The macroeconomic effects of AI-based innovation

Andrea Gazzani* Filippo Natoli†

February 24, 2025

Abstract

What are the macroeconomic consequences of AI-based technological progress? We use data on the artificial intelligence content of each patent granted by the US Patent Office between 1980 and 2019 to identify aggregate shocks to the AI *intensity* in US innovation. An increase in AI intensity leads to a delayed surge in industrial production and a decline in consumer prices in line with the transmission of a positive supply shock. Such effects descend from a gradual increase in total factor productivity. Employment, hours, and wages *increase* even in highly AI-exposed sectors, underscoring the crucial role of general equilibrium effects. Overall, shocks to the composition of US innovation in favor of AI technologies act as an additional force amplifying the already expansionary impact coming from positive (average) technology shocks. The expansionary impact of higher AI intensity in innovation comes at the cost of an increase in inequality across income and wealth distributions.

JEL: E32, O33, O34, C36

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

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

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

This paper benefited from feedbacks in seminars at the Bank of Italy, the ECB - Banco de España conference "The macroeconomics effects of AI" and the 12th IMF Statistical Forum. The views expressed in this paper do not necessarily reflect those of the Bank of Italy or ESCB.

1 Introduction

General-purpose technologies (GPTs) such as steam engines, electricity, computers, and the internet have been recognized to drive eras of technological progress and growth (Bresnahan and Trajtenberg, 1995). Recent works on artificial intelligence (AI), defined as "the science and engineering of making intelligent machines" (McCarthy et al., 2007), have argued that given its effectiveness in spurring innovation across different economic sectors, such technology might potentially play the role of a GPT (Cockburn et al., 2019; Goldfarb et al., 2023).

Such evidence and the recent AI breakthroughs – the release of AI-based chatbots like *Chat GPT* – have put the economic effects of AI adoption at the forefront of academic and policy debates. Although nascent empirical literature has focused on the partial equilibrium implications of the AI diffusion, such as those on firm-level innovation and performance (Babina et al., 2024; Chen et al., 2024) or sectoral hiring and wages in the labor market (Acemoglu et al., 2022), its full macroeconomic impact is yet under-investigated.[1] Indeed, artificial intelligence might have repercussions on multiple dimensions such as productivity, economic growth, inflation, labor market, and inequality both at the business cycle frequency and in the long run. Therefore, these innovations may have profound implications also for the conduct of monetary and fiscal policies.[2]

This work investigates empirically the macroeconomic consequences of AI innovation in the United States. AI's role in US innovation is widespread, and its weight in technological processes has grown considerably over time. We measure AI's role in US innovation by relying on granular data on US patents filed at the US Patent and Trademark Office (USPTO). In particular, we construct an aggregate measure of *AI intensity in innovation* by averaging, at the monthly frequency, patent-level AI scores proxying the artificial intelligence content in the underlying technology, which is available for all patents filed at the USPTO since the late 1970s. Scores are averaged based on the filing month of each patent as in Miranda-Agrippino

[1] For a review of the literature on the economics of AI, see Lu and Zhou (2021); Comunale and Manera (2024).

[2] On the impact of AI on firms' productivity, opposite views dominate the academic debate, and empirical evidence is still scant. Views go from skeptical (Gordon, 2018) to positive (Czarnitzki et al., 2023). Brynjolfsson et al. (2018) discusses possible reasons for the lack of productivity effect of AI in the short run.

et al. (2020). The obtained indicator is used to proxy an anticipated shock to the AI *intensity* of US innovation: in other words, a shock to the *composition* of future technology sources in favor of artificial intelligence.

