weekly hours, private final consumption expenditure serves as the measure of consumption, average weekly earnings per person are used as the measure of wages, and gross fixed capital formation is used to proxy investment.

[35] The TFP series is sourced from the ONS and is partially utilization-adjusted, accounting only for lower capital utilization during the COVID-19 pandemic. In the robustness section, we consider an alternative specification using fully utilization-adjusted TFP which is, however, limited to the period 1998:Q1–2020:Q1. The baseline results remain largely unchanged.

Figure 4: Responses to an AI patent-based news shock, used as an instrument.

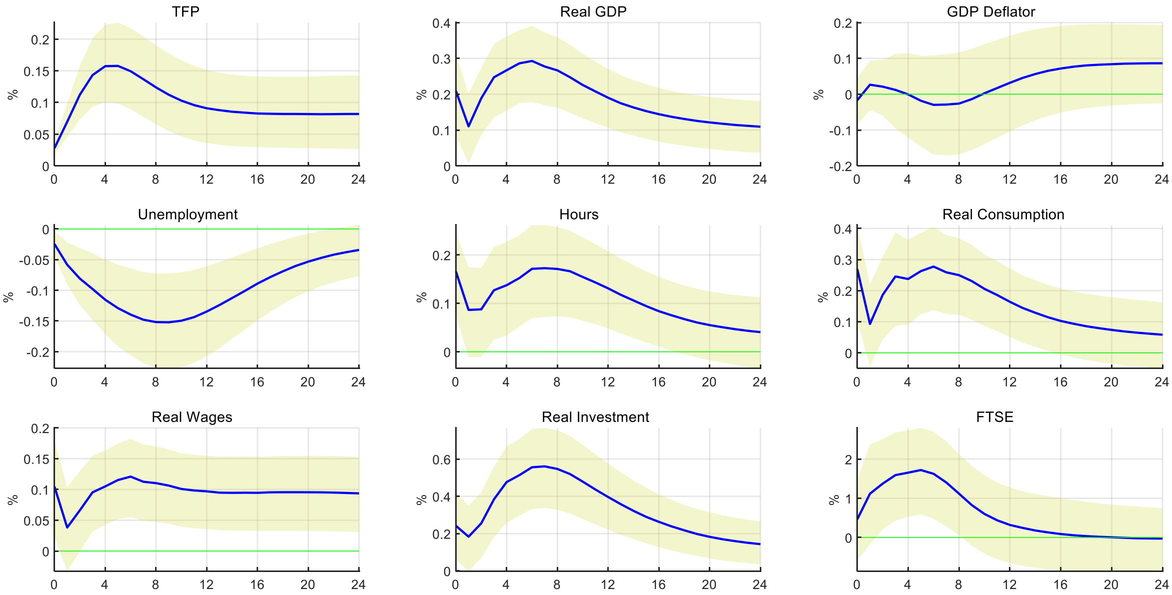


*Note:* The figure shows the impulse response functions (IRFs) to a one standard deviation AI patent-based news shock. Median responses are depicted as solid blue lines, with shaded areas representing 68% posterior bands. The horizon is 24 quarters (6 years).

significant throughout the forecast horizon. This suggests that news derived from AI patent activity convey essential information about aggregate productivity levels both in the present and over an extended period into the future. To put the magnitude of our estimate into perspective, the peak effect of broad technological news shocks on TFP reported in Cascaldi-Garcia and Vukotić (2022) is around 0.25%, while Kurmann and Sims (2021) and Barsky et al. (2015) report peak estimates of approximately 0.40% and 0.30%, respectively.

The positive response of TFP is consistent with a growing micro-level literature highlighting the productivity-enhancing effects of AI. At the firm level, Czarnitzki et al. (2023) find a positive and significant association between AI adoption and firm productivity. Similarly, Yang (2022) report firm-level evidence from Taiwan's electronics sector showing that AI innovations are positively linked to productivity, suggesting that AI adoption alters production methods in ways that enhance efficiency. Complementing these findings, Babina et al. (2024) show that AI investments foster firm growth primarily through product innovation, which expands output and market share, thereby reinforcing productivity gains.

Importantly, we find that TFP responds contemporaneously to AI-news shocks, that is, within three months after the shock.[36] This contrasts with the findings of previous studies in the broader technology news literature that typically find a delayed effect on TFP, often characterized by an S-shaped trajectory (Barsky et al., 2015; Miranda-Agrippino, 2023; Cascaldi-Garcia and Vukotić, 2022).[37] In addition, as Barsky et al. (2015); Kilian et al. (2023) and Kurmann and Sims (2021) highlight, there is no reason to *a priori* assume that technology news shocks are disconnected from current productivity. By contrast, producers may adjust their expectations about future fundamentals once tangible evidence emerges that some firms have successfully implemented the new technology.

The different timing of the TFP response between AI and general technology news shocks should also not be surprising, given the distinctive nature of the former technologies, which can be viewed as a form of intangible capital (Corrado et al., 2021; Brynjolfsson et al., 2017). AI-related patents, such as ML algorithm or LLMs, can be integrated into production processes much more rapidly than broader tangible technologies such as semiconductors or pharmaceuticals, which face higher adoption frictions and fixed costs. While earlier contributions in the AI literature adopted a more modest view on how quickly AI-driven productivity improvements could appear in official statistics (Brynjolfsson et al., 2017; Gordon, 2018), more recent sectoral evidence suggests otherwise. Brynjolfsson et al. (2025) find that access to AI increases worker productivity by around 15%. Importantly for our analysis, these gains arise immediately upon deployment, with AI augmenting existing workflows rather than requiring major organizational restructuring, thereby corroborating our finding that AI innovations affect TFP on impact.[38]

[36] As we discuss in the following section, this finding remains highly robust across various specifications.

[37] These studies, like ours, do not impose restrictions on the immediate TFP response.

[38] Note that in Appendix A2, we report the TFP response from a specification in which we replace our AI patent-based measure with the non-AI patent-based measure, which reflects broader technological categories and is more closely aligned with previous studies of technological news shocks. By plugging this index into our proxy SVAR, the TFP response follows the pattern observed in Cascaldi-Garcia and Vukotić (2022) and Miranda-Agrippino (2023), reflecting the typical slow diffusion of technological news. Indeed, the non-AI technology news shocks does not significantly affect TFP for the first four quarters; the response begins to increase thereafter and eventually becomes statistically significant. This reinforces the view that the immediate TFP response to the news shock observed in our baseline results is specific

Turning to the broader economic effects of the shock, the macroeconomic responses depicted in Figure 4 are consistent with a 'news-driven' business perspective, where the anticipation of forthcoming technological advancements in AI stimulates a business cycle expansion. An AIPNS triggers a noticeable co-movement among real output, real consumption, and real investment. All responses exhibit immediate increases, and show a persistent positive hump-shaped reaction, reaching their peak after approximately six quarters. The positive effects of these responses are consistent with previous literature on broader technological news shocks (Cascaldi-Garcia and Vukotić, 2022; Miranda-Agrippino, 2023) and with the recent study by Aldasoro et al. (2024), which finds positive long-run effects of AI adoption on output, consumption, and investment. The magnitude of these responses is economically substantial, with output and consumption rising by 0.3 percentage points at the peak, and investment increasing by 0.6 points. Subsequently, the responses of output and investment gradually slow down to their new long-run level. Both effects are long-lasting as the responses remain significant throughout the forecast horizon of 24 quarters. The response of consumption is sustained even approximately after four years (i.e., 16 quarters) but the error bands do not rule out a zero long-run effect. Unsurprisingly, the stock market promptly incorporates positive news regarding future AI developments, exhibiting an immediate upward response upon impact, consistent with the news-driven literature (Barsky et al., 2015).

Moving to the labour markets, the response of hours worked is positive on impact and rises significantly at the two-year horizon (i.e., 8 quarters), then gradually declines over time. For longer horizons, the effect is not statistically different from zero. Unemployment declines, reaching its lowest point around eight quarters following the AIPNS, and then gradually reverts to its initial impact level after the shock. Real wages also rise on impact and then remain robustly elevated over time. Combined with the subdued initial response of inflation, this indicates that news about future improvements in AI lead to a short-term increase in aggregate nominal wages. Taken together, the responses of unemployment, hours, and real wages point to an immediate and long-lived improvement

to AI-related technological innovations.

in labor market outcomes in response to AIPNS. Our results on labour markets provide also link to theoretical predictions on AI deployment (Acemoglu and Restrepo, 2018*b*, 2020, 2022). In this context, we view our findings as early evidence that productivity-enhancing mechanisms dominate displacement effects, as reflected in higher hours worked, lower unemployment, and rising real wages.[39]

Last, the response of inflation is quite muted, as indicated by the median which fluctuates around zero and the confidence bands indicating no statistically significant effect. The VAR literature presents mixed findings on the impact of technology news shocks on inflation: some studies document a sharp initial decline followed by a persistently negative response in prices (Barsky and Sims, 2011; Görtz and Tsoukalas, 2018; Kurmann and Sims, 2021), commonly referred to as the disinflation puzzle, while more recent studies document no significant response (Di Casola and Sichlimiris, 2018; Cascaldi-Garcia and Vukotić, 2022; Bouakez and Kemoe, 2023). Our result aligns with the latter group, suggesting that the disinflation puzzle disappears and the response of inflation is not significant. The response subsequently slowly builds up over time, consistent with Cascaldi-Garcia and Vukotić (2022) and predictions of standard New Keynesian models (Barsky and Sims, 2011; Jinnai, 2014) that technology news shocks are inflationary. However, the effect is not statistically different from zero at any horizon.

## 3.2 Forecast Error Variance Decompositions

To understand the extent to which AIPNS explain movements in all variables in our VAR, we report the forecast error variance decompositions (FEVD) in Figure 5. FEVD show the percentage share of variance accounted for by the AIPNS over a 24-quarter horizon. Several interesting observations emerge. The shock explains approximately 9 percent of the variation in TFP in the short run and around 7 percent at longer horizons, indicating that news concerning future advancements in AI carries relevant information about future productivity changes. Furthermore, the shock is a significant driver of real GDP and

[39]We run an alternative specification where total employment replaces the unemployment rate. The results are depicted in Figure A9. Employment responds positively and significantly to the AIPNS, and the remaining responses remain largely unchanged, confirming the baseline implications that the productivity-enhancing effects of AI dominate its potential displacement effects.

investment, as it accounts for approximately 16% and 14% of their overall variation, respectively, after 24 quarters. In both cases, the news shock explains a smaller fraction of variations in the short term, with an increasing effect over time. Additionally, the shock is responsible for approximately 10% of the fluctuations in real wages, consumption, and stock prices after six years, and it accounts for about 12% of changes in unemployment and hours over the same period.

Figure 5: Forecast Error Variance Decomposition.

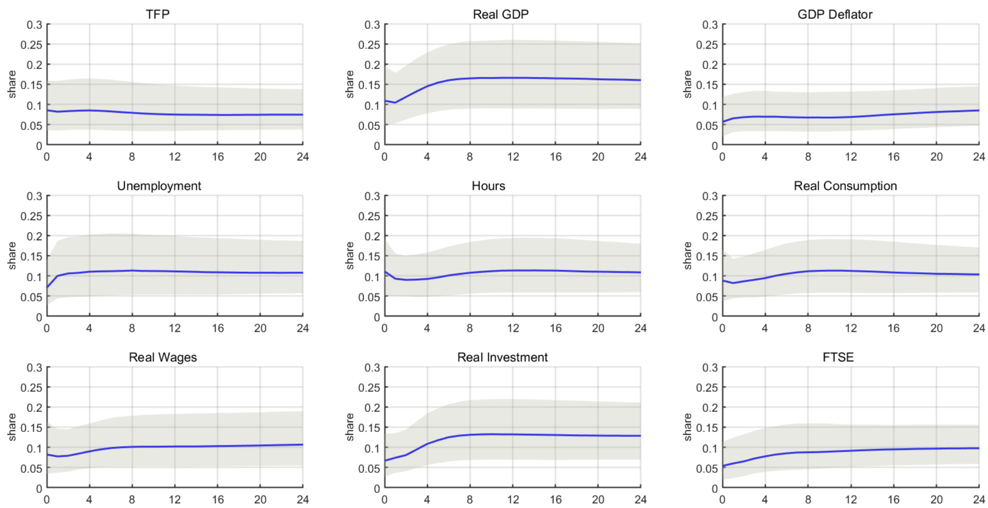


*Note:* The figure depicts the forecast error variance decompositions (FEVD) for key macroeconomic variables over a 24-quarter horizon. The results highlight the proportion of variance in each variable explained by the identified AI patent-based news shock. Median estimates are shown, with shaded areas representing the 68% posterior credible intervals.

The finding that AIPNS explains a substantial share of the variation in key macroeconomic variables reinforces our baseline results, as it suggests that the shocks capture meaningful information about future productivity, overall macroeconomic conditions, and labour market dynamics.

Our findings should be viewed as conservative estimates that understate the effects of AIPNS, given the long sample used in the VAR estimation relative to the IV sample and the fact that our patent-based measure captures only direct technological innovation activity, without accounting for indirect effects such as knowledge spillovers.

### 3.3 Alternative Identification of the AI News Shocks

To bolster the credibility of our findings and ensure they are not driven by our IV design, we further implement an alternative identification strategy. We identify the AIPNS following the maximum forecast error variance (MFV) approach proposed by Uhlig (2005), Barsky and Sims (2011), and Neville Francis and DiCecio (2014). This method identifies the structural shock that best explains the immediate, unforeseen variation in our AI patent-based measure. A key benefit of this approach is that it allows us to exploit the full informational content of the AI patent series, extending our sample back to 1982Q1.

