BANCA D'ITALIA
EUROSISTEMA

Temi di discussione

(Working Papers)

The macroeconomic effects of AI technology shocks

by Andrea Gazzani and Filippo Natoli

Number 1542

July 2026

BANCA D'ITALIA
EUROSISTEMA

Temi di discussione

(Working Papers)

The macroeconomic effects of AI technology shocks

by Andrea Gazzani and Filippo Natoli

Number 1542 - July 2026

BANCA D'ITALIA

EUROSISTEMA

ChaMP

This paper contains research conducted within the network "Challenges for Monetary Policy Transmission in a Changing World Network" (ChaMP). It consists of economists from the European Central Bank (ECB) and the national central banks (NCBs) of the European System of Central Banks (ESCB).

ChaMP is coordinated by a team chaired by Philipp Hartmann (ECB), and consisting of Diana Bonfim (Banco de Portugal), Margherita Bottero (Banca d'Italia), Emmanuel Dhyne (Nationale Bank van België/Banque Nationale de Belgique) and Maria T. Valderrama (Oesterreichische Nationalbank), who are supported by Gonzalo Paz-Pardo and Jean-David Sigaux (both ECB), 7 central bank advisers and 8 academic consultants.

ChaMP seeks to revisit our knowledge of monetary transmission channels in the euro area in the context of unprecedented shocks, multiple ongoing structural changes and the extension of the monetary policy toolkit over the last decade and a half as well as the recent steep inflation wave and its reversal. More information is provided on its website.

ISSN 2281-3950 (online)

Designed by the Printing and Publishing Division of the Bank of Italy

THE MACROECONOMIC EFFECTS OF AI TECHNOLOGY SHOCKS

by Andrea Gazzani* and Filippo Natoli**

Abstract

Using detailed data on the artificial intelligence (AI) content of U.S. patents from 1980 to 2019, we construct a novel monthly measure of AI intensity in innovation and identify AI related technology shocks. These shocks generate delayed increases in total factor productivity, output, employment, hours, and wages, alongside persistent declines in consumer prices, consistent with a positive supply shock. Although AI-intensive patents largely fall within the broader ICT and automation domains, they stand out as higher-quality innovations, attracting more citations and exhibiting greater technological and market value. AI shocks consistently have substantially larger aggregate effects than broader ICT or automation shocks, suggesting that AI has had a particularly high-impact role in the ICT revolution. Unlike general technology shocks, however, AI shocks reduce the labor share and increase wealth inequality, indicating that their gains are not distributed evenly across the economy. To connect these historical findings to the most recent wave of innovation, we develop a new time series of Generative-AI patents within the U.S. patent universe. A shock based exclusively on Gen-AI patents produces the same qualitative supply-side effects as broader AI shocks, linking the macroeconomic consequences of earlier AI innovation to those of the emerging Generative-AI era.

JEL Classification: E32, C32, C36, O33, O34.

Keywords: artificial intelligence, technology shocks, local projections, business cycle, patents.

DOI: 10.32057/0.TD.2026.1542

* Bank of Italy, Economics, Statistics and Research Department; andreagiovanni.gazzani@bancaditalia.it

** Bank of Italy, Economics, Statistics and Research Department; filippo.natoli@bancaditalia.it

1 Introduction[1]

Artificial intelligence (AI), famously defined by John McCarthy as "the science of making intelligent machines" (McCarthy, 2007), has recently moved to the center of economic and policy debates following breakthroughs in generative AI and large language models. For macroeconomists, these developments recast a classic question about how technology shocks affect aggregate productivity, output, prices, and the labor market over time. A large empirical literature has studied technology shocks using aggregate time series (Galí, 1999; Fisher, 2006; Beaudry and Portier, 2006; Fernald, 2012; Miranda-Agrippino et al., 2025), but much less is known about the macroeconomic effects of specific technological waves such as AI. Existing studies often compare AI with previous waves of innovation, including robotics and earlier digital technologies, emphasizing both similarities and structural differences (Aghion et al., 2017). While a growing body of evidence studies AI adoption and exposure at the firm, occupation, or task level, there is still limited evidence on whether AI generates aggregate disturbances that resemble standard technology shocks, and whether the gains from such disturbances are distributed evenly across workers, firms, and households.

We contribute by studying the aggregate implications of past AI innovation through the lens of patenting activity. Using detailed data on the AI content of every U.S. patent since 1980 (Giczy et al., 2022), we construct a novel monthly aggregate measure of AI intensity in U.S. innovation. This measure is designed to capture distinctly AI-related developments, rather than general technological progress or broad sectoral trends. Examples of AI patents include systems for detecting defects and anomalies in industrial and medical images, predicting outcomes from large datasets (e.g., online consumer preferences or oil reserves), automating fraud detection and credit-risk assessment in financial services, improving firms' customer support and optimizing the operation of robots and autonomous vehicles.

To assess the macroeconomic effects of AI technology shocks, we use our measure of AI

[1]We wish to thank Francesco Bianchi, Dario Caldara, Luisa Carpinelli, Francesco Filippucci, Stefano Eusepi, Jorgo Georgiadis, Mishel Ghassibe, Silvia Miranda-Agrippino, Joseba Martinez, Pablo Ottonello, Ayşegül Şahin, Paolo Surico, and Giovanni Veronese for fruitful discussions. This paper also benefited from feedback in seminars at the Bank of Italy, EUI, ECB, the ECB - Banco de España conference "The macroeconomics effects of AI", the 12th IMF Statistical Forum, the BI Norwegian Business School, Bank of Italy and Bank of Canada workshop "Applied macroeconomics in a changing world", and the SUERF - Bundesbank "AI and the Future of Central Banking". The views expressed in this paper do not necessarily reflect those of the Bank of Italy or ESCB. First Draft: July 14, 2024.

intensity in U.S. innovation within an instrumental variables (IV) framework. Patenting activity may respond endogenously to current or expected business cycle conditions, potentially affecting both the level of innovation and its composition toward or away from AI. To address this concern, we construct an instrument for AI technology shocks following the logic in Miranda-Agrippino et al. (2025). The instrument isolates variation in AI innovation that is orthogonal to current and expected business cycle fluctuations and other macroeconomic disturbances. We then estimate the dynamic effects of AI technology shocks using a local projection instrumental variable (LP-IV) approach. Since the LP-IV responses are very similar to those obtained from a standard OLS LP with a rich set of lagged controls, we use the OLS specification for quantitative decompositions and for exercises that exploit more granular variation across patents, sectors, and technology classes.

Our analysis shows that increases in the AI intensity in innovation act as positive supply shocks. They generate a delayed but persistent boost in economic activity and total factor productivity (TFP), while reducing consumer prices. Total employment also rises persistently while hours worked rise only temporarily, suggesting that, over the medium run, labor market adjustment occurs primarily through job creation. AI technology shocks are also economically meaningful: in our sample, they account for about 25% of GDP and employment fluctuations over a five-year horizon. Nonetheless, the fact that they are not the primary driver of business cycle fluctuations serves as a reassuring sanity check.

A distinguishing feature of patents that rely heavily on artificial intelligence — which we label *AI-intensive* patents — is their systematically higher quality. This pattern holds relative to the average U.S. patent and, crucially, relative to patents in the broader ICT and automation categories to which they largely belong. AI-intensive patents receive more forward citations, exhibit higher technological importance (Kelly et al., 2021) and estimated economic value (Kogan et al., 2017). Consistent with this micro-level evidence, AI technology shocks generate stronger effects on GDP, employment, and TFP than shocks associated with ICT and automation more broadly. Taken together, these findings suggest that AI-driven innovation has been an important source of ICT-related productivity gains over the past decades.

The transmission of AI technology shocks partly resembles that of general technology shocks (Miranda-Agrippino et al., 2025), although it differs in one key respect: its distributional consequences. Whereas general technological shocks are typically found to be labor-neutral (e.g., Fisher, 2006), AI technology shocks reduce the labor share. Thus, despite generating aggregate gains in employment and wages, AI-driven innovations appear to shift income toward capital rather than labor. This result squares with the fact that the gains from AI innovation are not distributed equally across households. Using monthly data on income and wealth shares across different percentiles of the US household distribution, we find that expansionary AI technology shocks increase the wealth share of the top 10% and reduce that of the bottom 50%. These distributional consequences are consistent with the asset returns channel (Moll et al., 2022). They are also visible across different educational levels: the employment gains are widespread but do not extend to workers with less than a high school education, a pattern consistent with the concept of skill-biased technological change.

Our results on employment speak directly to the recent debate on the medium-term effects of AI on the labor market. We first document that AI technology shocks lead to substantial labor-market reallocation: layoffs surge, but this is more than offset by new job openings. We analyze this point further by means of a sectoral analysis that sheds light on whether innovation based on AI might have differential effects across industries. We find that employment gains are broad-based, occurring even in sectors most exposed to AI, consistent with Aghion et al. (2025b). This suggests that productivity gains and general equilibrium effects outweigh displacement effects. Our time-series approach is particularly well suited to capture this mechanism as it avoids the missing-intercept problem that can arise in micro-estimates (Sarto, 2018; Wolf, 2023; Matthes et al., 2024).

Our baseline evidence relies on a broad historical measure of AI patenting, which captures the evolution of AI well before the recent rise of generative AI and includes technologies beyond those directly related to GenAI. To establish a closer link between our historical evidence and the current technological wave, whose full macroeconomic consequences have yet to materialize, we complement our baseline measure with a narrower series that isolates patents directly related to Generative AI. We identify these patents through a text- and classification-based search of the U.S. patent corpus, following the spirit of the WIPO taxonomy of Generative-AI technologies. The resulting patent set is validated against landmark Generative-AI inventions. Admittedly, our identification strategy is deliberately conservative: while it effectively isolates recent Gen-AI patents, including the foundational patent behind modern large language models such as ChatGPT and Claude (see Appendix A.3 for details), it may exclude some earlier contributions that are not easily identifiable through existing classifications or textual markers. Nonetheless, it provides, to the best of our knowledge, the first time series of Generative-AI patents within the universe of U.S. patents. As expected, the number of Gen-AI patent filings remains close to zero throughout most of the sample and rises sharply from the second half of the 2010s onward, consistent with the emergence of modern Generative-AI architectures.

A shock to the share of Generative-AI patents produces the same qualitative pattern as our baseline AI shock. GDP, employment, total factor productivity, stock prices, and wages increase with a delay, while consumer prices decline over the medium run. The similarity between the responses to Generative-AI and earlier AI innovations conveys an important message. Despite its much broader visibility and public adoption, Generative AI appears to propagate through the economy in a manner remarkably similar to previous generations of AI technologies. In this sense, Generative AI – at least at the stage of development captured by our sample – appears more as an incremental evolution than a macroeconomic discontinuity. Rather, it represents the latest phase of a technological trajectory through which AI, often operating behind the scenes, has already played a significant role in the broader ICT revolution.

Related Literature. Our paper contributes first to the empirical macroeconomic literature that studies technology shocks and their dynamic effects using aggregate time-series methods. This literature has developed a range of approaches to identify technological disturbances and trace their effects on productivity, output, prices, and labor-market outcomes, including Galí (1999), Fisher (2006), Beaudry and Portier (2006), Fernald (2012), and Miranda-Agrippino et al. (2025). We are particularly close to work that uses patent-based information to measure technological change and identify innovation shocks, such as Cascaldi-Garcia and Vukotić (2022), Miranda-Agrippino et al. (2025), Ferriani et al. (2024), and Gazzani et al. (2025).

Relative to this literature, we construct a monthly patent-based measure of AI intensity in U.S. innovation and use it to estimate the aggregate effects of AI-specific technology shocks.

The paper also relates to recent advances in empirical macroeconomic methods for estimating dynamic responses to shocks. We use local projections, following Jordà (2005), because the diffusion of AI technologies is likely to involve long and variable lags. Our identification strategy combines an external-instrument approach with a conditional local-projection specification, and we assess whether the implied shocks are orthogonal to standard macroeconomic shocks and professional forecasts using sufficient-information tests in the spirit of Forni and Gambetti (2014). This places the analysis within the broader time-series literature on shock proxies, local projections, and the propagation of medium-run technological disturbances.

We then speak to the growing literature on the macroeconomic effects of artificial intelligence. Existing work offers sharply different views on the aggregate implications of AI: some contributions emphasize limited productivity gains and adverse distributional effects (Acemoglu, 2024), whereas others highlight the potential for large productivity improvements, new tasks, and broad-based labor market gains (Aghion et al., 2017, 2025b). Model-based and policy-oriented studies also analyze how AI may affect output, inflation, and productivity dynamics (Aldasoro et al., 2024; Guerrón-Quintana et al., 2025). We contribute to this debate by providing direct time-series evidence on the dynamic effects of AI innovation shocks on aggregate productivity, output, prices, employment, wages, and asset prices.

Finally, our findings relate to the literature on automation, labor demand, and inequality. A large body of work studies AI and automation using firm-, occupation-, or sector-level variation (Felten et al., 2018; Webb, 2019; Babina et al., 2024; Brynjolfsson et al., 2025). We complement this evidence by estimating general-equilibrium responses at the aggregate level. Our results are consistent with mechanisms in which productivity and reinstatement effects can dominate displacement effects in aggregate employment (Autor and Salomons, 2018; Mann and Püttmann, 2023; Acemoglu and Restrepo, 2019), while the gains from technology remain unevenly distributed. In particular, the decline in the labor share and the increase in wealth concentration connect our evidence to work on capital-biased technical change and the asset-valuation channel of inequality (Karabarbounis and Neiman, 2013; Autor et al., 2020; Moll et al., 2022).

Our findings speak directly to the debate on the labor market effects of AI and automation. We find that AI technology shocks increase aggregate employment, consistent with evidence that automation can raise labor demand when productivity and general equilibrium effects dominate displacement effects (Autor and Salomons, 2018; Mann and Püttmann, 2023). The expansionary effects on output and productivity are also consistent with model-based analyses in which AI raises aggregate efficiency (Aldasoro et al., 2024; Guerrón-Quintana et al., 2025). More broadly, our time-series evidence emphasizes the importance of general-equilibrium effects, which can be missed by microeconomic estimates that recover only relative or partial-equilibrium effects (Sarto, 2018; Wolf, 2023; Matthes et al., 2024).

Our paper also relates to the large empirical literature on technology shocks and aggregate productivity, including Galí (1999), Beaudry and Portier (2006), Fernald (2012), Bloom et al. (2013), and Aghion et al. (2025a). We are particularly close to studies that use patent data to measure innovation and study its aggregate implications (Kogan et al., 2017; Kelly et al., 2021; Cascaldi-Garcia and Vukotić, 2022; Miranda-Agrippino et al., 2025; Ferriani et al., 2024; Gazzani et al., 2025). Relative to this literature, we isolate AI intensity in innovation as a distinct source of technological development. This allows us to compare AI technology shocks with broader technology shocks and with shocks to ICT and automation innovation. A key distinction is distributional: whereas general technology shocks are often found to be approximately labor-neutral (e.g., Fisher, 2006), AI technology shocks generate a significant decline in the labor share. This pattern is consistent with mechanisms such as the declining relative price of investment goods (Karabarbounis and Neiman, 2013) or the rise of superstar firms (Autor et al., 2020), and suggests that AI-driven growth can raise labor demand while redistributing income toward capital.

Finally, our work contributes to the literature on inequality and automation (Acemoglu and Restrepo, 2018; Prettner and Strulik, 2020; Moll et al., 2022). We show that AI-intensive patents are of higher quality than non-AI patents in the broader automation and ICT categories: they receive more citations, are more *important*, and are more economically valuable. In addition, we document that AI technology shocks increase the wealth share of the top 10% and reduce that of the bottom 50%. These findings provide empirical support for the asset-valuation channel proposed by Moll et al. (2022), whereby technological change raises the value of capital and equity claims held disproportionately by wealthier households.

Structure of the paper. The remainder of the paper is organized as follows. Section 2 describes the construction of our AI-intensity measure and presents the econometric framework used to identify AI technology shocks. Section 3 reports the main aggregate results, beginning with the LP-IV baseline and then turning to OLS local projections and quantitative implications. Section 4 dissects AI innovation by comparing it with general technology shocks, benchmarking it against ICT and automation innovation, and decomposing its effects into extensive and intensive margins. Section 5 studies labor-market transmission, focusing on labor-market flows, sectoral heterogeneity, and differences across educational groups, and examines the distributional consequences of AI technology shocks for income and wealth. Section 6 connects the historical evidence to the recent generative-AI wave by constructing a narrower series of Generative-AI patents and estimating its macroeconomic effects. Finally, Section 7 concludes.

2 Data and Econometric Framework

While AI advances are partly based on open-source platforms, AI patenting is considered a valuable source of information on the development and diffusion of AI technology (Webb, 2019, Grimm and Gathmann, 2022, Chen et al., 2024). We employ the Artificial Intelligence Patent Dataset (AIPD) from the United States Patent and Trademark Office (USPTO), a database constructed for research purposes by Giczy et al. (2022).[2] This dataset covers the whole patent database of the USPTO and analyzes it through machine learning models – then validated using manual review by patent examiners – to identify the AI content in all patents granted between 1976 and 2020. The AI content of patents can stem from innovations that either develop AI technologies or use AI and the application of AI technologies as a key instrument in other technological developments that firms wish to patent. The dataset contains information on a broad spectrum of AI fields divided into eight categories: machine

[2] A more synthetic description of the dataset and several example and interesting remarks are available at `https://www.uspto.gov/sites/default/files/documents/OCE-DH-AI.png`.

learning, natural language processing, computer vision, speech technology, knowledge processing, AI hardware, evolutionary computation, and planning and control systems. To each patent, the AIPD assigns a score ranging from 0 to 1 for each AI category, where the scores proxy the artificial intelligence content of each category in the patent, which we average into a single AI score for each patent.

2.1 AI intensity in US innovation

We use this granular patent-level database to construct an aggregate index of AI intensity in technological progress. Precisely, we measure the economy-wide *AI intensity of innovation* as the average AI intensity of patents filed each month t (then eventually granted):

$$AIint_t = \frac{1}{8N_t} \sum_{i=1}^{N_t} \sum_{m=1}^{8} AIint_{i,t}^{(m)} \tag{1}$$

where $AIint_{i,t}$ is the average AI score (ranging between 0 and 1) of patent i in month t across the 8 categories m, and N_t is the total number of (granted) patents filed in month t. Crucially, we aggregate the AI score of patents in the month each patent is *filed* (and not granted), consistently with the empirical literature, which uses the dynamics of patent applications to proxy for innovation shocks (Miranda-Agrippino et al., 2025; Cascaldi-Garcia and Vukotić, 2022; Ferriani et al., 2024; Gazzani et al., 2025).

Although AI patenting has accelerated markedly since the 2000s (see Baruffaldi et al., 2020, for a review), foundational advances in AI date back much earlier. Indeed, the diffusion of AI in the United States had already begun in the late 1970s (USPTO, 2020). In particular, AI-intensive patents filed in the 1980s contained the seeds of technologies that have become popular today, such as large language models and speech recognition. Appendix B illustrates notable examples of historically significant AI-intensive patents.

Figure 1 (upper panel) displays the dynamics of $AIint$ in our sample. The series starts from virtually 0 and exhibits an upward trend since 1980, which has intensified in the second half of the 1990s and peaks at about 0.06. Appendix A displays the sectoral breakdown of $AIint$ – finding that the aggregate trend is mainly driven by AI patents in the“computer software” and “communication” sectors – and the dynamics of patenting across the eight AI fields.

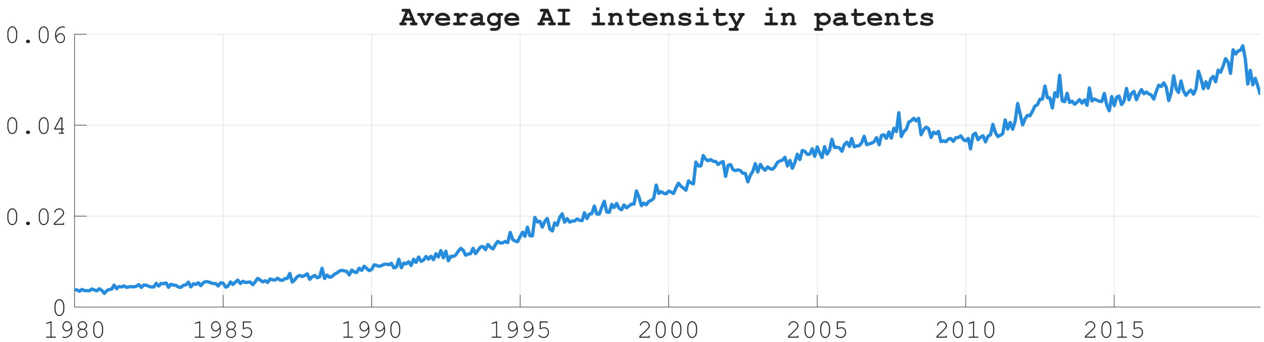

Figure 1: AI intensity in US patenting.

Note. The panel displays the average AI intensity in US patents ($AIint$). Sample: 1980m1-2019m12. Source: authors' calculation on the AIPD dataset.