Our monthly indicator is then used to estimate the dynamic effects of recomposition shocks towards AI technology between January 1980 and December 2019 using local projections (Jordà, 2005).[3] We find that an increase in the AI intensity of US innovation leads to a delayed surge in industrial production and a concomitant – albeit more muted – decline in consumer prices, suggesting the shock transmitting as a *positive supply shock*. Such positive effects descend from gradual improvement in total factor productivity (TFP), which peaks five years after the shock. The mildly deflationary effect is driven by cheaper, high-tech-intense durable goods, while services prices do not respond significantly to the shock. The expansionary effect on industrial production induces an upward shift of the government bonds yield curve, with a stronger impact on its short end, suggesting an endogenous contractionary response of the monetary authority. Overall, these effects hold even when contemporaneous fluctuations in the whole technological progress, proxied by variations in the total number of filed patents or TFP, are controlled for. For this reason, the expansionary effects of AI-tech recomposition shocks can be interpreted as an additional force amplifying the already expansionary impact coming from positive (average) technology shocks.

The overall expansionary macroeconomic effect relates to the implications of AI innovation to the labor market, which is the linchpin of its aggregate transmission and where the overall effects are not yet fully understood.[4] Indeed, while an increase in the AI content of US innovation seems to trigger recomposition effects within the workforce, these turn out to be overall employment-enhancing: following a positive AI tech shock, employers' demand for new workers surges (the number of job openings rises), more than offsetting the contemporary increase in layoffs. Overall, positive shocks lead to a rise in the number of employed people, the number of hours worked, and wages (measured by hourly earnings).

These results underscore the importance of considering the general equilibrium effects of

[3]Due to data availability, our analysis does not include the latest developments of generative AI that may mark a discontinuity in AI technology and its economic consequences.

[4]Survey evidence shows that firms see artificial intelligence as growth-enhancing but they are split about their prevailing labor market effects, see Morikawa (2016) and Ilzetzki and Jain (2023).

AI-driven technological developments, which can even reverse the negative partial-equilibrium repercussions documented in some micro-level estimates. To stress this point, we repeat our estimates on labor market variables at the sectoral level, i.e., by running local projections separately for 13 sectors available at the NAICS 2-digit level. We plot our results by ranking sectors based on their AI exposure, from the most to the least exposed, where sectoral exposure rests on recent assessments of the overlap between the main types of occupations in the sector and the tasks potentially be performed by AI (Felten et al., 2018; Acemoglu et al., 2022). We find that the effect of AI intensity shocks on employment and wages, among other variables, does not vary significantly between highly AI-exposed and lowly AI-exposed sectors (consistently with the conclusions in Albanesi et al., 2023). Such homogeneity underscores the power of AI technology's overall (general equilibrium) expansionary push in driving the results and, potentially, a certain degree of complementarity (rather than simply substitutability) between AI and human work. Our findings are consistent with Hampole et al. (2025) who also highlight the positive consequences on the labor market steaming from increased firms' productivity. The dominance of AI general equilibrium effects over the partial equilibrium effects is a finding that echoes those in Aldasoro et al. (2024) within a macroeconomic multi-sector theoretical model. Indeed, our empirical findings can help to discipline theoretical models that investigate the economic consequence of AI.

Aggregate and sector-level evidence may neglect potential distributional consequences. We take up this point by investigating the heterogeneous impact of our shock on the *shares* of total income, wealth, and labor income for specific percentiles of the respective distributions (data come from Blanchet et al., 2022). An increase in the AI intensity of US innovation leads to a rise in the labor income share and wealth share of the top 10% of the distribution (more than what happens to the top 1%) and a fall in the wealth share owned by the bottom 50%. In particular, the negative wealth effect on the bottom 50% is not just a matter of shares: additional estimates on the *level* of wealth across different percentiles of the distribution suggest that an increase in AI intensity anticipates a wealth loss for the poor. These results speak to the predictions of a wider gap between labor and capital income due to the diffusion of AI (Acemoglu, 2024) and inform the literature on the inequality effects of automation and technology more generally (Acemoglu and Restrepo, 2020; Prettner and Strulik, 2020;

Moll et al., 2022). AI innovation also appears to drive a substitution within the workforce as the employment of the least educated workers barely moves while it rises for all other groups. Consequently, the impact of AI-driven advances seems coherent with that predicted by the skill-biased technological change theory (as emphasized by Albanesi et al., 2023) and, historically, may have contributed to the secular decline in the demand for low-skilled workers in developed economies (Berman et al., 1998).