Recall from Section 2.3 the linear mapping of $u_t$ and the structural shocks $\varepsilon_t$, $u_t = A_0\varepsilon_t$. The $A_0$ impact matrix makes $\Sigma_u = \mathbb{E}[u_t u_t'] = \mathbb{E}[A_0\varepsilon_t(A_0\varepsilon_t)'] = A_0\mathbb{E}[\varepsilon_t\varepsilon_t']A_0' = A_0\Sigma_\varepsilon A_0' = A_0A_0'$, where $\Sigma_u$ is the variance-covariance matrix of the reduced-form innovations. Typically, identification is achieved by choosing $A_0$ as the Cholesky factor of $\Sigma_u$, implying a recursive ordering of the variables as in Sims (1986). In our context, we focus solely on identifying the innovation associated with the AIPNS; other structural shocks are not of interest. We can rewrite $A_0$ as $\hat{A}_0M$, where $\hat{A}_0$ is the lower triangular Cholesky factor of $\Sigma_u$ and $M$ is an $n \times n$ matrix satisfying $MM' = I$.

Let $\Omega_{i,j}(h)$ be the share of the forecast error variance of variable $i$ explained by structural shock $j$ at horizon $h$, and that the AI patent-based measure is driven only by one exogenous shock, that is, the AIPNS; it follows that $\Omega_{1,1}(0)_{AI} = 1$, where $i = 1$ refers to the AI measure (ordered first in the VAR), $j = 1$ is the AIPNS, and $h = 0$ denotes the effect of the shock on impact. The share of the forecast error variance of the AIPNS is defined as:

$$\Omega_{1,1}(0)_{AI_{news}} = \frac{e_1'\left(B_T\hat{A}_0Me_2e_2'M'\hat{A}_0'B_T'\right)e_1}{e_1'\left(B_T\Sigma_uB_T'\right)e_1} = \frac{B_{1,T}\hat{A}_0\gamma\gamma'\hat{A}_0'B_{1,T}'}{B_{1,T}\Sigma_uB_{1,T}'}.$$

where both $e_1$ and $e_2$ are selection vectors with 1 in the position $i = 1$ and zeros elsewhere, while $B_T$ is the matrix of moving average coefficients measured at each period until $T$. Since we are interested in the contemporaneous effect of the shock, this simplifies to $T = 0$. The combination of $e_1$ and $e_2$ with the proper column of $M$ can be written as

$\gamma$, which makes $\hat{A}_0\gamma$ the impact of AIPNS on all variables in the VAR. The AIPNS is identified by solving the optimization problem:

$$\gamma = \arg\max \Omega_{1,1}(0)_{AI_{news}} \quad \text{s.t.} \quad \gamma'\gamma = 1$$

This restriction imposes that $\gamma$ is orthonormal.

Intuitively, the AIPNS is identified as the shock that best explains unpredictable movements in the AI patents-based measure, on impact. The information set remains the same as described in the baseline, except that the AI series is now included in the VAR. This allows us to extend the sample back to 1982Q1, corresponding to the full availability of our AI patent series, which had previously been constrained by the unavailability of the BoE forecasts used in our IV. Figure 6 presents the IRFs obtained using the MFV identification scheme, confirming the robustness of our results to this alternative approach.

The responses are qualitatively similar to the benchmark model shown in Figure 4. The AIPNS triggers an immediate response and a persistent long-term increase in TFP. Consumption, investment, and consequently output experience an immediate rise in response to anticipated advancements in AI, while unemployment gradually falls. Similarly, hours worked and real wages also show a positive initial response and gradually adjust to their new lower long-term levels. The only difference compared to the baseline is that prices fall on impact in response to the shock, suggesting that the inflation response obtained from this identification scheme is in line with the disinflation puzzle observed in earlier studies of broader technological news shocks (Barsky and Sims, 2011; Görtz and Tsoukalas, 2018; Kurmann and Sims, 2021). As expected, the stock market index, being a forward-looking variable, quickly prices in the positive news and jumps on impact. In summary, our findings further confirm the significant impact of AI news shocks on the UK labor market and the broader macroeconomy in line with our benchmark specification.

Figure 6: Impulse Response Functions obtained through the MFV identification scheme.

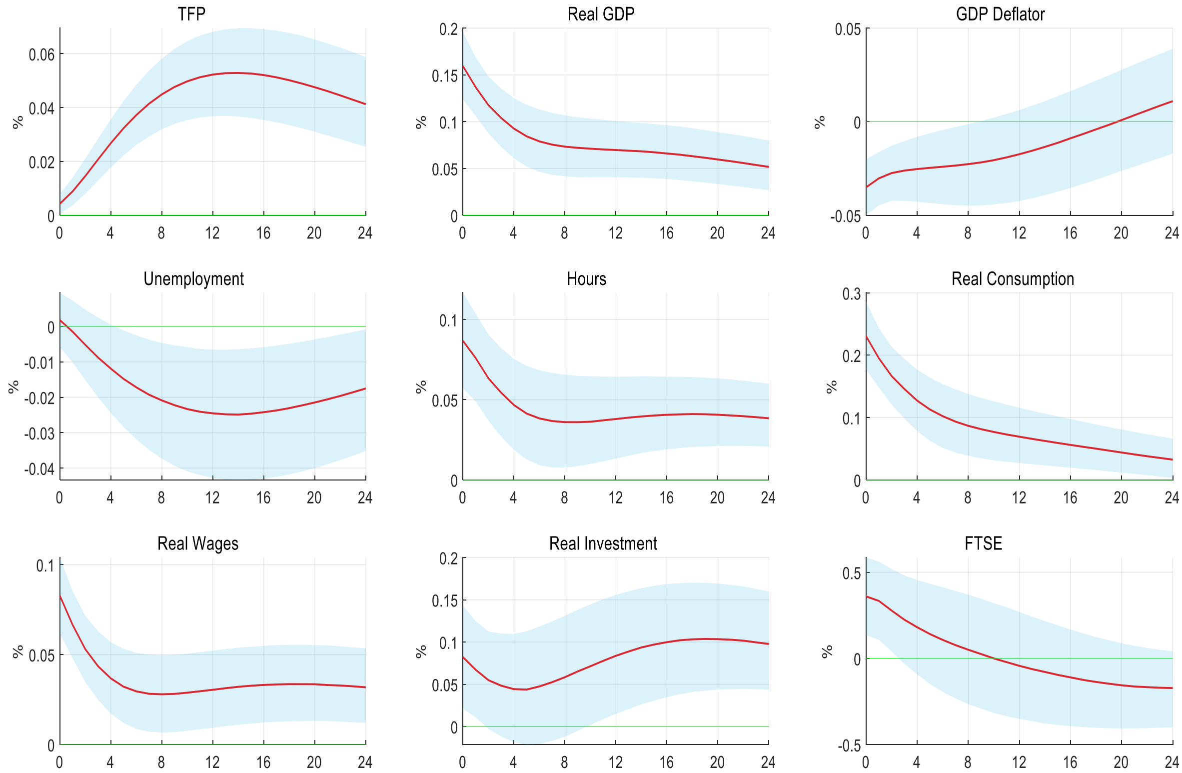


*Note:* The figure depicts the impulse response functions (IRFs) for the AI patent-based news shock using the MFV identification scheme. Median responses are shown with shaded areas representing 68% posterior credible intervals. The horizon is 24 quarters.

# 4 Robustness and Extensions

In this section, we replicate the baseline findings by exploring several extensions. We control for additional confounding factors to ensure the validity of our instrument, by augmenting Equation 1 with monetary and fiscal policy shocks, broader technological developments, and public R&D expenditures. We also consider an alternative construction of the patent measure based on forward citation weights, as well as an alternative measure of TFP. Finally, we compare our VAR results with estimates obtained from local projections.

## 4.1 Alternative IVs

A potential concern with the instrument construction in Section 2.2 arises from the possibility that AI patent filings may be correlated with other contemporaneous policy shocks beyond changes in expectations about the UK macroeconomy, as captured by BoE forecasts. In particular, monetary and fiscal policy may shape the macroeconomic environment, which in turn could influence AI patent activity. We address this concern by introducing the following two series as controls: a measure for UK monetary policy shocks proposed by Miranda-Agrippino (2023) which projects market surprises on central banks' forecasts and forecast revisions of key variables expected to enter the central bank's reaction function.[40] We also incorporate the narrative tax shocks series estimated by Cloyne (2013) and extended by Hayo and Mierzwa (2023).[41]

The results of the IV regression are reported in Table A1, which extends the baseline IV specification by incorporating alternative instruments described in this section. Columns (2) ('policy shocks') and (3) ('$\mathbb{E}_t$ + policy shocks') report results from augmenting equation 1 with the aforementioned policy shock series. In both specifications, whether or not we control for the BoE's macroeconomic forecasts, we find that fiscal policy shocks are statistically significant predictors of AI patenting activity, suggesting they may directly influence firms' innovation behavior. In contrast, monetary policy shocks do not exhibit a statistically significant association with AI patent filings.

The IRFs are reported in Appendix G, Figure A3.[42] Despite the shorter IV sample ending in 2015, it is reassuring that all responses are qualitatively similar to those in the benchmark model shown in Figure 4 with the exception of a decline in inflation. While some responses are delayed and exhibit less pronounced effects, they are consistent with positive news regarding future AI developments prompting an increase in aggregate productivity and a broad-based expansionary business cycle phase.

Another potential source of endogeneity arises from the possibility that AI patent

[40]Note that we also consider an alternative specification where we utilize the monetary surprises series constructed by Cesa-Bianchi et al. (2020), covering almost the same period, that is, from 1996:Q2 to 2015:Q1. The responses (not reported here) are almost identical to the results in Figure A3 (see below).

[41]The tax shock series is available up to 2018:Q2, whereas the monetary policy shock series spans an even narrower period, up to 2015:Q1. Owing to these data limitations, equation 1 is estimated over a correspondingly narrower IV sample relative to our baseline specification, covering the period from 1997:Q3 to 2015:Q1. The VAR sample remains unchanged from the baseline.

[42]The figure depicts the IRFs from the specification that includes the lagged AI values and the narrative policy shocks, that is, the specification corresponding to the regression output in column (2) in Table A1. The results are largely unchanged when replacing it with the IV from the specification in column (3) in Table A1.

applications are not fully exogenous to broader technological activity seemingly unrelated to AI. Since AI innovation is often intertwined with complementary technologies such as semiconductors, our baseline IV may capture not only expected developments in AI but also expectations about broader technological trends. To address this, we construct an alternative IV by augmenting the regression with a control for broader techonological trends, proxied by our non-AI patent-based measure. Indeed, the results in Table A1 in Appendix G reveal that non-AI patenting activity contains information about future AI patent production, as reflected in the significance of its current and lagged coefficient values.

The IRFs derived from this specification, presented in Figure A4 in Appendix G, support the results of our baseline model. Positive news about future AI activity raise TFP and initiate GDP expansion and improvement of labor market outcomes. A subtle short-run difference emerges: real GDP, consumption, real wages, and hours worked exhibit temporary initial declines, while unemployment remains flat. These effects likely reflect transitional dynamics from the task-displacing nature of AI, as firms reorganize production and delay hiring before productivity gains materialize. This is consistent with Acemoglu and Restrepo (2020), who document short-run declines in labor demand and consumption following automation and AI adoption. Nevertheless, these effects are temporary. Within a few quarters, the anticipated productivity improvements begin to materialize as firms adjust to the expected advancement of AI, giving rise to a broad-based expansion in macroeconomic and labor market indicators, closely mirroring the medium- and long-run dynamics observed in the baseline model.

Another potential endogeneity concern is that public innovation policies may directly influence innovation decisions (Bloom et al., 2019). In particular, government R&D spending could affect firms' incentives to invest in AI, confounding the identification of exogenous technology news shocks. To address this, we estimate an alternative specification that controls for UK public R&D expenditure.[43] The regression results, presented

[43]The data on government R&D allocations are sourced from the OECD. They are reported at an annual frequency. We interpolate these into quarterly observations using cubic spline interpolation to match the frequency of the other variables in our VAR.

in the last column of Table A1, indicate that public R&D does not exert a statistically significant effect on AI patenting activity. Consistent with this finding, the IRFs shown in Figure A5, which incorporate public R&D and its lags as controls, remain largely unchanged relative to our benchmark specification.

## 4.2 Weighting Patents by Scientific Contribution

The use of patent counts is often criticized because this measure may either underestimate or overestimate the true scientific and economic impact of innovations (Kelly et al., 2021; Verendel, 2023). A potential remedy is the use of patent citation counts. Patent citations also have limitations, but they remain a well-established proxy for knowledge flows in innovation economics (Criscuolo and Verspagen, 2008; Jaffe et al., 1993; Jaffe and Trajtenberg, 1999) and are increasingly used to identify major innovation breakthroughs in the economics literature (Kelly et al., 2021). In this robustness test, we weight patents by their forward citations, a measure widely regarded as capturing a patent's scientific relevance (Miranda-Agrippino, 2023; Kelly et al., 2021).

PATSTAT reports the citations each patent receives (i.e., which patents cite a given patent). We use this information to construct a weighted AI patent series. Specifically, each patent is multiplied by the number of its citations plus one: $Patents_{adj} = Patents \times (Citations + 1)$, where $Patents_{adj}$ denotes the adjusted number of patents scaled by their citations.[44] Figure 7 compares the weighted AI patent series with the unweighted baseline. Note that the citation-weighted series begins to diverge noticeably from the baseline from 2018 onward, with a pronounced decline toward the end of the sample. This occurs because citation weights aggravate the truncation bias (Lerner and Seru, 2021). In other words, it can take many years after a patent is filed before its influence is fed to forward citations.[45]

Figure A6 in Appendix G plots the responses. Relative to the baseline specification, the results are qualitatively similar: TFP and real GDP rise significantly, unemployment

[44]The unit ensures that patents without citations are not excluded (Miguelez and Moreno, 2018).

[45]As noted in Lerner and Seru (2021) and Miranda-Agrippino (2023), citations for patents can only accumulate over time, leading to more recent patents being underweighted due to having fewer citations, regardless of their actual scientific value.

Figure 7: Comparison of weighted and unweighted AI patent series.

