The Macroeconomic Effects of AI Technology Shocks *

Andrea Gazzani† Filippo Natoli‡

First Draft: July 14, 2024
Current Draft: March 26, 2026

Abstract

Using detailed data on the artificial intelligence (AI) content of all U.S. patents granted between 1980 and 2019, we construct a novel measure of AI intensity in technological innovation, which we use to identify unanticipated shifts in AI technology. AI technology shocks generate a delayed rise in productivity and economic activity, alongside a decline in consumer prices, consistent with a positive supply shock. Employment, hours worked, and wages also increase, even in sectors most exposed to AI, suggesting that general equilibrium effects dominate displacement. Yet, unlike general technology shocks, AI tech shocks compress the labor share. The expansionary effects of AI-intensive innovation are also accompanied by growing inequality, particularly stark in the wealth distribution.

JEL: E32, O33, O34, C36

Keywords: artificial intelligence, patents, innovation, technology shocks, business cycle

*We wish to thank Francesco Bianchi, Dario Caldara, Luisa Carpinelli, Jorgo Georgiadis, Silvia Miranda-Agrippino, Joseba Martinez, Pablo Ottonello, Paolo Surico, and Giovanni Veronese for fruitful discussions. This paper also benefited from feedbacks 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.

†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

John McCarthy, a founding father of Artificial Intelligence (AI), defined it as *"the science and engineering of making intelligent machines, especially intelligent computer programs"* (McCarthy, 2007). Recent breakthroughs such as generative AI and large language models have brought the economic implications of AI to the forefront of academic and policy debates. Like the steam engine, electricity, and the internet before it, AI is increasingly seen as a general-purpose technology (GPT; Bresnahan and Trajtenberg, 1995) with the potential to catalyze a new era of innovation and economic growth given its widespread sectoral effects (Cockburn et al., 2019; Goldfarb et al., 2023). While a nascent empirical literature has begun to explore the firm-level and sectoral impacts of AI adoption, its full macroeconomic effects remain largely uncharted.

In this paper, we assess the macroeconomic consequences of AI technology shocks in the United States. Leveraging the granular data in Giczy et al. (2022) on the AI content of every patent granted by the U.S. Patent and Trademark Office since 1980, we construct a novel aggregate measure of *AI intensity in U.S. innovation*, reflecting the salience of artificial intelligence in patented technologies. We use this measure to identify AI technological shocks orthogonal to fluctuations in the business cycle, changes in financing conditions, equity prices, or expectations regarding key economic aggregates. Our primary identification strategy constructs an instrument for AI technology following Miranda-Agrippino et al. (2025) and estimates dynamic effects via local projection instrumental variables (LP-IV). We also apply our measure within a local projection framework incorporating a comprehensive set of macro-financial controls; both approaches yield consistent results.

Our analysis reveals that an increase in the AI intensity of innovation acts as a positive supply shock, leading to a delayed but persistent surge in economic activity and total factor productivity. These expansionary effects lead to a notable increase in total employment, average hours worked, and wages. Importantly, these phenomena are broad-based, occurring even in sectors where jobs are most exposed to AI, consistently with Aghion et al. (2025b). This underscores the powerful role of productivity gains and general equilibrium forces, which appear to outweigh displacement effects (*the missing intercept problem*; Wolf, 2023; Sarto, 2018; Matthes et al., 2024). This net positive outcome occurs even as the labor force un-

dergoes a recomposition; an analysis of occupational flows shows that while layoffs increase, they are more than offset by a surge in new job openings.

AI technology shocks compress overall consumer prices. A sectoral analysis reveals that most components of the PCE price index decline. The fall is particularly stark in high-tech goods such as computers, hardware and software, audio and video equipment. Energy prices constitute a notable exception in line with the energy intensity of AI technology.

Contrasting AI technology shocks with general technological development (Miranda-Agrippino et al., 2025), we find that, despite the positive effects on the labor market, AI technology shocks compress the labor share. Conversely, general technological shocks are labor-neutral (Fisher, 2006) and do not significantly affect the labor share.

We then examine the association of AI-intensive patents with automation and ICT patents. We find that nearly 80% of AI-intensive patents are ICT patents, and half are classified as automation patents Mann and Püttmann (2023). Conversely, AI-intensive patents constitute only a modest share of the broader class of ICT and automation patents. Using leading measures of patent quality such as forward citations, importance Kelly et al. (2021), and value Kogan et al. (2017), we document that AI-intensive patents are higher-quality patents within both the ICT and automation categories. This pattern squares with a comparison of the macroeconomic effects of these innovations. AI technology shocks generate larger expansionary effects on GDP, TFP, and employment, but also depress the labor share to a larger extent.

Despite the aggregate benefits, the gains from AI innovation are not distributed equally across the population. The expansionary shock leads to a rise in the wealth and labor-income shares of the top 10%, while the wealth share of the bottom 50% declines persistently. These distributional consequences can be linked to the asset returns channel (Moll et al., 2022) but they are also apparent across different education levels. The employment gains are widespread, except for workers with less than a high school education, consistent with the concept of skill-biased technological change.

Related Literature. A growing body of research studies the implications of AI diffusion at the firm level or occupation-level (Felten et al., 2018; Webb, 2019; Babina et al., 2024;

Brynjolfsson et al., 2025) or with a macroeconomic perspective (Aghion et al., 2018; Lu and Zhou, 2021; Comunale and Manera, 2024). On the potential macroeconomic effects, views are polarized (Acemoglu, 2024; Aghion et al., 2025b). We present a novel empirical investigation on the aggregate effects of AI intensity in the innovation process. We emphasize the crucial role played by general equilibrium effects (Wolf, 2023; Sarto, 2018; Matthes et al., 2024) and find an overall positive response of employment as in Autor and Salomons (2018) and Mann and Püttmann (2023) and expansionary effects on output and productivity as in the model by Aldasoro et al. (2024).

A broad empirical literature, exemplified among many others by Galí (1999), Beaudry and Portier (2006), Fernald (2012), Bloom et al. (2013), Aghion et al. (2025a) has studied the relationship between technology shocks and total factor productivity. Our paper is particularly related to works that exploit patent data to study the aggregate implications of innovation, such as Kogan et al. (2017), Kelly et al. (2021), Cascaldi-Garcia and Vukotić (2022), Miranda-Agrippino et al. (2025), Ferriani et al. (2024), and Gazzani et al. (2025). Our work examines the intensity of AI technology shocks as a distinct engine of technological development. Unlike general technology shocks, which are labor-neutral and leave the labor share unaffected (Fisher, 2006), AI technology shocks generate a persistent decline in the labor share, highlighting their distinct transmission mechanism.

Finally, our work contributes to the literature on inequality, particularly regarding automation (Acemoglu and Restrepo, 2020a; Prettner and Strulik, 2020; Moll et al., 2022). We show that AI intensity in innovation significantly increases the wealth share of the top 1% while eroding that of the bottom 50%, providing empirical support for the asset valuation channel proposed by Moll et al. (2022).

Structure of the paper. The remainder of the paper is organized as follows. Section 2 details the construction of our novel AI intensity measure and outlines the econometric framework used for identification. Section 3 presents the main results on aggregate macroeconomic dynamics, beginning with the LP-IV baseline, followed by OLS LP estimates and quantitative implications. Section 4 dissects AI innovation by comparing it with general technology shocks, decomposing it into extensive and intensive margins, and benchmarking it against

automation and ICT innovation. Section 5 investigates the transmission through the labor market, analyzing occupational flows, sectoral heterogeneity, and the role of education. Section 6 discusses the distributional consequences for income and wealth. 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 here make use of 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).[1] 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 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 across the whole universe of filed patents each month

[1] 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.

t:

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

where $AIint_{i,t}$ is the average AI score 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 A 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. The lower panel exhibits the share of AI (intensive) patents over total patents and $AIint$ within this sub-group. These two variables represent, respectively, the extensive and intensive margins of AI patent intensity. 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. Appendix B discusses the decomposition of the overall $AIint$ into the intensive and extensive margin components. We also report how $AIint$ has varied by sectoral breakdown – finding the aggregate trend driven mainly by "computer software" and "communication" – and specific AI fields. In Appendix F we describe exercises aimed at ruling out that our findings are spuriously driven by $AIint$ capturing general technological trends or automation.

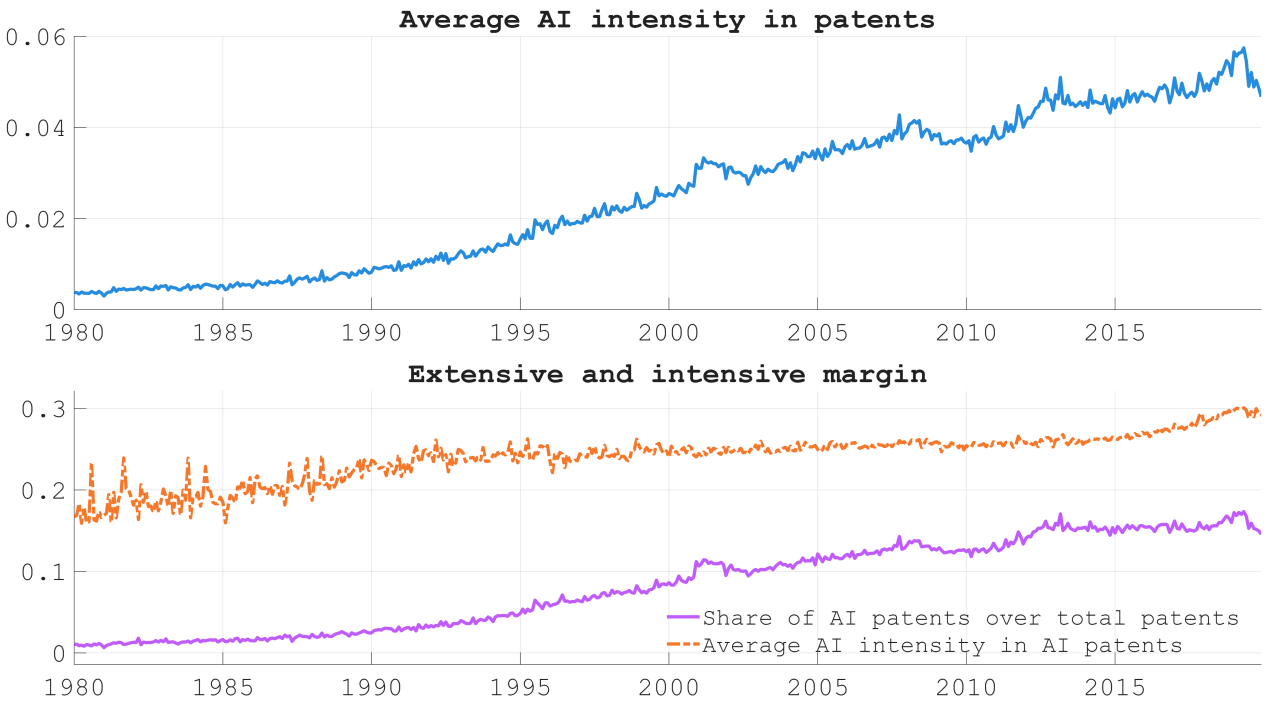


Figure 1: AI INTENSITY IN US PATENTING.

Note. The upper panel displays the average AI intensity in US patents (AIint). The lower panel reports the decomposition of AIint into the contributions of the extensive (share of AI patents) and intensive margins (average AI intensity in AI patents). 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). We choose LPs over VAR methods 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 in favor of LPs (Li et al., 2024). Our baseline identification relies on an instrumental variable (IV) strategy following Miranda-Agrippino et al. (2025), which isolates exogenous variation in *AIint* orthogonal to macroeconomic expectations and monetary policy. We complement this with an internal-instrument approach 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:

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

where y_{t+h} represents the outcome variable in log-levels as (Montiel Olea et al., 2025), α denotes the constant, and $AIint_t$ denotes AI patent intensity. The coefficient β_h captures the dynamic response at horizon h or Impulse Response Functions (IRFs); the inference is based on Newey and West (1994) standard errors. The control vector $\mathbf{w}_t$ includes 12 lags of y_t and $AIint_t$, as well as industrial production, the headline Personal Consumption Expenditures (PCE) price index, employment, weekly hours worked, hourly earnings, the 1- and 10-year Treasury yields, and the S&P500. These controls capture variations due to business cycles, expectations of future economic conditions, and monetary policy, among other factors. Given the time lag inherent in innovation development, the comprehensive set of lagged control variables allows us to isolate unexpected fluctuations in $AIint_t$.