Our empirical results contribute to the literature by shedding light on the aggregate implications of AI technology diffusion, whose overall economic impact depends not only on the partial equilibrium effects – substitutability or complementarity of working tasks with AI – but also on general equilibrium effects. In this regard, the findings obtained from micro-level estimates must be interpreted as partial equilibrium results as they suffer from the "missing intercept" problem (Wolf, 2023; Sarto, 2018; Matthes et al., 2024).[5] To the best of our knowledge, this work constitutes the first empirical analysis of the aggregate economic implications of AI-driven technological progress that considers general equilibrium effects. According to our findings, AI diffusion within US innovation acts as an additional, expansionary positive supply shock, which amplifies the effects of the underlying technological developments underway. Its aggregate economic effects come from the interplay of two major forces. On the one hand, aggregate productivity gains dominate the potential displacement effects of such innovation; on the other hand, AI technology diffusion transforms the skills demanded by firms. The findings contained in this paper also help to make sense of the future potential consequences of the latest technological advancements of generative AI, to the extent of whether they will feed into the whole stream of future technological progress of the United States.

[5]Results in this literature are mixed. For instance, Bonfiglioli et al. (2023) find a negative displacement effect for low-skill and production workers, while the effect turns positive for workers at the top of the wage distribution. Hui et al. (2023) and Grennan and Michaely (2020) document a negative impact of AI on the employment of freelancers and financial analysts, respectively. Abis and Veldkamp (2024) suggest that higher data intensity led to a decline in the labor share in the investment management industry. Conversely, Brynjolfsson et al. (2023) highlight instead the complementarity of AI to workers by showing that the productivity of customer support agents has surged.

2 Transmission channels

The macroeconomic effects of AI primarily descend from its effects within the labor market. Indeed, AI adoption in workplaces might influence overall employment levels and wages via several direct mechanisms. Firstly, 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 (2020) 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 using AI to improve services that have already been automated. Moreover, new tasks may emerge as a result of AI, which could raise the demand for labor also in non-automated areas and influence the productivity of the entire production process. Thirdly, 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*).

Importantly, various indirect channels can contribute to the aggregate effect of AI adoption. A clear example is the presence of spillover effects, where productivity gains are shared across sectors via intermediate goods, or where rising incomes lead to greater overall demand. Overall, AI is expected to enhance productivity growth by automating not just the creation of goods and services, but also the generation of ideas (Aghion et al., 2017).

3 Data

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).[6] This dataset, obtained from the whole patent database of the USPTO with machine learning models and validated using manual review by patent examiners, identifies 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. Search into each patent encompasses 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 and then average these scores to get a single AI score for each patent.

3.1 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. In formulas,

$$AIint_t = \sum_{i=1}^{N_t} AIscore_{i,t}$$

where $AIscore_{i,t}$ is the average AI score of patent i in month t, 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 using the dynamics of patent applications to proxy news shocks over future technology (Miranda-Agrippino et al., 2020; Cascaldi-Garcia and Vukotić, 2022; Ferriani et al., 2024). The intuition about the news effect of patenting for future technology is that filing a patent discloses information to the public – granting can come years after filing – on future technology developments.

[6] 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.

Figure 1 displays the dynamics of our measure of $AIint$. The series displays an upward trend since the beginning of our sample in 1980, which has steepened since the second half of the 1990s. In 2019 (the end of our estimation sample), the AI intensity in patenting was around 5% and about 16% of patents filed each month were AI-intensive patents. In Section 5 we describe exercises that we run to rule out that our findings are spuriously driven by $AIint$ capturing general technological trends.