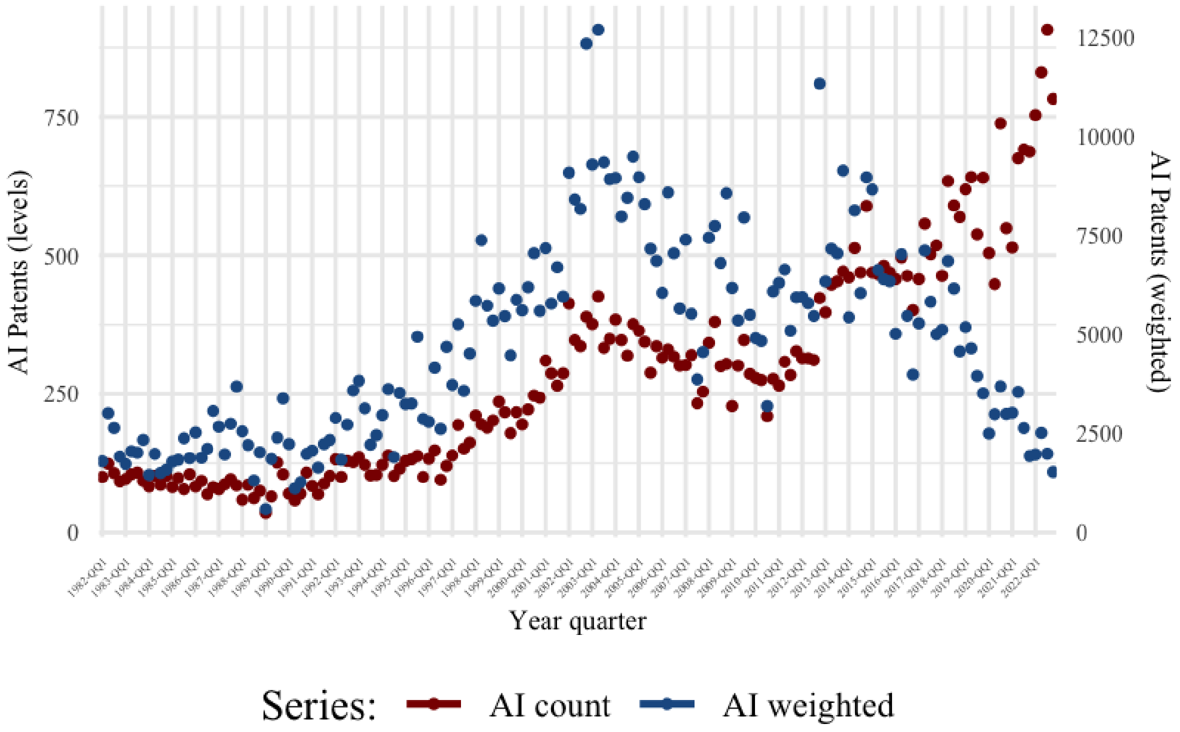


*Note:* The red dotted series is our baseline AI patent-based index (AI count). The blue dotted series (AI weighted) scales each patent by the number of forward citations it has received (plus one), thus accounting for their scientific relevance.

falls, and both real wages and investment increase. The main difference is that the positive effects implied by the citation-weighted AIPNS are generally smaller in magnitude. In some cases, such as hours worked and consumption, their significance fades more quickly, reflecting truncation bias, since patents granted near the end of the sample have fewer citations and their true impact is underestimated. Taken together with the limitations discussed above, this warrants caution when relying on the citation-weighted scheme as an instrument for AI news shocks.[46]

## 4.3 Utilization adjusted TFP Series

The VAR technology news literature typically utilizes utilization-adjusted TFPs to identify technology news shocks.[47] As Fernald (2014) shows, adjusting for these utilization

[46]To partially address truncation bias, we estimate an alternative specification excluding observations from 2018:Q2 onward. The resulting IRFs, not reported here, are broadly similar to the baseline.

[47]Utilization-adjusted TFP accounts for variations in factor utilization, e.g. hours worked per employee, labor effort, and capital workweeks.

margins helps strip out cyclical non-technological fluctuations, yielding a cleaner measure of underlying technological change. Unfortunately, apart from the US case, there are currently no other utilization-adjusted quarterly TFP estimates covering a long sample period. In our main specification, we use the closest available series sourced from the ONS, which accounts for lower capital utilization during the COVID-19 pandemic. Consequently, our baseline TFP measure remains only partially adjusted.

Nevertheless, a salient feature of our identification strategy is its departure from imposing assumptions regarding the short- and long-term behavior of TFP. Thus, our findings should remain robust regardless of whether productivity is adjusted for capacity utilization. To verify this, we incorporate the sole UK-adjusted TFP measure available at the time of writing, as designed in Kamber et al. (2017) and extended by Mouratidis et al. (2024). Its coverage is restricted to 1998Q1–2020Q1, thereby excluding the period marked by the diffusion of recent advances in generative AI. Consequently, the VAR sample is smaller than in the baseline due to limited TFP data availability. Figure 8 suggests that the response of the adjusted TFP series to an AIPNS closely mirrors our baseline scenario, characterized by an initial jump in response to the shock, followed by a reversion to its new, higher long-run level of about 0.05%. Therefore, our results remain unaffected regardless of whether productivity is adjusted for capacity utilization. [48]

[48] The remaining IRFs remain largely unchanged. The full figure is available upon request.

Figure 8: Impulse response functions (IRFs) for utilization-adjusted TFP.

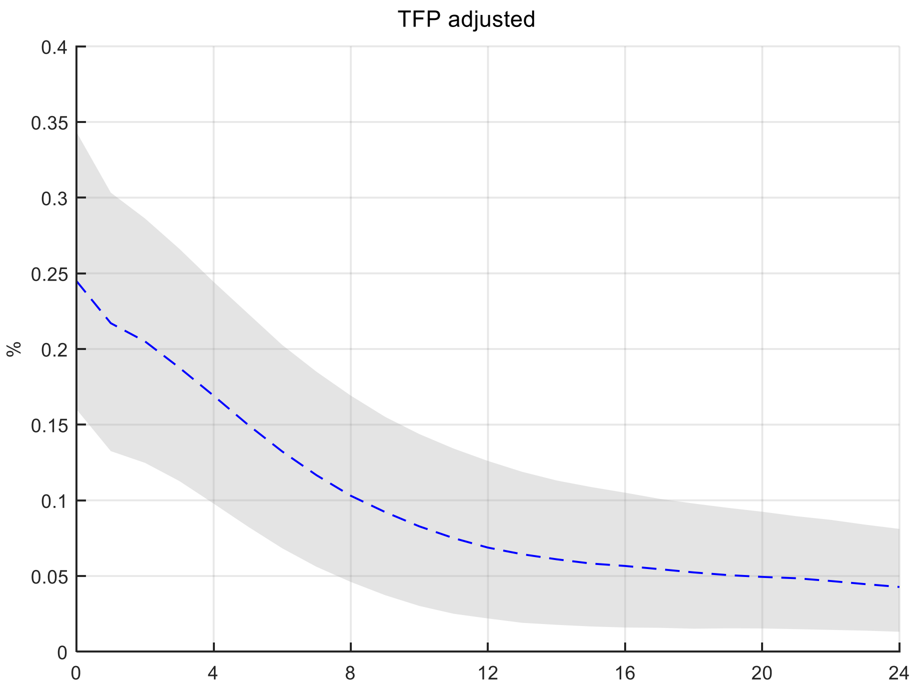


*Note:* The figure shows the IRFs from a VAR using the utilization-adjusted TFP series from Kamber et al. (2017). Estimation sample 1998Q1–2020Q1, IV sample 1998Q1-2020Q1.

### 4.4 Local Projections

It is often argued that lag truncation in VAR models can bias medium- and long-term impulse responses, whereas local projections (LPs) are less susceptible to this issue (Jordà et al., 2024). In our case, the risk of lag truncation bias is limited, since we include a relatively large number of lags in the VAR. As Jordà et al. (2024) point out, using more than one lag helps mitigate the adverse effects of truncation when capturing the impact of shocks at medium horizons.

Still, to check robustness, we also estimate the impulse responses to an AI news shock using LPs à la Jordà (2005) that estimate the IRFs directly from the following linear regression:

$$y_{t+h} = c_0^h + \sum_{j=1}^{p} B_j^h y_{t-j} + u_{t+h}, \tag{7}$$

where $y_t$ is the vector of the same endogenous variables used in our VAR, $h$ is the IRFs horizon, and $u_{t+h}$ denotes residuals.[49] The IRFs for the AI news shock at horizon $h$ can be computed as $B_1^h A_0^{(1)}$, where $A_0^{(1)}$ denotes the contemporaneous effect of the AI news shock obtained using the MFV VAR.[50,51]

Overall, the results depicted in Figure 9 are qualitatively consistent with those from the MFV VAR: under the LP approach, TFP, real GDP, hours worked, consumption, wages, investment, and the stock market increase in response to the shock, while unemployment declines. The magnitudes of most LP estimates are comparable to the average effects in Figure 6. However, the LP responses are more erratic and display considerably wider confidence bands, rendering them statistically insignificant over longer periods. This wider variability of LP estimates is consistent with the findings of Kilian et al. (2023), Miranda-Agrippino (2023), Budnik and Rünstler (2023), and Plagborg-Møller and Wolf (2021) on the relative efficiency of structural VARs versus local projections.

Figure 9: Local projections.

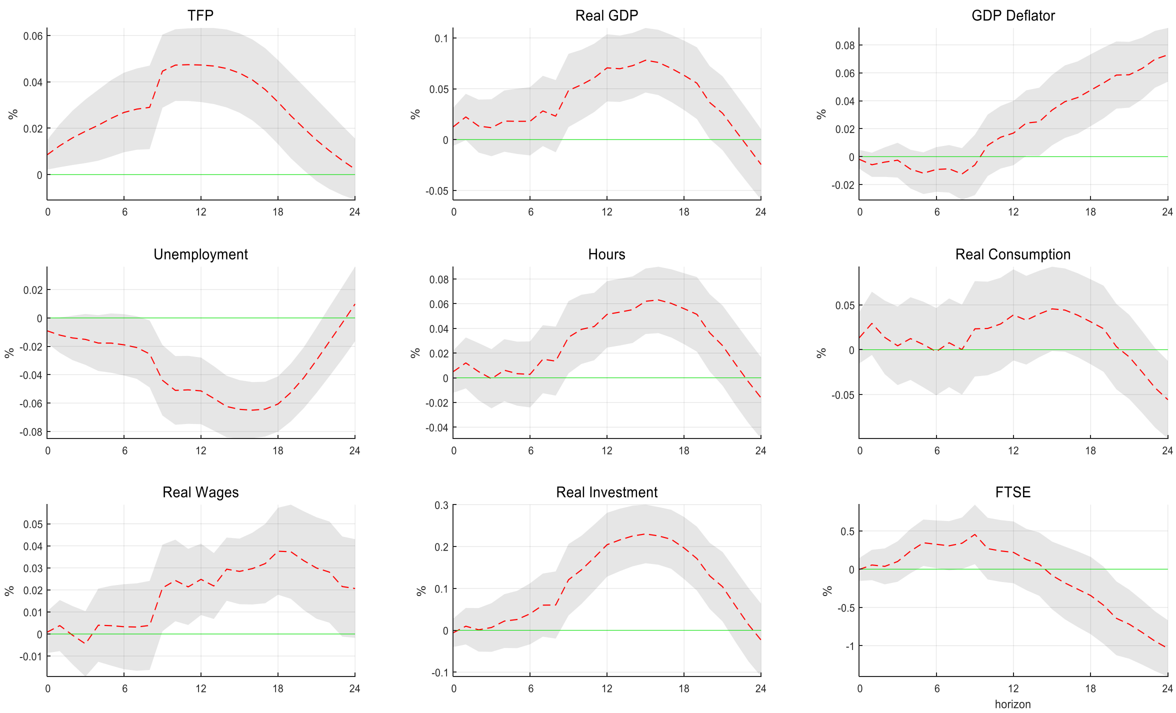


*Note:* The figure shows the IRFs from Local Projections following Jordà (2005). Median responses are shown in red dotted lines, with grey shaded areas representing 68% posterior credible intervals. The horizon is 24 quarters.

[49] The lag order $p$ depends on $h$.

[50] Note that at horizon 0, the VAR and local projection coincide.

[51] We use the MFV-identified shock for LP because it incorporates the AI patent series directly into the VAR, enabling estimation on an extended sample (1982Q1–2022Q4) compared to the IV sample. Using the proxy-SVAR news shock would restrict the LP estimation to the shorter IV sample.

### 4.5 Additional robustness

We perform additional robustness checks, including a pre-crisis sample, alternative inflation and productivity measures, total employment in place of the unemployment rate, and an IV excluding 8-quarter-ahead CPI and GDP forecasts. A detailed description of these specifications and results is provided in Appendix H.

## 5 Industry-Level Evidence of AI Adoption

In this section, we aim to better understand which industries drive the propagation of news regarding future technological advancements in AI. To this end, we account for sectoral heterogeneity in the generation of AI innovations by constructing sectoral indices of AI-related patents, allocating each patent to its industry of origin. We follow Dorner and Harhoff (2018), which allows us to map International Patent Classification (IPC) codes to NACE sectors.[52,53]

After matching our AI measure with the concordance tables, we obtain a quarterly dataset of AI patents that originates in 272 sub-sectors, which correspond to three digit NACE numerical codes (01.1 to 99.0), over the period 1982-2022. Table A3 in the Appendix reports the share of AI patents as a proportion of total AI patents across the broad NACE industrial sectors during the sample period. The table includes only those sectors in which at least one AI patent has been recorded.

We observe that 59% of AI patents originate in the manufacturing sector. Within

[52]Specifically, we identify the industry of origin for each patent by linking its International Patent Classification (IPC) codes to NACE sector codes, following the concordance developed by Dorner and Harhoff (2018). This mapping combines inventor–employee information from PATSTAT with administrative employment data, enabling the assignment of patents to the industries where inventors are employed. For instance, consider patents classified under IPC `G06F`, which pertains to electronic data processing technologies (e.g., computer hardware and software). Dorner and Harhoff (2018) show that many inventors associated with `G06F` patents are employed in firms classified under NACE 26.20 (Manufacture of Computers and Peripheral Equipment) and NACE 62.01 (Computer Programming Activities).