2.3 Instrumental variable construction

Following Miranda-Agrippino et al. (2025), we construct an external instrument for AI technology shocks by isolating unpredictable movements in $AIint$ that are orthogonal to macroeconomic expectations and monetary policy. Specifically, we regress the change in $AIint$ on its own lags, Romer and Romer (2004) monetary policy shocks, and Blue Chip forecasts:

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

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 D.

Our analysis proceeds in five steps. First, we focus on the macroeconomic effects on the US economy, presenting LP-IV baseline responses and confirming them with OLS LP estimates. Second, we dissect AI innovation by comparing it with general technology shocks,

decomposing it into extensive and intensive margins, and benchmarking it against automation and ICT innovation. Third, we delve into the labor market by studying occupational flows, sectoral heterogeneity, and the role of education. Fourth, we investigate the consequences of AI technology diffusion for inequality.

3 Main Results

This section presents the macroeconomic effects of AI technology shocks. We begin with the LP-IV baseline, which provides a clean identification of the causal effects of AI innovation by instrumenting $AIint$ with the residual from equation (3). 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 LP-IV Baseline Effects

Figure 2 displays the aggregate responses to an AI technology shock identified via LP-IV, normalized to generate a 1% peak effect on GDP. The instrument isolates exogenous variation in $AIint$ that is orthogonal to macroeconomic expectations and monetary policy (Section 2.3). The dynamics are consistent with an expansionary technology shock with diffusion lags: (i) all variables respond with delay and (ii) TFP and output increase while consumer prices fall (Portier, 2014; Miranda-Agrippino et al., 2025). Stock prices rise significantly and earlier than real variables, reflecting forward-looking expectations of future productivity gains.

Labor market quantities—employment, hours worked, and wages—increase, consistent with the view that general equilibrium effects dominate displacement, as emphasized in 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 C. 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.

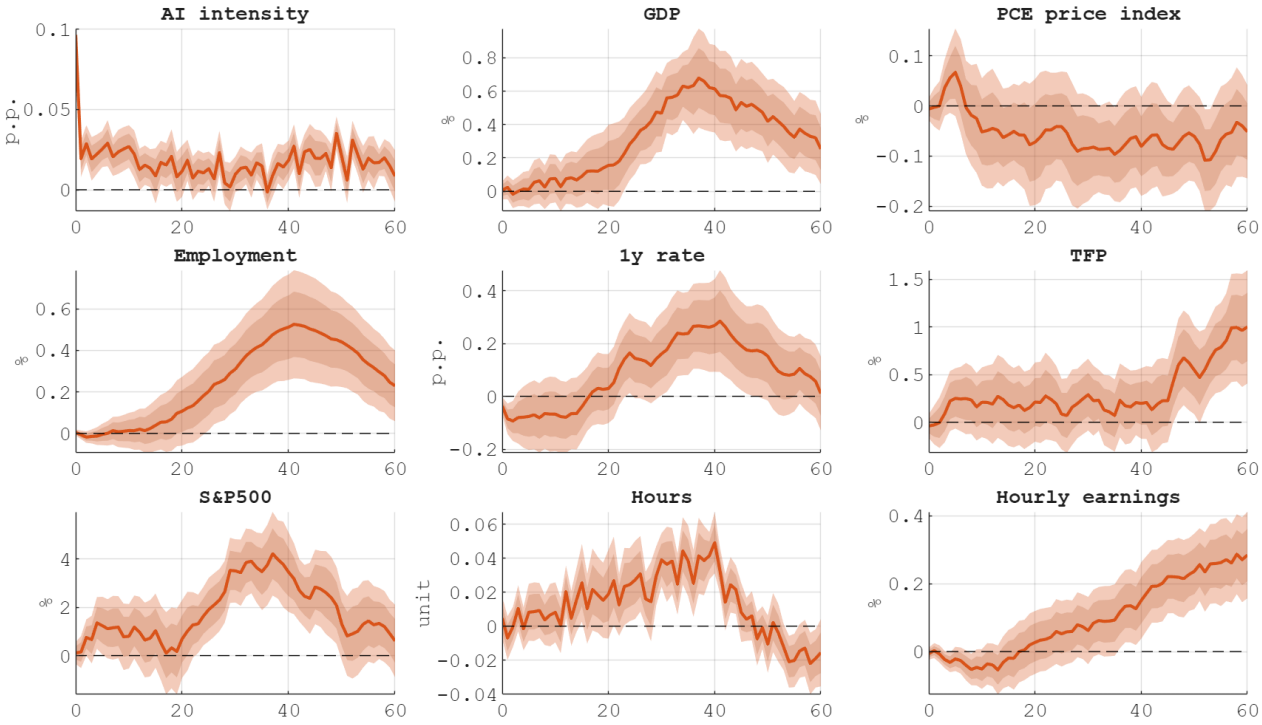


Figure 2: LP-IV: IRFs to an AI technology shock.

Note. The figure displays the IRFs to an AIint shock estimated by LP-IV, where the instrument is constructed following Miranda-Agrippino et al. (2025) as described 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.

3.2 OLS LP and Alternative Specifications

Given the consistency between the LP-IV results and an internal-instrument approach using OLS local projections with extensive controls, in this section we present the OLS LP estimates

and demonstrate the stability of our findings across a range of alternative specifications. The OLS LP approach is particularly useful for the remainder of the paper, as it allows us to easily compute the forecast error variance decomposition (Section 3.3) and compare AI shocks with other technological developments (Section 4) over a longer sample.

3.2.1 OLS LP results

Figure 3 displays the OLS LP impulse responses estimated from equation (2) with the full set of macro-financial controls. The results closely mirror the LP-IV estimates in Figure 2, confirming that the controls in the LP specification successfully absorb potential endogeneity in $AIint$. The same qualitative patterns emerge: a delayed but persistent rise in GDP and TFP, a decline in consumer prices, increases in employment, hours worked, and hourly earnings, and a significant early response of stock prices. The close alignment between the two identification strategies reinforces confidence in a causal interpretation of the estimated effects.

Table 1 provides further support for the validity of the internal-instrument approach. Using the approach 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 $x_t \perp \mathbf{w_t}$ do not exhibit any linear dependence on macroeconomic shocks or projections from professional forecasters.

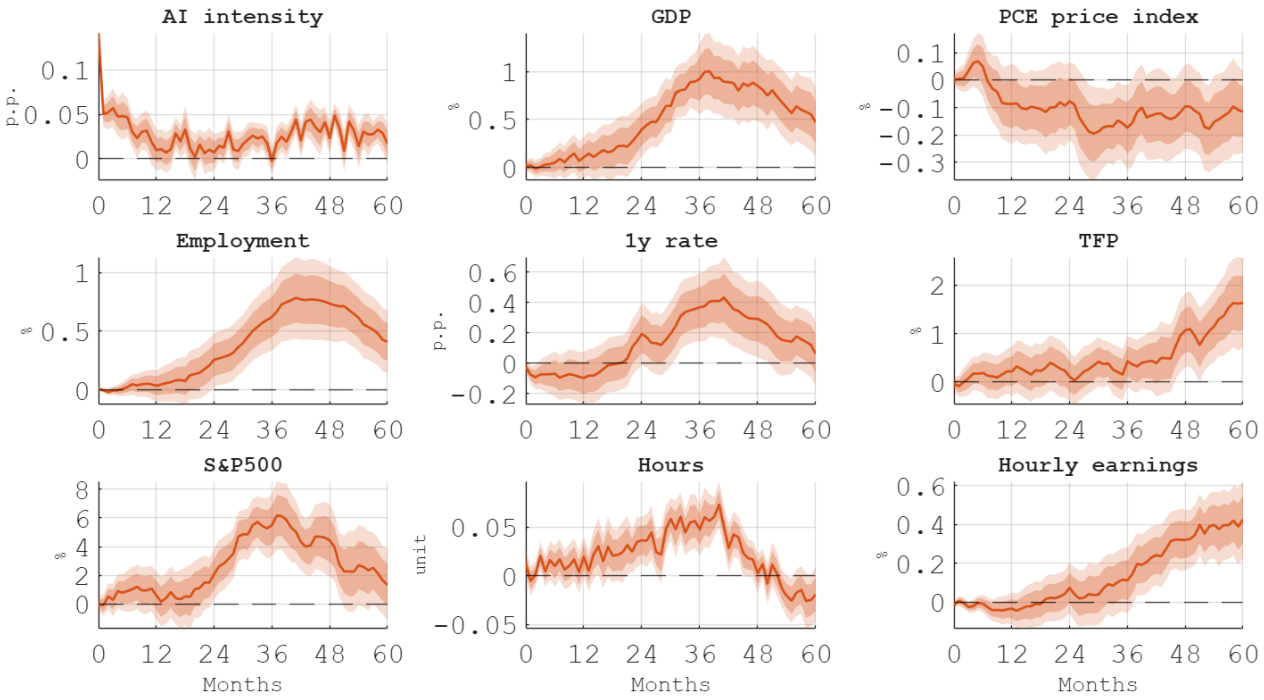


Figure 3: OLS LP: IRFs to an AI technology shock.

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

3.2.2 Alternative specifications

We also assess the sensitivity of our findings to several perturbations of the LP specification. First, similar results hold if we include the number of AI-intensive patents rather than the share $AIint$ (see Section 4.3 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 4.2 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.

Table 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		*Quarterly Shocks*	
Shock Series	p-value	Shock Series	p-value
Känzig (2021) oil supply news	0.611	Romer and Romer (2010) tax	0.257
Baumeister and Hamilton (2019)		Ramey (2011) fiscal	0.132
oil supply	0.246	Mertens and Ravn (2013) priv. tax	0.337
Gilchrist and Zakrajšek (2012) credit	0.387	Mertens and Ravn (2013) corp. tax	0.532
Gertler and Karadi (2015) mon pol.	0.927	Basu et al. (2006) TFP	0.567
Romer and Romer (2004) mon pol.	0.677	Barsky and Sims (2011) news	0.368
Miranda-Agrippino et al. (2025)		Kurmann and Otrok (2013) news	0.273
tech. news	0.387	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)-(v) 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.

Appendix F shows that our results are not driven by spurious factors, such as controlling for general technological advancements and robotics innovation, and by filtering out low-frequency components in the data, such as detrending or transforming the variables into growth rates.

3.3 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 themselves 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^*:

$$e_{t+h^*} = \sum_{h=0}^{h} \theta_h \varepsilon_{t+h} + u_{t+h} \tag{7}$$

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 4 highlight that overall, AI technology shocks have not been a major driver of the US business cycle. Nonetheless, they possess non-negligible explanatory power for key aggregates such as GDP and employment. These results contrast with the findings in the literature on technology shocks that typically find large explanatory power for macroeconomic variables (e.g., Miranda-Agrippino et al., 2025).