2.2 Econometric Framework

Our aim is to assess the macroeconomic consequences of sudden surges in AI innovation. We employ the local projections (LPs) introduced by Jordà (2005) due to the long and variable lags that are likely to characterize the diffusion of AI technology (Brynjolfsson et al., 2018), which tilt the bias-variance trade-off between VAR models and LPs in favor of the latter (Li et al., 2024). Our baseline identification relies on an instrumental variable (IV) strategy inspired by that used by Miranda-Agrippino et al. (2025) for general technology shocks. In particular, we isolate exogenous variation in $AIint$ orthogonal to macroeconomic expectations and monetary policy shocks. We complement this with a standard LP OLS estimation in which a rich set of controls addresses potential endogeneity between patenting activity and economic conditions. Both strategies yield consistent results throughout the paper.

Specifically, our LP specification estimates the following set of regressions for each forecast horizon h:

$$y_{t+h} = \alpha_h + \beta_h AIint_t + \boldsymbol{\gamma}_h \mathbf{w}_t + v_{h,t}, \qquad h = 0, 1, ..., H \tag{2}$$

where y_{t+h} represents the outcome variable in log-levels as suggested in Montiel Olea et al. (2025), α denotes the constant, and $AIint_t$ denotes AI patent intensity, which is instrumented by z_t defined in eq.(3). The coefficient β_h captures the dynamic response at horizon h or

Impulse Response Functions (IRFs), which in our baseline is estimated via a local-projection instrumental variable approach; the inference is based on Newey and West (1994) standard errors. The main outcome variables of interest are GDP and TFP (the monthly aggregates constructed by Miranda-Agrippino et al., 2025), the headline Personal Consumption Expenditures (PCE) price index, total employment, weekly hours worked (not logged), hourly earnings, the 1-year Treasury yield (not logged), and the S&P500 price index. The control vector $\mathbf{w}_t$ includes 12 lags of all outcome variables and of $AIint_t$.

2.3 Instrumental variable construction

Our baseline strategy relies on an instrumental variable (IV) to identify shocks to AI technology. The rationale is that patent filings may incorporate information about technological progress, but they may also respond to business-cycle conditions or previously released news that affect both the level and the composition of innovation. To mitigate this concern, we construct an IV for AI technology shocks based on patent data, applying, in our setting, the methodology proposed by Miranda-Agrippino et al. (2025) for general technology shocks. Specifically, we use the residual variation in changes in AI intensity after removing the predictable component from lagged AI-intensity changes, Blue Chip macroeconomic forecasts, and monetary policy shocks. The identifying assumption is that the remaining variation in AI intensity is not systematically correlated with other contemporaneous shocks that affect the outcomes except through AI-related technology shocks.

Specifically, we regress the change in $AIint$ on its own lags, Romer and Romer (2004) monetary policy shocks, and Blue Chip forecasts:

$$\Delta AIint_t = c + \gamma(L)\Delta AIint_t + \sum_{h\in\{3,12\}} \beta_h \mathbb{E}_t[x_{t+h}] + \delta\eta_t + z_t, \tag{3}$$

where $\Delta AIint_t$ denotes the monthly change in $AIint$, and $\gamma(L) = \sum_{j=1}^{12} \gamma_j L^j$. $\mathbb{E}_t[x_{t+h}]$ are the Blue Chip forecasts one quarter and one year ahead for real GDP, CPI, and the unemployment rate. η_t is the monetary policy control. The residual z_t serves as our external instrument, available for the period 1982 to 2014 due to data limitations. Additional details on the IV construction are reported in Appendix C.

3 Main Results

This section presents the macroeconomic effects of AI technology shocks. We begin with the LP-IV baseline, which estimates the dynamic effects of AI technology shocks by instrumenting $AIint$ in eq.(2) with z_t. We then confirm the results using OLS local projections with extensive macro-financial controls, assess the quantitative importance of AI shocks via a forecast error variance decomposition, and discuss a battery of robustness exercises.

3.1 The Dynamic Effects of AI Technology Shocks

Figure 2 displays the aggregate responses to an AI technology shock identified via LP-IV, normalized to generate a 1% peak effect on TFP. The instrument isolates the variation in $AIint$ that is orthogonal to macroeconomic expectations and monetary policy (Section 2.3). The first-stage F-statistic is comfortably above conventional thresholds (above 100), indicating that weak-instrument concerns are unlikely to drive the LP-IV results. The dynamic effects are consistent with an expansionary technology shock with diffusion lags. First, all variables respond with a delay; second, TFP and output increase while consumer prices fall (Portier, 2014; Miranda-Agrippino et al., 2025).

Labor market quantities — employment, hours worked, and wages — rise, consistent with general equilibrium effects dominating the direct displacement effect, as emphasized in parts of the automation literature (Autor and Salomons, 2018; Mann and Püttmann, 2023). This pattern aligns with the "missing intercept" insight that partial-equilibrium displacement estimates can understate aggregate gains (Wolf, 2023; Sarto, 2018; Matthes et al., 2024), and with models in which AI raises productivity through task reallocation and complementarity (Aghion and Bunel, 2024; Aghion et al., 2025b). Employment adjusts primarily through job creation rather than hours worked, and effects remain significant at the 60-month horizon.

Regarding price effects, the mild and progressive decrease in the PCE price index masks heterogeneous dynamics across consumption categories, which we document in detail in Appendix D. While most subcategories exhibit gradual and persistent price declines consistent with standard technology-driven cost reductions, the strongest disinflationary effects are concentrated in durable goods, notably high-tech products. In contrast, energy prices increase on impact, suggesting that AI-related innovations are energy-intensive in the short run despite generating broader efficiency gains.

Finally, despite these positive labor market outcomes, AI technology shocks compress the labor share, as we document in Section 4.1, consistent with a shift in income toward capital.

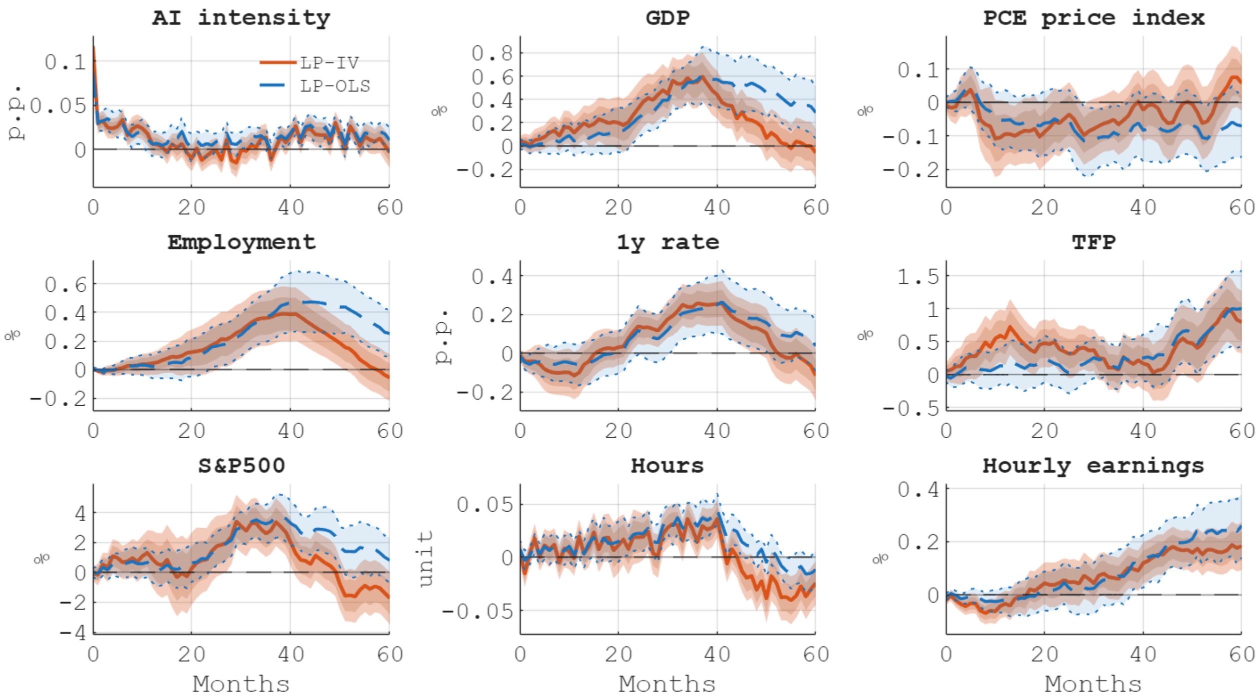

Figure 2: IRFs to an AI technology shock.

Note. The figure displays the IRFs to an AIint shock estimated by LP-IV, where the instrument is constructed as in Section 2.3, and the LP specification follows eq.(2). Sample 1982m1-2014m12. The estimates are based on local projections with Newey-West standard errors. Point estimates and 68%-90% confidence bands.

We also report OLS local projection estimates based on the same specification in equation (2), with the full set of macro-financial controls. These estimates closely mirror the LP-IV responses in Figure 2. The same patterns emerge: GDP and TFP rise with delay, consumer prices fall, employment and hourly earnings increase, and stock prices respond earlier than real variables. The OLS responses are somewhat more persistent than the LP-IV estimates, especially at longer horizons, but the overall dynamics are very similar.[3]

This close alignment between the LP-IV and OLS estimates is reassuring and suggests that spurious factors are unlikely to be driving the main results. The OLS LP approach is

[3] A one-standard-deviation innovation in $AIint$ in the OLS specification generates a peak TFP response of about one half of the response shown under our baseline normalization.

particularly useful for the remainder of the paper, as it allows us to compute the forecast error variance decomposition in Section 3.2 and to compare AI shocks with other technological developments in Section 4 over a longer sample. We describe the OLS specification, the associated orthogonality tests, and additional robustness exercises in Appendix D.4.

Appendix D.1 provides further support for the validity of the OLS-LP approach. Using the test proposed by Forni and Gambetti (2014), we find that our set of controls does not have explanatory power for $AIint_t$ in addition to $AIint_t$ own lags, suggesting that $AIint_t$ can be considered as conditionally exogenous to business cycle fluctuations. Additionally, the implicit shocks of interest $AIint_t \perp \mathbf{w_t}$ do not exhibit any linear dependence on macroeconomic shocks or projections from professional forecasters.

Alternative specifications We also assess the sensitivity of our findings to several perturbations of the LP specification (Appendix D.5). First, similar results hold if we include the number of AI-intensive patents rather than the share $AIint$ (see Section 4.2 for a discussion of the role of the intensive and extensive margin in $AIint$). Second, we address low-frequency movements in the data by either detrending $AIint$ or by directly including a trend in our LP specification. Third, we include stock prices as the ratio of high-tech to industrial stocks to capture sector expectations that may drive AI technology. Fourth, we include patent counts to control for general technological development (total) or for robotics automation. Section 4.1 explicitly compares AI technology shocks with automation shocks. Fifth, we include additional controls, such as the Economic Policy Uncertainty (EPU; Baker et al., 2016) and the Excess Bond Premium (EBP; Gilchrist and Zakrajšek, 2012). The results are consistent across all these alternatives, suggesting that specific identification or modeling choices are inconsequential for our economic conclusions.

3.2 Quantitative implications

The impulse responses provide insights into the qualitative implications of $AIint$ shocks. To gauge the quantitative consequences of $AIint$ shocks for the macroeconomy, we estimate the forecast error variance decomposition using the OLS LP framework. We follow Gorodnichenko and Lee (2020) and apply their R^2 approach, which yields the best performance in finite samples among other alternatives. The method amounts to extracting the shocks to $AIint$, by regressing them on the information set at $t-1$ and employing them in a regression on the forecast error of the endogenous variable of interest. Formally,

$$AIint_t = \sum_{j=1}^{12} \boldsymbol{\phi_j}\, \mathbf{y}_{t-j} + \varepsilon_t \tag{4}$$

Given any variable of interest $y^{(e)}$ and forecast horizon h such that

$$\Delta_h y_t^{(e)} := y_{t+h}^{(e)} - y_{t-1}^{(e)} \tag{5}$$

We regress $\Delta_h y_t^{(e)}$ on lagged $\mathbf{y_t^*} = [\mathbf{y_t}\, \varepsilon_t]$ that contains the observables and the shock itself to construct the forecast error of the baseline model:

$$\Delta_h y_t^{(e)} = \alpha + \sum_{j=1}^{12} \boldsymbol{\phi_j^*} \mathbf{y^*}_{t-j} + e_{t+h} \tag{6}$$

Then we regress the forecast error on the shocks realized between t and the forecast horizon h:

$$e_{t+h} = \sum_{l=0}^{h} \theta_{h,l} \varepsilon_{t+l} + u_{t+h} \tag{7}$$

The R^2 of this regression for each horizon h is the average variance contribution of shocks to $AIint$ to the endogenous variable of interest. The inference is based on 10000 bootstrap replications.

Results in Figure 3 highlight that AI technology shocks have not been a dominant driver of the US business cycle. Nonetheless, they possess non-negligible explanatory power (roughly 25% in the medium run) for key aggregates such as GDP and employment.

Their lower explanatory power for TFP is not inconsistent with the TFP impulse responses documented above. Our shocks are identified at the innovation margin and can be interpreted as news about future AI-driven productivity. Such productivity gains may materialize only gradually, through diffusion, adoption, complementary investment, and organizational adjustment (Brynjolfsson et al., 2021), and may therefore extend beyond the five-year horizon considered in the variance decomposition.

An additional possibility is that AI innovation shocks contain a noisy signal about future productivity, as in the news-shock literature (Forni et al., 2017). Under this interpretation, the shock can move economic activity and asset prices by changing expectations, while explaining only a lower share of realized TFP volatility over medium horizons.

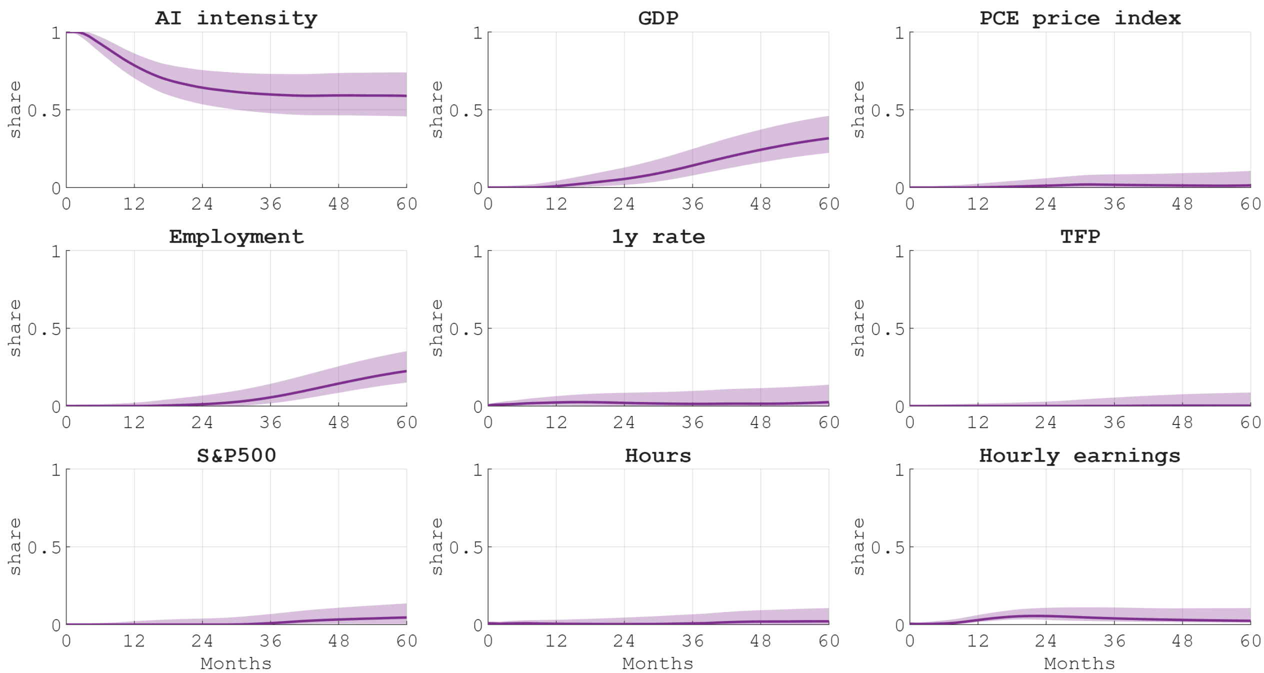

Figure 3: Variance contribution of AI technology shocks.

Note. The figure displays the variance contribution of $\mathcal{AI}$int shock. Sample 1980-2019. The estimates are based on local projections with 12 lags using the R^2 of Gorodnichenko and Lee (2020). Point estimates and 90% confidence bands from 10000 bootstrap replications.

4 Dissecting AI Innovation

This section deepens the analysis of AI-driven technological change by comparing it with the broader technology landscape, combining patent-level evidence with macroeconomic comparisons against general technology, ICT, and automation. We then decompose AI innovation into its extensive and intensive margins to identify the source of its aggregate effects.

4.1 AI in the broader technology landscape

We begin by asking how AI innovation relates to other technological domains and whether it displays distinctive characteristics relative to the broader innovation process. To this end, we compare AI-intensive patents with general technology, ICT, automation, and robotics. The evidence combines descriptive patent-level facts with comparisons of estimated macroeconomic effects.

Consistent with our patent-based approach to identifying AI technology, we use patent classifications to proxy broader technological domains: automation is defined using the taxonomy in Mann and Püttmann (2023), robotics is defined using official USPTO CPC robotics codes, and ICT is defined using the CPC-based definition in Inaba and Squicciarini (2017).[4]

Table 1 summarizes the relationship between AI-intensive patents and broader technology classes in a compact way. Panel A reports, first, the share of AI-intensive patents that belong to each technology class and, second, the share of patents within each class that are AI-intensive. Two facts stand out. First, AI innovation is overwhelmingly concentrated in ICT: nearly 80% of AI-intensive patents are ICT patents, whereas more than half also fall in the automation category. By contrast, robotics accounts for a negligible fraction of AI-intensive patents. Second, the overlap is asymmetric: while AI patents are often ICT- or automation-related, only a minority of ICT and automation patents are AI-intensive. This indicates that AI represents a specific and narrow frontier within these broader technology classes, rather than simply mirroring them.

Panel B turns to patent quality. AI-intensive patents consistently outperform non-AI patents, not only relative to the average patent, but also within ICT and automation. They receive more forward citations (by other patents), exhibit higher importance scores (Kelly et al., 2021), and produce higher estimated patent value (Kogan et al., 2017). These differences suggest that AI-intensive innovation is not merely a relabeling of broader digital or automation-related innovation, but a particularly impactful and economically valuable subset of it. This pattern provides a natural explanation for why AI shocks may generate stronger macroeconomic effects than shocks associated with other technology groups.

[4] Mann and Püttmann (2023) database has full coverage only through 2014.

Table 1: AI-intensive patents and broader technology classes

Panel A: AI and broader technology classes

	All patents	ICT	Automation	Robotics
Share of AI patents in category	100.0%	79.5%	55.7%	0.5%
Share of patents in category that are AI patents	9.4%	21.6%	16.1%	13.4%

Panel B: Patent characteristics by AI-intensive status

	All patents		ICT patents		Automation patents	
	Not AI	AI	Not AI	AI	Not AI	AI
Importance	0.21	0.35*	0.27	0.37*	0.24	0.35*
Forward citations	12.09	19.05*	11.93	18.73*	19.32	31.27*
Value	10.51	14.73*	8.33	15.35*	9.98	14.96*
Value (imputed)	0.69	1.18*	0.33	1.21*	0.62	1.16*

*Notes: AI patents stand for "AI-intensive" patents determined as an indicator equal to one when Giczy et al. (2022) classifies the patent as AI-intensive (above the 50% threshold in at least one domain). ICT patents are identified using CPC codes from Inaba and Squicciarini (2017). Robotics patents are classified according to official USPTO CPC robotics codes. Automation patents follow the Mann and Püttmann (2023) classification (with reliable coverage through 2014). Panel A reports, first, the share of AI-intensive patents that belong to each broader technology class and, second, the share of patents within each class that are AI-intensive. The value for "All patents" in the second row is the unconditional AI-intensive share. Panel B reports mean patent characteristics; * indicates the difference between AI and non-AI patents is statistically significant at the 1% level. Importance follows Kelly et al. (2021); patent value follows Kogan et al. (2017).*

Comparison with general technology. We next compare the macroeconomic effect of AI technology shocks with general technology shocks. To address this question, we use the IV-based approach of Miranda-Agrippino et al. (2025) (hereafter MAHHB) as a benchmark. Specifically, we contrast shocks identified using the original MAHHB monthly instrument—which captures general technology news—with those obtained from our AI-specific IV constructed analogously (Section 2.3).

The correlation between our AI IV and the MAHHB IV is 0.08, which is low and not statistically significant at conventional levels. This indicates that the AI IV captures shocks specific to AI innovation while remaining largely orthogonal to general technological advancements. We then compare the aggregate effects of the two shocks. To maximize comparability and avoid imposing arbitrary choices on which variable to instrument within an LP-IV framework, we use the IVs directly as shocks in an internal-instrument approach, setting them to zero when unavailable (Plagborg-Møller and Wolf, 2021; Noh, 2018).

Figure 4 reports the impulse response functions, normalizing both shocks to generate a 1% peak effect on TFP. We focus on four key variables: GDP, the PCE price index, TFP, and the labor share. While GDP, prices, and TFP respond similarly to AI and general technology shocks, their labor-share implications differ markedly. Consistent with the interpretation of labor-neutral technology shocks (Fisher, 2006), general technology shocks do not produce persistent movements in the labor share. In contrast, AI technology shocks generate a larger and more persistent decline, consistent with automation-driven dynamics (Autor and Salomons, 2018; Acemoglu and Restrepo, 2020). The effect gradually dissipates at longer horizons, potentially reflecting the reinstatement mechanism emphasized in Acemoglu and Restrepo (2019). This comparison suggests that AI shocks resemble positive technology shocks along standard macroeconomic dimensions, but remain distinct in their distributional implications between labor and capital.