[53]The literature distinguishes between two main approaches to patent–industry concordance: the industry-of-origin (IOO) approach (Dorner and Harhoff, 2018), and the industry-of-use (IOU) approach captured through a probabilistic matching method (Lybbert and Zolas, 2014). While IOU concordance tables could, in principle, be useful for capturing the end-use of AI technologies, existing implementations are confined to the manufacturing sector. In contrast, the IOO approach, adopted in this paper, enables us to map IPC codes to a much broader range of industries such as IT, research and software development, highly relevant for AI.

manufacturing, the majority of these patents, 76% of the total AI patents in this sector, are concentrated in the production of *computer and electronic products* (NACE code C26). The majority of AI patents within C26 are concentrated in the following categories: *computers and peripheral equipment* (C26.2) account for 28%, *communication equipment* (C26.3) for 22%, *instruments for measuring, testing, and navigation* (C26.5) for 16%, and *electronic components and boards* (C26.1) for (7%), of all AI patents in manufacturing.

Outside the manufacturing sector, the two service-related sectors with the lion's share of AI patents are: a) *computer programming, data processing, and related activities* (NACE code J62) within the broader industry class J: *Information and Communication*, and, b) *research and development* (NACE code M72) within the broader industry class M: *Professional, Scientific and Technical Activities.*[54,55] Together, these two service sectors dominate the AI patent landscape within the broader services category (NACE codes J and M), accounting for a combined share of 96%. Last, less technologically intensive sectors, such as agriculture, mining, real estate, education, and arts and entertainment, have unsurprisingly generated only a minimal number of AI patent filings.

Figure 10 displays the effects of three shocks on selected macroeconomic variables: (i) the baseline (aggregate) AIPNS; (ii) the AIPNS based on patents originating in the manufacturing sector (industry C in NACE); and (iii) the AIPNS based on patents originating in professional, scientific, and information services sectors (the combined contribution of industries J and M in NACE, hereafter referred to as the 'services' news shock for brevity). Each panel reports the FEVD at horizons of 0, 6, 12, 18, and 24 quarters.

Starting with TFP, the services news shock accounts for an average of about 8% of its fluctuations, a share very close to that implied by all patents in the sample. The manufacturing sector contributes a lower share, around 5% of the FEVD of TFP. This pattern suggests that AI-related patenting activity in services is more relevant for future movements in TFP than patenting activity in manufacturing. The stronger impact

[54] J62 includes technologies related to software development, modification, testing, and support; the planning and design of integrated computer systems; and the on-site management and operation of clients' computer systems and data processing facilities.

[55] Our estimates are in line with Vukotić (2019) and Cascaldi-Garcia and Vukotić (2022), who also estimate technology news shocks at the level of similar sub-industries within the manufacturing and services sectors.

of the services shock on TFP, relative to manufacturing, reflects the nature of patents in this industry. Specifically, the most related subindustry is *computers and peripheral equipment*, which accounts for 28% of all AI patents in manufacturing. This subindustry develops and deploys core digital and AI-related technologies, such as cloud computing, data analytics, algorithm design, and AI consultancy. Patents in this sector are intangible and software-driven, facilitating easier integration into existing processes and technologies. Consequently, their productivity-enhancing effects tend to be more pronounced than those of patents in manufacturing.

Turning to the macroeconomic variables, a clear pattern emerges. The manufacturing AIPNS is more influential than the baseline AIPNS, particularly for real GDP: it explains 23% of output fluctuations after 24 quarters, compared to 16% under the baseline AIPNS. Similar results hold for real wages and investment, where the manufacturing shock accounts for about 14% of the variance, equal to or exceeding the baseline effect. Thus, patenting activity in manufacturing appears important for explaining movements in key UK macroeconomic indicators.[56] The services news shock also matters, though less pronounced, explaining 12% of GDP fluctuations, 9% of real wage variation, and 11% of investment variation at the 24-quarter horizon.

Figure 11 reports the FEVDs for the key sub-industries identified earlier. The first three columns correspond to the main manufacturing sub-industries, while the final two columns show the dominant service sub-industries. The results indicate that the largest contributions to TFP fluctuations come from the *research and development* sub-sector in the short term (about 8% on average after 6 quarters) and the *instruments for measuring, testing, and navigation* sub-sector in the longer term (18–24 quarters), which accounts for roughly 7% of the variation. When looking at the FEVD of GDP, AI patenting activity in the manufacturing sub-sectors of *computer* and *communication equipment*, accounts for sizable future movements, about 22% and 19% after 24 quarters, respectively. Among the two dominant service sub-sectors, *computer programming* is especially relevant, explaining on average 15% of GDP fluctuations across the forecast horizon. For real investment,

[56] The manufacturing news shock also explains 17% of the variation in both hours worked and consumption after 24 quarters. Full FEVD results for all variables and shocks are available upon request.

the manufacture of *communication equipment* is most important, accounting for 15% of fluctuations in the short term (up to 6 quarters) and 13% at longer horizons. Finally, the *computer* and *communication equipment* industries also dominate in explaining real wage dynamics, each contributing around 10% of the FEVD after 24 quarters.

Figure 10: FEVDs explained by industry AI patent based news shock.

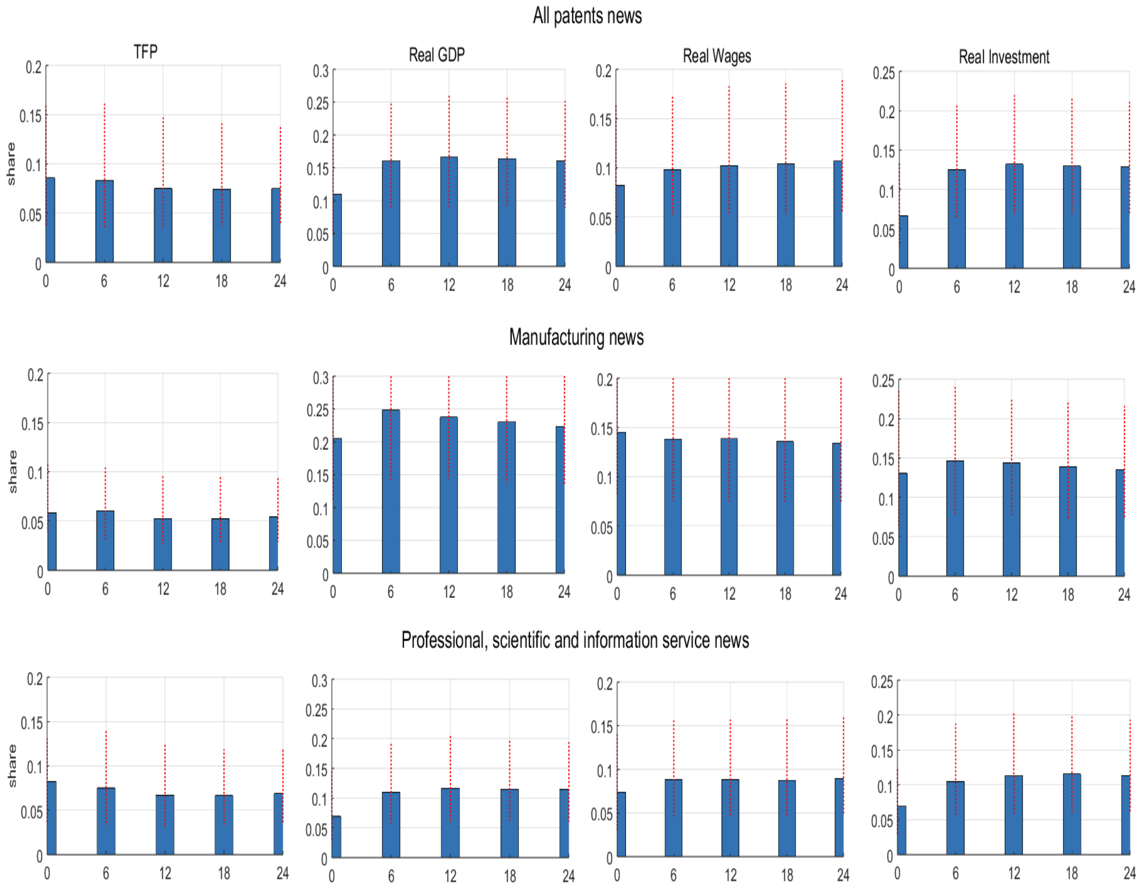


*Note:* The figure reports the FEVDs of TFP, GDP, wages, and investment explained by all AI patent-based news shock, the manufacturing sector AI patent-based news shock, and the services sector patent-based news shock, at different horizons (0,6,12,18,24).

Figure 11: FEVDs explained by sub-industries AI patent based news shock.

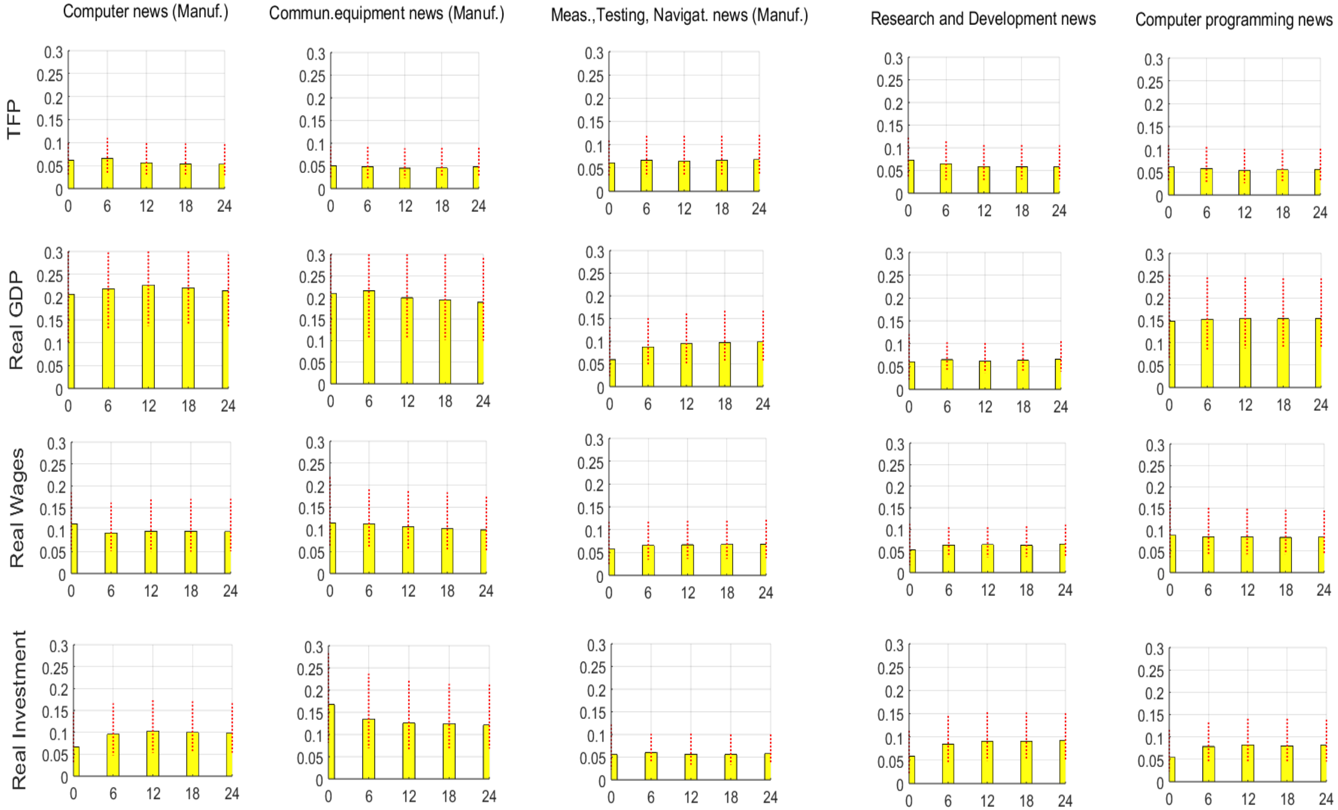


*Note:* The figure reports the FEVDs of TFP, GDP, wages, and investment explained by AI patent-based news shocks from the most relevant subindustries, at different horizons (0,6,12,18,24).

# 6 Conclusion

Recent advances in artificial intelligence have generated both optimism and concern. While some fear that rapid, AI-driven automation will displace workers, others highlight its potential to boost productivity and create new forms of employment. This ambiguity makes early, forward-looking indicators, such as AI-related patent filings, particularly valuable for understanding how expectations about AI shape productivity, labour markets, and the broader economy.

Motivated by this, we ask: How does the UK, an advanced economy and a forerunner in technological innovation, respond to shocks that convey public information about future improvements in AI?

To address this question, we construct a novel measure of AI inventive activity by classifying UK patent text data into AI and non-AI technologies over the past four decades. Classifying AI patents poses a significant challenge, as the scope of AI has evolved substantially over time, rendering traditional keyword searches and technology codes often inconsistent across long periods. We classify over 550 million patent texts using a machine learning algorithm and validate our estimates both internally and with the assistance of an external domain expert.

We exploit this measure of AI innovation activity to build an instrument that is independent of macroeconomic expectations, fiscal and monetary policy, broader trends in non-AI technological innovation, and public R&D spending, factors that could otherwise affect patenting activity. This instrument is then used within a Bayesian proxy VAR to identify AI news shocks, which capture information about future technological developments in AI.

We find that TFP rises immediately following an AI news shock. This immediate jump contrasts with the delayed diffusion pattern typically documented for broad technological-news shocks. Unlike traditional innovations in IT and manufacturing, advances in AI can be integrated more rapidly into existing production processes. Although this interpretation should be viewed with some caution, since many recent AI technologies are only beginning to be applied, our findings of rapid productivity gains are consistent with the emerging literature documenting the rapid adoption of generative AI in the workplace (Bick et al., 2024; Brynjolfsson et al., 2025).

The AI news shocks identified here also trigger a broad-based and persistent macroeconomic expansion, as shown by increases in output, investment, and consumption. This expansion operates primarily through the labour market, which emerges as the central channel of aggregate transmission. Specifically, positive AI news shocks are employment-enhancing: they reduce unemployment while increasing wages and hours worked.