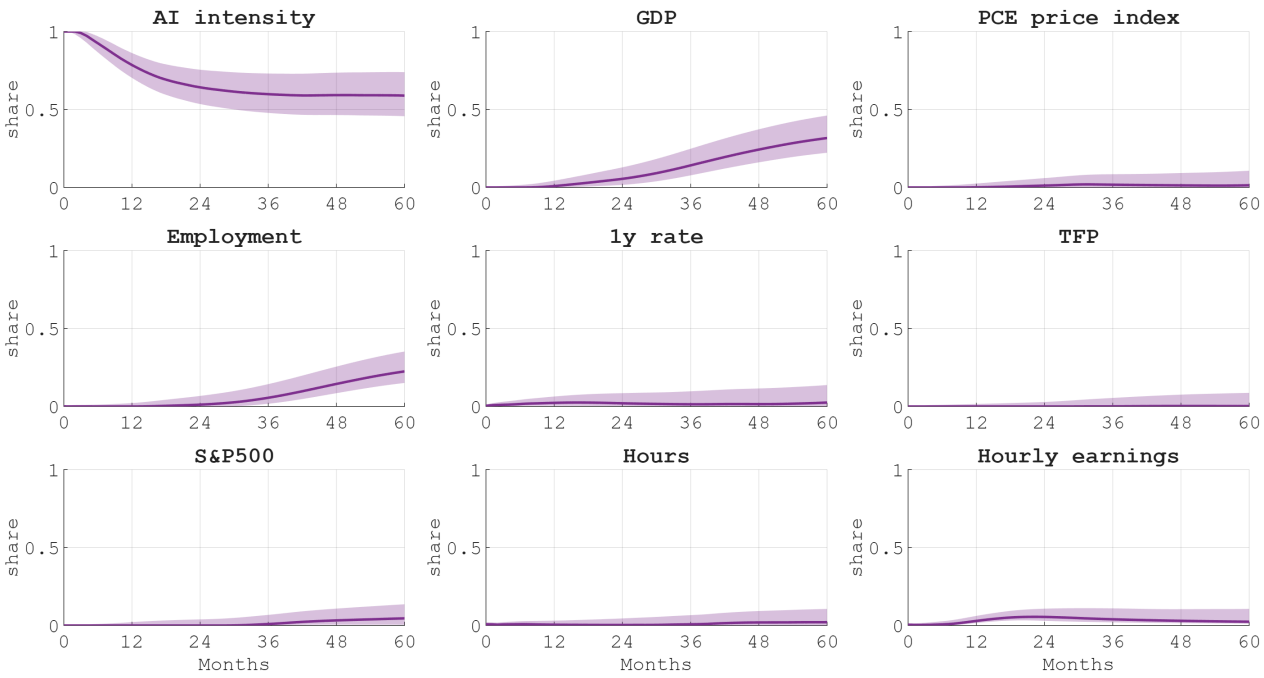


Figure 4: VARIANCE CONTRIBUTION OF AI TECHNOLOGY SHOCKS.

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 unpacking its underlying dimensions and benchmarking it against other forms of innovation. We first compare AI technology shocks with general technology shocks to highlight the distinct transmission mechanism of AI innovation. We then decompose AI innovation into its extensive and intensive margins to identify the sources of its aggregate effects. Finally, we compare AI shocks with automation and ICT innovations, combining macroeconomic evidence with granular patent characteristics.

4.1 AI innovation versus general technology shocks

How do the effects of unexpected increases in AI innovation compare with those of 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).

First, the correlation between our AI IV and the MAHHB IV is 0.08, which is low and not statistically significant at any standard threshold (p-value = 0.104). This indicates that the AI IV captures shocks specific to AI innovation while being orthogonal to general technological advancements.

Second, we explicitly compare the aggregate effects of the two shocks. To maximize comparability and avoid imposing arbitrary instrument choices 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 5 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 in a similar fashion 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 sizable and persistent decline—lasting up to 40 quarters—consistent with automation-driven dynamics (Autor and Salomons, 2018; Acemoglu and Restrepo, 2020b). The effect gradually dissipates at longer horizons, potentially reflecting the reinstatement mechanism emphasized in Acemoglu and Restrepo (2019).

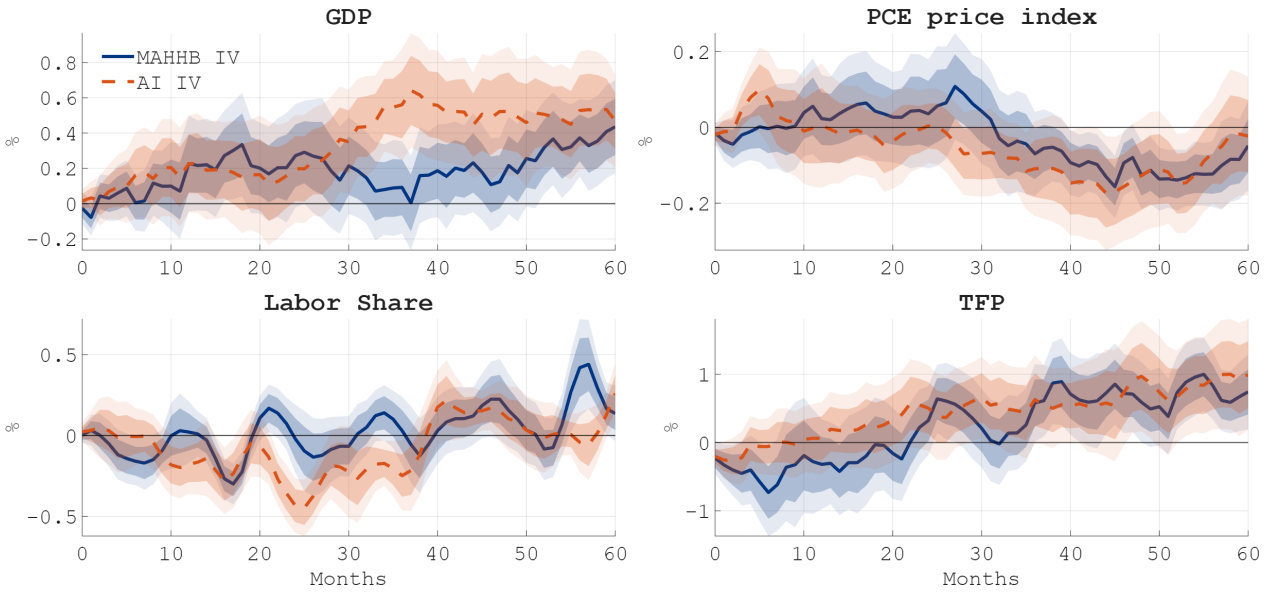


Figure 5: AI INNOVATION VERSUS GENERAL INNOVATION.

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.

4.2 Relationship with Automation and ICT

Why is AI innovation particularly powerful? What distinguishes it from the broader pool of technological innovation? Consistent with our patent-based approach to identifying AI technology, we also use patent classifications to proxy broader technological domains: automation is defined using the taxonomy in Mann and Püttmann (2023), robotics using official USPTO CPC robotics codes, and ICT using the CPC-based definition in Inaba and Squicciarini (2017).[2]

Table 2 summarizes descriptive evidence on the overlap between AI-intensive patents and these technology groups. Panel A reports the share of patents classified as AI-intensive within each category. AI intensity is particularly concentrated in ICT and automation: 21.6% of ICT patents and 16.1% of automation patents are AI-intensive, compared with 3.3% and 6.8% among non-ICT and non-automation patents, respectively. For robotics, the AI-intensive share is 13.4%, only modestly higher than the 10.0% observed outside robotics.

Panel B reverses the exercise and reports how frequently each patent type appears among

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

AI-intensive patents. AI-intensive patents strongly overlap with ICT (79.5%) and, within the sample for which the automation taxonomy is available, with automation (55.7%), whereas overlap with robotics is very small (0.5%). This indicates that patented AI innovation is predominantly ICT-related rather than robotics-related. At the same time, the overlap is asymmetric: while a majority of AI-intensive patents carry the automation tag, only a minority of automation patents (16.1%, Panel A) are classified as AI-intensive.

Panel C reports mean AI intensity scores by patent categories. ICT patents exhibit the highest average AI intensity (0.065 versus 0.011 outside ICT), followed by automation (0.048) and robotics (0.041). The substantial gap between ICT and non-ICT patents supports the interpretation that AI innovation — at least as captured by these patent-based measures — is concentrated in ICT-related technologies.

Finally, Panel D compares patent characteristics within ICT and automation, conditional on AI-intensive status. We consider leading measures of patent quality proposed in the literature as importance and forward citations (Kelly et al., 2021) and patent value (Kogan et al., 2017). Forward citations count the number of times a patent is cited by subsequent patents, providing a standard measure of technological impact. Importance is a text-based similarity measure constructed by Kelly et al. (2021) that weights positively (negatively) how closely a patent's content resembles that of future (previous) patents, thereby identifying patents that are different from past innovations while influencing subsequent innovations. Patent value, from Kogan et al. (2017), estimates the private economic value of a patent based on the stock market reaction to its grant announcement. The imputed patent value refines this measure by addressing the attribution problem that arises when multiple patents are granted to the same firm on the same day.

AI-intensive patents consistently display higher quality and value. Within ICT, AI-intensive patents receive 57% more forward citations (18.73 vs. 11.93), exhibit higher importance scores (0.37 vs. 0.27), and have two to four times higher estimated patent value (15.35 vs. 8.33 raw; 1.21 vs. 0.33 imputed). Within automation, AI-intensive patents receive 62% more forward citations (31.27 vs. 19.32), score higher on importance (0.35 vs. 0.24), and have about 50% higher estimated patent value (14.96 vs. 9.98 raw) and nearly double the imputed patent value (1.16 vs. 0.62). Across all quality metrics—forward citations, importance, and

patent value—AI-intensive patents significantly outperform their non-AI counterparts within both ICT and automation categories, suggesting that AI represents a high-impact frontier within these broader technology classes. Overall, these patterns indicate that AI-intensive patents possess unique features and constitute a high-impact, valuable subset of broader ICT (and, to a lesser extent, automation) innovation. These granular patent characteristics align well with their broader macroeconomic effects, which we document next.

Table 2: AI-intensive patents: overlap with ICT, robotics, and automation tags

Panel A: Share of AI-intensive patents by technology classification

	Outside category		Inside category	
	N	AI share (%)	*N*	AI share (%)
ICT	4,420,459	3.3	2,693,595	21.6
Robotics	7,083,929	10.0	30,125	13.4
Automation	4,815,450	6.8	2,298,604	16.1

Panel B: Technology composition of patents by AI-intensive status

	Not AI-intensive		AI-intensive	
	N	Share (%)	*N*	Share (%)
ICT	6,447,949	37.8	666,105	79.5
Robotics	6,447,949	0.4	666,105	0.5
Automation	6,447,949	32.3	666,105	55.7

Panel C: Mean AI intensity score by technology classification

	Outside category		Inside category	
	N	Mean score	*N*	Mean score
ICT	4,420,459	0.011	2,693,595	0.065
Robotics	7,083,929	0.030	30,125	0.041
Automation	4,815,450	0.021	2,298,604	0.048

Panel D: Patent characteristics by AI-intensive status within technology

	ICT patents		Automation patents	
	Not AI	AI	Not AI	AI
Importance	0.27	0.37*	0.24	0.35*
Forward citations	11.93	18.73*	19.32	31.27*
Value	8.33	15.35*	9.98	14.96*
Value (imputed)	0.33	1.21*	0.62	1.16*