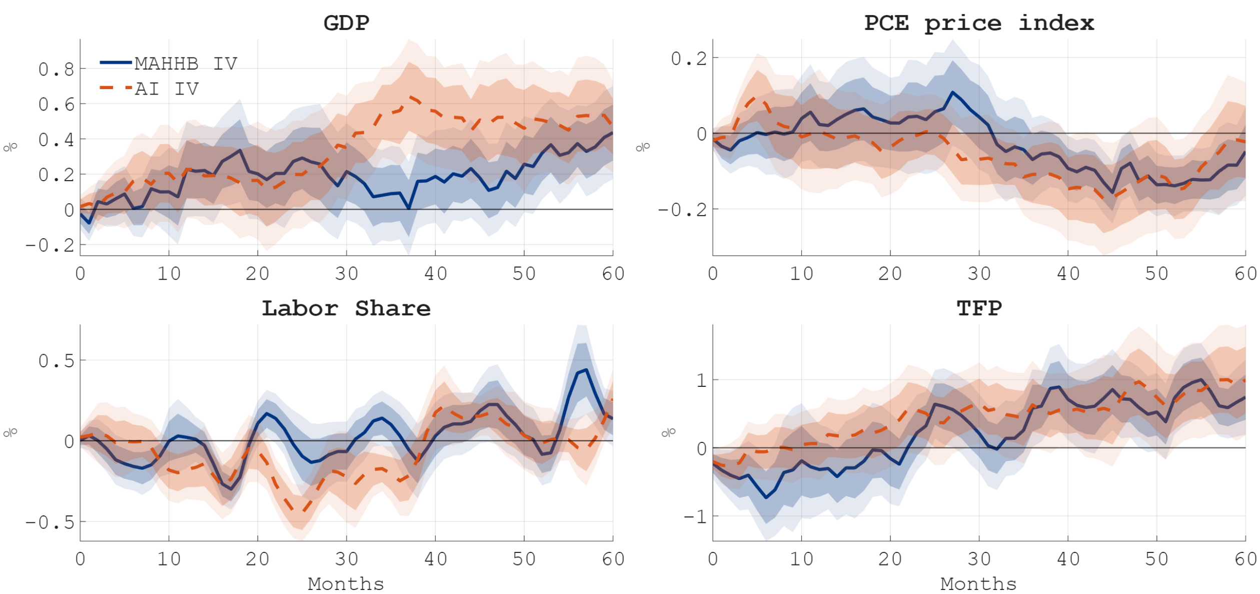

Figure 4: AI innovation versus General innovation.

Note. The figure compares the IRFs to an AI technology shock (AI IV) with those from a general technology innovation (MAHHB IV). Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimates and 68%-90% confidence bands.

Comparison with ICT and automation. We then turn to a narrower comparison with ICT and automation. Since robotics patents are both relatively few and only weakly overlapping with AI-intensive patents, we exclude them from the macroeconomic comparison. For this exercise, we estimate local projections for each patent category, expressed as a share of total patents, to ensure maximum comparability over time.

Figure 5 contrasts the macroeconomic effects of AI, automation, and ICT technology shocks. The IRFs depict the response of GDP and TFP (panel a) and employment and the labor share (panel b) to an unanticipated 1 percentage point increase in each patent share. The results reveal a clear ranking. AI technology shocks produce larger and more persistent positive effects on GDP, TFP, and employment than both automation and ICT shocks, with ICT in turn dominating automation. This ranking is fully consistent with the patent-level evidence in Table 1: AI-intensive patents are not only concentrated in ICT and automation, but also display systematically higher quality than non-AI patents within those categories. At the same time, AI technology shocks generate a more pronounced and persistent decline in the labor share relative to other forms of innovation. Taken together, these findings suggest that AI-driven productivity gains are particularly powerful, raising labor demand in absolute terms while shifting the distribution of income toward capital more strongly than either ICT or automation more broadly.

4.2 Extensive and intensive margins

Fluctuations in $AIint$ combine variations in two distinct dimensions of the AI innovation flow, i.e. the intensive margin (the AI intensity of each innovation) and the extensive margin (the number of AI-intensive patents produced). The extensive margin reflects entry of new innovations and an expansion in variety of technologies, while the intensive margin captures improvements in quality and a broadening in the range of uses of existing technologies. Variations in the two margins may map into potentially different effects on productivity and the labor market: identifying the dominant margin helps clarify whether AI acts mainly as an expansion of the innovation set towards AI or as a sectoral-specific technological shock.

We provide a formal decomposition of our indicator of AI intensity in US innovation (AIint) in two groups, i.e. that of AI-intensive patents (with score in at least one category greater than 0.5, as defined in Giczy et al., 2022) and non-AI-intensive ones (all the remaining

Figure 5: The Dynamic Effects of Innovation in AI, Automation and ICT

(a) GDP & TFP

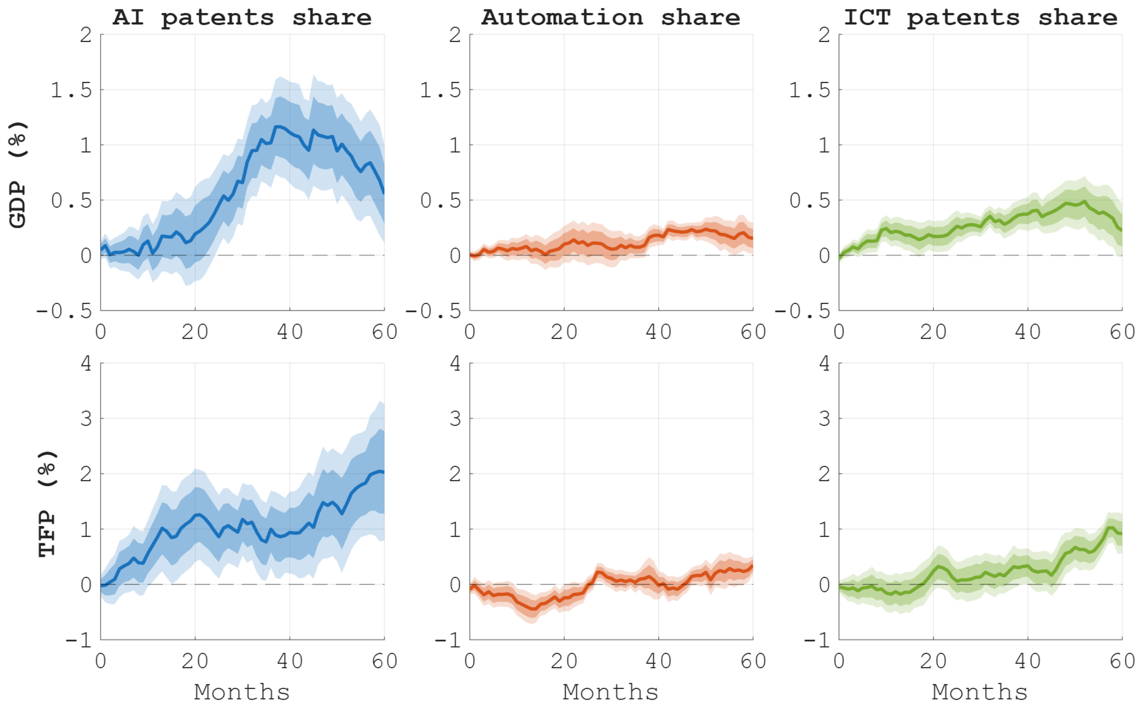

(b) Employment & Labor Share

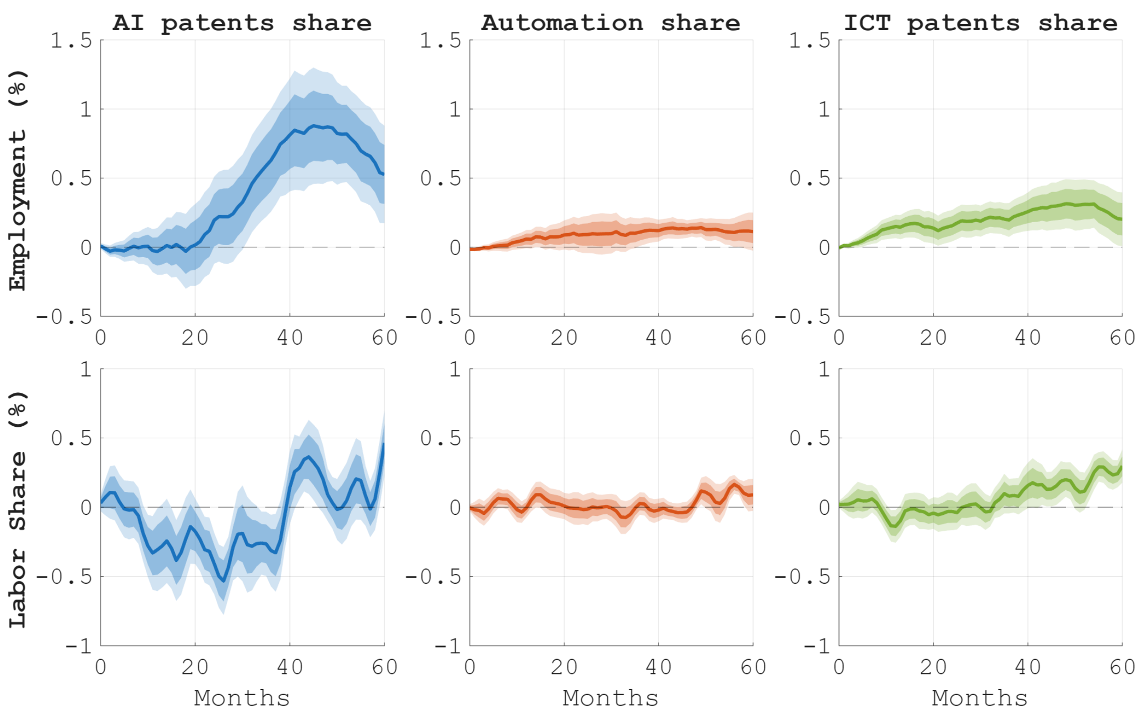

Note. The figure compares the dynamic responses to innovation shocks in AI, automation, and ICT patents. Each shock is defined as an unanticipated 1 percentage point increase in the corresponding patent share (relative to total patents), estimated via local projections as in (2) with the baseline set of controls. Point estimates are shown with 68% and 90% confidence bands based on Newey–West standard errors. Sample: 1980–2019.

Figure 6: Decomposition of AI Macroeconomic Effects.

(a) Intensive versus extensive margin comparison.

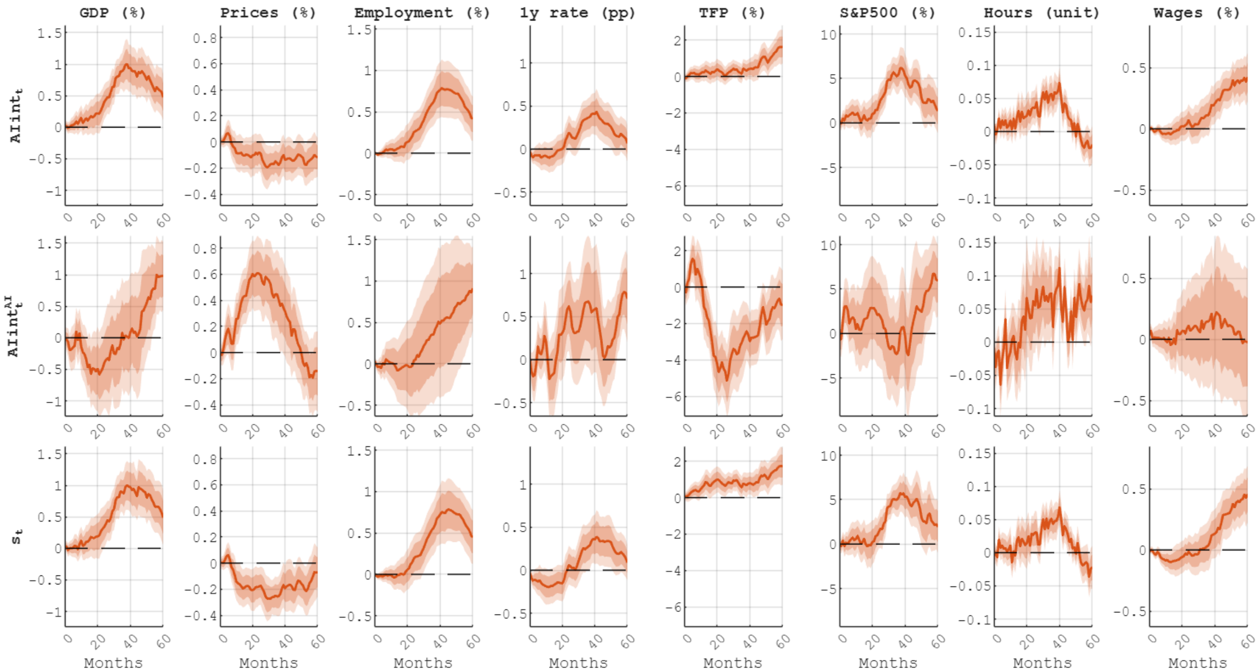

(b) Dispersion versus extensive margin comparison.

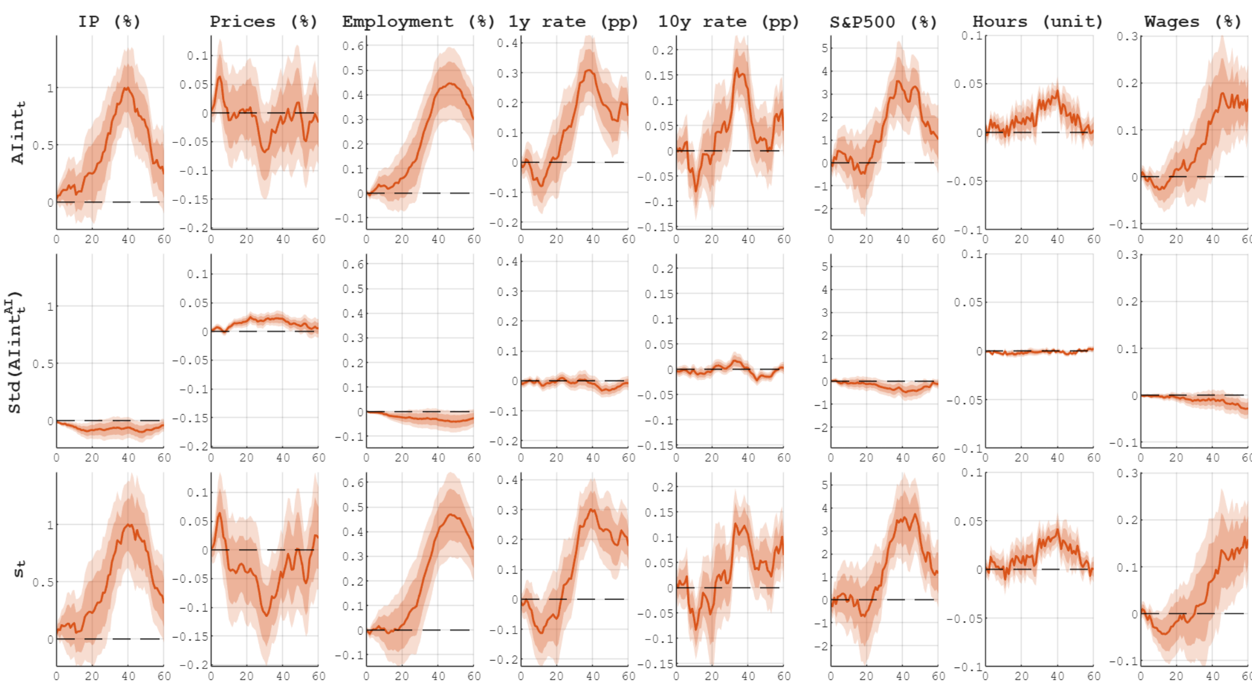

Note. The figure displays the IRFs to an AI technology shock. Panel (a) compares the effect of the intensive margin (average score in the AI-intensive group) with the extensive margin (share of AI patents). Panel (b) compares the effect of the dispersion of AI scores within the AI-intensive category with the extensive margin. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimates and 68%-90% confidence bands.

patents). In formulas, AI-intensive patents are:

$$a_{i,t} \;=\; \mathbf{1}\left\{ \max_{m\in\{1,\ldots,8\}} AIint^{(m)}_{i,t} > 0.5 \right\}. \tag{8}$$

Let $N_t^* = \sum_{i=1}^{N_t} a_{i,t}$ denote the number of AI-intensive patents in month t, and define the extensive margin (the share of AI-intensive patents) as

$$s_t \;=\; \frac{N_t^*}{N_t}. \tag{9}$$

Details on the decomposition, as well as descriptive statistics of the two groups, are available in Appendix A. Appendix Figure A.1 exhibits the share of AI (intensive) patents over total patents and $AIint$ within this sub-group, representing the extensive and intensive margins of AI patent intensity, respectively. The share of AI-intensive patents fluctuated between zero at the beginning of the sample and about 15% in the latest years. $AIint$ within AI-intensive patents evolved from about 0.18 to 0.3.

We then explore whether fluctuations in the intensive or extensive margin drive the overall macroeconomic effects. We repeat our baseline estimation strategy to estimate an unexpected increase in the average score in the AI-intensive group (intensive-$AIint^{AI}$) and in the share of AI-intensive patents relative to total patents (extensive-s). Figure 6a compares the baseline impacts of $AIint$ with those stemming from the intensive or the extensive margin. The overall effect is very sharply explained by the extensive margin, whereas variations in the intensive margin within the group of AI-intensive patents is inconsequential. These results strongly suggest that the increasing diffusion of AI within the overall pool of innovation generates the expansionary effects documented in the baseline analysis.

The previous analysis leaves open the question of whether a different intensive margin dimension of AI innovation matters for its overall macroeconomic effects – namely, whether concentrating highly AI-intensive patents within a cluster of patents has different aggregate implications than a lower level of AI intensity dispersed across the broader pool of AI-based innovations. If a high AI intensity is concentrated in a few patents, aggregate effects may be limited and driven by rents of few highly innovative firms; if dispersion of AI intensity does not matter, and the effects of AI tech mainly passes through newly created AI intense patents, AI is more likely to generate spillovers and wider macroeconomic impacts. This distinction clarifies whether AI shocks stem from a few breakthroughs or from generalized technological upgrading.

For this purpose, we focus on the cross-sectional dispersion of patent-level AI scores. Results in Figure 6b show that innovations in the cross-sectional dispersion of AI scores within the AI-intensive category is also inconsequential for macroeconomic outcomes. Overall, the evidence points to the importance of diffusion across a broader pool of innovations, rather than further increases in AI intensity within patents that are already highly AI-intensive.

5 Labor Market and Inequality

The aggregate results show that AI technology shocks raise output, productivity, employment, hours, and wages, while reducing consumer prices. These effects point to an expansionary supply-side shock, but they do not by themselves reveal how the gains from AI innovation are transmitted across workers, sectors, and households. This section examines these distributional margins.

We proceed in two steps. First, we study the labor-market transmission of AI technology shocks. We analyze labor-market flows, sectoral employment responses, and heterogeneity by educational attainment to assess whether the aggregate employment gains mask displacement or reallocation across workers and industries. This evidence helps distinguish between the direct displacement effects of automation and the productivity, spillover, and reinstatement effects through which AI may raise labor demand.

Second, we move from labor-market outcomes to the distribution of income and wealth. Even if AI shocks raise aggregate employment and wages, their gains may not be shared evenly across households. In particular, the decline in the labor share documented above suggests that AI-driven productivity gains may accrue disproportionately to capital owners. We therefore examine the responses of income, labor-income, and wealth shares across the U.S. household distribution. The evidence shows that AI shocks are expansionary but distributionally non-neutral: they generate broad labor-market gains, while increasing wealth concentration through an asset-valuation channel.

Among the direct mechanisms, technological advancements can lead to job losses by automating certain tasks (*displacement effect*). Secondly, these technologies may enhance human labor, allowing for a more adaptable distribution of tasks and boosting productivity (*productivity effect*). According to Acemoglu and Restrepo (2018) and Acemoglu (2024), such productivity effects might come from different sources: AI systems can reduce the cost of automated job tasks and enhance productivity in tasks that are not entirely automated (for example enabling workers to specialize) thereby increasing the marginal product of labor; AI adoption can raise the productivity of capital by improving services that have already been automated. Moreover, new tasks may emerge from AI, which could also raise labor demand in non-automated areas and influence the productivity of the entire production process. Additionally, displacement and productivity effects may combine: while some roles are eliminated, new ones emerge, either due to innovation or because older technologies become so affordable that their demand rises (*reinstatement effect, see* Acemoglu and Restrepo, 2019).

Various indirect channels can contribute to the aggregate effect of AI technology shocks. A clear example is spillover effects, in which productivity gains are shared across sectors via intermediate goods, or in which rising incomes lead to greater overall demand. Overall, AI is expected to enhance productivity growth by automating not only the production of goods and services but also the generation of ideas (Aghion et al., 2017).