This study, and the identification of AI news shocks based on our proposed AI patent-based series, provides useful insights into measuring the effects of anticipated future technological developments in AI on the overall macroeconomic performance of the UK. Future research could expand by examining the heterogeneous impacts of AI innovations across sectors and different classes of AI technologies. As firm deployment data on AI become available, future research will be able to provide evidence about the impact of AI innovations as they materialize. Yet our evidence already shows that expectations about AI can exert sizeable and positive effects.

# References

Abood, A. and Feltenberger, D. (2018), 'Automated patent landscaping', *Artificial Intelligence and Law* **26**(2), 103–125.

Acemoglu, D. (2025), 'The simple macroeconomics of ai', *Economic Policy* **40**(121), 13–58.

Acemoglu, D., Autor, D., Hazell, J. and Restrepo, P. (2022*a*), 'Artificial intelligence and jobs: Evidence from online vacancies', *Journal of Labor Economics* **40**(S1), S293–S340.

Acemoglu, D., Autor, D., Hazell, J. and Restrepo, P. (2022*b*), 'Artificial intelligence and jobs: Evidence from online vacancies', *Journal of Labor Economics* **40**(S1), S293–S340.

Acemoglu, D. and Restrepo, P. (2018*a*), Artificial intelligence, automation and work, NBER Working Paper 24196, National Bureau of Economic Research.

Acemoglu, D. and Restrepo, P. (2018*b*), 'The race between man and machine: Implications of technology for growth, factor shares, and employment', *American economic review* **108**(6), 1488–1542.

Acemoglu, D. and Restrepo, P. (2020), 'Robots and jobs: Evidence from us labor markets', *Journal of Political Economy* **128**(6), 2188–2244.

Acemoglu, D. and Restrepo, P. (2022), 'Tasks, automation, and the rise in us wage inequality', *Econometrica* **90**(5), 1973–2016.

Acs, Z. J. and Varga, A. (1994), 'High-technology employment and r&d in cities: a new data base of high-tech centers', *Economic Development Quarterly* **8**(2), 107–118.

Aghion, P., Bénabou, R., Martin, R. and Roulet, A. (2023), 'Environmental preferences and technological choices: Is market competition clean or dirty?', *American Economic Review: Insights* **5**(1), 1–19.

Aldasoro, I., Doerr, S., Gambacorta, L. and Rees, D. M. (2024), *The impact of artificial intelligence on output and inflation*, Bank for International Settlements, Monetary and Economic Department.

Alexopoulos, M. (2011), 'Read all about it!! what happens following a technology shock?', *American Economic Review* **101**(4), 1144–1179.

Babina, T., Fedyk, A., He, A. and Hodson, J. (2024), 'Artificial intelligence, firm growth, and product innovation', *Journal of Financial Economics* **151**, 103745.

Baron, J. and Schmidt, J. (2014), 'Technological standardization, endogenous productivity and transitory dynamics', (503).

Barsky, R. B., Basu, S. and Lee, K. (2015), 'Whither news shocks?', *NBER Macroeconomics Annual* **29**(1), 225–264.

Barsky, R. B. and Sims, E. R. (2011), 'News shocks and business cycles', *Journal of Monetary Economics* **58**, 273–289.

Baruffaldi, S., van Beuzekom, B., Dernis, H., Harhoff, D., Rao, N., Rosenfeld, D. and Squicciarini, M. (2020), 'Identifying and measuring developments in artificial intelligence: Making the impossible possible'.

Beaudry, P. and Lucke, B. (2010), 'Letting different views about business cycles compete', *NBER Macroeconomics Annual* **24**(1), 413–456.

Beaudry, P. and Portier, F. (2006*a*), 'Stock prices, news, and economic fluctuations', *American Economic Review* **96**(4), 1293–1307.

Beaudry, P. and Portier, F. (2006*b*), 'Stock prices, news, and economic fluctuations', *American Economic Review* **96**(4), 1293–1307.

Beaudry, P. and Portier, F. (2014), 'News-driven business cycles: Insights and challenges', *Journal of Economic Literature* **52**(4), 993–1074.

Bekkers, R., Bongard, R. and Nuvolari, A. (2011), 'An empirical study on the determinants of essential patent claims in compatibility standards', *Research Policy* **40**(7), 1001–1015.

Berger, F., Blind, K. and Thumm, N. (2012), 'Filing behaviour regarding essential patents in industry standards', *Research Policy* **41**(2), 216–225.

Bhattacharya, S. and Guriev, S. (2006), 'Patents vs. trade secrets: Knowledge licensing and spillover', *Journal of the European Economic Association* **4**(6), 1112–1147.

Bick, A., Blandin, A. and Deming, D. J. (2024), The rapid adoption of generative ai, Technical report, National Bureau of Economic Research.

Bloom, N., Jones, C. I., Van Reenen, J. and Webb, M. (2020), 'Are ideas getting harder to find?', *American Economic Review* **110**(4), 1104–1144.

Bloom, N., Van Reenen, J. and Williams, H. (2019), 'A toolkit of policies to promote innovation', *Journal of Economic Perspectives* **33**(3), 163–184.

Bluwstein, K., Hacioglu Hoke, S. and Miranda-Agrippino, S. (2020), 'When creativity strikes: News shocks and business cycle fluctuations'.

Bonfiglioli, A., Crinò, R., Gancia, G. and Papadakis, I. (2025), 'Artificial intelligence and jobs: evidence from us commuting zones', *Economic Policy* **40**(121), 145–194. Special Issue on Artificial Intelligence and the Economy.

Bouakez, H. and Kemoe, L. (2023), 'News shocks, business cycles, and the disinflation puzzle', *Journal of Money, Credit and Banking* **55**(8), 2115–2151.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I. and Amodei, D. (2020), 'Language models are few-shot learners', *Advances in Neural Information Processing Systems (NeurIPS)* **33**, 1877–1901.

Brynjolfsson, E., Li, D. and Raymond, L. (2025), 'Generative AI at work', *The Quarterly Journal of Economics* **140**(2), 889–942.

Brynjolfsson, E., Rock, D. and Syverson, C. (2017), Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics, Technical report, National Bureau of Economic Research.

Buarque, B. S., Davies, R. B., Hynes, R. M. and Kogler, D. F. (2020), 'Ok computer: the creation and integration of ai in europe', *Cambridge Journal of Regions, Economy and Society* **13**(1), 175–192.

Budnik, K. and Rünstler, G. (2023), 'Identifying structural vars from sparse narrative instruments: Dynamic effects of us macroprudential policies', *Journal of Applied Econometrics* **38**(2), 186–201.

Buesa, M., Heijs, J., Pellitero, M. and Baumert, T. (2010), 'The determinants of regional innovation in europe: A combined factorial and regression knowledge production function approach', *Research Policy* **39**(6), 722–735.

Caldara, D. and Herbst, E. (2019), 'Monetary policy, real activity, and credit spreads: Evidence from bayesian proxy svars', *American Economic Journal: Macroeconomics* **11**(1), 157–192.

Cascaldi-Garcia, D. and Vukotić, M. (2022), 'Patent-based news shocks', *The Review of Economics and Statistics* **104**.

Cesa-Bianchi, A., Thwaites, G. and Vicondoa, A. (2020), 'Monetary policy transmission in the united kingdom: A high frequency identification approach', *European Economic Review* **123**, 103375.

Chawla, N. V. (2005), Data mining for imbalanced datasets: An overview, *in* O. Maimon and L. Rokach, eds, 'Data Mining and Knowledge Discovery Handbook', Springer, pp. 853–867.

Choudhury, P., Wang, D., Carlson, N. A. and Khanna, T. (2019), 'Machine learning approaches to facial and text analysis: Discovering ceo oral communication styles', *Strategic Management Journal* **40**(11), 1705–1732.

Christiansen, L. E. (2008), Do technology shocks lead to productivity slowdowns? evidence from patent data, Working Paper 2008/024, International Monetary Fund. Accessed May 21, 2025.

Cloyne, J. (2013), 'Discretionary tax changes and the macroeconomy: New narrative evidence from the united kingdom', *American Economic Review* **103**(4), 1507–1528.

Cockburn, I. M., Henderson, R. and Stern, S. (2018), *The impact of artificial intelligence on innovation*, Vol. 24449, National bureau of economic research Cambridge, MA, USA.

Corrado, C., Haskel, J. and Jona-Lasinio, C. (2021), 'Artificial intelligence and productivity: an intangible assets approach', *Oxford Review of Economic Policy* **37**(3), 435–458.

Criscuolo, P. and Verspagen, B. (2008), 'Does it matter where patent citations come from? inventor vs. examiner citations in european patents', *Research policy* **37**(10), 1892–1908.

Cunha, W., Rocha, L. and Gonçalves, M. A. (2025), 'A thorough benchmark of automatic text classification: From traditional approaches to large language models', *arXiv preprint arXiv:2504.01930* .

Czarnitzki, D., Fernández, G. P. and Rammer, C. (2023), 'Artificial intelligence and firm-level productivity', *Journal of Economic Behavior & Organization* **211**, 188–205.

Dahlke, J., Beck, M., Kinne, J., Lenz, D., Dehghan, R., Wörter, M. and Ebersberger, B. (2024), 'Epidemic effects in the diffusion of emerging digital technologies: evidence from artificial intelligence adoption', *Research Policy* **53**(2), 104917.

Di Casola, P. and Sichlimiris, S. (2018), Towards technology-news-driven business cycles, Technical report, Sveriges Riksbank Working Paper Series.

Domingos, P. and Pazzani, M. (1997), 'On the optimality of the simple bayesian classifier under zero-one loss', *Machine Learning* **29**(2).

Dorner, M. and Harhoff, D. (2018), 'A novel technology-industry concordance table based on linked inventor-establishment data', *Research Policy* **47**(4), 768–781.

European Patent Office (2017), Patents and the fourth industrial revolution, Technical report, European Patent Office. Accessed: February 2025.

Fernald, J. (2014), A quarterly, utilization-adjusted series on total factor productivity, Federal Reserve Bank of San Francisco.

Filippucci, F. et al. (2024), The impact of artificial intelligence on productivity, distribution and growth: Key mechanisms, initial evidence and policy challenges, OECD Artificial Intelligence Papers 15, OECD Publishing, Paris.
**URL:** *https://doi.org/10.1787/8d900037-en*

Forni, M., Gambetti, L. and Sala, L. (2014), 'No news in business cycles', *Economic Journal* **124**, 1168–1191.

Frey, C. B. and Osborne, M. A. (2017), 'The future of employment: How susceptible are jobs to computerisation?', *Technological Forecasting and Social Change* **114**, 254–280.

Fujii, H. and Managi, S. (2018), 'Trends and priority shifts in artificial intelligence technology invention: A global patent analysis', *Economic Analysis and Policy* **58**, 60–69.

Giczy, A. V., Pairolero, N. A. and Toole, A. A. (2022), 'Identifying artificial intelligence (ai) invention: A novel ai patent dataset', *The Journal of Technology Transfer* **47**(2), 476–505.

Gordon, R. J. (2018), Why has economic growth slowed when innovation appears to be accelerating?, Technical report, National Bureau of Economic Research.

Görtz, C. and Tsoukalas, J. D. (2018), 'Sectoral tfp news shocks', *Economics Letters* **168**, 31–36.

Gross, D. P. and Sampat, B. N. (2023), 'America, jump-started: World war ii r&d and the takeoff of the us innovation system', *American Economic Review* **113**(12), 3323–3356.

Hayo, B. and Mierzwa, S. (2023), 'The effect of legislated tax changes on the trade balance: Empirical evidence for the united states, germany, and the united kingdom', *Journal of Macroeconomics* **78**.

Jaffe, A. B. and Trajtenberg, M. (1999), 'International knowledge flows: Evidence from patent citations', *Economics of innovation and new technology* **8**(1-2), 105–136.

Jaffe, A. B., Trajtenberg, M. and Henderson, R. (1993), 'Geographic localization of knowledge spillovers as evidenced by patent citations', *the Quarterly journal of Economics* **108**(3), 577–598.

Jaimovich, N. and Rebelo, S. (2009), 'Can news about the future drive the business cycle?', *American Economic Review* **99**(4), 1097–1118.

Jamasb, T. and Pollitt, M. G. (2011), 'Electricity sector liberalisation and innovation: An analysis of the uk's patenting activities', *Research Policy* **40**(2), 309–324.

Jinnai, R. (2014), 'R&d shocks, news shocks, and price levels', *Available at SSRN 2439059* .

Jordà, Ò., Singh, S. R. and Taylor, A. M. (2024), 'The long-run effects of monetary policy', *Review of Economics and Statistics* pp. 1–49.

Jordà, O. (2005), 'Estimation and inference of impulse responses by local projections', *American Economic Review* **95**(1), 161–182.

Kamber, G., Theodoridis, K. and Thoenissen, C. (2017), 'News-driven business cycles in small open economies', *Journal of International Economics* **105**, 77–89.

Kelly, B., Papanikolaou, D., Seru, A. and Taddy, M. (2021), 'Measuring technological innovation over the long run', *American Economic Review: Insights* **3**(3), 303–320.

Kilian, L., Plante, M. D. and Richter, A. W. (2023), Estimating macroeconomic news and surprise shocks, Dallas Fed Working Paper 2304, Federal Reserve Bank of Dallas. Revised March 2024.
**URL:** *https://www.dallasfed.org/research/papers/2023/wp2304*

Kogan, L., Papanikolaou, D., Seru, A. and Stoffman, N. (2017), 'Technological innovation, resource allocation, and growth', *The Quarterly Journal of Economics* **132**(2), 665–712.

Kumar, P. and Sharma, M. (2022), 'Data, machine learning, and human domain experts: none is better than their collaboration', *International Journal of Human–Computer Interaction* **38**(14), 1307–1320.

Kurmann, A. and Sims, E. (2021), 'Revisions in utilization-adjusted tfp and robust identification of news shocks', *Review of Economics and Statistics* **103**(2), 216–235.