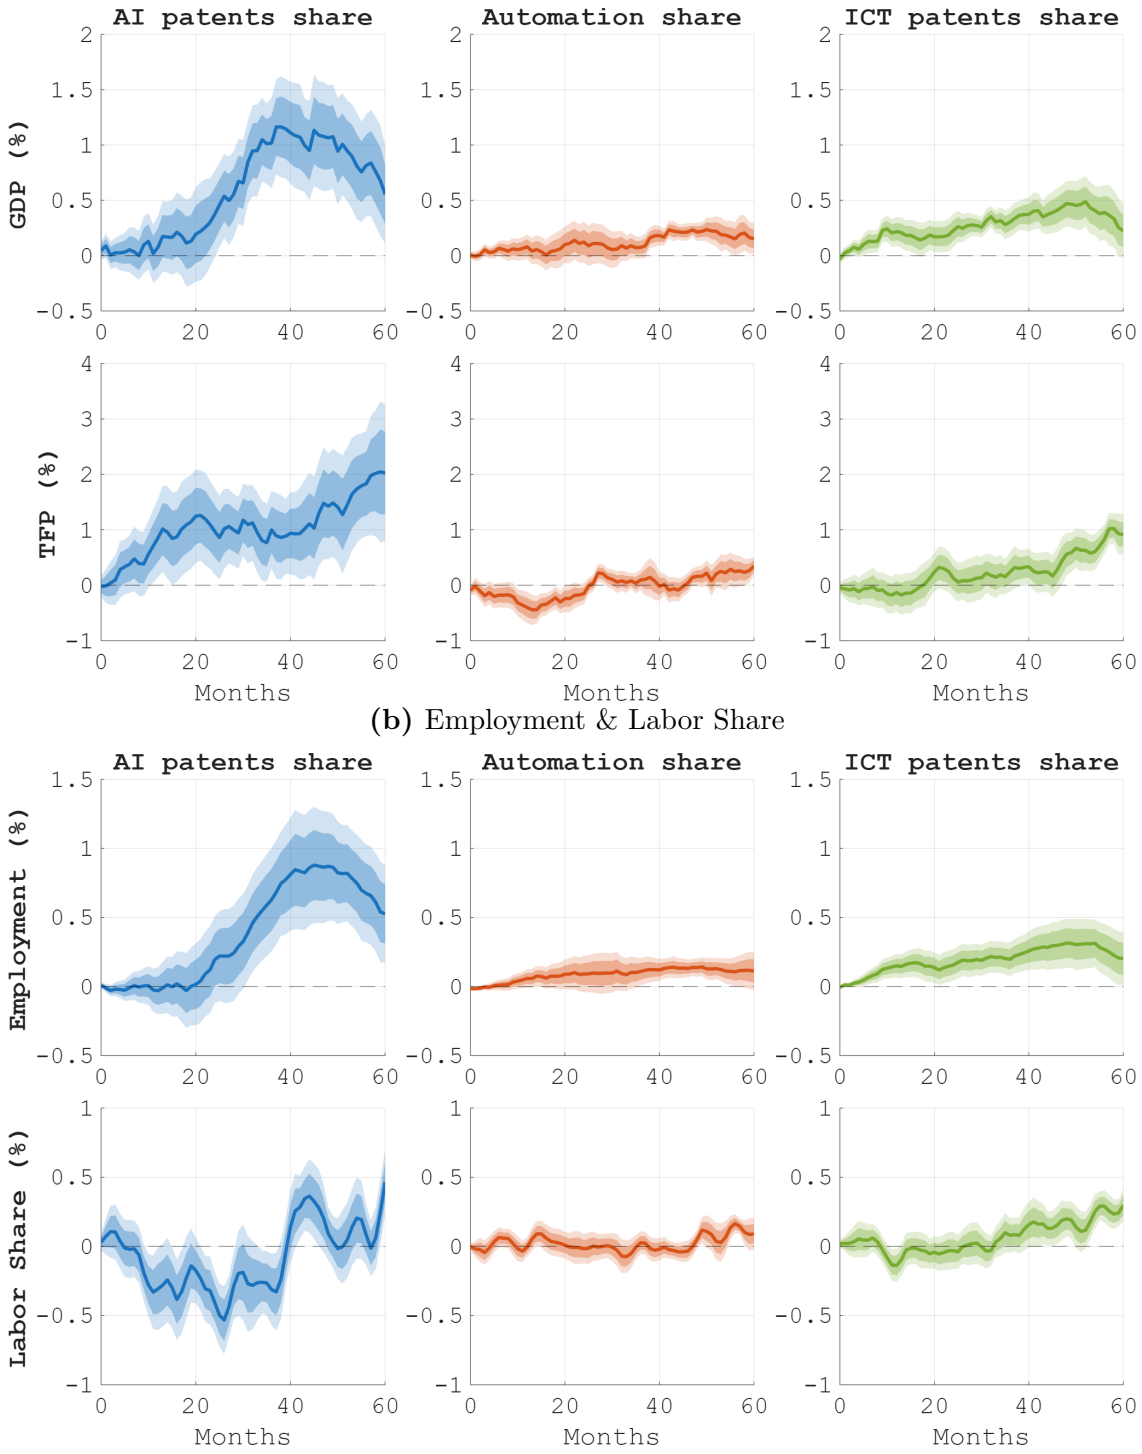
*Notes: "AI-intensive" is an indicator equal to one when Giczy et al. (2022) classifies the patent as AI-intensive (above the 50% threshold in at least one domain). "AI share" denotes the percentage of patents classified as AI-intensive. ICT patents are identified using CPC codes from Inaba and Squicciarini (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 the share of AI-intensive patents inside and outside each technology category. Panel B reports the share of each technology category within AI-intensive and non-AI-intensive patents. Panel C reports mean AI intensity scores. Panel D reports mean patent characteristics; * indicates the difference between AI and non-AI patents is statistically significant at the 1% level.*

Comparison of Aggregate Effects. We compare the aggregate effects of AI technology shocks with those arising from automation and ICT innovation. We exclude robotics patents from this analysis: as documented in Table 2, robotics patents are relatively few, making them poorly suited for time-series estimation, and exhibit limited overlap with AI-intensive patents. For this exercise, we estimate local projections for each patent category, expressed as a share of total patents, as in our extensive-margin specification, to ensure maximum comparability. Figure 6 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 compared to both automation and ICT shocks, with ICT dominating automation. This ranking is consistent with the patent quality differentials documented in Table 2, where AI-intensive patents exhibit higher forward citations, importance scores, and estimated patent values. At the same time, AI technology shocks generate a more pronounced and persistent decline in the labor share relative to other types of innovation. This pattern suggests that AI-driven productivity gains are sufficiently large to increase labor demand in absolute terms, even as capital captures a growing share of the expanding output. The decline in the labor share thus reflects a shift in the distribution of income toward capital rather than an absolute deterioration in labor market outcomes.

4.3 Extensive and intensive margins

$AIint$ captures both the intensive and extensive margins as an aggregate measure of AI intensity in the innovation process (see Appendix B.1.1). In this section, we 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 7a 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, where the intensive margin is inconsequential. These results strongly suggest that the increasing diffusion of AI in innovation generates

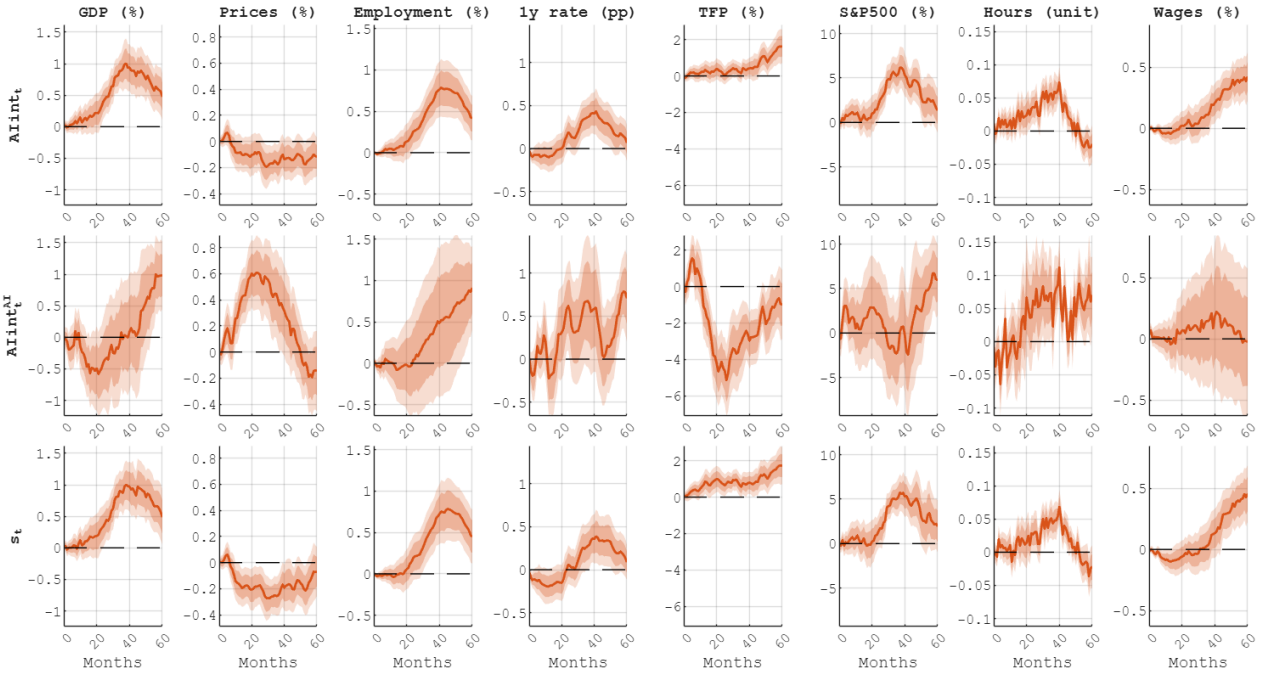
Figure 6: The Dynamic Effects of Innovation in AI, Automation and ICT



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 estimatess are shown with 68% and 90% confidence bands based on Newey–West standard errors. Sample: 1980–2019.

expansionary effects. Conversely, an increase in AI intensity among patents already in the AI category is inconsequential at the aggregate level. To test this hypothesis in more detail, we also consider the standard deviation of AI patent scores. We discuss analytically in Appendix B.1.2 why, given the empirical distribution of $AIint$, fluctuations in the overall standard deviation of $AIint$ effectively reflect changes in the extensive margin. Conversely, we focus on the standard deviation of AI scores within the AI-intensive category. Results in Figure 7b show that the dispersion of AI scores within the AI-intensive category is also inconsequential for macroeconomic outcomes. These findings support the mechanism proposed in Martinez (2021). Automation by laggard firms raises wages and the labor share (extensive margin), whereas additional automation by firms already at the frontier mainly boosts rents (intensive margin).

(a) Intensive versus extensive margin comparison.



(b) Dispersion versus extensive margin comparison.

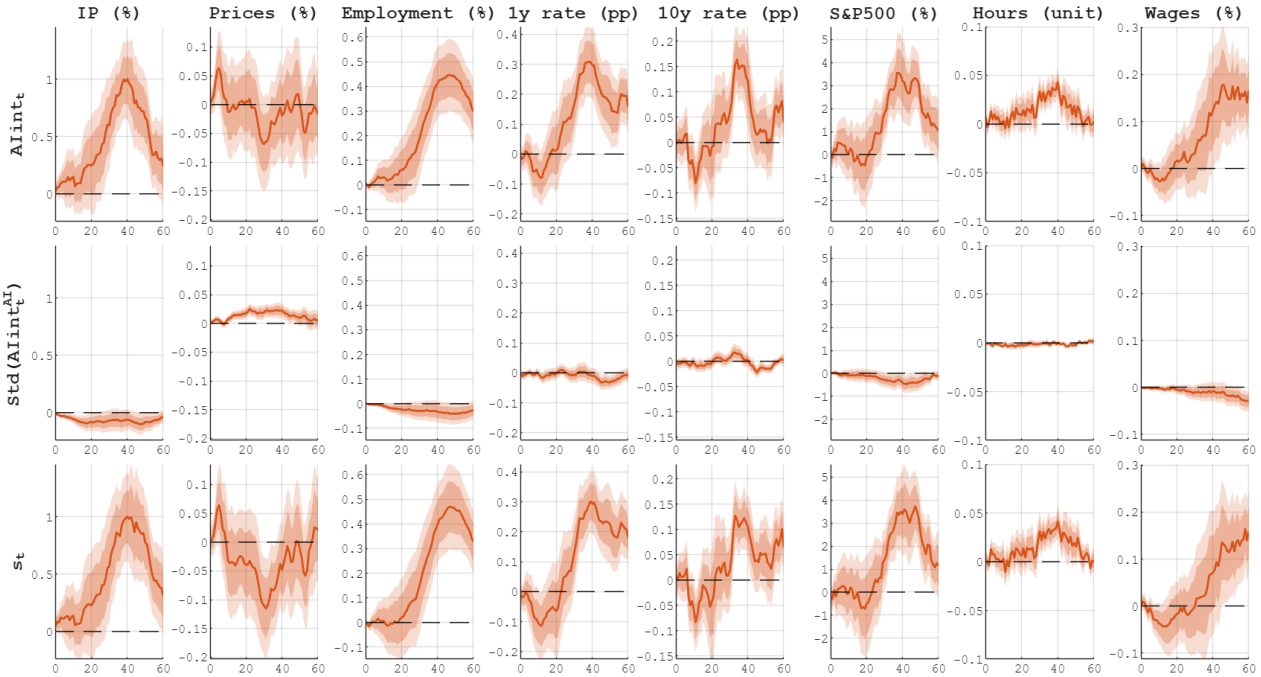


Figure 7: Decomposition of AI Macroeconomic Effects.