We examine in further detail the consequences of $AIint$ for labor market flows and sectoral quantities. Figure 7 displays the aggregate responses of occupational flows. Despite increasing employment, hours, and wages, $AIint$ shocks transform the US labor market by altering the composition of jobs and workers. Both job openings and separations surge after the shock hits, but the former dominates the latter consistently with the response of employment in Figure 2. By examining the dynamics of outflows in detail, we find that separations rise with higher layoff rates, while the number of quits falls. This pattern is consistent with the *reinstatement effect*, as firms shift their labor force composition toward workers who are more complementary to AI-based progress, thereby exploiting new technological advancements.

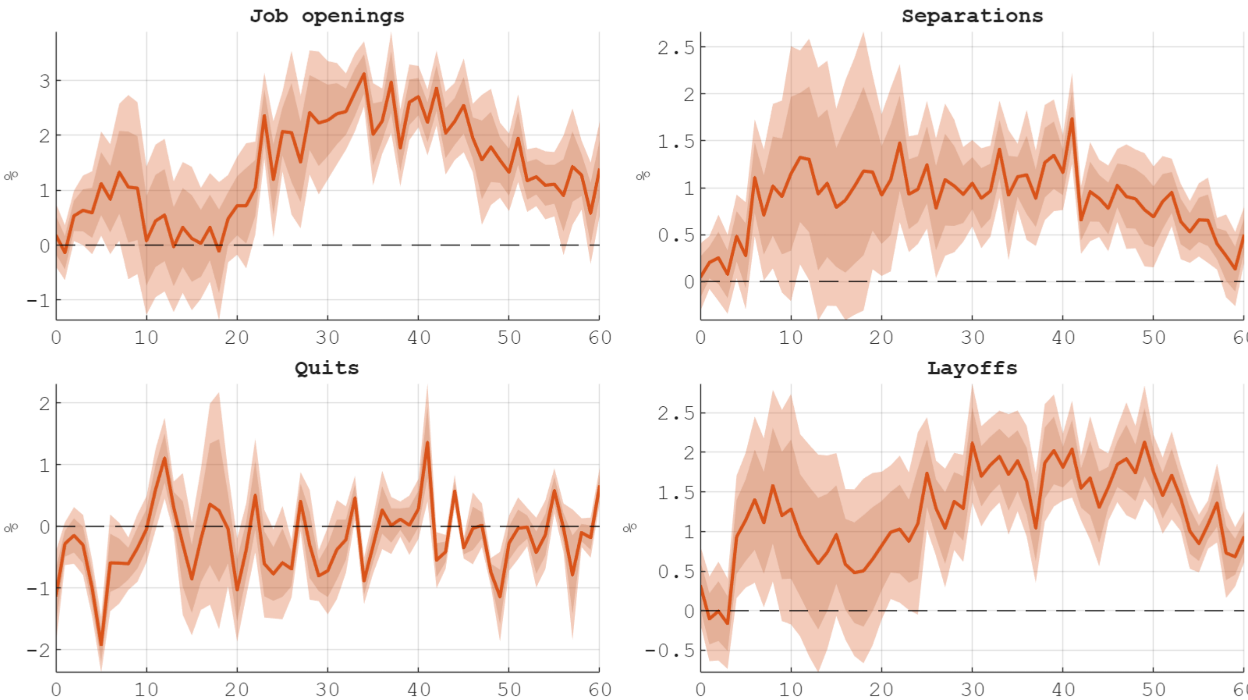

Figure 7: Response of labor market flows to an AI technology shock.

Note. The figure displays the IRFs to an AIint shock. Sample 2006-2019. The estimates are based on local projections with Newey-West standard errors. Point estimates and 68%-90% confidence bands.

5.1 Sectoral Heterogeneity and AI Exposure

To determine whether the aggregate expansion documented above masks significant reallocation across industries, we disaggregate the labor market response by sector. We stratify sectors based on their exposure to AI technologies, employing the metric constructed by Acemoglu et al. (2022). This measure aggregates occupation-level exposure scores from Felten et al. (2018), weighted by the vacancy structure of each industry, to identify sectors where the overlap between workforce skills and AI capabilities is most pronounced.[5]

Appendix Figure D.5 presents the cumulative impulse responses of labor quantities and wages over a 60-month horizon, with sectors ordered from lowest to highest AI exposure. Regarding labor quantities, we observe a remarkably homogeneous response as employment gains are broad-based, occurring even in sectors with the highest AI exposure. This uniformity underscores the dominance of general equilibrium forces: the aggregate demand and

[5]Specifically, Acemoglu et al. (2022) calculate sector-level exposure as the weighted average of occupation-level scores, using the share of vacancies posted by each sector in each occupation as weights. Alternative mappings of occupations to AI tasks include Webb (2019) and Eloundou et al. (2023).

productivity spillovers generated by AI appear to overwhelm the direct sectoral displacement effects. The response of real wages and job openings, while being slightly more sector-specific – e.g., some sectors with low exposure to AI experience stronger wage growth, suggesting some local displacement effect – broadly confirm this pattern.

5.2 The Role of Education

Education may constitute a critical source of heterogeneity in the transmission of $AIint$ shocks. To assess whether AI innovation favors specific skill groups, we disaggregate the labor market response by educational attainment, ranging from workers with less than a high school diploma to those with a bachelor's degree or higher.

Figure 8a presents the impulse responses for employment. The results uncover a distinct non-monotonic pattern: employment gains are most pronounced for the intermediate groups (High School and Some College), which exhibit a robust and persistent expansion. In contrast, workers with the lowest level of education (Less than High School) appear to be excluded from the expansion, with their employment response statistically indistinguishable from zero. The most educated group (Bachelor's and above) also benefits, though the magnitude of the response is slightly smaller than in the "Some College" category. Consequently, while AI stimulates labor demand for the majority of the workforce, it appears to reinforce the secular decline in relative demand for low-skilled workers, consistent with the hypothesis of skill-biased technological change (Berman et al., 1998).

The dynamics of earnings, displayed in Figure 8b, paint a different picture. Unlike the employment response, the positive impact on real wages is remarkably homogeneous across the skill distribution. All education groups, including the lowest-skilled, experience a persistent rise in earnings. This suggests that while AI shocks create barriers to entry for low-skilled workers (the extensive margin), those who remain employed share in the aggregate productivity gains through higher wages (the intensive margin).

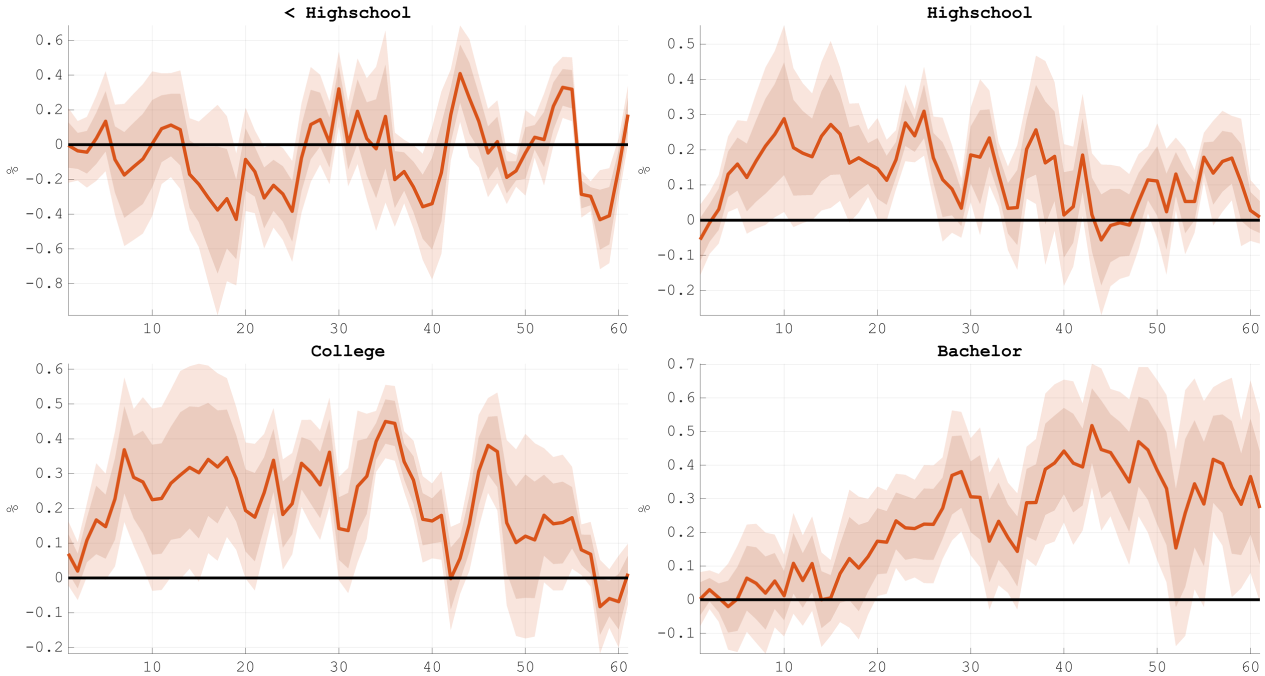

(a) Employment by educational attainment

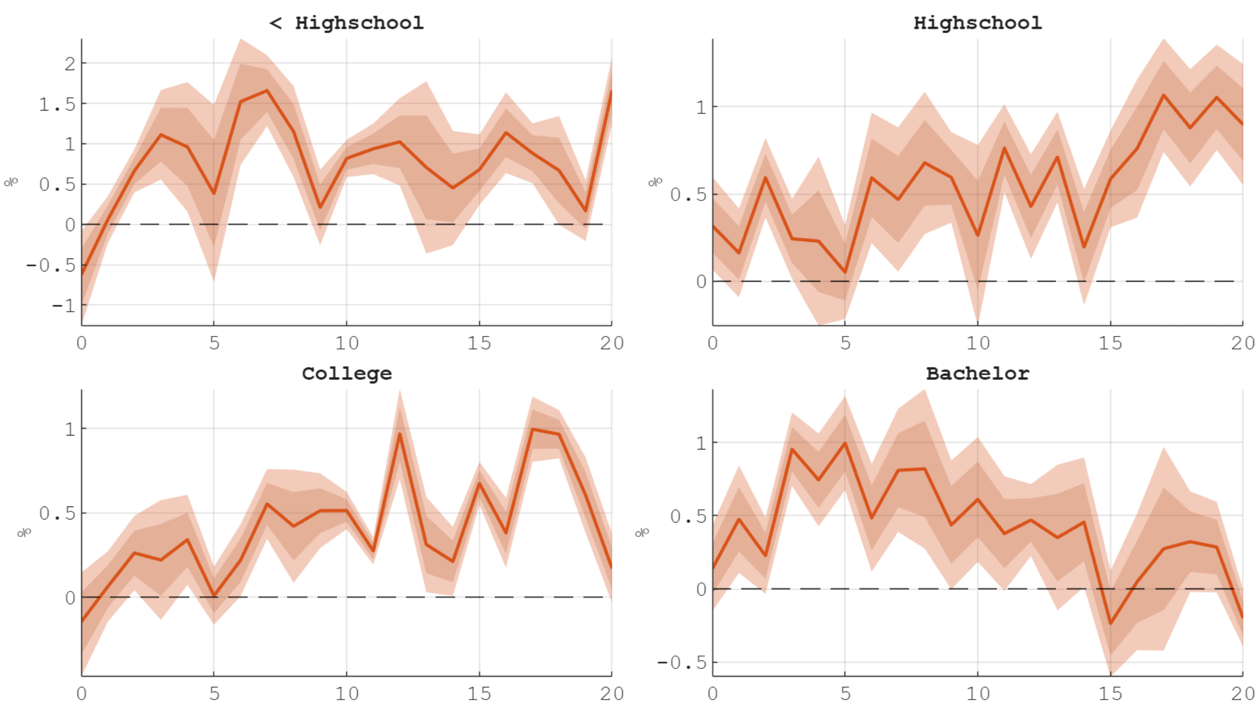

(b) Earnings by educational attainment

Figure 8: Responses of employment and earnings by educational attainment to an AI technology shock.

Note. The figure displays the IRFs to an AIint shock. Panel (a) reports employment responses using the 2006–2019 sample; Panel (b) reports earnings responses using the quarterly 2000–2019 sample. The estimates are based on local projections with Newey–West standard errors. Shaded areas denote 68% and 90% confidence bands.

5.3 Implications for Inequality

The aggregate results so far show that AI technology shocks generate broad-based gains: TFP, output, employment, hours, and wages all rise, while consumer prices fall. We now ask whether these gains are distributed evenly across households. To address this question, we use high-frequency measures of income, wealth, and labor-income shares across the U.S. distribution from Blanchet et al. (2022). We estimate the responses of these variables across selected groups of the distribution to an AI technology shock in our baseline LP framework (Figure 9).

Three main patterns emerge. First, the response of income shares is limited. The shares of the top 1% and top 10% drift mildly upward after an AI shock, but the responses are generally not statistically distinguishable from zero. Second, in sharp contrast, AI shocks generate large and persistent changes in the wealth distribution. The wealth share of the top 10% rises significantly and persistently, while the wealth share of the bottom 50% declines. Third, the labor-income distribution displays a similar but more transitory polarization: labor-income shares rise at the top and fall at the bottom, but the responses revert toward baseline within roughly forty months.

The wealth response is not merely a redistribution of relative shares. Appendix Figure D.9 reports the response of real income and real wealth in levels across the distribution. AI shocks raise real income across the distribution, consistent with the expansionary aggregate and labor-market effects documented above. By contrast, the real wealth responses are sharply heterogeneous. Real wealth rises at the top of the distribution, while the bottom 50% experiences an absolute decline in real wealth. The combination of broad-based income gains and absolute wealth losses at the bottom is the most striking distributional feature of AI shocks.

These dynamics are consistent with the asset-valuation channel emphasized by Moll et al. (2022). This interpretation is supported by the strong response of equity prices documented in Figure 2. Since financial and business assets are held disproportionately by households at the top of the wealth distribution, these valuation gains translate into higher wealth concentration. Households at the bottom, which hold few claims to capital and rely more heavily on labor income and low-return assets, benefit less from the valuation gains and experience declines in real wealth.

The evidence in this section complements the labor-share result documented in Section 4.1. Taken together, the results indicate that AI shocks are non-neutral along two dimensions. They are non-neutral across factors of production, as the labor share falls despite rising employment and wages. They are also non-neutral across households, as wealth gains are concentrated at the top and real wealth falls at the bottom of the distribution.

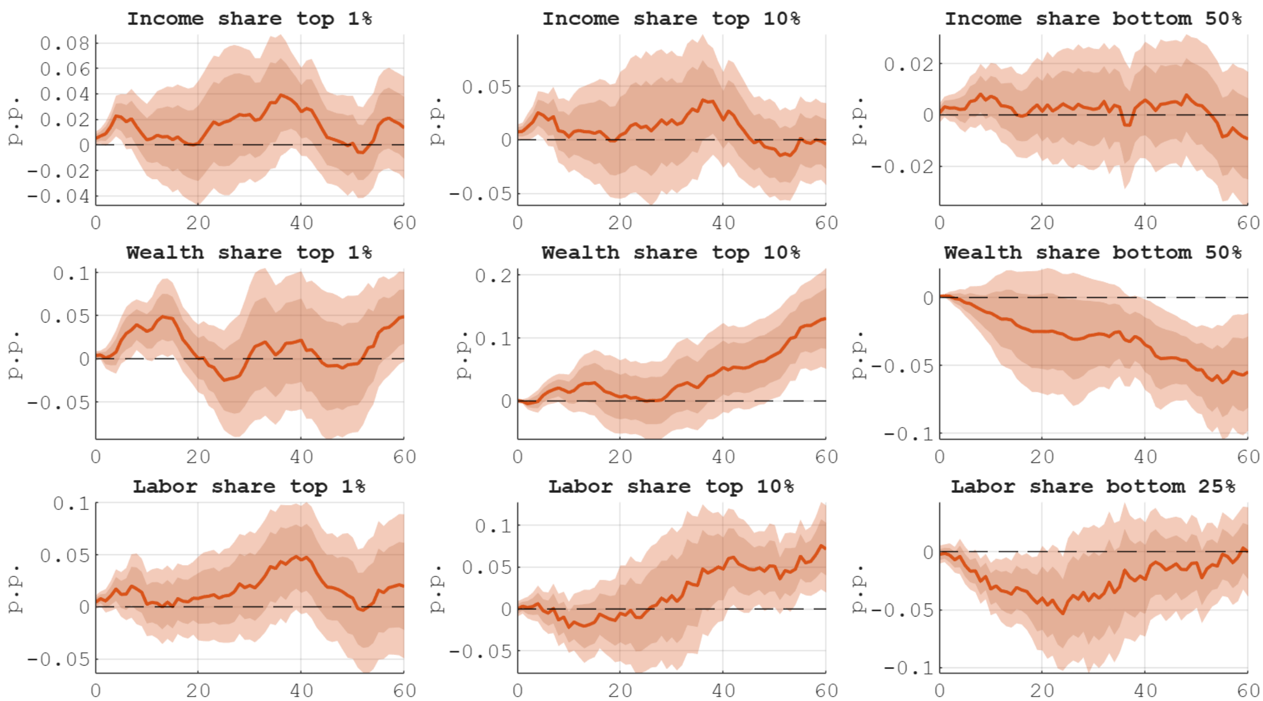

Figure 9: RESPONSE OF INEQUALITY TO AN AI TECHNOLOGY SHOCK.

Note. The figure displays the IRFs to an AIint shock. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimates and 68%-90% confidence bands.

6 Generative-AI patents

The results presented thus far are based on a broad universe of historical AI patents and should therefore not be interpreted as evidence on Generative AI alone. Nevertheless, this broader set includes many of the foundational innovations that underpin today's generative technologies, including key advances in machine learning and neural-network architectures that ultimately enabled modern tools for *generating* text, images, audio, and other content. To strengthen the connection between this historical evidence and the current technological frontier, we complement our baseline measure with a narrower series focused exclusively on patents related to Generative AI. We identify Gen-AI patents through a text- and classification-based search of the U.S. patent corpus using the Google Patents Public Datasets on BigQuery, in the spirit of the WIPO taxonomy of Generative-AI technologies (WIPO, 2024). The retrieved set includes several pioneering Generative-AI patents. Most notably, it contains Google's Transformer patent,[6] widely regarded as the foundational architecture behind modern large language models such as ChatGPT and Claude. It also includes early patents on generating music, images, and new drug candidates with deep-learning models, highlighting the technological roots of today's systems capable of creating text, images, audio, and other forms of content.

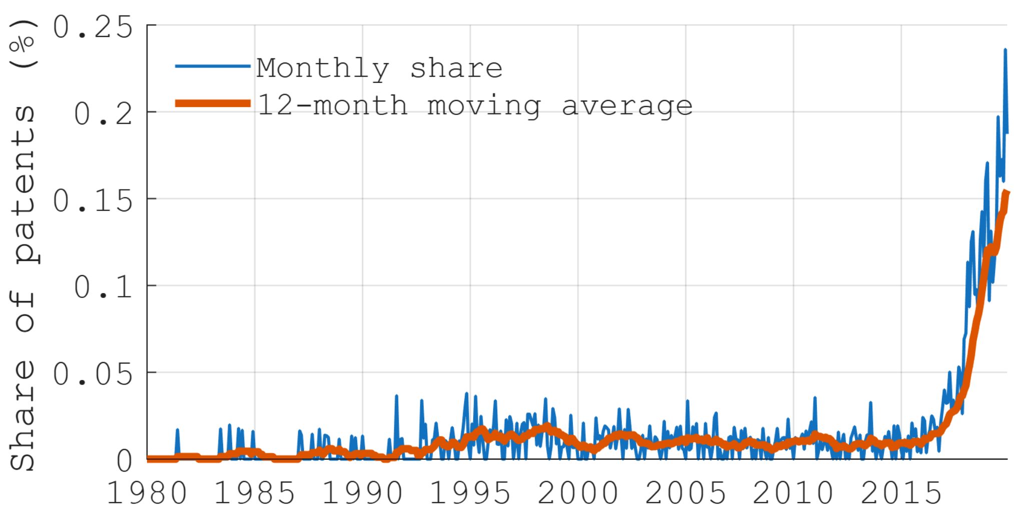

Figure 10: Generative-AI patent share.

Note. The figure displays the monthly share of generative-AI patents in total patents. Patents are dated by filing month. Sample: 1980m1–2019m12.

As measure of Gen-AI incidence in US innovation, we consider the share of generative AI patents out of total patents, which abstracts from fluctuations in overall patenting activity, in analogy with the extensive-margin specification of Section 4.2. Figure 10 displays the monthly series from 1980 to 2019. As expected, the share of Gen-AI patents is close to zero for most of the early sample and rises sharply in the second half of the 2010s, consistent with the emergence of modern generative-AI architectures. While the series is constructed to favor precision over coverage, so it may miss some less explicitly labeled generative-AI innovations

[6]"Attention-based sequence transduction neural networks" patent (US10452978B2).

in the earlier part of the sample, it represents, as far as we know, the first attempt to construct a historical series of Gen-AI patents in the US landscape.

We then use Gen-AI patent shares to construct impulse responses of macroeconomic variables. In particular, we estimate equation (2) replacing $AIint_t$ with the standardized generative-AI share, using the same baseline set of controls. Figure 11 reports the impulse responses, normalized so that the peak GDP response equals 1 percent. The dynamics qualitatively resemble those of the baseline AI shock: GDP, employment, TFP, equity prices, hours, and hourly earnings rise over the medium run, while consumer prices eventually decline—the pattern of a technology-driven expansion in which productivity gains materialize gradually.