Lerner, J. and Seru, A. (2021), 'The use and misuse of patent data: Issues for finance and beyond', *The Review of Financial Studies* **35**(6), 2667–2704.

Lybbert, T. J. and Zolas, N. J. (2014), 'Getting patents and economic data to speak to each other: An 'algorithmic links with probabilities' approach for joint analyses of patenting and economic activity', *Research Policy* **43**(3), 530–542.

Manning, C. D., Raghavan, P. and Schütze, H. (2009), *An Introduction to Information Retrieval*, Cambridge University Press, Cambridge.

McCarthy, J. (1997), 'What is artificial intelligence?'.

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.

Miguelez, E. and Moreno, R. (2018), 'Relatedness, external linkages and regional innovation in europe', *Regional studies* **52**(5), 688–701.

Minsky, M. and Papert, S. (1969), *Perceptrons: An Introduction to Computational Geometry*, MIT Press.

Miranda-Agrippino, S. (2023), Unsurprising shocks: information, premia, and the monetary transmission, Technical report, Bank of England Staff Working Paper No. 626.

Miric, M., Jia, N. and Huang, K. G. (2023), 'Using supervised machine learning for large-scale classification in management research: The case for identifying artificial intelligence patents', *Strategic Management Journal* **44**(2), 491–519.

Mouratidis, K., Papapanagiotou, G. and Thoenissen, C. (2024), 'Uk productivity: the long and the short of it', *The Sheffield Economic Research Paper Series (SERPS)* **2024005**(2024005).

Murphy, K. P. (2012), *Machine Learning: A Probabilistic Perspective*, MIT Press, Cambridge, MA.

Negro, M. D. and Schorfheide, F. (2011), Bayesian macroeconometrics, *in* G. K. H. van Dijk and J. Geweke, eds, 'Handbook of Bayesian Econometrics', Oxford University Press, pp. 293–389.

Neville Francis, Michael T. Owyang, J. E. R. and DiCecio, R. (2014), ‘A flexible finite-horizon alternative to long-run restrictions with an application to technology shocks’, *The Review of Economics and Statistics* **96**(4), 638–647.

OECD (2022), ‘Oecd framework for the classification of ai systems’, *OECD Digital Economy Papers, No. 323* .

Plagborg-Møller, M. and Wolf, C. K. (2021), ‘Local projections and vars estimate the same impulse responses’, *Econometrica* **89**(2), 955–980.

Qiu, Z. and Wang, Z. (2023), ‘Technological origination and evolution analysis by combining patent claims and citations: A case of surgical robot domain’, *Advanced Engineering Informatics* **58**, 102145.

Shea, J. (1999), What do technology shocks do?, *in* B. S. Bernanke and J. J. Rotemberg, eds, ‘NBER Macroeconomics Annual 1998, Volume 13’, MIT Press, Cambridge, MA, pp. 275–322.

Sims, C. A. (1986), ‘Are forecasting models usable for policy analysis?’, *Federal Reserve Bank of Minneapolis Quarterly Review* **Winter**, 2–16.

Stock, J. H. and Watson, M. W. (2012), ‘Disentangling the channels of the 2007-09 recession’, *Brookings Papers on Economic Activity* **44**(1), 81–156.

Stock, J. H. and Watson, M. W. (2018), ‘Identification and estimation of dynamic causal effects in macroeconomics using external instruments’, *The Economic Journal* **128**(610), 917–948.

Straccamore, M., Bruno, M., Monechi, B. and Loreto, V. (2023), ‘Urban economic fitness and complexity from patent data’, *Scientific Reports* **13**(1), 3655.

Syverson, C. (2017), ‘Challenges to mismeasurement explanations for the us productivity slowdown’, *Journal of Economic Perspectives* **31**(2), 165–186.

Tong, X. and Frame, J. D. (1994), ‘Measuring national technological performance with patent claims data’, *Research Policy* **23**(2), 133–141.

Uhlig, H. (2005), 'What are the effects of monetary policy on output? results from an agnostic identification procedure', *Journal of Monetary Economics* **52**, 381–419.

UNCTAD (2025), Chapter I: AI at the technology frontier, *in* 'Technology and Innovation Report 2025: Inclusive Artificial Intelligence for Development', UNCTAD, pp. 1–30.
**URL:** *https://unctad.org/system/files/official-document/tir2025ch1˙en.pdf*

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, and Polosukhin, I. (2017), Attention is all you need, *in* 'Advances in Neural Information Processing Systems (NeurIPS)'.

Verendel, V. (2023), 'Tracking artificial intelligence in climate inventions with patent data', *Nature Climate Change* **13**(1), 40–47.

Vukotić, M. (2019), 'Sectoral effects of news shocks', *Oxford Bulletin of Economics and Statistics* **81**(1), 215–249.

Williams, S., Glass, A. J., Matos, M., Elder, T. and Arnett, D. (2025), 'The uk productivity puzzle: A survey of the literature and expert views', *International Journal of the Economics of Business* **32**(1), 31–65.

Winfield, A. (2020), 'Intelligence is not one thing', *Journal of Artificial General Intelligence* **11**(2), 97–100.

Xiao, D., Kokas, A. and Zhang, J. (2023), 'The emergence and diffusion of artificial intelligence', *Research Policy* **52**(6).

Yang, C.-H. (2022), 'How artificial intelligence technology affects productivity and employment: Firm-level evidence from taiwan', *Research Policy* **51**(6), 104536.

# A Definitions of AI components

This paper considers eight broad components of AI technology following Giczy et al. (2022): machine learning, AI hardware, natural language processing, speech recognition, computer vision, knowledge processing, planning/control, and evolutionary computation. A patent may belong to more than one of these categories. For example, an invention in computer vision could utilize a machine learning method. Below, we provide a definition for each AI technology.

1. Machine learning: Machine learning is a field concerned with the development and study of statistical algorithms that can learn patterns from data and generalize to previously unseen observations, enabling them to perform tasks without explicit programming. Examples of machine learning models include decision trees, support vector machines, and neural networks.

2. AI hardware: This field refers to computer hardware specifically designed to execute AI algorithms efficiently. For example, graphics processing units (GPUs) are commonly used to train deep neural networks.

3. Natural Language Processing (NLP): NLP is a subfield of AI concerned with enabling computers to recognize, understand, and generate human language in both text and speech. NLP is closely related to linguistics and draws on methods from computational linguistics to model language. One of the earliest NLP programs is ELIZA, a computer program developed to explore communication between humans and machines.

4. Speech recognition: Speech recognition focuses on the development of methodologies and techniques that enable the conversion of human speech into text. For example, the Hidden Markov Model (HMM) represents speech as a sequence of hidden states corresponding to phonemes and estimates the most likely word sequences using dynamic programming algorithms.

5. Computer vision: Computer vision is a subfield of AI that develops methods enabling computers and systems to interpret and process information from visual inputs. For example, blob detection can be used to identify neighboring regions in a digital image that differ in properties such as color.

6. Knowledge processing: This is a subfield of AI concerned with representing information in a form that enables a system to deduce new knowledge. Specialized hardware and software are often required to process this information effectively. An inference engine is an example of software commonly used in knowledge processing.

7. Planning and control: Planning and control refers to AI technologies focused on the design and execution of strategies to achieve specified goals. Given an initial state, a desired outcome, and a set of possible actions, the planning problem involves determining a sequence of actions that ensures the desired outcome. Combinatorial optimization techniques are often employed, particularly in environments that are dynamically uncertain.

8. Evolutionary computation: Evolutionary computation is a family of algorithms for global optimization that are inspired by biological evolution. An initial set of potential solutions is defined and iteratively updated. In each iteration, less desirable solutions are discarded, while small modifications are introduced to the remaining candidates. Evolutionary algorithms include, among others, genetic algorithms and programming, differential evolution, and neuroevolution.

# B Manual Classification Example

Figure A1 presents two representative examples of both AI-related and non-AI-related patents, illustrating the criteria employed in our manual classification process.

In the *Non-AI* case (left; GB2429930), the independent claim describes a program that detects sender IDs via a wireless interface and triggers predetermined game actions. The inventive step concerns protocol-driven communication and deterministic control flow: conditional on receipt of an identifier, the system executes fixed routines. There is clearly no estimation or prediction from data (no learning, perception, or planning) and no specialized AI hardware. In our terms, this is an information transmission rather than information extraction; accordingly, it does not map to any of the eight AI components and is classified as non-AI.

By contrast, the AI example (right; GB2207049) discloses biometric authentication using an optical-flow sensor configured to capture time-varying skin motion, optionally fused with EMG/proximity signals, to detect facial muscle activity. Implementing this claim requires extracting spatiotemporal features from raw sensor data and making a decision about an unobserved state (muscle activation) for authentication. Consequently, this is a measurement, inference pipeline, information extraction and prediction, which maps to computer vision and knowledge processing. Consistent with our framework, we classify this patent as AI-related.

Figure A1: Examples of AI and Non-AI Patent Classification

**Non-AI Classification**

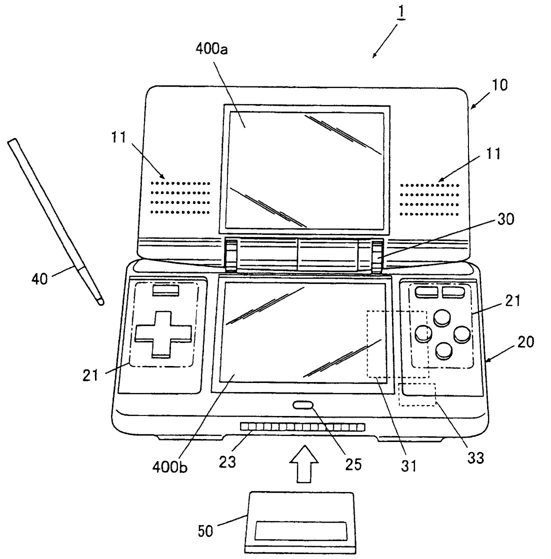


**Non-AI Example:**
GB2429930: Wireless gaming character generator
*Claims:* I. A program for causing a portable computer provided with wireless communication section which performs wireless communication according to a specific communication standard to execute a game in which a player enjoys collecting characters, the program causing the computer to execute: a reception detection step of detecting that the wireless communication section has received a sender ID transmitted from another communication instrument which is located in a current wireless communication area and performs wireless communication ...

**AI Classification**

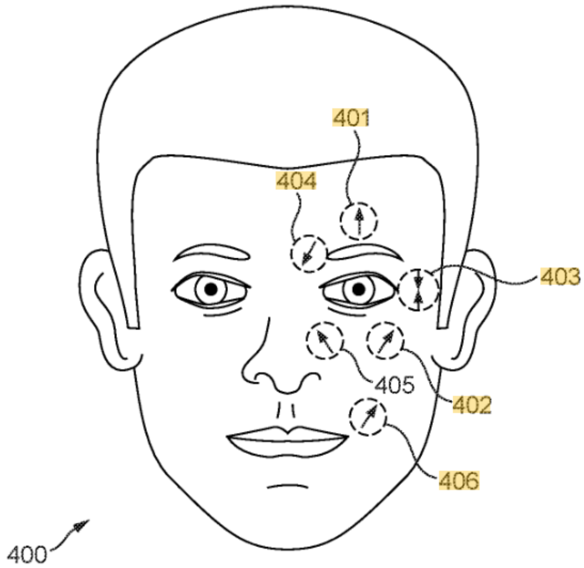


**AI Example:**
GB2207049: User authentication using biometric data, e.g. fingerprints, iris scans or voiceprints.
*Claims:* A system for detecting facial muscle activity comprising: wearable apparatus comprising; an optical flow sensor located so as to, in use when the apparatus is worn, image an area of skin associated with one or more facial muscles, the optical flow sensor being configured to capture time-varying data describing movement of the respective area of skin substantially in the plane of the imaged area of skin; and an EMG sensor and/or a proximity sensor arranged to detect the activity of a first muscle along with the optical flow sensor; and...

*Note:* Patent examples and their images were retrieved from EPO's Espacenet online query database and correspond to observations in our classification sample. The images are only used for illustratory purposes in this study and were not employed in the classification process which relies only on textual inputs (abstracts, descriptions, claims, etc).