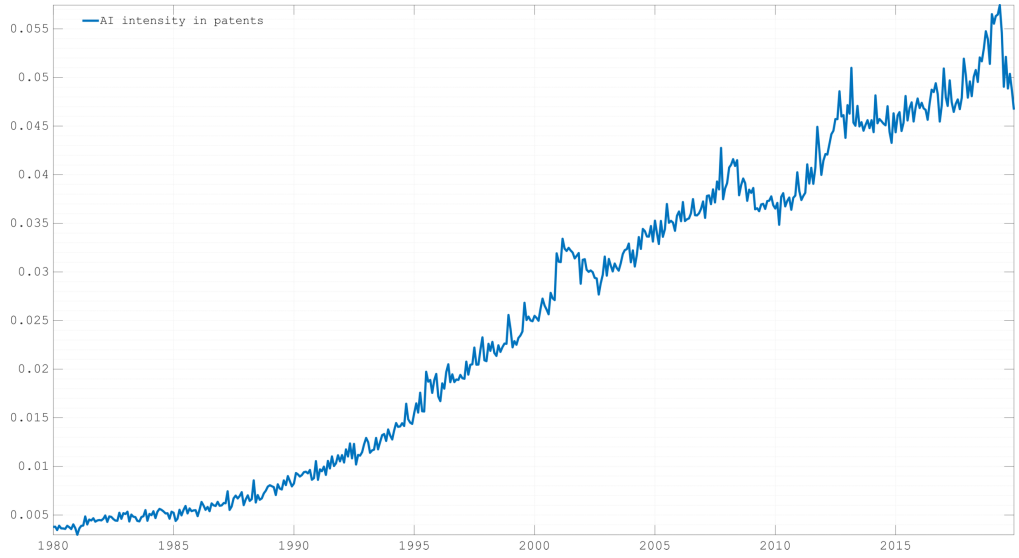


Figure 1: AI INTENSITY IN US PATENTING.

Note. The figure displays AIint at the monthly frequency between 1980 and 2019. Source: authors' calculation on the AIPD dataset

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

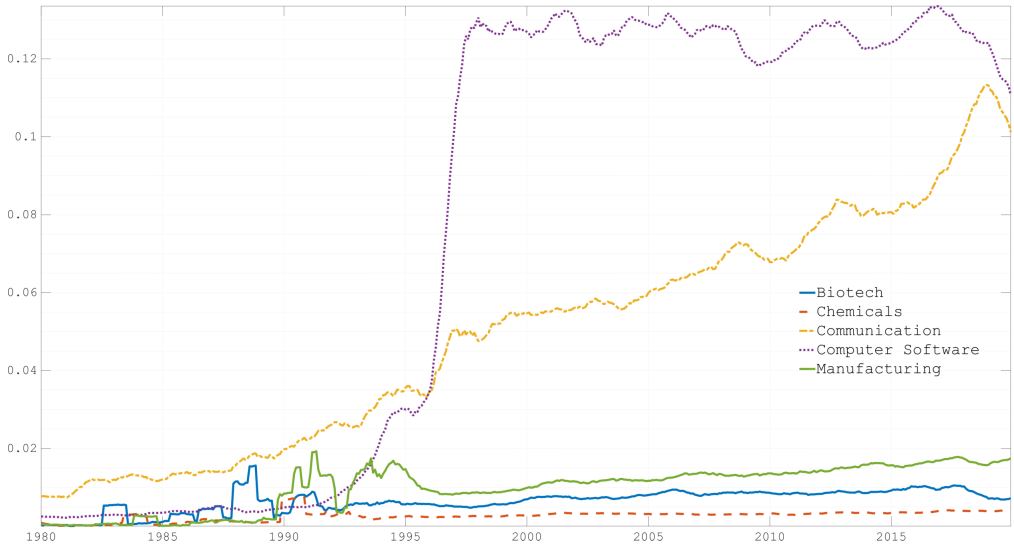


Figure 2: Number of AI-intensive patents by sector.