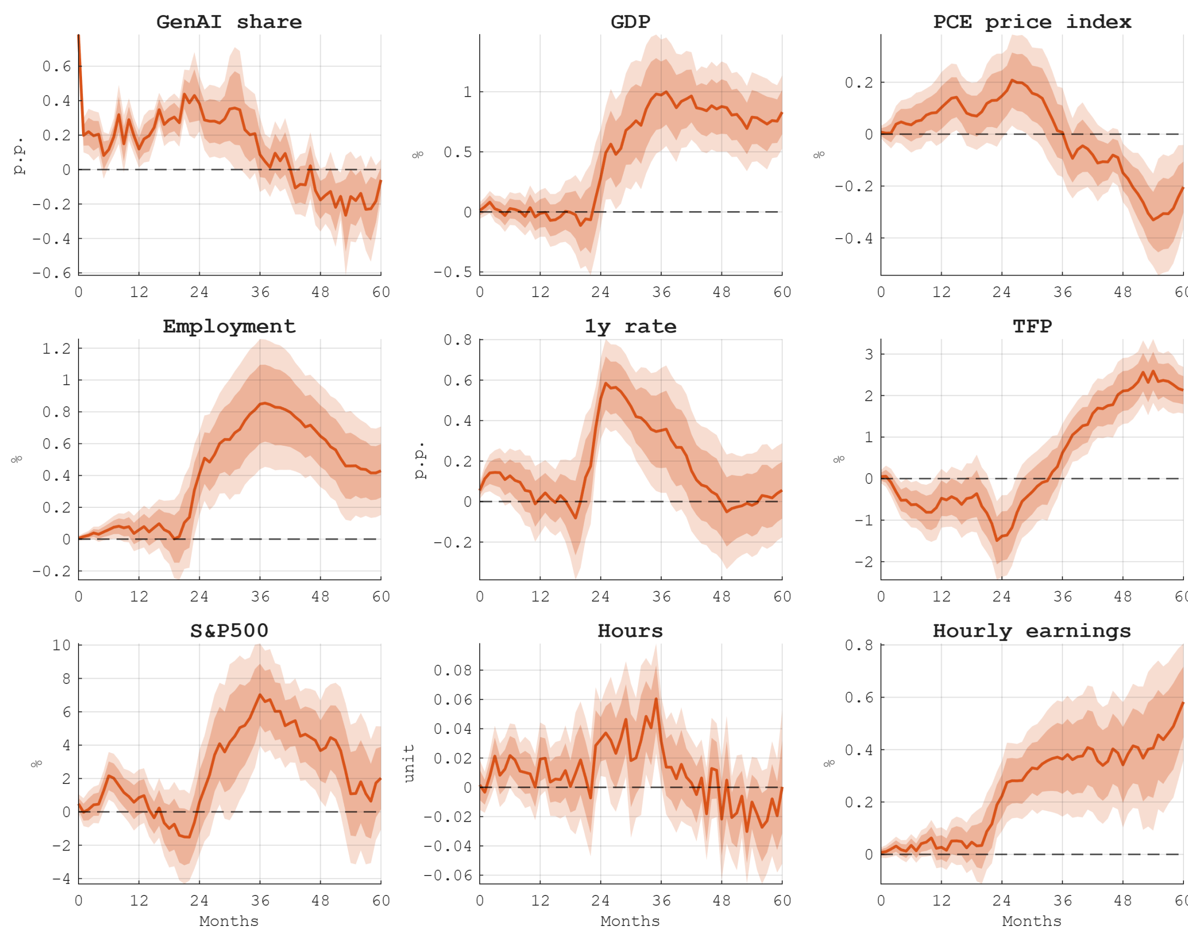

Figure 11: IRFs to a generative-AI technology shock.

Note. The figure displays the IRFs to a generative-AI technology shock, defined as a shock to the share of generative-AI patents in total patents. Sample: 1980m1–2019m12. The estimates are based on local projections with Newey–West standard errors. Point estimates and 68%–90% confidence bands.

The fact that Generative-AI and earlier AI innovations elicit similar economic responses carries an important implication. Despite being far more visible and directly accessible to the public, the Generative-AI technologies in our sample appear to propagate through the economy in a manner remarkably similar to previous generations of AI. Earlier AI innovations powered a wide range of applications across sectors – from e-commerce optimization and resource exploration to anomaly detection in industrial settings – often operating behind the scenes. Generative AI – at least at the stage of development captured by our sample – can therefore be interpreted less as a sharp discontinuity than as the latest stage in a longer process of AI-driven technological change, whose origins predate the recent mass diffusion of generative tools.

7 Conclusions

The recent breakthroughs in generative AI have placed the macroeconomic implications of artificial intelligence at the center of academic and policy debate. This paper exploits four decades of patent-level data on the AI content of U.S. innovation to provide systematic empirical evidence on the dynamic effects of AI technology shocks on the U.S. economy.

Three findings stand out. First, AI innovation behaves as an expansionary supply-side shock with long diffusion lags. It raises TFP, output, employment, hours, and wages, while lowering consumer prices. The labor-market gains are broad-based across sectors and educational groups, although they do not extend equally to workers with the lowest educational attainment. These findings highlight the importance of general-equilibrium spillovers that partial-equilibrium analyses focused on displacement may miss. Second, AI innovation appears to be a particularly powerful component of digital and automation-related technological change. AI-intensive patents are largely concentrated in ICT and automation, but they exhibit higher citations, technological importance, and market value than other patents in these categories. Consistent with this micro evidence, AI shocks generate stronger aggregate responses than broader ICT or automation shocks. A narrower measure based on patents directly related to generative AI delivers the same qualitative supply-side dynamics, suggesting that the baseline AI shock captures a broader technological force that includes, but predates, the recent GenAI wave. Third, unlike general technology shocks, AI shocks are distributionally non-neutral. They reduce the labor share and raise wealth inequality, with gains concentrated toward the top of the wealth distribution, consistent with an asset-valuation channel through which AI raises returns to capital held disproportionately by wealthy households.

These results should be interpreted with two caveats. First, our sample ends in 2019 and therefore predates the mass diffusion of the most recent Generative-AI applications. However, patents most directly related to Generative AI, including the foundational ones, already displayed macroeconomic dynamics similar to those of the broader series of AI shocks occurred between 1980 and 2019. To the extent that the mechanisms documented in the paper—diffusion lags, general-equilibrium spillovers, and asset-valuation effects—are features of AI as a general-purpose technology rather than of a specific technological wave, the evidence remains informative for the current episode. Whether the magnitudes will be similar, larger, or smaller is ultimately an empirical question that future work with post-2019 data will need to address. Second, our patent-based identification captures shocks to the innovation margin rather than directly measuring the diffusion or adoption of AI technologies. The impulse responses estimated via local projections, however, naturally incorporate the average diffusion and adoption dynamics that followed AI innovation in our sample, consistent with the long lags and persistent effects we document. Disentangling these margins and studying their interactions with firms' organizational choices and workers' tasks is an important avenue for future research.

References

Acemoglu, D. (2024): “The Simple Macroeconomics of AI,” Working Paper w32487, NBER.

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

Acemoglu, D. and P. Restrepo (2018): “Artificial Intelligence, Automation and Work,” Boston University - Department of Economics - The Institute for Economic Development Working Papers Series dp-298, Boston University - Department of Economics.

——— (2019): “Automation and new tasks: How technology displaces and reinstates labor,” *Journal of economic perspectives*, 33, 3–30.

——— (2020): “Robots and Jobs: Evidence from US Labor Markets,” *Journal of Political Economy*, 128, 2188–2244.

Aghion, P., A. Bergeaud, T. Boppart, and J.-F. Brouillette (2025a): “Resetting the Innovation Clock: Endogenous Growth through Technological Turnover,” Tech. rep.

Aghion, P. and S. Bunel (2024): “AI and Growth: Where Do We Stand?” Policy note, Federal Reserve Bank of San Francisco.

Aghion, P., S. Bunel, X. Jaravel, T. Mikaelsen, A. Roulet, and J. Søgaard (2025b): “How different uses of AI shape labor demand: evidence from France,” in *AEA Papers and Proceedings*, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, vol. 115, 62–67.

Aghion, P., B. F. Jones, and C. I. Jones (2017): “Artificial Intelligence and Economic Growth,” Working Paper 23928, National Bureau of Economic Research.

Aldasoro, I., S. Doerr, L. Gambacorta, and D. Rees (2024): “The impact of artificial intelligence on output and inflation,” Tech. rep., Bank for International Settlements.

Autor, D., D. Dorn, L. F. Katz, C. Patterson, and J. Van Reenen (2020): "The Fall of the Labor Share and the Rise of Superstar Firms*," *The Quarterly Journal of Economics*, 135, 645–709.

Autor, D. and A. Salomons (2018): "Is Automation Labor Share–Displacing? Productivity Growth, Employment, and the Labor Share," *Brookings Papers on Economic Activity*, 1–63.

Babina, T., A. Fedyk, A. He, and J. Hodson (2024): "Artificial intelligence, firm growth, and product innovation," *Journal of Financial Economics*, 151, 103745.

Baker, S. R., N. Bloom, and S. J. Davis (2016): "Measuring economic policy uncertainty," *The quarterly journal of economics*, 131, 1593–1636.

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

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

Basu, S., J. G. Fernald, and M. S. Kimball (2006): "Are technology improvements contractionary?" *American Economic Review*, 96, 1418–1448.

Baumeister, C. and J. D. Hamilton (2019): "Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks," *American Economic Review*, 109, 1873–1910.

Beaudry, P. and F. Portier (2006): "Stock Prices, News, and Economic Fluctuations," *American Economic Review*, 96, 1293–1307.

——— (2014): "News-driven business cycles: Insights and challenges," *Journal of Economic Literature*, 52, 993–1074.

Berman, E., J. Bound, and S. Machin (1998): "Implications of Skill-Biased Technological Change: International Evidence*," *The Quarterly Journal of Economics*, 113, 1245–1279.

Blanchet, T., E. Saez, and G. Zucman (2022): "Real-time inequality," Tech. rep., National Bureau of Economic Research.

Bloom, N., M. Schankerman, and J. Van Reenen (2013): "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 81, 1347–1393.

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

Brynjolfsson, E., D. Rock, and C. Syverson (2018): *Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics*, University of Chicago Press, 23–57.

——— (2021): "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies," *American Economic Journal: Macroeconomics*, 13, 333–372.

Cascaldi-Garcia, D. and M. Vukotić (2022): "Patent-Based News Shocks," *The Review of Economics and Statistics*, 104, 51–66.

Chen, W., S. Xinyuan, Terrence Tianshuo, and S. Srinivasan (2024): "The Value of AI Innovations," *Harvard Business School Working Paper No. 24-069.*

Eloundou, T., S. Manning, P. Mishkin, and D. Rock (2023): "GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models," Papers 2303.10130, arXiv.org.

Felten, E. W., M. Raj, and R. Seamans (2018): "A Method to Link Advances in Artificial Intelligence to Occupational Abilities," *AEA Papers and Proceedings*, 108, 54–57.

Fernald, J. G. (2012): "A quarterly, utilization-adjusted series on total factor productivity," Tech. rep.

Ferriani, F., A. G. Gazzani, and F. Natoli (2024): "The Macroeconomic Effects of Green Technology Shocks," *Available at SSRN: https://ssrn.com/abstract=4787604*.

Fisher, J. D. M. (2006): "The Dynamic Effects of Neutral and Investment-Specific Technology Shocks," *Journal of Political Economy*, 114, 413–451.

Forni, M. and L. Gambetti (2014): "Sufficient information in structural VARs," *Journal of Monetary Economics*, 66, 124–136.

Forni, M., L. Gambetti, M. Lippi, and L. Sala (2017): "Noisy News in Business Cycles," *American Economic Journal: Macroeconomics*, 9, 122–152.

Galí, J. (1999): "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review*, 89, 249–271.

Gazzani, A., J. Martinez, F. Natoli, and P. Surico (2025): "The Public Origins of American Innovation," CEPR Discussion Paper 20788, CEPR Press, Paris & London.

Gertler, M. and P. Karadi (2015): "Monetary policy surprises, credit costs, and economic activity," *American Economic Journal: Macroeconomics*, 7, 44–76.

Giczy, A. V., N. A. Pairolero, and A. A. Toole (2022): "Identifying artificial intelligence (AI) invention: a novel AI patent dataset," *The Journal of Technology Transfer*, 47, 476—505.

Gilchrist, S. and E. Zakrajšek (2012): "Credit spreads and business cycle fluctuations," *American economic review*, 102, 1692–1720.

Gorodnichenko, Y. and B. Lee (2020): "Forecast error variance decompositions with local projections," *Journal of Business & Economic Statistics*, 38, 921–933.

Grimm, F. and C. Gathmann (2022): "The Diffusion of Digital Technologies and its Consequences in the Labor Market," Tech. rep.

Guerrón-Quintana, P., T. Mikami, and J. Nosal (2025): "The Macroeconomic Implications of the Gen-AI Economy," SSRN Working Paper, last revised October 2025.

Inaba, T. and M. Squicciarini (2017): “ICT: A new taxonomy based on the international patent classification,” *OECD Science, Technology and Industry Working Papers*, 2017, 1.

Jordà, Ò. (2005): “Estimation and inference of impulse responses by local projections,” *American economic review*, 95, 161–182.

Känzig, D. R. (2021): “The macroeconomic effects of oil supply news: Evidence from OPEC announcements,” *American Economic Review*, 111, 1092–1125.

Känzig, D. R. and C. T. Williamson (2025): “Unraveling the drivers of energy-saving technical change,” Tech. rep., National Bureau of Economic Research.

Karabarbounis, L. and B. Neiman (2013): “The Global Decline of the Labor Share*,” *The Quarterly Journal of Economics*, 129, 61–103.

Kelly, B., D. Papanikolaou, A. Seru, and M. Taddy (2021): “Measuring Technological Innovation over the Long Run,” *American Economic Review: Insights*, 3, 303–20.

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

Kurmann, A. and C. Otrok (2013): “News shocks and the slope of the term structure of interest rates,” *American Economic Review*, 103, 2612–2632.

Li, D., M. Plagborg-Møller, and C. K. Wolf (2024): “Local projections vs. VARs: Lessons from thousands of DGPs,” *Journal of Econometrics*, 105722.

Mann, K. and L. Püttmann (2023): “Benign effects of automation: New evidence from patent texts,” *Review of Economics and Statistics*, 105, 562–579.

Matthes, C., N. Nagasaka, and F. Schwartzman (2024): “Estimating The Missing Intercept,” *Mimeo*.

McCarthy, J. (2007): “What is artificial intelligence,” *Mimeo*.

Mertens, K. and M. O. Ravn (2013): "The dynamic effects of personal and corporate income tax changes in the United States," *American economic review*, 103, 1212–1247.

Miranda-Agrippino, S., S. Hacioğlu-Hoke, and K. Bluwstein (2025): "Patents, News, and Business Cycles," *Review of Economic Studies, forthcoming.*

Moll, B., L. Rachel, and P. Restrepo (2022): "Uneven growth: automation's impact on income and wealth inequality," *Econometrica*, 90, 2645–2683.

Montiel Olea, J. L., M. Plagborg-Møller, E. Qian, and C. K. Wolf (2025): "Local Projections or VARs? A Primer for Macroeconomists," Working Paper 33871, National Bureau of Economic Research.

Newey, W. K. and K. D. West (1994): "Automatic lag selection in covariance matrix estimation," *The Review of Economic Studies*, 61, 631–653.

Noh, E. (2018): "Impulse-response analysis with proxy variables," *Mimeo.*

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

Portier, F. (2014): "Comment on" Whither News Shocks?"," Tech. rep., National Bureau of Economic Research.

Prettner, K. and H. Strulik (2020): "Innovation, automation, and inequality: Policy challenges in the race against the machine," *Journal of Monetary Economics*, 116, 249–265.

Ramey, V. A. (2011): "Identifying government spending shocks: It's all in the timing," *The Quarterly Journal of Economics*, 126, 1–50.

Romer, C. D. and D. H. Romer (2004): "A new measure of monetary shocks: Derivation and implications," *American economic review*, 94, 1055–1084.

——— (2010): "The macroeconomic effects of tax changes: estimates based on a new measure of fiscal shocks," *American Economic Review*, 100, 763–801.

Sarto, A. (2018): "Recovering macro elasticities from regional data," *Mimeo.*

Smets, F. and R. Wouters (2007): "Shocks and frictions in US business cycles: A Bayesian DSGE approach," *American economic review*, 97, 586–606.

USPTO (2020): "Inventing AI Tracing the diffusion of artificial intelligence with U.S. patents," IP DATA HIGHLIGHTS 5, U.S. Patent and Trademark Office • Office of the Chief Economist.

Webb, M. (2019): "The Impact of Artificial Intelligence on the Labor Market," *Available at SSRN: https://ssrn.com/abstract=3482150*.

WIPO (2024): "Patent Landscape Report: Generative Artificial Intelligence," Tech. rep., World Intellectual Property Organization, Geneva, wIPO Patent Landscape Report.

Wolf, C. K. (2023): "The missing intercept: A demand equivalence approach," *American Economic Review*, 113, 2232–2269.

A Data Definitions and Details

This appendix provides additional details on the data used in the paper and on the construction of the AI-related patent measures. Section A.1 describes the Artificial Intelligence Patent Dataset, the construction of our aggregate AI-intensity measure, and the decomposition of this measure into extensive and intensive margins. Section A.2 documents the auxiliary macroeconomic, financial, labor-market, patent-quality, and distributional data used in the empirical analysis. Section A.3 describes the construction of the generative-AI patent series used in the main-text section on Generative-AI patents.

A.1 Artificial Intelligence Patent Dataset

The Artificial Intelligence Patent Dataset (AIPD) classifies U.S. patent documents as AI-related and provides domain-specific AI scores using machine-learning methods documented by Giczy et al. (2022). For each patent document, the dataset reports predicted AI content in eight domains: machine learning, natural language processing, computer vision, speech technology, knowledge processing, AI hardware, evolutionary computation, and planning/control. Scores lie in $[0, 1]$ and can be interpreted as the predicted relevance of each domain to the patent document.

In the main analysis, we (i) restrict attention to granted patents, (ii) time patents by application filing date, and (iii) aggregate to the monthly frequency. Timing by filing date is standard in the macro-patent literature because applications better align with the innovation decision and mitigate grant/publication timing distortions (Miranda-Agrippino et al., 2025; Cascaldi-Garcia and Vukotić, 2022; Ferriani et al., 2024). To limit end-of-sample biases induced by publication lags, we end the sample in 2019.

A.1.1 Intensive and extensive margins

Denote patents by $i = 1, \ldots, N_t$, where N_t is the total number of patents filed in month t. Let $AIint_{i,t}^{(m)} \in [0, 1]$ be the AIPD score of patent i in AI subfield $m = 1, \ldots, 8$. We define the overall AI intensity score of patent i as the average across subfields,

$$AIint_{i,t} \;=\; \frac{1}{8}\sum_{m=1}^{8} AIint_{i,t}^{(m)} \;\in\; [0,1]. \tag{10}$$

Following Giczy et al. (2022), we classify a patent as *AI-intensive* if at least one subfield score exceeds the threshold $\tau = 0.5$:

$$a_{i,t} \;=\; \mathbf{1}\left\{\max_{m\in\{1,\ldots,8\}} AIint_{i,t}^{(m)} > \tau\right\}. \tag{11}$$

Let $N_t^* = \sum_{i=1}^{N_t} a_{i,t}$ denote the number of AI-intensive patents in month t, and define the extensive margin (the share of AI-intensive patents) as

$$s_t \;=\; \frac{N_t^*}{N_t}. \tag{12}$$

The aggregate AI intensity of innovation is the cross-sectional average of patent scores in month t,

$$AIint_t \;=\; \frac{1}{N_t}\sum_{i=1}^{N_t} AIint_{i,t}. \tag{13}$$

Define the group-specific average AI intensity within the AI-intensive and non-AI-intensive groups as

$$AIint_t^{AI} \;=\; \frac{1}{N_t^*}\sum_{i=1}^{N_t} a_{i,t}\, AIint_{i,t}, \tag{14}$$

$$AIint_t^{non} \;=\; \frac{1}{N_t - N_t^*}\sum_{i=1}^{N_t} (1-a_{i,t})\, AIint_{i,t}. \tag{15}$$

By construction, the aggregate AI intensity can be written as the mixture of the intensive margins weighted by the extensive margin:

$$AIint_t \;=\; s_t\, AIint_t^{AI} \;+\; (1-s_t)\, AIint_t^{non}. \tag{16}$$

Given that AI intensity among non-AI-intensive patents is empirically close to zero (Table A.1), a useful approximation is

$$AIint_t \;\approx\; s_t\, AIint_t^{AI}. \tag{17}$$

This decomposition motivates the empirical exercise in Section 4.2 of the main text, where we assess whether the aggregate macroeconomic effects are primarily driven by fluctuations in the extensive margin s_t or by changes in the intensive margin $AIint_t^{AI}$.