# C Instrument Construction

Table A1: Regression results based on Equation 1

| **Variable** | **Baseline** (1) | **Policy shocks** (2) | $\mathbb{E}_t$ + **Policy shocks** (3) | $\mathbb{E}_t$ + **Broader tech** (4) | $\mathbb{E}_t$ + **Public R&D** (5) |
|---|---|---|---|---|---|
| Intercept | -12.662*** (0.007) | 1.552* (0.0569) | -4.262 (0.554) | -3.491 (3.465) | 34.314 (0.334) |
| $\text{AI}_{t-1}$ | -0.264*** (0.000) | -0.474*** (0.0000) | -0.618*** (0.000) | -0.502*** (0.110) | -0.651*** (0.000) |
| $\text{AI}_{t-2}$ | -0.042 (0.645) | -0.388*** (0.0040) | -0.559*** (0.000) | -0.366*** (0.086) | -0.520*** (0.000) |
| $\text{AI}_{t-3}$ | -0.343*** (0.002) | -0.180 (0.1519) | -0.224** (0.013) | -0.042 (0.091) | -0.148 (0.144) |
| $\text{AI}_{t-4}$ | -0.413*** (0.000) | -0.038 (0.7618) | -0.022 (0.865) | 0.008 (0.068) | 0.002 (0.980) |
| $\mathbb{E}_t[\pi_{t+1}]$ | 0.275 (0.779) | | 0.057 (0.974) | -0.607 (1.117) | -0.327 (0.731) |
| $\mathbb{E}_t[\pi_{t+4}]$ | 0.569 (0.564) | | 3.306 (0.146) | 0.939 (1.044) | 1.096 (0.305) |
| $\mathbb{E}_t[\pi_{t+8}]$ | 6.539*** (0.007) | | 1.889 (0.528) | 2.498* (1.496) | 2.639 (0.262) |
| $\mathbb{E}_t[Y_{t+1}]$ | -0.847 (0.213) | | -0.541 (0.483) | -0.065 (0.452) | -0.377 (0.491) |
| $\mathbb{E}_t[Y_{t+4}]$ | 1.090** (0.043) | | 0.508 (0.703) | -0.486 (0.536) | 0.738 (0.332) |
| $\mathbb{E}_t[Y_{t+8}]$ | -0.612 (0.507) | | -2.192 (0.236) | 0.102 (0.925) | -1.872 (0.179) |
| monetary | | -17.128 (0.5158) | -10.611 (0.736) | | |
| $\text{monetary}_{t-1}$ | | 15.757 (0.6657) | -4.212 (0.893) | | |
| $\text{monetary}_{t-2}$ | | -10.720 (0.7386) | -33.053 (0.254) | | |
| fiscal | | 0.431** (0.0464) | 0.731** (0.014) | | |
| $\text{fiscal}_{t-1}$ | | 0.012 (0.9759) | 0.243 (0.438) | | |
| $\text{fiscal}_{t-2}$ | | 0.627 (0.1596) | 0.790** (0.048) | | |
| $\text{Non-AI}_t$ | | | | 0.943*** (0.147) | |
| $\text{Non-AI}_{t-1}$ | | | | 0.437* (0.229) | |
| $\text{Non-AI}_{t-2}$ | | | | 0.256* (0.137) | |
| Public R&D | | | | | -0.610 (0.613) |
| Public $\text{R\&D}_{t-1}$ | | | | | 0.859 (0.701) |
| Public $\text{R\&D}_{t-2}$ | | | | | -0.289 (0.788) |
| F-statistic | 6.402*** | 2.069*** | 2.284*** | 8.765*** | 3.988*** |
| Adjusted $R^2$ | 0.348 | 0.134 | 0.229 | 0.499 | 0.277 |
| Observations | 102 | 70 | 70 | 102 | 102 |
| *Wald tests for joint significance of controls* | | | | | |
| BoE forecasts | 6.120 (0.000) | | | | |
| Policy shocks | | 2.051 (0.045) | | | |
| BoE forecasts + policy shocks | | | 5.299 (0.000) | | |
| BoE forecasts + broader tech | | | | 12.001 (0.000) | |
| BoE forecasts + public R&D | | | | | 2.174 (0.031) |

*Notes:* Regression results are based on Eq.(1). The dependent variable is $AIPatents_t$, the annual growth rate of all AI patent applications. BoE forecasts are for the two core macroeconomic variables included in $x_t$, namely GDP growth $Y_t$ and CPI inflation $\pi_t$. Policy shocks include narrative *monetary* policy and narrative tax changes (*fiscal*). Additional controls include broader technological developments, as captured by the *non-AI* measure, and *public R&D* expenditures. Robust standard errors are reported in parentheses. The lower panel reports Wald test statistics for the joint significance of the controls, with associated p-values in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

# D Granger causality tests

Table A2: Instrument: Granger causality tests

| | Wald Test | p-value | Adjusted $R^2$ | $N$ |
|---|---|---|---|---|
| $E[\pi_{T+1}], E[Y_{T+1}]$ | 0.094 | 0.910 | -0.017 | 98 |
| $E[\pi_{T+4}], E[Y_{T+4}]$ | 0.039 | 0.961 | -0.018 | 98 |
| $E[\pi_{T+8}], E[Y_{T+8}]$ | 0.021 | 0.978 | -0.018 | 98 |

*Notes*: The dependent variable is the residual of Eq.(1). The numbers reported are Wald test statistics for joint significance of the BoE's forecasts of GDP and CPI, at different horizons. All the regressions include 4 lags of the dependent variable and a constant.

# E TFP response to a non-AI technological news shock

Figure A2: TFP response to a broad technological news shock that excludes AI patents.

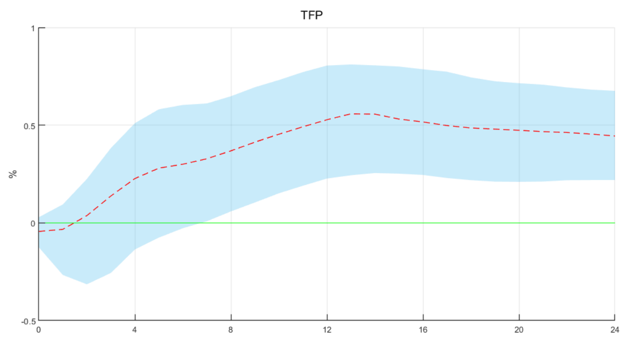


*Note:* The figure shows the response of TFP to a one standard deviation broad technological patent-based news shock that excludes AI. The median response is depicted as dotted red line, with shaded areas representing 68% posterior bands. The horizon is 24 quarters (6 years).

# F IRFs with Different Instruments

Figure A3: IRFs with alternative IV proxies for monetary and fiscal shocks.

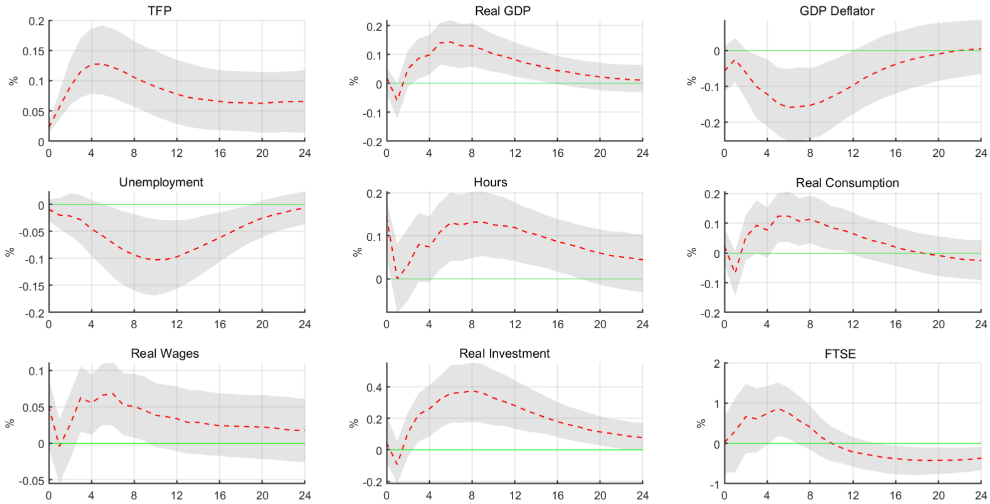


*Note:* The figure reports IRFs identified with an IV that controls for policy shocks (fiscal and monetary shocks). Estimation sample as in baseline, IV sample 1997Q3-2015Q1).

Figure A4: IRFs with alternative IV proxy for broader technology news shocks.

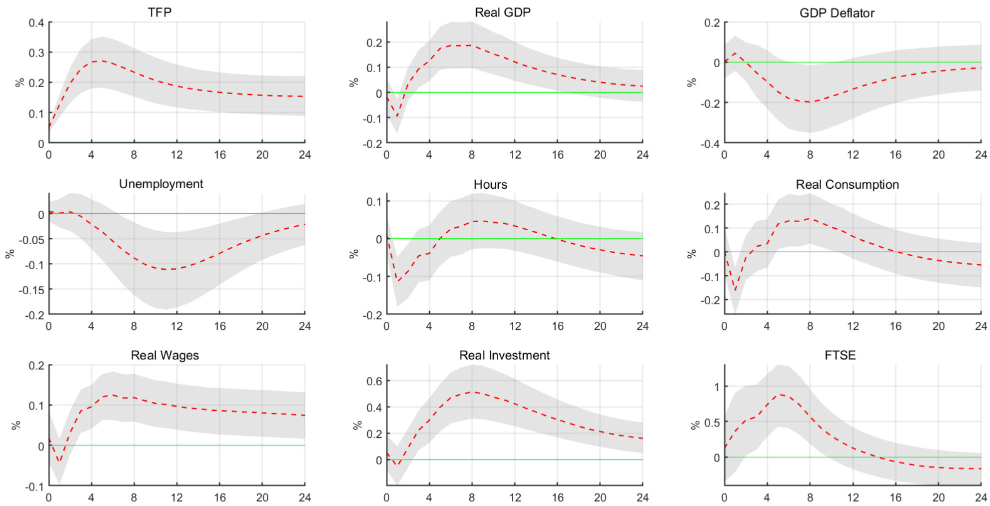


*Note:* The figure reports IRFs identified with an IV that controls for broader technology news shocks. Estimation sample as in baseline, IV sample 1997Q3-2022Q4.

Figure A5: IRFs with alternative IV proxy for public R&D allocation.

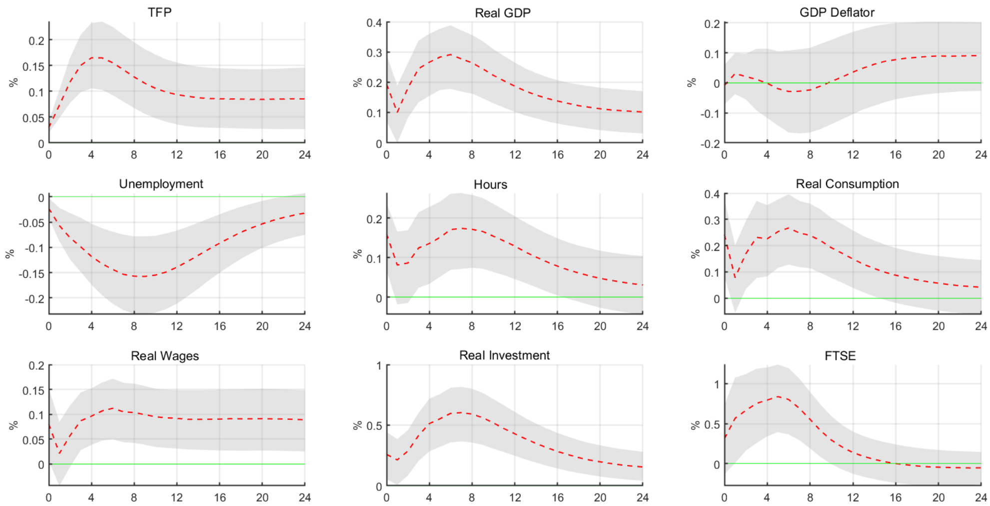


*Note:* The figure reports IRFs identified with an IV that controls for public R&D allocation. Estimation sample as in baseline, IV sample 1997Q3-2022Q4.

# G IRFs based on the weighted AI patent measure

Figure A6: IRFs with alternative weighted AI patents.

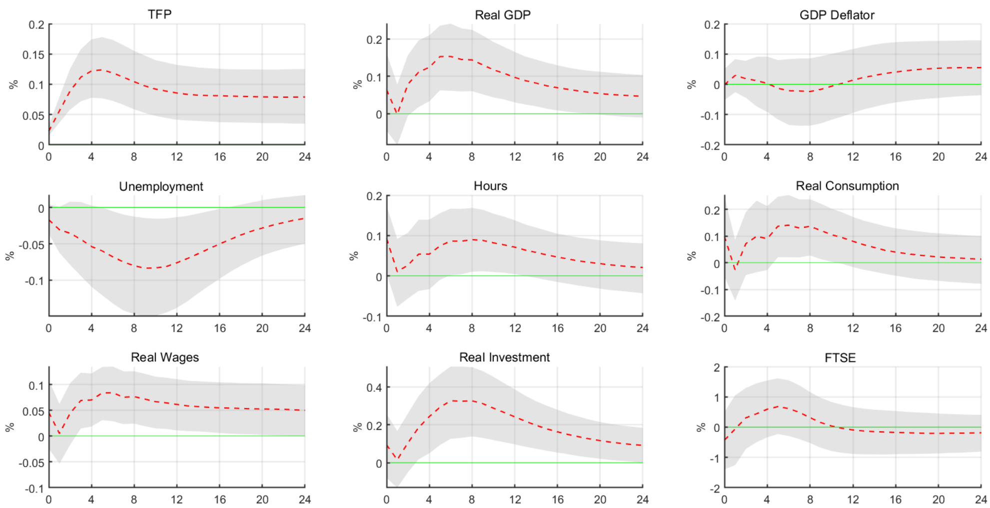


*Note:* The figure reports IRFs based on AI patent applications weighted by their forward citation counts. Estimation sample as in baseline, IV sample 1997Q3-2022Q4.

# H Additional Robustness Chesks

Figure A7 reports IRFs estimated over a sample that excludes the 2008 financial crisis (estimation sample: 1971Q1–2007Q4). The IRFs are broadly similar in the short and medium run, though some differences emerge. The impact response of hours, output, and investment are muted, while consumption and real wages decline. The subsequent fall in inflation likely reflects this delayed response of output, consumption, and investment. These differences may be attributable to the limited VAR sample, as well as the IV spanning only 1997Q3–2007Q4, compared with the baseline, which may reduce the precision of the estimated responses. Accordingly, caution is warranted when comparing these results with the baseline.

Additionally, we run the following three alternative specifications. First, Figure A8 reports IRFs for a proxy SVAR that includes CPI inflation instead of the GDP deflator in (log) levels. Estimation and identification samples are the same as in the baseline. The responses to an AIPNS are largely unchanged.

The same holds for an alternative specification in which total employment, sourced from the ONS, replaces the unemployment rate, shown in Figure A9. Employment responds positively and significantly to the AI news shock, and the remaining responses remain largely unchanged, reinforcing the baseline implications that the productivity-enhancing effects of AI dominate its potential displacement effects.

Figure A10 reports the IRFs from another VAR specification where TFP is replaced by output per worker (or average labor productivity) taken by the ONS, representing an alternative measure of productivity (Christiansen, 2008; Neville Francis and DiCecio, 2014; Acemoglu and Restrepo, 2018*b*). The results are very similar, with all the responses qualitatively matching and, in most cases, showing quantitative similarities, with those obtained in the baseline, thereby confirming the robustness of our findings when using the alternative productivity measure. One quantitative difference is the larger initial response of labor productivity relative to TFP, which then gradually increases to a permanently higher level that is about twice as high as that for TFP. This could potentially reflect capital deepening, which is consistent with the immediate rise in investment following the AI news shock. This interpretation also aligns with Kurmann and Sims (2021), who find that labor productivity increases considerably more than TFP on impact, due to short-run capital deepening effects.