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.

5 Zooming on the Labor Market

Shocks to AI-based technological diffusion are crucially transmitted through the labor market via several direct and indirect mechanisms.

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 (2020a) 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 adoption. 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 occupational flows and sectoral quantities. Figure 8 displays the aggregate responses of occupational flows. Despite improving 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.

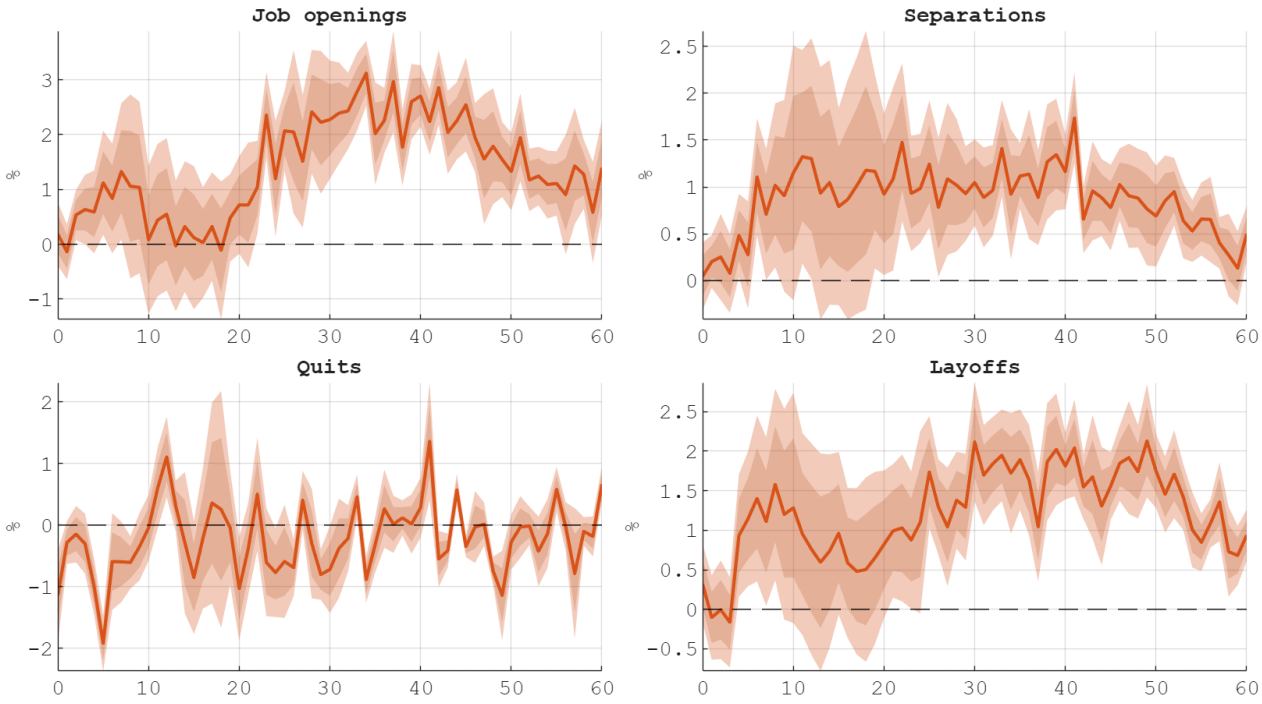


Figure 8: Response of labor market flows to an AI technology shock.

Note. The figure displays the IRFs to a 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.[3]

Figure B1 presents the cumulative impulse responses of labor quantities and wages over a 60-month horizon, with sectors ordered from lowest to highest AI exposure. Two key patterns emerge. First, regarding labor quantities, we observe a remarkably homogeneous

[3]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).

response. Contrary to the hypothesis that AI would trigger mass layoffs in highly exposed industries, 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 productivity spillovers generated by AI appear to overwhelm the direct sectoral displacement effects. Second, distinct wage heterogeneity emerges in our estimates. While real wages rise on average, the magnitude of this increase is inversely related to AI exposure. Sectors with *low* exposure to AI experience significantly stronger wage growth compared to their highly exposed counterparts. A similar pattern characterizes job openings.

We interpret these results as the manifestation of the *displacement effect* in partial equilibrium. In highly exposed sectors, the threat of automation could exert downward pressure on wage bargaining power, dampening the gains from the aggregate expansion. Conversely, in low-exposure sectors, workers benefit from the surge in aggregate demand without facing the direct substitution threat of the new technology.

5.2 The Role of Education

Education constitutes 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 9 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 10, paint a different picture. Unlike the em-

ployment 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).

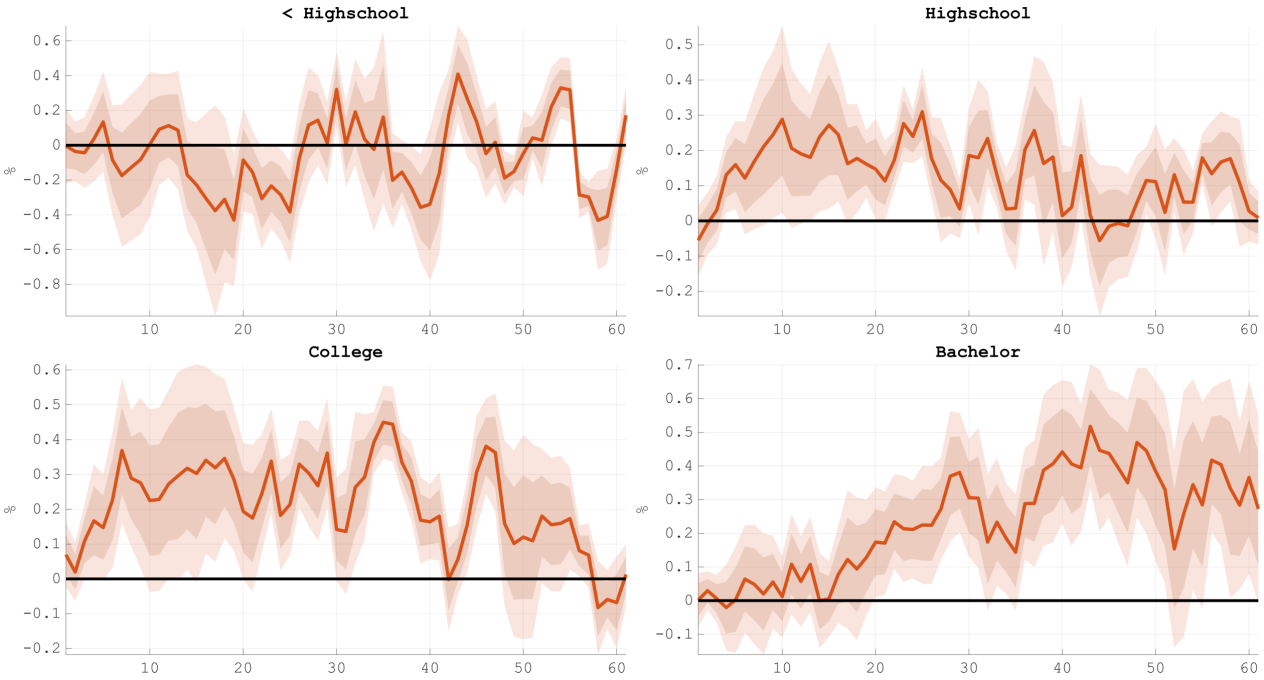


Figure 9: Response of employment by educational attainment to $AIint$.

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

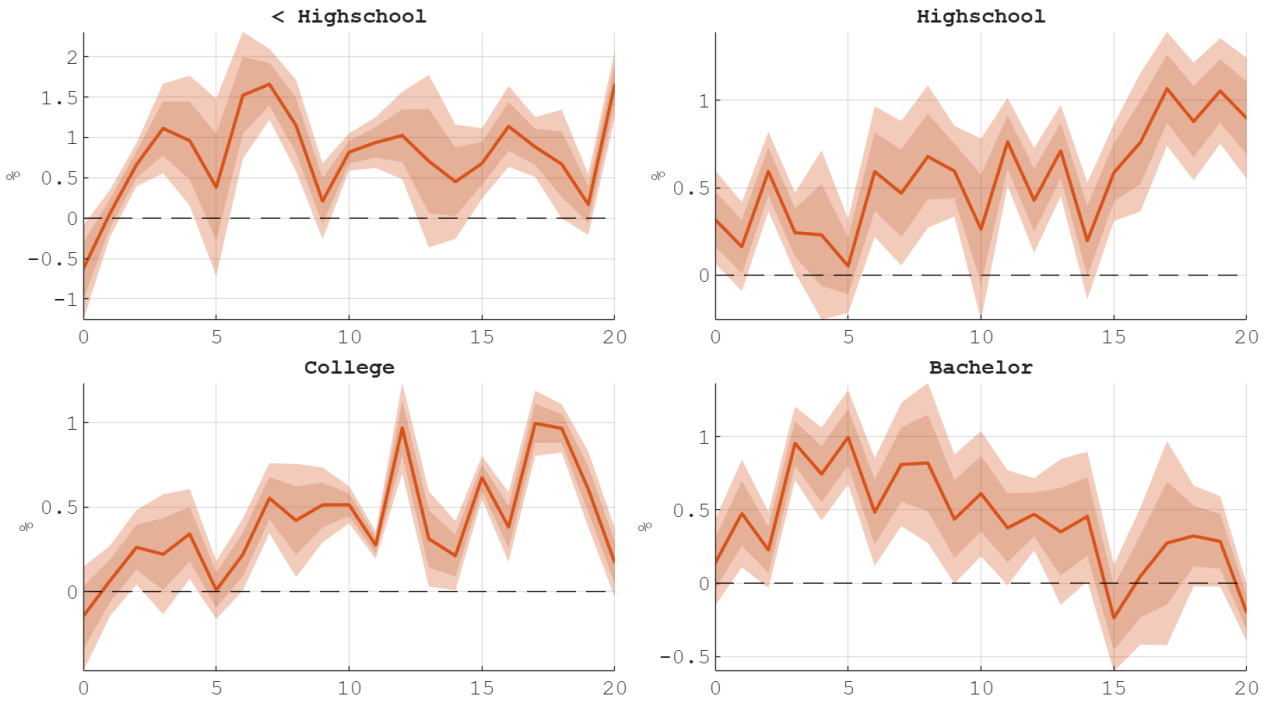


Figure 10: Response of earnings by educational attainment to an AI technology shock.

Note. The figure displays the IRFs to a AIint shock. Sample 2000-2019. The estimates are based on local projections with Newey-West standard errors. Point estimates and 68%-90% confidence bands.

6 Implications for Inequality

We conclude our empirical analysis by examining the distributional consequences of AI innovation. To explore this dimension, we leverage on the distributional accounts data compiled by Blanchet et al. (2022).[4]

Figure 11 reports the response of total income, wealth, and labor income shares to an *AIint* shock. The impact on the total income distribution (top row) is relatively muted; while the shares of the top 1% and 10% exhibit a mild upward drift, the response lacks statistical precision. A far more striking divergence emerges in the wealth distribution (middle row). An AI intensity shock triggers a significant and persistent rise in the wealth share held by the top 1% and, most notably, the top 10%. In sharp contrast, the wealth share of the bottom 50% declines persistently. Crucially, this is not merely a relative decline: additional estimates

[4]The data is available at https://realtimeinequality.org/.

on the *level* of real wealth, displayed in Figure C3, indicate that an increase in AI intensity anticipates an absolute wealth *loss* for the bottom 50%. Conversely, AI innovation tends to lift absolute real *incomes* for all groups.

These dynamics provide empirical support for the mechanism theorized by Moll et al. (2022), who argue that automation increases inequality primarily through asset returns. Since AI innovation boosts the productivity of capital and equity valuations (as shown in Figure 2), it disproportionately benefits households at the top of the distribution who hold high-return assets, while eroding the relative position of those relying on labor income or low-return savings.

A similar polarization characterizes the labor income share (bottom row). The labor share of the top 10% rises while that of the bottom 25% falls.[5] However, unlike the persistent divergence in wealth, the effects on labor income shares appear to be transitory, reverting to the baseline after approximately 40 months.