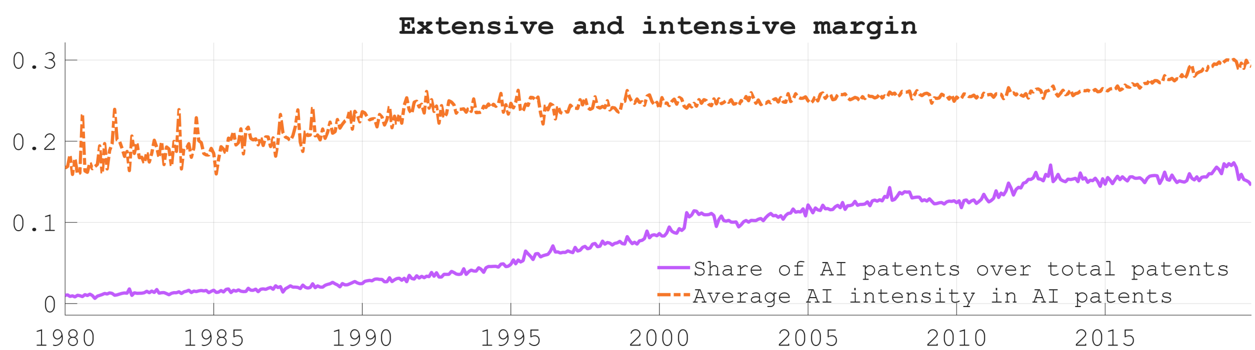

Figure A.1: Intensive and extensive margins in AI intensity.

Note. The panel displays the intensive and extensive margins of AI patent intensity. Sample: 1980m1-2019m12. Source: authors' calculation on the AIPD dataset.

A.1.2 Cross-sectional dispersion of AI scores

Within month t, let $\sigma_{1,t}^2 := \mathrm{Var}(AIint_{i,t} \mid a_{i,t} = 1)$ and $\sigma_{0,t}^2 := \mathrm{Var}(AIint_{i,t} \mid a_{i,t} = 0)$. Using the notation in Section A.1.1, the cross-sectional variance of $AIint_{i,t}$ satisfies the law of total variance

$$\mathrm{Var}(AIint_{i,t}) = s_t\,\sigma_{1,t}^2 \;+\; (1-s_t)\,\sigma_{0,t}^2 \;+\; s_t(1-s_t)\left(AIint_t^{AI} - AIint_t^{non}\right)^2. \tag{18}$$

Table A.1 shows that $AIint_t^{non}$ is close to zero and $\sigma_{0,t}^2$ is negligible; under this empirically valid approximation,

$$\mathrm{Var}(AIint_{i,t}) \;\approx\; s_t\,\sigma_{1,t}^2 \;+\; s_t(1-s_t)\left(AIint_t^{AI}\right)^2. \tag{19}$$

Since $AIint_t = s_t\, AIint_t^{AI}$ by Section A.1.1, substituting into (19) yields

$$\operatorname{Var}(AIint_{i,t}) \;\approx\; s_t\,\sigma_{1,t}^2 + \frac{1-s_t}{s_t}\,\big(AIint_t\big)^2, \qquad \operatorname{sd}(AIint_{i,t}) \;\approx\; \sqrt{s_t\,\sigma_{1,t}^2 \;+\; \frac{1-s_t}{s_t}\,\big(AIint_t\big)^2}. \tag{20}$$

If the within–AI dispersion is modest relative to the squared mean (i.e., $\sigma_{1,t}^2$ is small compared to $\big(AIint_t^{AI}\big)^2$), a convenient simplification is

$$\operatorname{sd}(AIint_{i,t}) \;\approx\; AIint_t^{AI}\,\sqrt{s_t(1-s_t)} \;=\; AIint_t\,\sqrt{\frac{1-s_t}{s_t}}. \tag{21}$$

Equations (18)–(21) make clear that, given the distributional features in Table A.1, the cross-sectional standard deviation of $AIint_{i,t}$ is primarily governed by the extensive margin s_t through the mixture term, rather than by polarization within the intensive margin of AI–intensive patents.

Using Table A.1, the share of AI-intensive patents in the pooled sample is

$$s \;=\; \frac{1{,}023{,}744}{9{,}059{,}463} \;\approx\; 0.113,$$

the AI–group mean and variance are $\mu_1 \approx 0.26$ and $\sigma_1^2 \approx (0.17)^2 \approx 0.03$, and for non–AI patents $\mu_0 \approx 0.01$ and $\sigma_0^2 \approx (0.01)^2 = 0.0001$. Plugging these values into (18) gives

$$Var(AIint) \;=\; s\,\sigma_1^2 \;+\; (1-s)\,\sigma_0^2 \;+\; s(1-s)(\mu_1-\mu_0)^2.$$

Evaluating the three components,

$$s\,\sigma_1^2 \approx 0.00339, \qquad (1-s)\,\sigma_0^2 \approx 0.00009, \qquad s(1-s)(\mu_1-\mu_0)^2 \approx 0.00627,$$

which sum to

$$Var(AIint) \approx 0.01, \qquad \operatorname{sd}(AIint) \approx 0.10,$$

matching the unconditional variance and standard deviation reported in the table.

The mixture (between–group) term $s(1-s)(\mu_1-\mu_0)^2$ accounts for approximately

$$\frac{0.00627}{0.00975} \approx 64\% \quad \text{of the total variance,}$$

the within–AI term $s\,\sigma_1^2$ for about 35%, and the within–non–AI term $(1-s)\,\sigma_0^2$ is negligible. These calculations corroborate the interpretation following (19)–(21): the cross–sectional dispersion of $AIint$ is primarily driven by the extensive margin s_t through the mixture term, rather than by polarization within the intensive margin.

Indeed, we display in Figure D.4 the effects of a shock to the overall standard deviation of $AIint$ that are nearly identical to a shock to $AIint$, and in particular s_t itself.

Table A.1: AI intensity distribution across groups

Statistic	AI intensive group	Unconditional	Non AI intensive group
Percentiles			
1%	0.07	0.00	0.00
5%	0.08	0.00	0.00
10%	0.10	0.00	0.00
25%	0.13	0.00	0.00
50%	0.22	0.00	0.00
75%	0.34	0.01	0.00
90%	0.51	0.11	0.01
95%	0.62	0.25	0.03
99%	0.80	0.53	0.07
Moments			
Mean	0.26	0.03	0.01
Std. Dev.	0.17	0.10	0.01
Variance	0.03	0.01	0.00
Skewness	1.29	4.22	4.97
Kurtosis	4.39	23.57	34.62
Observations	1,023,744	9,059,463	8,035,719

A.1.3 Breakdown by sector and AI field

Figure A.2 reports the share of AI-intensive patents by sector. “Computer software” and “communication” are by far the most AI-intensive sectors. “Communication” has steadily risen during our sample, whereas “computer software” spiked in the mid-1990s. Towards the end of our sample, both sectors feature about 10% of AI-intensive patents out of total patents. Figure A.3 reports the average AI intensity value across AI fields instead. “Planning/control” and “Knowledge processing” lie at the top of the distribution. “AI hardware” and “Vision” follow in a second cluster. Finally, “Machine learning” has spiked in recent years, leaving behind the residual categories.

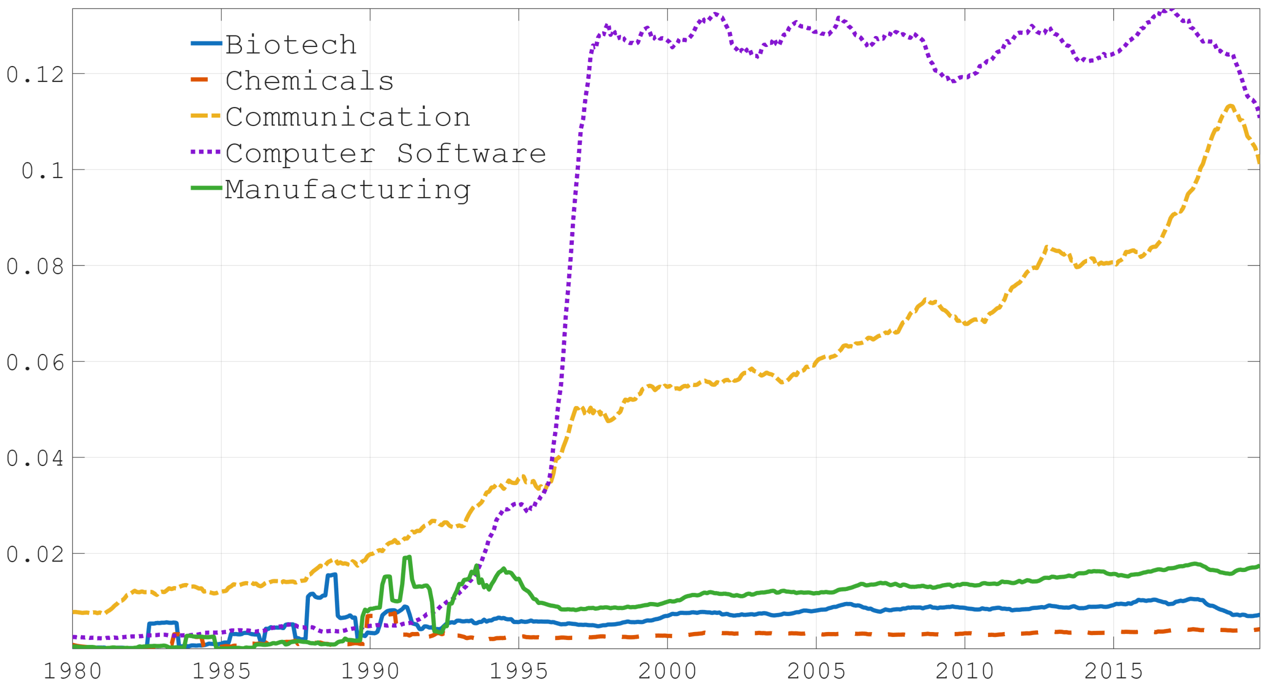

Figure A.2: Share of AI-intensive patents by sector.

Note. The figure displays AIint by sector between 1980 and 2019. Source: authors' calculation on the AIPD dataset

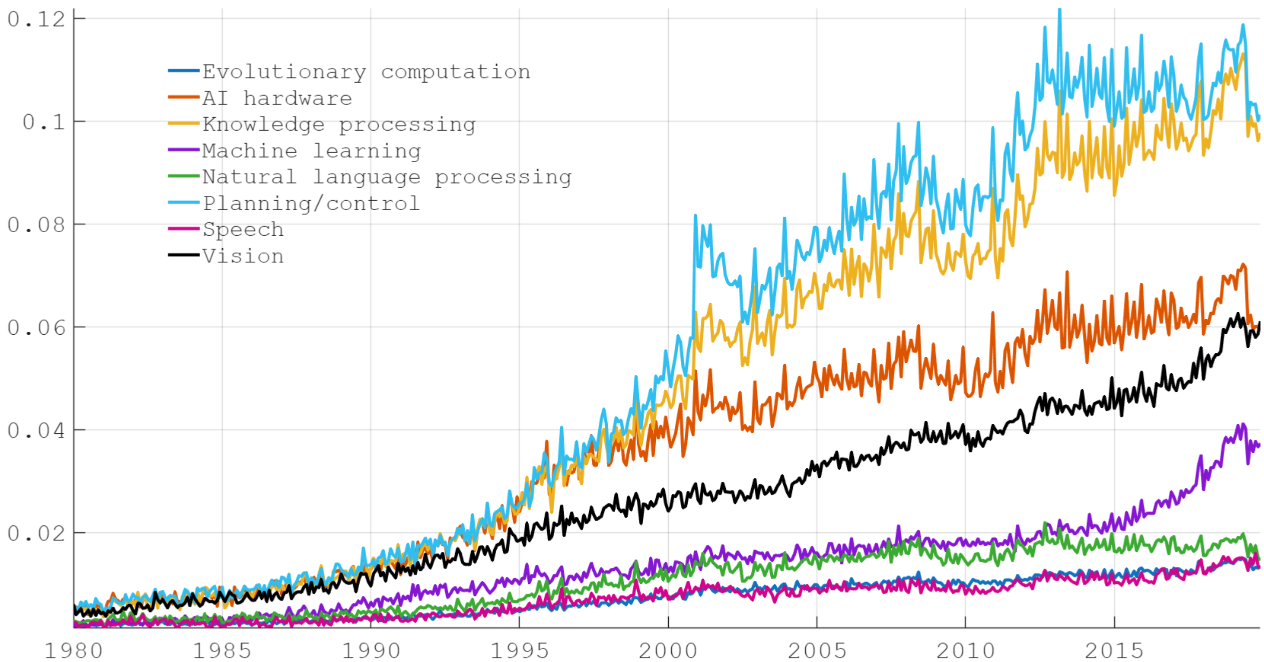

Figure A.3: Share of AI-intensive patents by field

Note. The figure displays AIint by AI field between 1980 and 2019. Source: authors' calculation on the AIPD dataset

A.2 Other data sources and variable definitions

This section documents the auxiliary datasets used in the empirical analysis and provides definitions consistent with the baseline information set in Section 2 of the main text.

Macro aggregates. GDP and TFP are the monthly aggregate constructed by Miranda-Agrippino et al. (2025). Industrial production is the Federal Reserve Board index for total industrial production (FRED:INDPRO). Prices are measured using the headline Personal Consumption Expenditures (PCE) price index (FRED:PCEPI) and its core counterpart excluding food and energy (FRED:PCEPILFE). Employment is total nonfarm payroll employment (FRED:PAYEMS). Average weekly hours are average weekly hours of production and nonsupervisory employees, total private (FRED:AWHNONAG). Hourly earnings are average hourly earnings of production and nonsupervisory employees, total private (FRED:AHETPI). In robustness exercises, we also use CPI measures: headline CPI (FRED:CPIAUCSL) and core CPI excluding food and energy (FRED:CPILFESL).

Financial variables. The short- and long-term interest rates are the 1-year and 10-year U.S. rates (FRED:GS1 and FRED:GS10). In robustness exercises, we also use the NASDAQ Composite (FRED:NASDAQCOM) and the Dow Jones Industrial Average (FRED:DJIA).

Uncertainty and credit conditions (robustness controls). Economic policy uncertainty is measured using the U.S. Economic Policy Uncertainty index (FRED:USEPUINDXM). When we proxy credit conditions using corporate bond spreads, we use the ICE BofA U.S. corporate option-adjusted spread (FRED:BAMLC0A0CM).

Labor market flows. Job openings, separations, layoffs and discharges, quits, and hires are taken from the Job Openings and Labor Turnover Survey (JOLTS): job openings (FRED:JTSJOL), total separations (FRED:JTSTSL), layoffs and discharges (FRED:JTSLDL), quits (FRED:JTSQUL), and total hires (FRED:JTSHIL).

Sectoral outcomes and AI exposure. Sectoral employment, hours, wages, and labor market flows are mapped to a consistent industry classification. Sector-level exposure to AI follows Acemoglu et al. (2022), who aggregate occupation-level AI exposure measures to industries using vacancy-based weights. Sectors are ordered by this exposure measure in Figure D.5.

Patent classifications beyond AI. ICT patents are identified via CPC-based ICT definitions following Inaba and Squicciarini (2017). Robotics patents are identified using official USPTO CPC robotics codes. Automation patents follow the taxonomy constructed by Mann and Püttmann (2023) (available with reliable coverage through 2014). When comparing across technology classes, we express each group as a share of total patents to maximize comparability across time.

Patent quality measures. Forward citations are the count of subsequent patent citations received by each patent. Importance is the text-based novelty/impact measure of Kelly et al. (2021). Patent value follows Kogan et al. (2017) and is based on the stock market reaction to patent grant events; when using the imputed value measure, we rely on the version that addresses same-day multiple-grant attribution.

Distributional accounts. Income and wealth shares are taken from the distributional national accounts compiled by Blanchet et al. (2022). We use the publicly available series for top and bottom groups and study both shares and (when available) real levels to distinguish redistribution from absolute gains.

A.3 Generative-AI patents

This appendix describes the construction of the generative-AI patent series used in Section 6 of the main text. The purpose of the exercise is to connect the historical evidence based on $AIint$ to the recent generative-AI wave. The series is not meant to replace the baseline AI measure; rather, it isolates a narrower set of patents most directly connected to generative AI and asks whether this subset produces similar macroeconomic dynamics.

We identify generative-AI patents through a text- and classification-based search of the U.S. patent corpus in Google Patents BigQuery. The search follows the spirit of the WIPO *Patent Landscape Report: Generative Artificial Intelligence* (WIPO, 2024), which defines generative AI through a combination of model classes, data modes, and application areas rather than through a single patent class. In particular, we build a high-precision query around explicit generative-AI signals in patent titles, abstracts, and claims, together with relevant patent classifications. The query targets both general references to generative AI and specific model families emphasized in the WIPO taxonomy, including generative adversarial networks, variational autoencoders, diffusion models, autoregressive models, large language models, transformers, and text, image, video, code, and speech generation.

This procedure should be interpreted as a WIPO-style search rather than a replication of the full WIPO pipeline. WIPO combines patent searches, AI-assisted prompts, and a classifier to construct its global GenAI patent landscape. We instead use the search component of that taxonomy to build a transparent and reproducible U.S. patent seed. We clean and deduplicate the resulting records at the patent-document level and validate the set against landmark generative-AI patents, including the foundational transformer patent (US10452978B2) as well as early generative adversarial network and diffusion-model patents. The search identifies 1,008 generative-AI patent documents in the 1980m1–2019m12 estimation sample.

Let B_t denote the number of matched generative-AI patents filed in month t and N_t the total number of AIPD-scored patents filed in the same month. Our preferred measure is the monthly share

$$GenAI_t \;=\; 100 \times \frac{B_t}{N_t}. \tag{22}$$

This normalization abstracts from fluctuations in overall patenting activity, in analogy with the extensive-margin specification in the main text; accordingly, the share specification does not include additional controls for total patenting.

The local-projection exercise in the main text replaces $AIint_t$ with the standardized generative-AI share in (22), using the same baseline set of controls. For comparability across exercises, the impulse responses are normalized so that the peak GDP response equals 1 percent. As a robustness check, Section D.7 repeats the exercise using the raw number of matched generative-AI patents, controlling for total patenting.

Overall, the generative-AI patent series should be interpreted as a narrow, high-precision validation exercise. Its advantage is that it is anchored in patents directly identified as generative-AI related and then matched to AIPD. We use the series to assess whether patents most directly connected to generative AI display macroeconomic dynamics similar to those of the broader historical AI measure.

Validation using landmark patents. As a qualitative check on the search, we verify that it retrieves patent documents associated with recognizable milestones and applications of generative AI. Beyond the Transformer patent already mentioned in the main text (US10452978B2), the retrieved set includes Google's "Generating music with deep neural networks" (US10068557B1), eBay's "Generating a digital image using a generative adversarial network" (US10552714B2), and Preferred Networks' "Generative machine learning systems for drug design" (US10776712B2). As summarized in Table A.2, these examples span sequence transduction, audio and music generation, image generation with generative adversarial networks, and scientific applications of latent-variable generative models. They are reported only as a qualitative validation: the empirical measure uses all matched generative-

AI patent documents and does not rely on any individual patent.

Table A.2: Examples of generative-AI patent documents retrieved by the search

Technology family	**Example patent document**	**Applicant**	**Validation role**
Transformer / sequence generation	"Attention-based sequence transduction neural networks" (US10452978B2; 2018)	Google LLC	Benchmark patent for the Transformer architecture underlying modern large language models.
Neural music generation	"Generating music with deep neural networks" (US10068557B1; 2017)	Google LLC	Neural generation of musical audio; links the search to generative audio.
GAN-based image generation	"Generating a digital image using a generative adversarial network" (US10552714B2; 2018)	eBay Inc.	Image generation via generator–discriminator (GAN) architectures.
Generative models for drug design	"Generative machine learning systems for drug design" (US10776712B2; 2016)	Preferred Networks Inc.	Scientific application of latent-variable generative models.

Note. The table reports examples of generative-AI patent documents retrieved by the WIPO-style search and contained in the seed used in the empirical exercise. The examples serve only as a qualitative validation and do not define the series; the empirical measure uses all matched generative-AI patent documents, dated by filing month. Filing years in parentheses.

B Examples of AI-intensive patents

This appendix reports examples of historically significant AI-intensive patents contained in the AIPD-based measure. Early examples include IBM's 1980 patent titled "Parallel pattern verifier with dynamic time warping" (US4348553A), which consisted of a speech recognition system employing a network of elementary local decision modules to match observed time-varying speech patterns against stored prototypes. Shortly thereafter, in 1982, Analis SA filed "Standard hardware-software interface for connecting any instrument which provides a digital output stream with any digital host computer" (US4485439A), relating to the connection of automated laboratory apparatus to clinical laboratory or hospital information systems. In 1984, a private individual affiliated with the University of Michigan filed "Adaptive computing system capable of learning and discovery" (US4697242A), describing an algorithm that employed a language of conditional statements, called "classifiers", to operate on fixed-length binary words. This was followed in 1985 by Standard Systems Corp.'s "Artificial intelligence system" (US4670848A) for accepting, understanding, and responding to a statement. The decade closed with advancements in linguistics, such as Texas Instruments Inc.'s 1989 patent "Chart parser for stochastic unification grammar" (US4984178A), which focused on spoken language processors incorporating rule and observation probabilities.

As the field matured into the 1990s and 2000s, the focus shifted toward classification and search. In 1995, Lucent Technologies Inc. filed "Method and apparatus for training a text classifier" (US5675710A) to classify text by employing a supervised learning algorithm, including applications for the Web. In 2003, IBM filed "Information search using knowledge" (US6636848B1) for automatic sourcing within a corpus of documents. Audio processing also saw refinement, as evidenced by 21CT Inc.'s 2009 patent, "Hidden Markov model for speech processing with training method" (US9020816B2), which describes a method for identifying non-language speech sounds in an audio signal.