Last, we run an alternative specification where our constructed IV does not control for expectations regarding CPI and GDP for 8 quarters ahead; specifically, forecasts for CPI and GDP are included for horizons of 1 and 4 quarters ahead. The IRFs in A11 are almost identical to the baseline. We also perform several robustness checks focusing on the textual analysis stages used to construct our AI patent-based news series. Results, not reported here for brevity but available upon request, show that the IRFs remain very similar to our baseline findings across all cases.[57]

Figure A7: IRFs excluding the 2008 financial crisis.

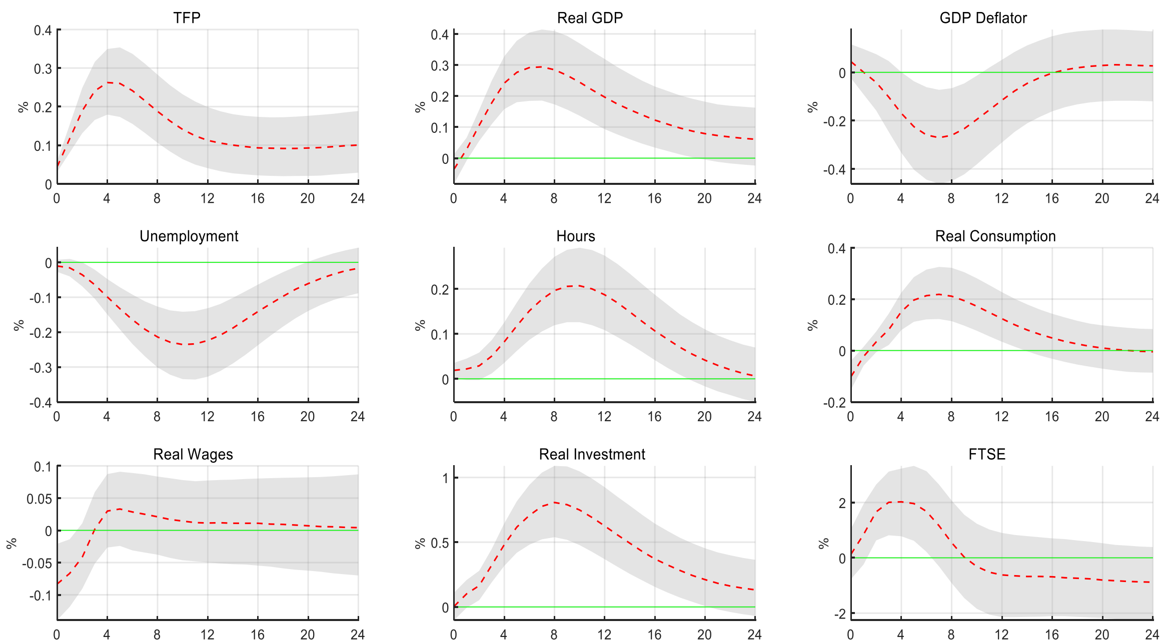


*Note:* The figure shows IRFs estimated over a pre-crisis sample that excludes the 2008 financial crisis (estimation sample 1971Q1 -2007Q4).

[57]Specifically, we employ the K-Nearest Neighbors (KNN) algorithm as an alternative classifier. Unlike the NBC, KNN does not assume feature independence and can capture more complex relationships, though at the cost of potentially excluding patents that are AI-related but not sufficiently similar to the training data. Second, we estimate an alternative BVAR using AI patent predictions derived from different feature selection criteria. In contrast to our baseline mutual information criterion, we test the Gini coefficient, gain-based weighting, and symmetrical uncertainty, each emphasizing different aspects of feature relevance. Third, we adopt a more conservative NBC classification by raising the cutoff ratio from 1 to 1.2. Finally, we experiment with alternative text-processing choices, such as using Term Frequency (TF) instead of TF-IDF weighting and varying the number of top features selected (40 and 130 features, versus the baseline of 105).

Figure A8: IRFs using CPI inflation instead of GDP deflator.

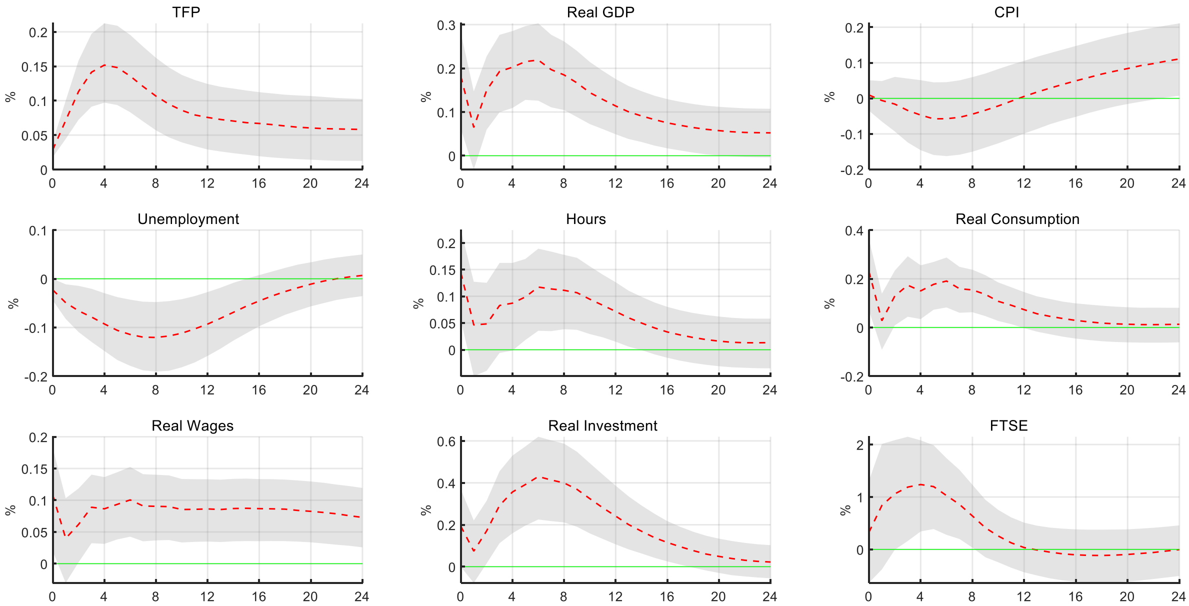


*Note:* The figure reports IRFs that uses CPI inflation instead of GDP deflator.

Figure A9: IRFs with employment.

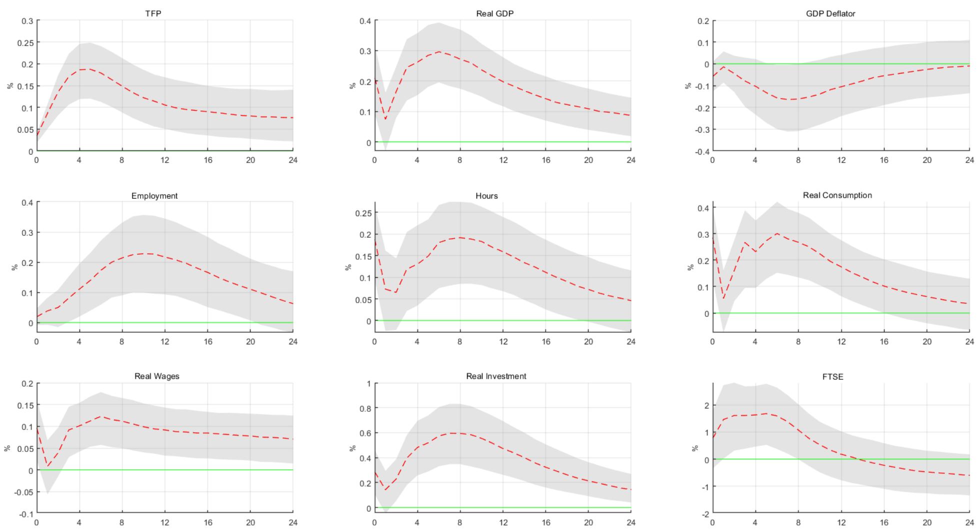


*Note:* The figure shows IRFs using the log of employment in place of the unemployment rate.

Figure A10: IRFs with alternative productivity measure (output per worker).

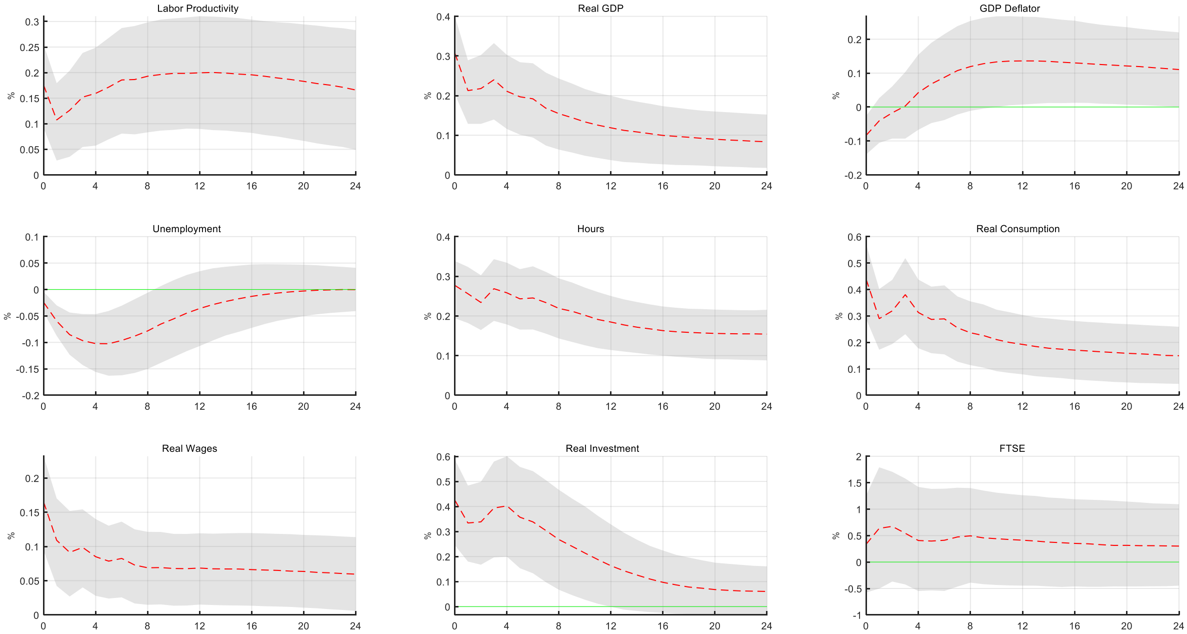


*Note:* The figure shows IRFs using output per worker as an alternative productivity measure.

Figure A11: IRFs excluding long-horizon expectations in IV construction.

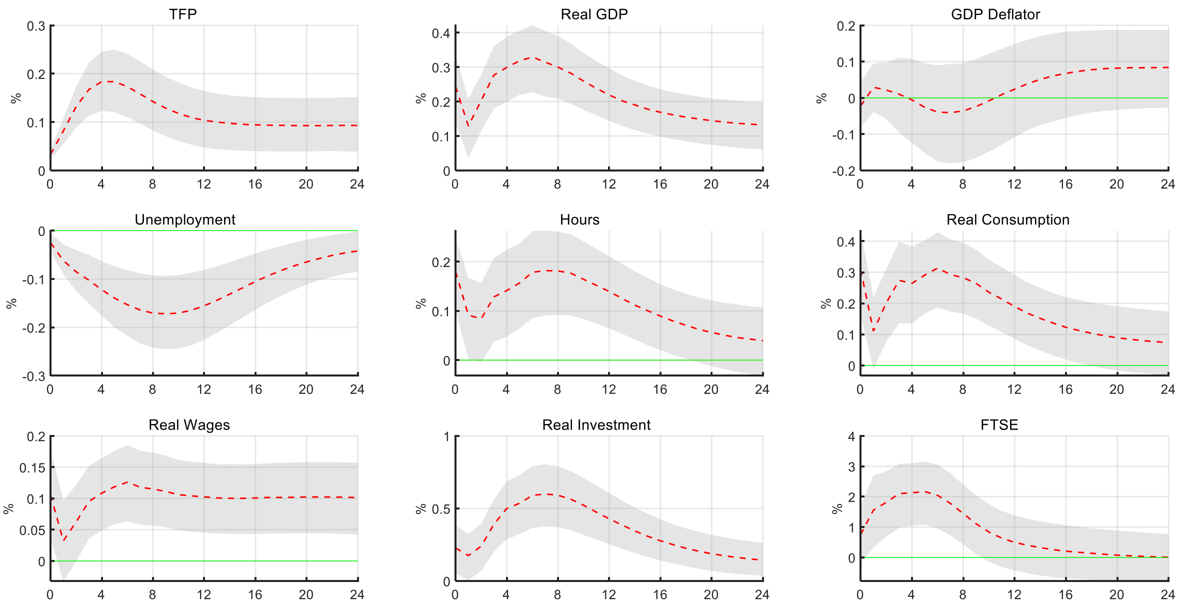


*Note:* The figure reports IRFs identified with an IV that does not control for beliefs regarding CPI and GDP for 8 quarters ahead (i.e. forecast of CPI and GDP are included for $h = 1, 4$ quarters ahead.

# I AI Sectoral Break-down

Table A3: Distribution of AI patents across industries in the UK

| Industry | NACE | Share of AI patents (%) |
|---|---|---|
| Manufacturing | C | 59.0 |
| Information and communication | J | 25.3 |
| Professional, scientific, and technical activities | M | 15.2 |
| Wholesale and retail trade | G | 0.31 |
| Electricity, gas, steam and air conditioning supply | D | 0.16 |
| Education | P | 0.03 |
| Transportation and Storage | H | 0.02 |
| Construction | F | 0.02 |
| Other | – | 1.1 |

*Note:* We report industries with at least one AI-related patent. Industries follow NACE Rev. 2 sections (codes A–U).