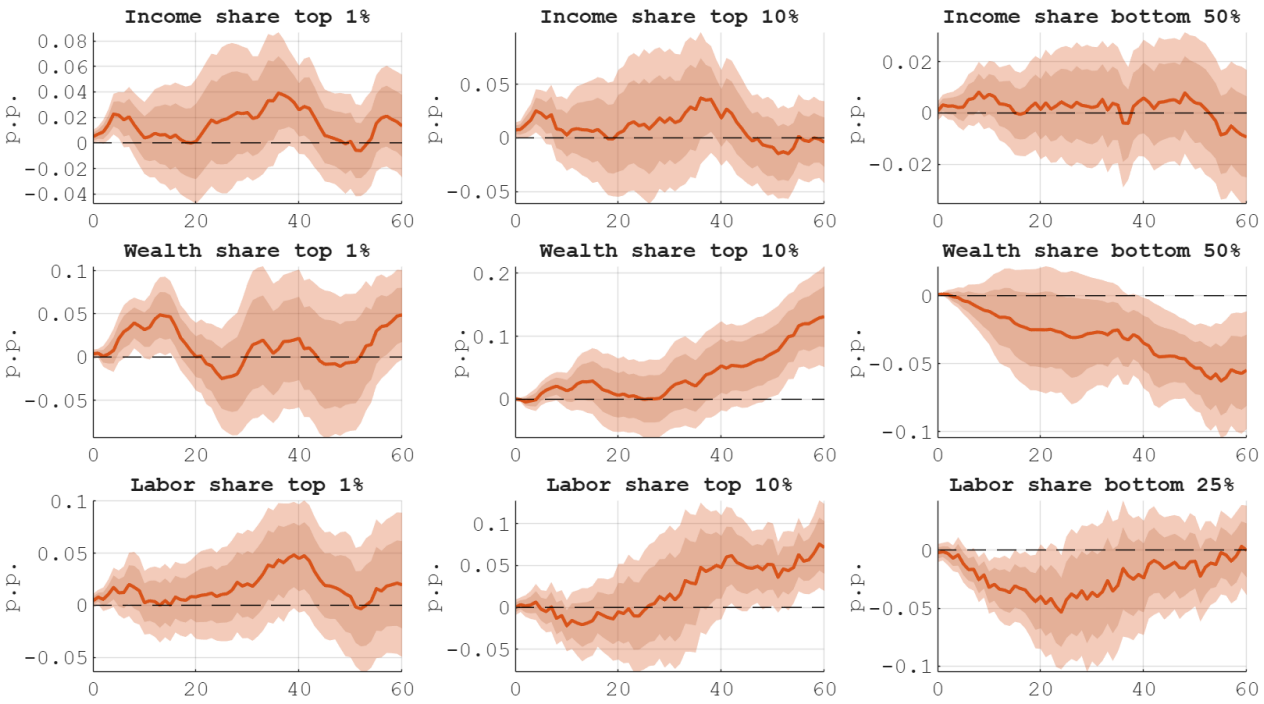


Figure 11: Response of inequality to an AI technology shock.

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

[5]Note that the demographic breakdown provided in https://realtimeinequality.org/ differs slightly between total income/wealth and labor income.

7 Conclusions

AI-based innovations such as *ChatGPT* have brought AI to the center of the academic and policy debate. The future socio-economic implications span a broad spectrum, from idyllic to apocalyptic scenarios. Although the future direction and pace of AI technology are difficult to anticipate, in this paper we have exploited data on the AI intensity of US patents from 1980 to 2019 to examine the macroeconomic consequences of AI-based innovation. Our findings can contribute to disciplining the theoretical modeling of AI's economic effects. The key insight is that the general equilibrium effects of AI innovation can be substantial and should be appropriately accounted for when drawing policy implications. In particular, partial equilibrium analyses fail to capture AI's contribution to aggregate TFP. In this paper, we document a positive nexus: AI innovation is found to increase industrial production, compress consumer prices, and boost the labor market, as reflected in employment, hours worked, and salaries. However, the more intensive use of AI technology in innovation produces adverse distributional effects. The share of wealth owned by the bottom 50% of the population falls, potentially calling for a policy response that should be explored in 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 (2019): "Automation and new tasks: How technology displaces and reinstates labor," *Journal of economic perspectives*, 33, 3–30.

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

——— (2020b): "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.

——— (2018): *Artificial Intelligence and Economic Growth*, 237–282, in Agrawal et al. (2019).

Agrawal, A., J. Gans, and A. Goldfarb (2019): *The Economics of Artificial Intelligence: An Agenda*, University of Chicago Press.

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. 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.

Bresnahan, T. F. and M. Trajtenberg (1995): "General purpose technologies 'Engines of growth'?" *Journal of Econometrics*, 65, 83–108.

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*, 23–57, in Agrawal et al. (2019).

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*.

Cockburn, I. M., R. Henderson, and S. Stern (2019): "The Impact of Artificial Intelligence on Innovation: An Exploratory Analysis," *The Economics of Artificial Intelligence: An Agenda*.

Comunale, M. and A. Manera (2024): "The economic impacts and the regulation of AI: A review of the academic literature and policy actions," .

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.

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.

Goldfarb, A., B. Taska, and F. Teodoridis (2023): "Could machine learning be a general purpose technology? A comparison of emerging technologies using data from online job postings," *Research Policy*, 52, 104653.

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.

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.

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.

Lu, Y. and Y. Zhou (2021): "A review on the economics of artificial intelligence," *Journal of Economic Surveys*, 35, 1045–1072.

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

Martinez, J. (2021): "Putty-Clay Automation," CEPR Discussion Paper 16022, Centre for Economic Policy Research, Paris & London.

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*.

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

A Examples of AI-intensive patents

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.

B Data Definitions and Details

This Appendix provides a more detailed description of the AI measure employed in the paper. Section B.1.1 describes analytically the decomposition of $AIint$ into the extensive margin and intensive margin components. Section B.1.3 contains further details on the sectoral and field breakdown of $AIint$.

B.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.

B.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,

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

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

$$a_{i,t} \;=\; \mathbf{1}\left\{\max_{m\in\{1,\ldots,8\}} AIint_{i,t}^{(m)} > \tau\right\}. \tag{9}$$

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

$$s_t \;=\; \frac{N_t^*}{N_t}. \tag{10}$$

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

$$AIint_t \;=\; \frac{1}{N_t}\sum_{i=1}^{N_t} AIint_{i,t}. \tag{11}$$

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

$$AIint_t^{AI} \;=\; \frac{1}{N_t^*}\sum_{i=1}^{N_t} a_{i,t}\, AIint_{i,t}, \tag{12}$$

$$AIint_t^{non} \;=\; \frac{1}{N_t - N_t^*}\sum_{i=1}^{N_t} (1-a_{i,t})\, AIint_{i,t}. \tag{13}$$

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

$$AIint_t \;=\; s_t\, AIint_t^{AI} \;+\; (1-s_t)\, AIint_t^{non}. \tag{14}$$

Given that AI intensity among non-AI-intensive patents is empirically close to zero (Table A1), a useful approximation is

$$AIint_t \;\approx\; s_t\, AIint_t^{AI}. \tag{15}$$

This decomposition motivates the empirical exercise in Section 4.3, 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 $AIint_t^{AI}$.

B.1.2 Standard deviation

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

$$\mathrm{Var}(AIint_{i,t}) = s_t\,\sigma^2_{1,t} \;+\; (1-s_t)\,\sigma^2_{0,t} \;+\; s_t(1-s_t)\left(AIint_t^{AI} - AIint_t^{non}\right)^2. \tag{16}$$

Table A1 shows that $AIint_t^{non}$ is close to zero and $\sigma^2_{0,t}$ is negligible; under this empirically valid approximation,

$$\mathrm{Var}(AIint_{i,t}) \;\approx\; s_t\,\sigma^2_{1,t} \;+\; s_t(1-s_t)\left(AIint_t^{AI}\right)^2. \tag{17}$$

Since $AIint_t = s_t\, AIint_t^{AI}$ by Section B.1.1, substituting into (17) yields

$$\mathrm{Var}(AIint_{i,t}) \;\approx\; s_t\,\sigma^2_{1,t} + \frac{1-s_t}{s_t}\left(AIint_t\right)^2, \qquad \mathrm{sd}(AIint_{i,t}) \;\approx\; \sqrt{s_t\,\sigma^2_{1,t} \;+\; \frac{1-s_t}{s_t}\left(AIint_t\right)^2}. \tag{18}$$

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

$$\mathrm{sd}(AIint_{i,t}) \;\approx\; AIint_t^{AI}\sqrt{s_t(1-s_t)} \;=\; AIint_t\sqrt{\frac{1-s_t}{s_t}}. \tag{19}$$

Equations (16)–(19) make clear that, given the distributional features in Table A1, 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 A1, 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 (16) 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

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

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 (17)–(19): the cross–sectional standard deviation 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 A3 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 A1: 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

B.1.3 Breakdown by sector and AI field

Figure A1 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 since in the mid-'90s. Towards the end of our sample, both sectors feature about 10% of AI-intensive patents out of total patents. Figure A2 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.

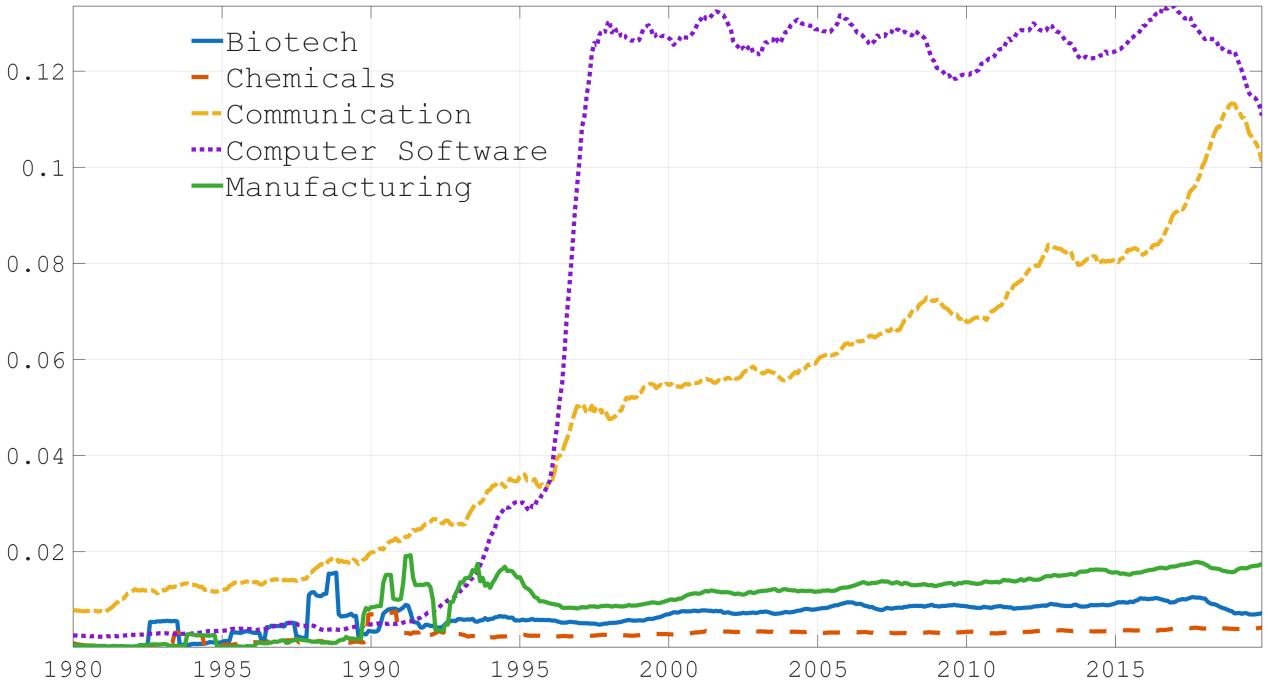


Figure A1: Share of AI-intensive patents by sector.