More recent patents reflect the era of big data and conversational AI. In 2013, Google filed "Data driven word pronunciation learning and scoring with crowd sourcing based on the word's phonemes pronunciation scores" (US9741339B2) for determining pronunciations for particular terms. Finally, illustrating the rise of modern language models, IBM filed a patent in 2018 for "Training data expansion for natural language classification" (US10726204B2) for training a natural language classifier associated with a computer system chat interface.

C Construction of the IV for AI Technology Shocks

This appendix provides additional details on the construction of the monthly external instrument introduced in Section 2.3 of the main text. Following Miranda-Agrippino et al. (2025), we isolate unexpected variation in our AI intensity measure by regressing $\Delta AIint_t$ on its own lags, monetary policy shocks, and survey expectations:

$$\Delta AIint_t = c + \sum_{j=1}^{12} \gamma_j \Delta AIint_{t-j} + \sum_{h \in \{1Q,1Y\}} \beta_h \mathbb{E}_t[x_{t+h}] + \delta \eta_t^{RR} + z_t, \tag{23}$$

where $\mathbb{E}_t[x_{t+h}]$ denotes Blue Chip forecasts for CPI, real GDP, and unemployment at horizons of one quarter and one year, and η_t^{RR} is the Romer and Romer (2004) monetary policy shock.

The residual z_t serves as our external instrument, capturing unpredictable movements in AI intensity that are orthogonal to macroeconomic expectations and monetary policy.

Table C.1 reports the estimation results. The first 12 lags of $\Delta AIint_t$ are jointly significant, indicating substantial persistence in AI intensity changes. The coefficients on CPI and GDP forecasts suggest some predictability from expectations, which the instrument purges. The monetary policy shock (RRMP) is not statistically significant, indicating limited contamination from monetary policy.

Table C.1: Instrument construction (monthly)

	Dependent variable: $\Delta AIint_t$
Constant	-0.014
	(0.06)
$\Delta AIint_{t-1}$	-0.668***
	(0.06)
$\Delta AIint_{t-2}$	-0.420***
	(0.07)
$\Delta AIint_{t-3}$	-0.266***
	(0.08)
$\Delta AIint_{t-4}$	-0.169**
	(0.07)
$\Delta AIint_{t-5}$	-0.075
	(0.08)
$\Delta AIint_{t-6}$	-0.006
	(0.08)
$\Delta AIint_{t-7}$	-0.008
	(0.08)
$\Delta AIint_{t-8}$	0.048
	(0.08)
$\Delta AIint_{t-9}$	0.112
	(0.08)
$\Delta AIint_{t-10}$	0.137*
	(0.08)
$\Delta AIint_{t-11}$	0.154*
	(0.08)
$\Delta AIint_{t-12}$	0.037
	(0.07)
η_t^{RR} (Romer-Romer)	-0.059
	(0.04)
$\mathbb{E}_t[CPI_{t+1Q}]$	0.040**
	(0.02)
$\mathbb{E}_t[CPI_{t+1Y}]$	-0.045***
	(0.02)
$\mathbb{E}_t[GDP_{t+1Q}]$	0.020***
	(0.01)
$\mathbb{E}_t[GDP_{t+1Y}]$	-0.005
	(0.01)
$\mathbb{E}_t[Unemp_{t+1Q}]$	-0.041
	(0.04)
$\mathbb{E}_t[Unemp_{t+1Y}]$	0.045
	(0.05)
F-statistic	8.20
p-value	0.000
Adjusted R^2	0.290
Observations	387

Note. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. *Sample: 1982m1–2014m12.*

D Additional Results

This appendix collects additional empirical results that complement the baseline evidence in the main text. The appendix serves two purposes. First, it provides further evidence on the mechanisms underlying the aggregate responses to AI technology shocks. Second, it reports validation and robustness exercises that support the use of the baseline OLS local-projection specification in the quantitative and heterogeneity analyses.

We begin by examining the response of consumer prices across expenditure categories, showing that the aggregate decline in prices masks substantial heterogeneity across goods and services. We then report impulse responses by AI field and compare the effects of innovations to different moments of the AI-intensity distribution, including the extensive margin and the dispersion of AI scores. Next, we present additional sectoral labor-market responses, which complement the discussion of labor-market transmission in Section 5 of the main text.

The final part of the appendix reports the OLS local-projection estimates, the sufficient-information tests, and a set of alternative specifications. These exercises assess whether the baseline dynamics are robust to different controls, transformations, and measures of AI-related innovation. We also report the absolute responses of income and wealth, which complement the distributional results in Section 5 of the main text, and the generative-AI raw-count robustness exercise.

D.1 Heterogeneity in Consumer Prices

This subsection analyzes the heterogeneous effects of AI technology shocks on consumer prices across expenditure categories, complementing the aggregate PCE results in Section 3.1 of the main text.

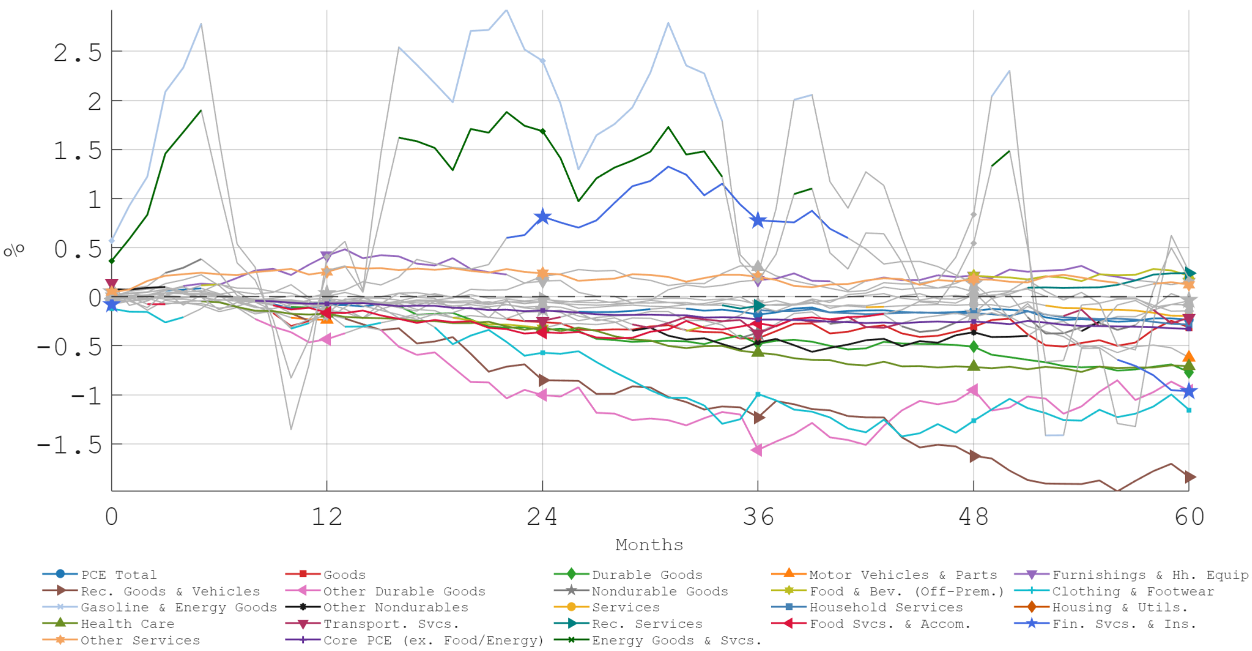

Figure D.1: Consumer prices response to an AI technology shock.

Note. The figure displays the IRFs to an AIint shock. Sample 1980–2019. The estimates are based on local projections with Newey–West standard errors. Point estimates. Colored lines and markers denote significance at 68% confidence.

Figure D.1 shows impulse responses of prices across consumption categories. For most items, prices decline gradually and persistently, consistent with the typical effects of technology shocks. The strongest response is concentrated in durable "recreational goods and services," a category that includes many high-tech products such as information-processing equipment (computers and software), video and audio equipment, and photographic equipment, as well as "other durable goods." A plausible mechanism is that AI-related innovations improve the efficiency of software and information-processing technologies, reducing production costs and lowering (quality-adjusted) prices for durables that rely on these inputs. Consistent with this interpretation, many AI-intensive patents are filed by large technology firms (e.g., IBM, Microsoft, Google), and efficiency gains may pass through relatively quickly to affected hardware and software products.

Conversely, energy prices surge in response to an AI technology shock. This is consistent with the energy intensity of digital technology and suggests that, at least in the short run, AI technology shocks are energy-intensive rather than energy-saving (Känzig and Williamson, 2025).

D.2 Responses by AI Field and AI-Intensity Moments

This subsection reports additional evidence on whether the aggregate responses differ across AI fields and across alternative dimensions of the AI-intensity distribution.

Figure D.2 reports the responses of GDP, consumer prices, employment, and TFP to shocks constructed from each of the eight AI fields in turn, holding the baseline specification fixed; the colored lines highlight individual fields against the grey envelope of the remaining categories. Two features stand out. First, the field-level responses are broadly consistent with one another and reproduce the aggregate pattern documented in the main text: a delayed expansion in GDP, employment, and TFP alongside a gradual decline in consumer prices. Second, the comovement across fields indicates that the aggregate $AIint$ result is not driven by any single AI domain, but reflects a supply-side dynamic common to the technological components of AI innovation. Some heterogeneity remains in the timing and, for individual fields such as machine learning, in the short-run sign of the response, but the medium-run dynamics align across fields.

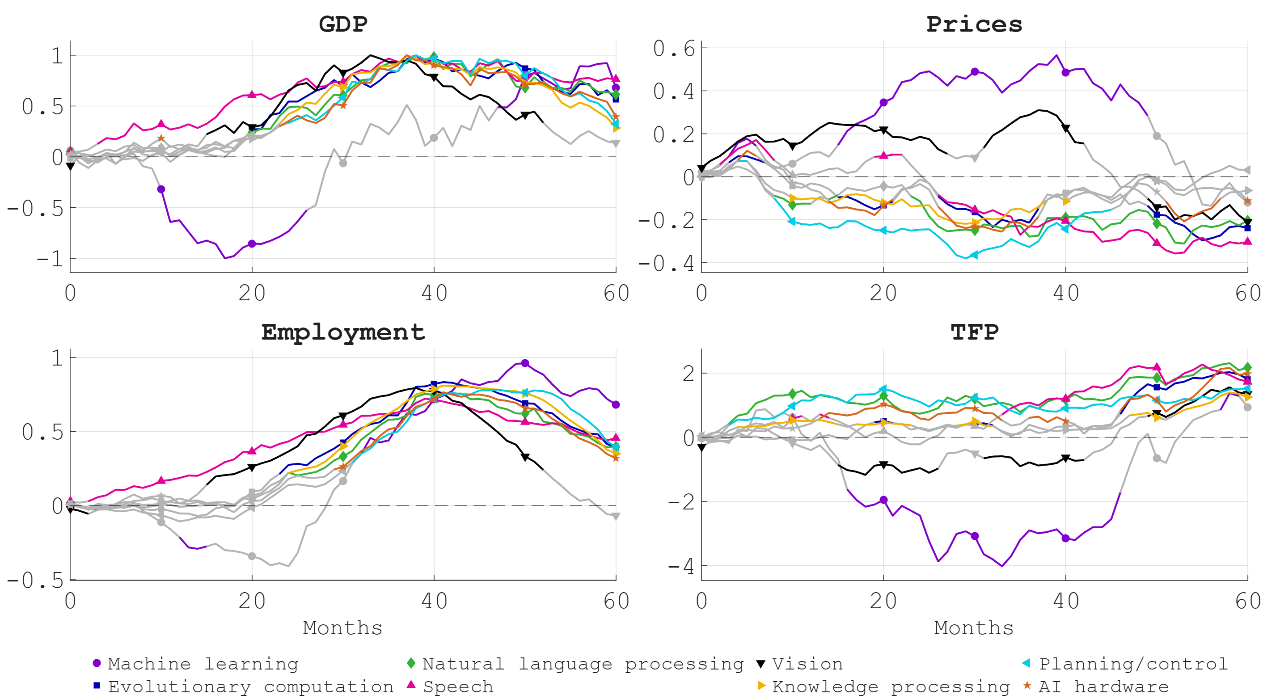

Figure D.2: IRFs by AI categories.

Note. The figure displays the IRFs to an AIint shock. Sample 1980–2019. The estimates are based on local projections with Newey–West standard errors. Point estimates and 68%–90% confidence bands.

Figure D.3 repeats the exercise replacing GDP with industrial production. The field-level responses again line up with the aggregate, confirming that the expansionary supply-side pattern is not specific to the choice of activity measure.

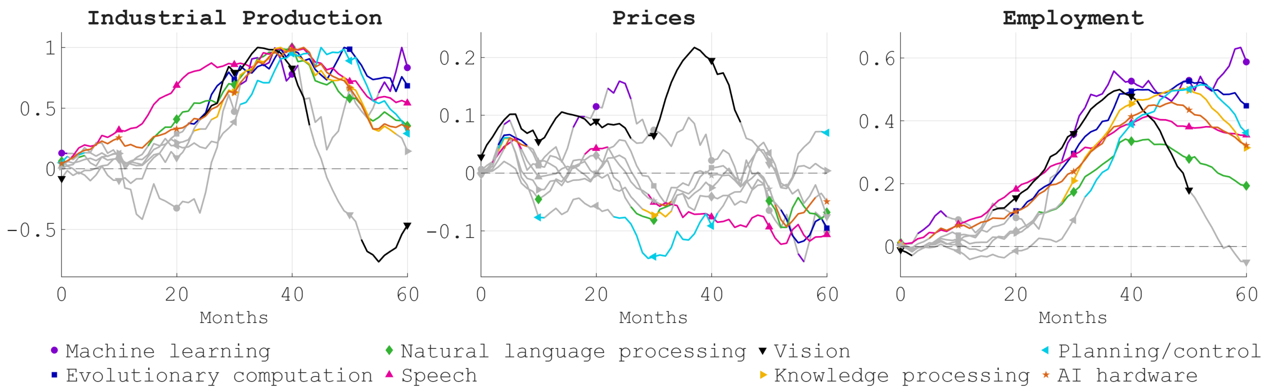

Figure D.3: IRFs by AI categories – Specification with Industrial Production.

Note. The figure displays the IRFs to an AIint shock. Sample 1980–2019. The estimates are based on local projections with Newey–West standard errors. Point estimates and 68%–90% confidence bands.

Figure D.4 compares the responses to three innovations: the aggregate measure $AIint_t$, the cross-sectional standard deviation of patent-level AI scores, and the extensive margin s_t. As anticipated by the variance decomposition in Section A.1.2, the three sets of responses are nearly identical. The dispersion of AI scores therefore carries essentially no information beyond the extensive margin: the aggregate effects of AI innovation operate through the diffusion of AI across the broader pool of patents, rather than through increasing polarization within already AI-intensive patents. This reinforces the extensive-margin interpretation of the decomposition in Section 4.2 of the main text.

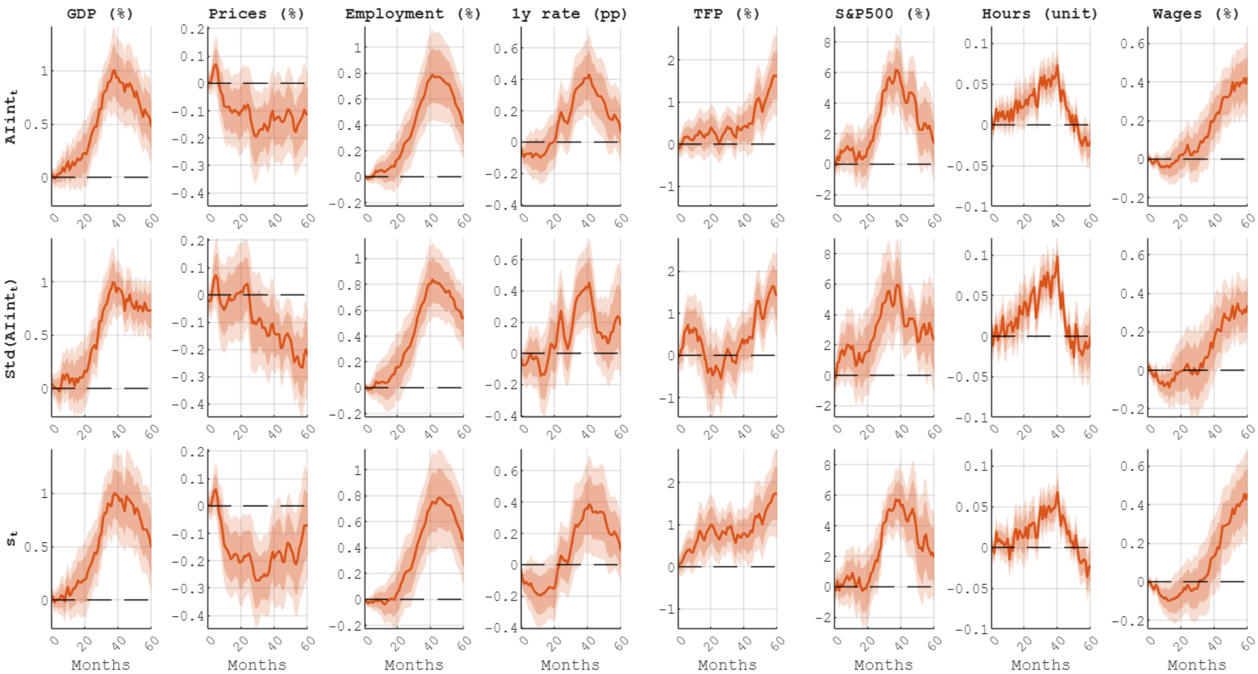

Figure D.4: Overall standard deviation versus extensive margin comparison.

Note. The figure compares IRFs to innovations in aggregate AI intensity ($AIint_t$), the cross-sectional standard deviation of patent-level AI scores, and the extensive margin s_t defined as the share of AI-intensive patents. Sample 1980–2019. The estimates are based on local projections with Newey–West standard errors. Point estimates and 68%–90% confidence bands.

D.3 Sectoral Labor-Market Responses

This subsection reports additional sectoral labor-market responses, complementing the evidence in Section 5.1 of the main text.

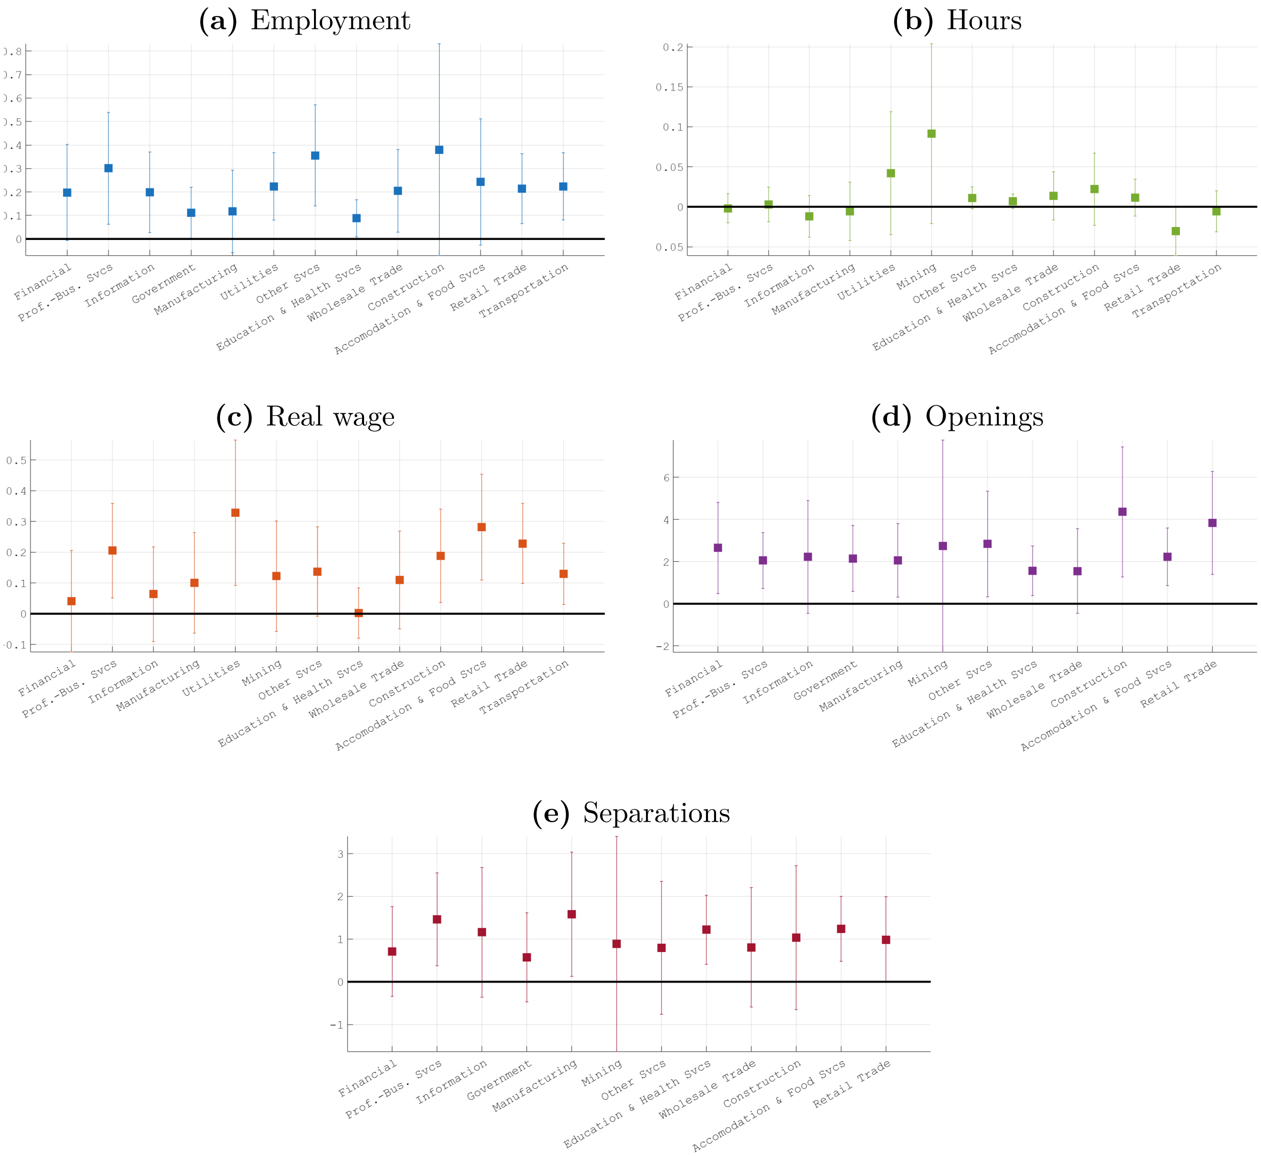

Figure D.5: Sectoral labor market responses – Average.

Note. The figure displays the cumulated IRFs over 60 months to an AIint shock. Sample 2006–2019. The estimates are based on local projections with Newey–West standard errors. Point estimates and 90% confidence bands.

Figure D.5 reports the cumulated 60-month responses of sectoral employment, hours, real wages, job openings, and separations, with sectors ordered by AI exposure. Consistent with Section 5.1 of the main text, the employment gains are broad-based: they are positive across essentially all sectors and do not decline with AI exposure, underscoring the dominance of general-equilibrium forces over sector-specific displacement. The responses of real wages and job openings are somewhat more dispersed across sectors—with some lower-exposure sectors experiencing stronger wage growth—but the overall picture confirms a widespread expansion rather than a reallocation concentrated in the least exposed industries.

D.4 OLS LP Results and Sufficient Information Tests

This subsection reports the OLS local-projection estimates and the sufficient-information tests used to assess whether the baseline information set captures the main predictable movements in $AIint$.

Figure D.6 reports the OLS local-projection responses estimated from the baseline specification in eq. (2) of the main text. They closely track the LP-IV responses of Figure 2 in the main text: GDP and TFP rise with a delay, consumer prices decline, and employment, hours, and hourly earnings increase, while stock prices respond earlier than real variables. The close correspondence between the OLS and LP-IV estimates supports the use of the OLS specification—on which the variance decomposition and the comparisons across technology classes rely—throughout the paper.

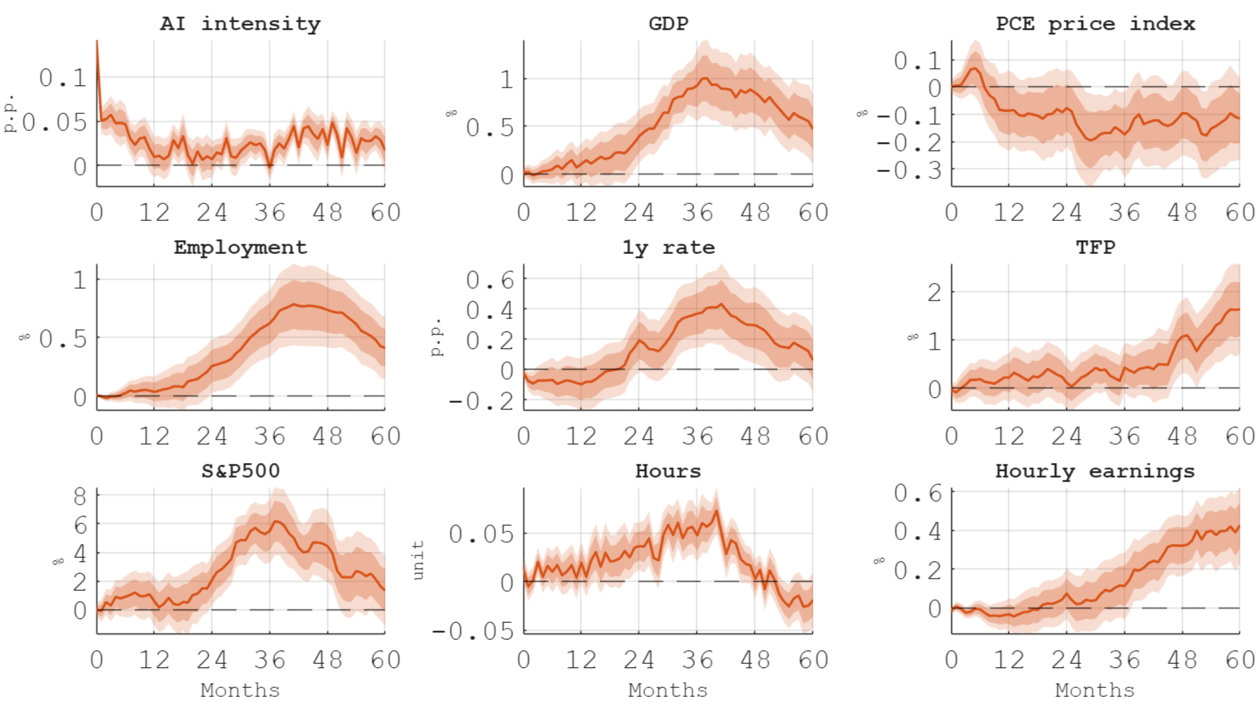

Figure D.6: OLS LP: IRFs to an AI technology shock.

Note. The figure displays the IRFs to an AIint shock estimated by OLS LP specified in eq. (2) of the main text. Sample 1980m1–2019m12. The estimates are based on local projections with Newey–West standard errors. Point estimates and 68%–90% confidence bands.

Table D.1: Sufficient Information Test

Panel A: Summary

Test Description	P-value (monthly)	P-value (quarterly)
Model Specification Tests		
(i) LP Specification	0.255	-
Orthogonality Tests		
(ii) All Shocks	0.422	0.662
(iii) SPF Forecasts	0.773	0.782
(iv) All Shocks + SPF	0.465	0.801

Panel B: Individual Shocks

Monthly Shocks

Shock Series	p-value
Känzig (2021) oil supply news	0.611
Baumeister and Hamilton (2019) oil supply	0.246
Gilchrist and Zakrajšek (2012) credit	0.387
Gertler and Karadi (2015) mon pol.	0.927
Romer and Romer (2004) mon pol.	0.677
Miranda-Agrippino et al. (2025) tech. news	0.387

Quarterly Shocks

Shock Series	p-value
Romer and Romer (2010) tax	0.257
Ramey (2011) fiscal	0.132
Mertens and Ravn (2013) priv. tax	0.337
Mertens and Ravn (2013) corp. tax	0.532
Basu et al. (2006) TFP	0.567
Barsky and Sims (2011) news	0.368
Kurmann and Otrok (2013) news	0.273
Beaudry and Portier (2014) news	0.757
Smets and Wouters (2007) TFP	0.192

Note. Panel A reports the summary results of the orthogonality test at the monthly and quarterly frequency (based on averaged shocks). Column (i) reports p-values for the F-test of joint significance on the coefficients associated with the set of controls from our baseline LP specification, excluding $AIint$ own lags and the constant. Rows (ii)–(iv) report p-values from an orthogonality test where the implied shocks to $AIint$ are regressed on external information. SPF stands for Survey of Professional Forecasters: one and four-quarter-ahead forecasts for the unemployment rate, the GDP deflator, real non-residential investment, and real corporate net profits, as in Miranda-Agrippino et al. (2025). Individual shocks are reported in Panel B.

Table D.1 reports the associated sufficient-information tests. The baseline controls have no joint explanatory power for $AIint_t$ beyond its own lags (row i), and the implied $AIint$ shocks are orthogonal to a broad set of external macroeconomic shocks and professional forecasts (rows ii–iv and Panel B). Taken together, these results support treating $AIint_t$, conditional on the baseline information set, as exogenous to business-cycle fluctuations.

D.5 Alternative Specifications

This subsection reports additional robustness checks aimed at ruling out spurious drivers of our results, including controls for broader innovation trends, automation and robotics innovation, and financial and uncertainty conditions.

Figure D.7 overlays the baseline responses (black) with those obtained under the battery of alternative specifications discussed in the main text—linear and quadratic detrending, an explicit trend, the ratio of high-tech to industrial stock prices, controls for total patent counts and for robotics patents, and additional uncertainty and credit controls (EPU and EBP). The baseline responses lie well within the range spanned by these alternatives at all horizons, indicating that the main dynamics are not driven by any particular modeling or identification choice.

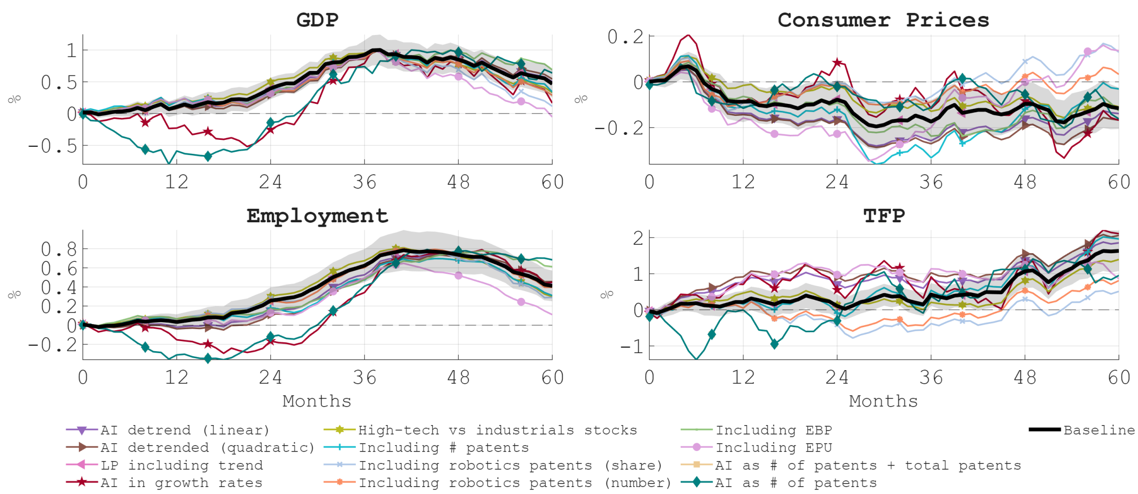

Figure D.7: Robustness Across Specifications

Note. The figure displays the IRFs to a shock to AIint across several alternative specifications. Sample 1980–2019. The estimates are based on local projections with Newey–West standard errors. Point estimates and 68%–90% confidence bands.

Figure D.8 re-estimates the baseline using only variables available at the monthly frequency, replacing GDP with industrial production. The responses are very similar to the baseline, confirming that the results do not hinge on the inclusion of interpolated monthly aggregates.

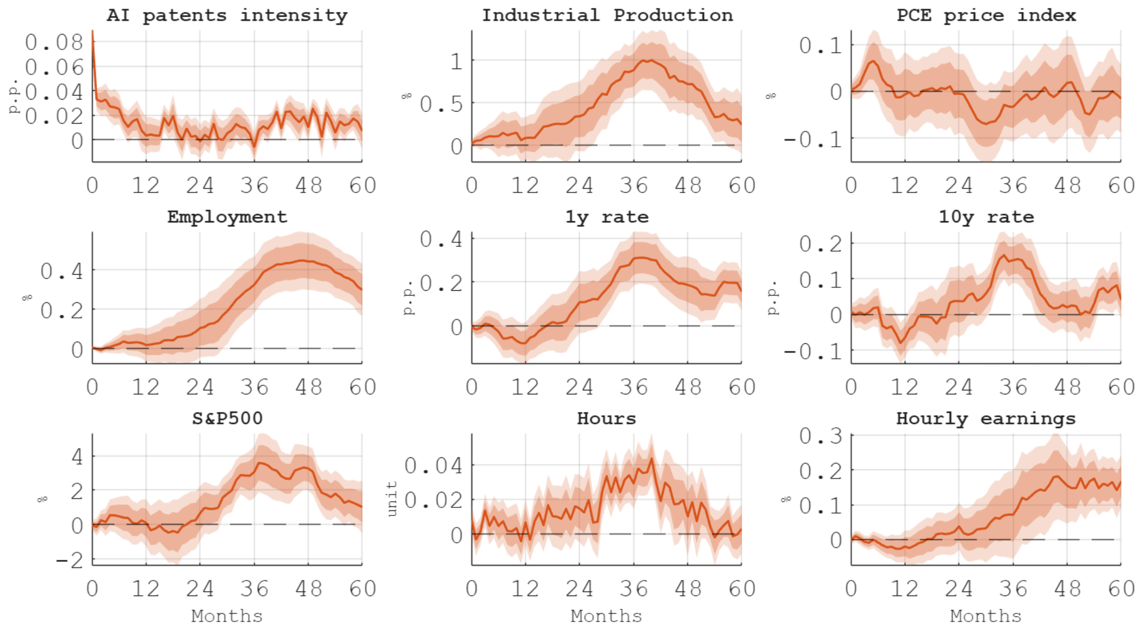

Figure D.8: IRFs with Monthly Observables Only

Note. The figure displays the IRFs to a shock to AIint including only observable variables at the monthly frequency. Sample 1980–2019. The estimates are based on local projections with Newey–West standard errors. Point estimates and 68%–90% confidence bands.

D.6 Absolute Responses of Income and Wealth

This subsection reports the responses of real income and real wealth in levels across the distribution, complementing the distributional share results in of the main text.

Figure D.9 shows that real incomes rise across the distribution, consistent with the aggregate expansion documented above. The real wealth responses, by contrast, are sharply heterogeneous: real wealth rises at the top of the distribution while the bottom 50% experiences an absolute decline. The combination of broad-based income gains and absolute wealth losses at the bottom is the distributional signature emphasized in the main text, and complements the share-based results of Section 5.

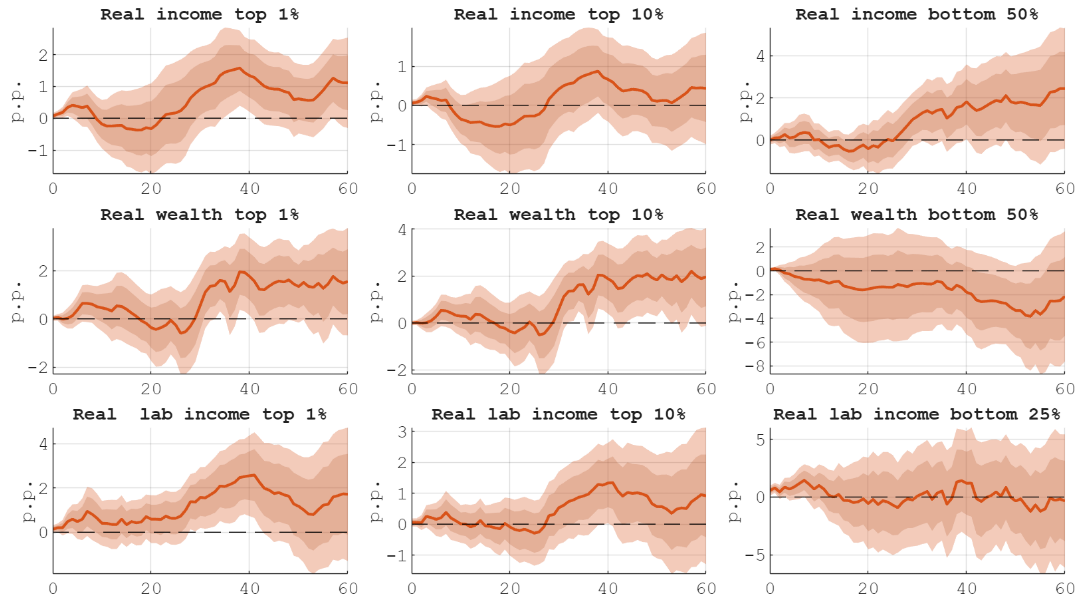

Figure D.9: Absolute Response of Income and Wealth

Note. The figure displays the IRFs to an AIint shock. Sample 1980–2019. The estimates are based on local projections with Newey–West standard errors. Point estimates and 68%–90% confidence bands.

D.7 Generative-AI patents: raw-count specification

The main text uses the generative-AI patent share in (22). As a robustness exercise, we repeat the local projection using the raw number of matched generative-AI patents, B_t. Since this variable is not expressed as a share, the specification includes twelve lags of total AIPD-scored patenting, measured as $100 \log(1 + N_t)$, to absorb fluctuations in the size of the patent universe while leaving the shock variable in raw-count units.

Figure D.10 reports the results. The responses are qualitatively similar to those obtained with the share measure in the main text, indicating that the macroeconomic pattern is not an artifact of the normalization by total patenting.

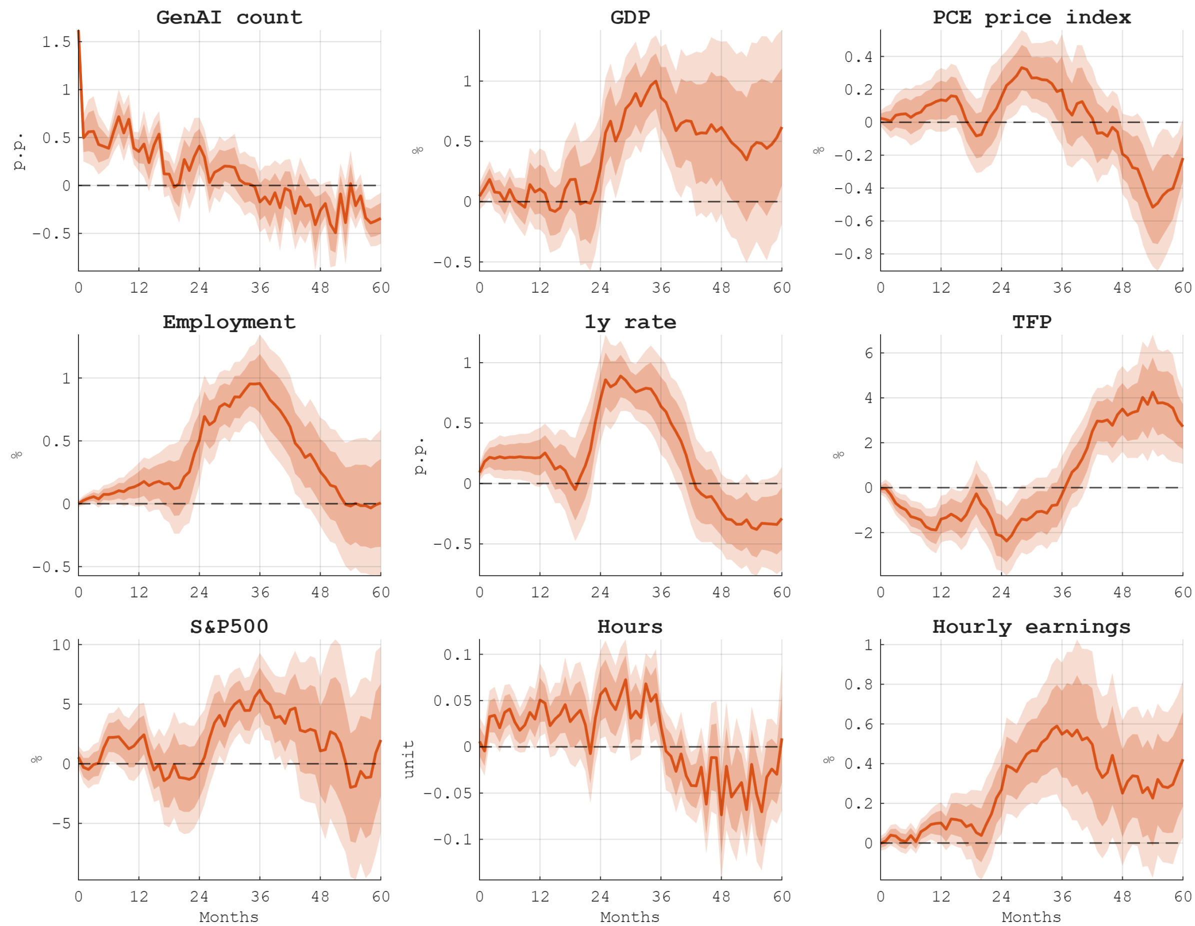

Figure D.10: IRFs to a generative-AI raw-count shock.

Note. The figure displays the IRFs to a shock to the raw number of matched generative-AI patents, B_t, dated by filing month and standardized before estimation. The specification includes the baseline monthly macroeconomic information set and twelve lags of total AIPD-scored patenting, measured as $100 \log(1 + N_t)$. Sample: 1980m1–2019m12. The estimates are based on local projections with Newey–West standard errors. Point estimates and 68%–90% confidence bands. IRFs are normalized so that the peak GDP response equals 1 percent.