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

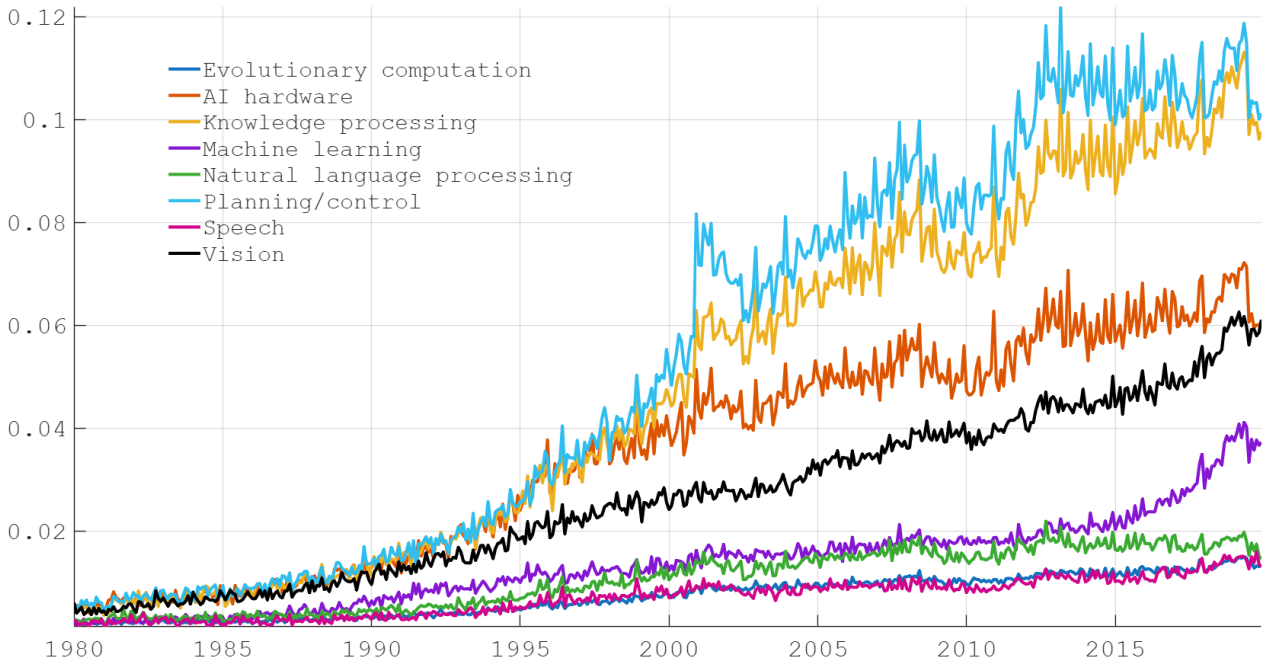


Figure A2: Share of AI-intensive patents by field

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

B.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.

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 risk-free rates are measured using the 1-year and 10-year U.S (FRED:GS1 and FRED:GS10). Treasury yields. (FRED:SP500). 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 B1.

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.

C Heterogeneity in Consumer Prices

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

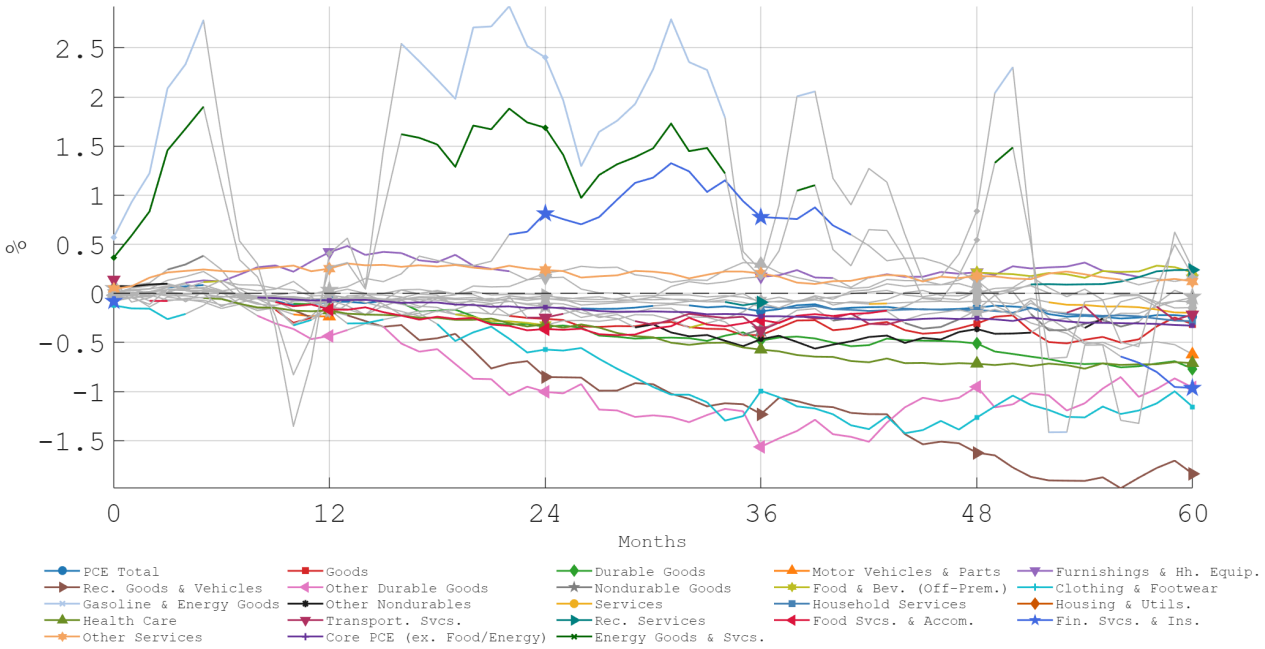


Figure A1: Consumer prices response to an AI technology shock.

Note. The figure displays the IRFs to a 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 A1 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 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. 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:

$$\Delta AIint_t = c + \sum_{j=1}^{12} \gamma_j \Delta AIint_{t-j} + \sum_{h \in \{1Q, 1Y\}} \beta_h \mathbb{E}_t[x_{t+h}] + \delta \eta_t^{RR} + z_t, \tag{20}$$

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 A2 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 A2: 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.*

E Additional Results

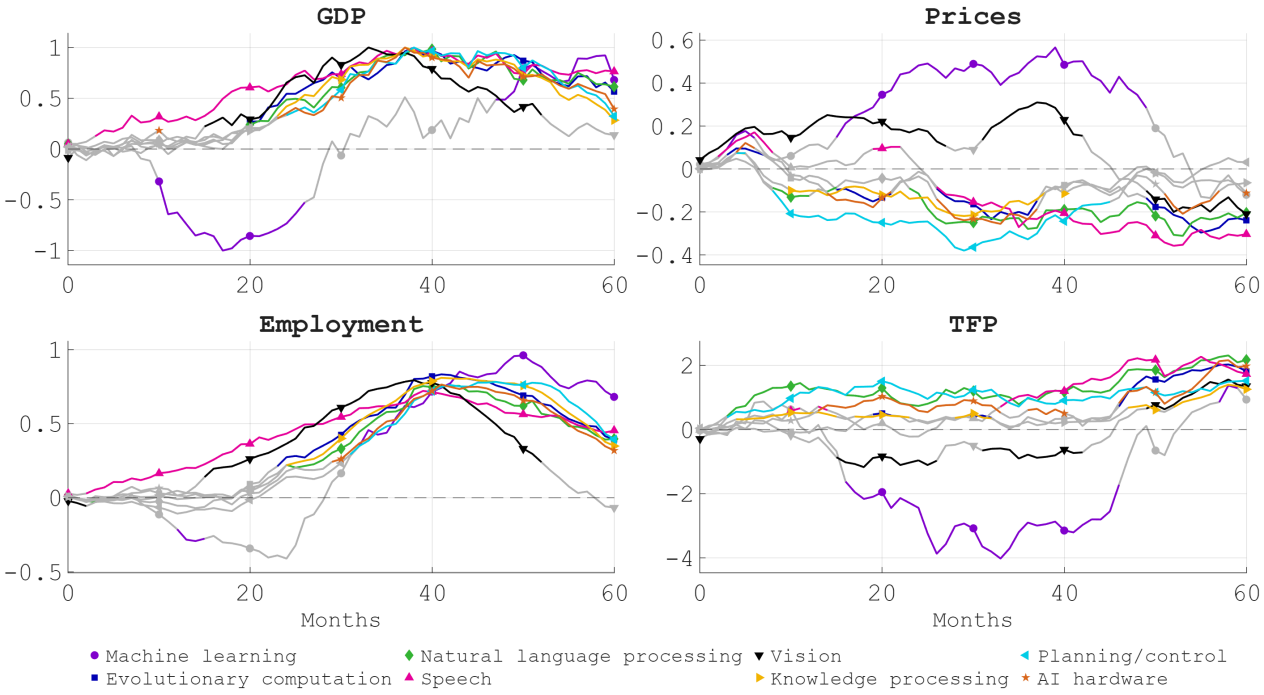


Figure A1: IRFs by AI categories

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

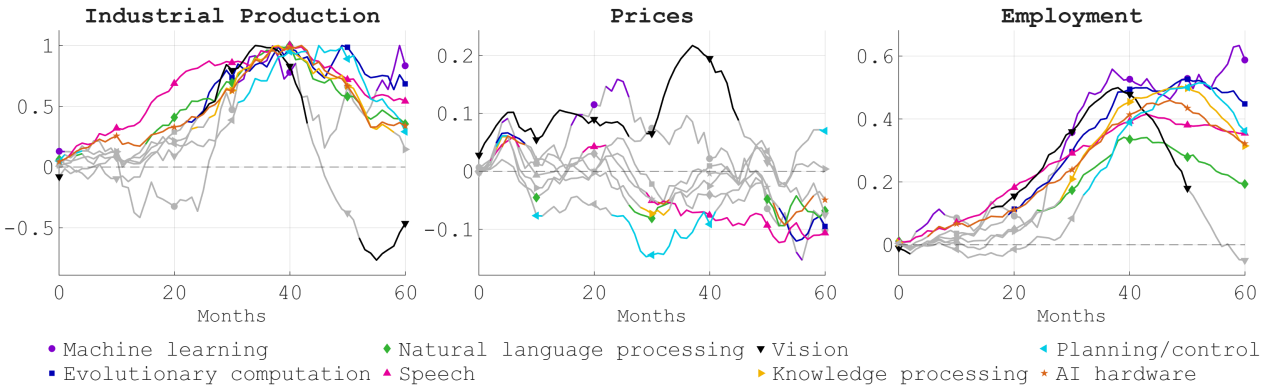


Figure A2: IRFs by AI categories - Specification with Industrial Production

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

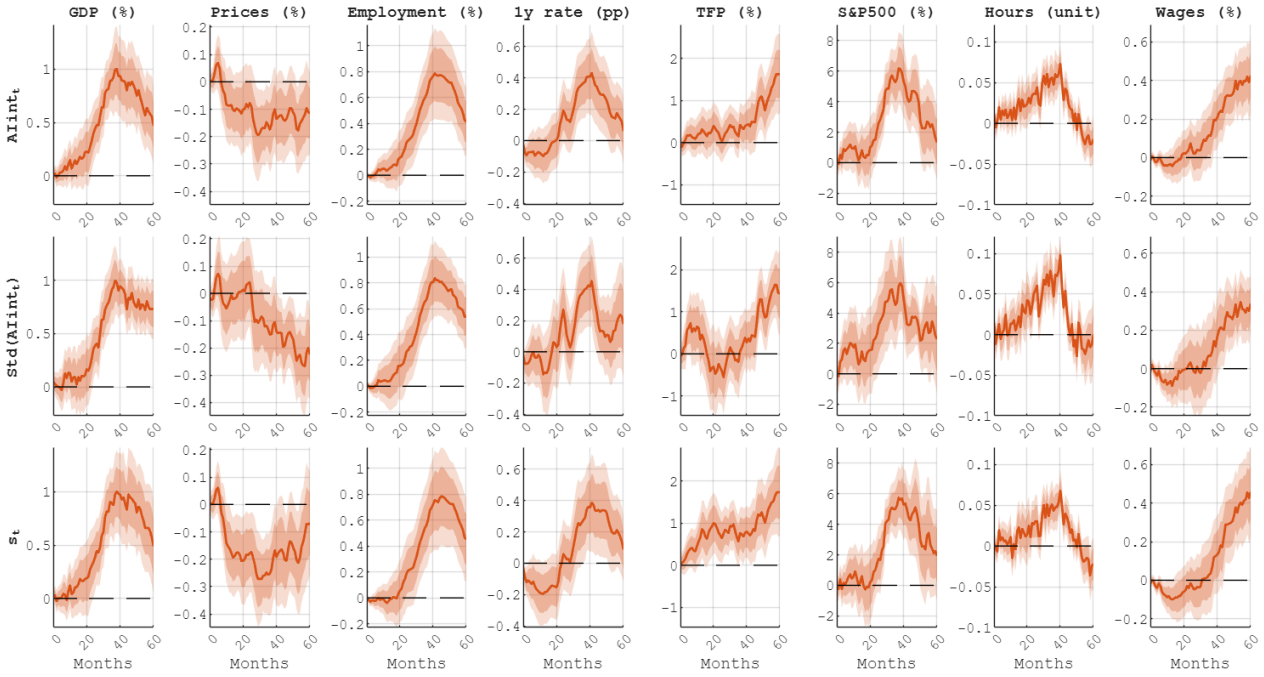


Figure A3: Overall standard deviation versus extensive margin comparison

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

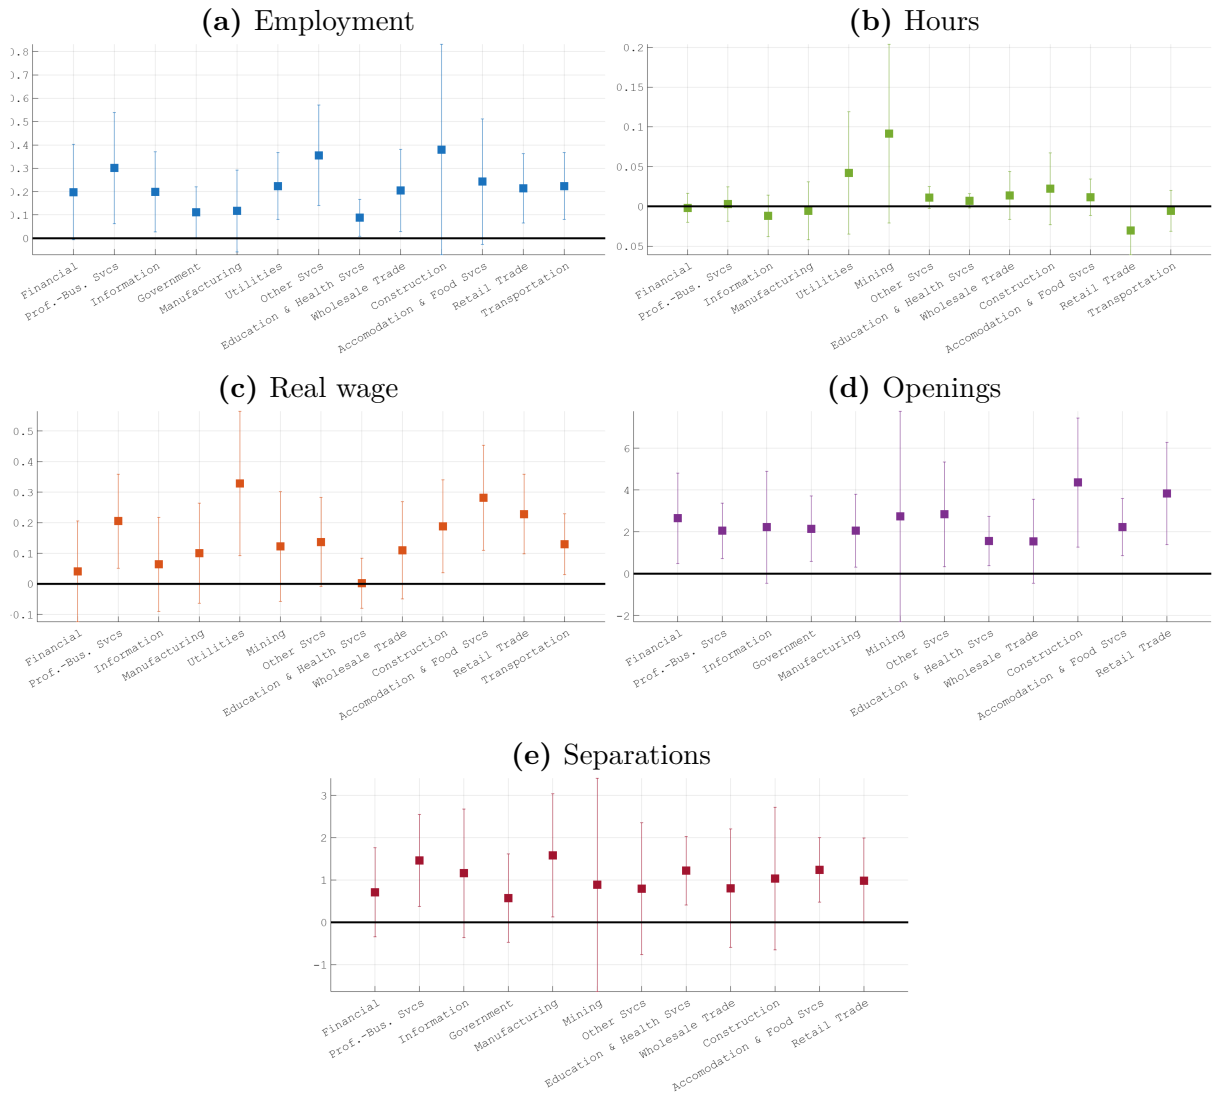


Figure B1: Sectoral labor market responses - Average

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

F Robustness Exercises

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

F.1 Robustness

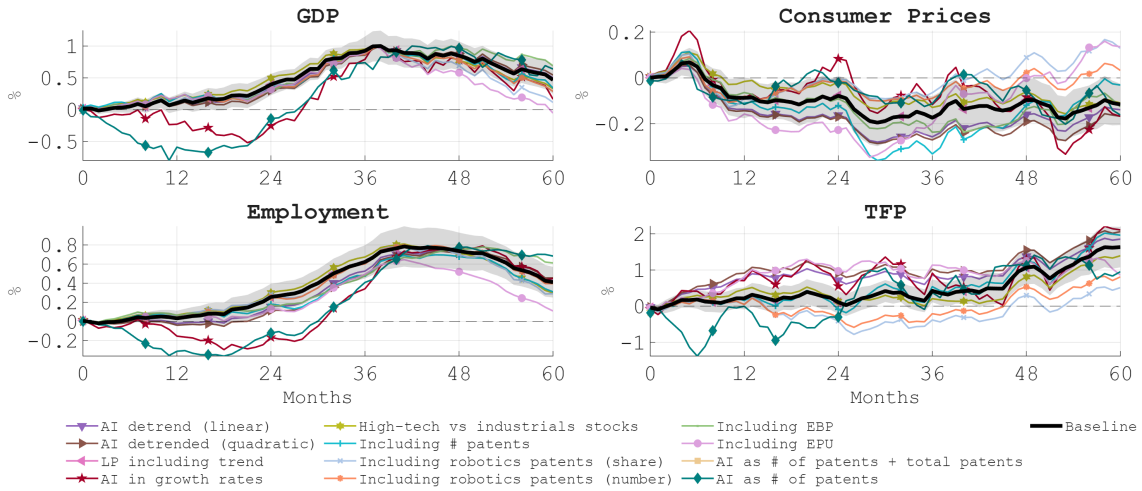


Figure C1: Robustness Across Specifications

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

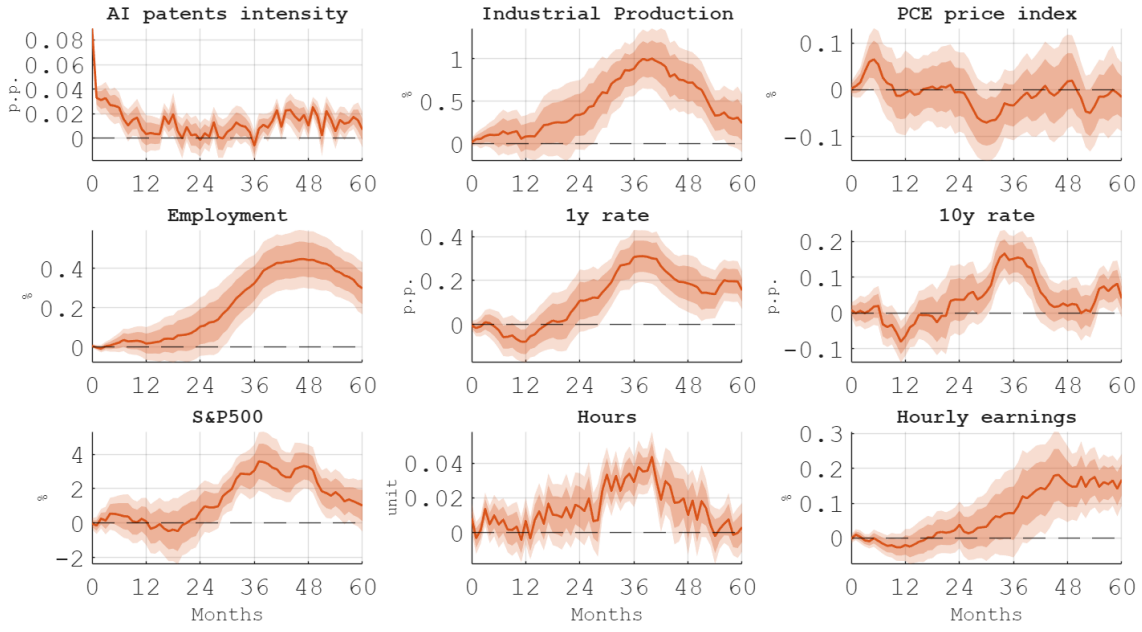


Figure C2: IRFs with Monthly Observables Only

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.

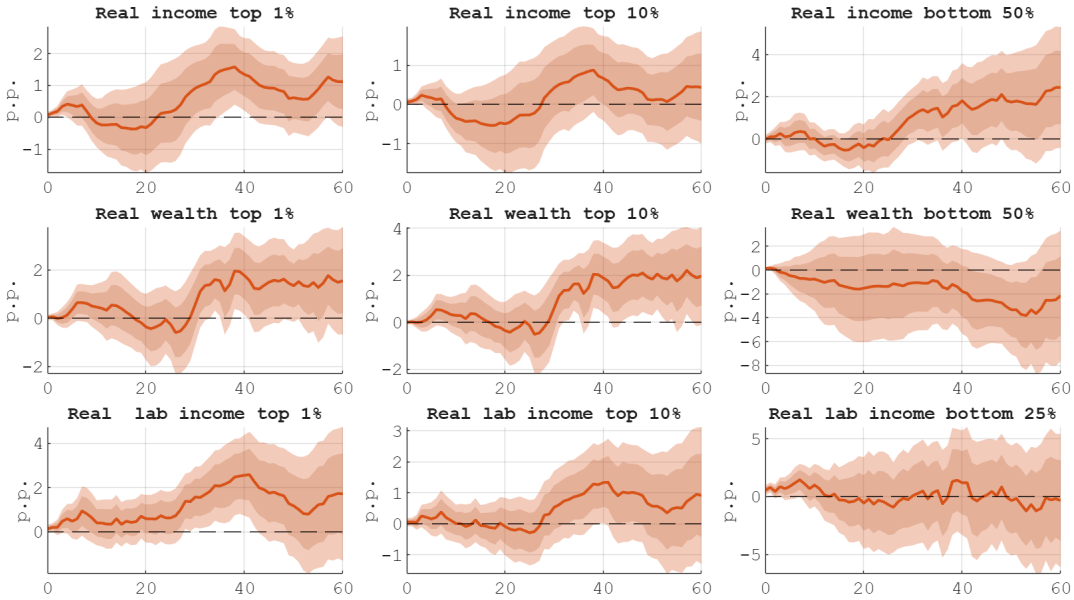


Figure C3: Absolute Response of Income and Wealth

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