Note. The figure displays AIint at the monthly frequency between 1980 and 2019. Source: authors' calculation on the AIPD dataset

3.2 Examples of AI intensive patents

While AI patenting has gained momentum since the 2000s (see Baruffaldi et al., 2020 for a review), seminal AI technology developments started much earlier. Indeed, the diffusion of AI within other technologies in the United States was already underway since the late 1970s (USPTO, 2020). In particular, in the 1980s, AI-intensive patents contained the seed of those technological developments that have become popular nowadays such as large language models or speech recognition. Below we provide some notable examples and more recent AI-related patents.[7]

- In 1980, IBM filed a patent titled "Parallel pattern verifier with dynamic time warping" (US4348553A) consisting of a speech recognition system that employs a network of elementary local decision modules for matching an observed time-varying speech pattern against all possible time warpings of the stored prototype patterns.

- In 1982, Analis SA filed a patent titled "Standard hardware-software interface for connecting any instrument which provides a digital output stream with any digital host

[7] All of the following patents can be found on Google Patents.

computer" (US4485439A) that relates to the connection of data acquisition apparatus to computer systems in general, and particularly to the connection of automated laboratory apparatus to clinical laboratory and/or hospital information systems.

- In 1984, a private individual related to the University of Michigan filed the patent "Adaptive computing system capable of learning and discovery" (US4697242A). The algorithm employed a language composed of conditional statements called "classifiers" which responded to and operate fixed-length binary words called "messages".

- In 1985, Standard Systems Corp. filed the patent "Artificial intelligence system" (US4670848A) for accepting a statement, understanding the statement and making a response to the statement based upon at least a partial understanding of the statement.

- In 1989, Texas Instruments Inc. filed the patent "Chart parser for stochastic unification grammar" (US4984178A) related to spoken language interfaces, and more particularly to a spoken language processor containing a chart parser that incorporates rule and observation 10 probabilities with stochastic unification grammars.

- In 1995, Lucent Technologies Inc. filed the patent "Method and apparatus for training a text classifier" (US5675710A) to classify text by employing a supervised learning algorithm, including in the Web.

- In 2003, IBM filed the patent "Information search using knowledge" (US6636848B1) for automatic sources within a corpus of documents, including in the Web.

- In 2009, 21CT Inc. filed the patent "Hidden Markov model for speech processing with training method" (US9020816B2) relative to a method for identifying non-language speech sounds in a speech or audio signal.

- In 2013, Google filed the patent "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.

- In 2018, IBM filed the patent "Training data expansion for natural language classification" (US10726204B2) for training a natural language classifier associated with a chat interface of a computer system provided.

3.3 Orthogonality of $AIint$

As a preliminary step, we assess whether $AIint$ can be considered exogenous with respect to macro-financial conditions. Similarly to Miranda-Agrippino et al. (2020) and Ferriani et al. (2024), we evaluate the orthogonality of $AIint$ with respect to macroeconomic aggregates and structural shocks from the literature (Table 1).

Regarding the former (panel a), we employ macroeconomic expectations from professional forecasters surveyed by *Long Term Consensus Forecast*. Whereas the literature typically employs short-term forecasts, this dataset provides us with forecasts for a longer horizon, which may be important for AI innovation. We focus on the 1, 3, and 5-year ahead expectations on GDP, CPI, and bond yields that are the series available for a more extended period. Next, we assess the orthogonality of $AIint$ with respect to the extensive information contained in the FRED-MD database (McCracken and Ng, 2016), from which we extract seven factors. Notice that this test is quite severe and is typically employed to test the invertibility of shocks identified from structural VAR models. In panel (b)-(c), we report the correlation of $AIint$ with several structural shocks in the literature at the monthly and quarterly frequency, respectively. Those include monetary and fiscal policy, oil supply, and financial and uncertainty shocks.

Overall, these tests diagnose extremely weak contamination of $AIint$ by other structural shocks identified in the literature, no linear dependence on past macro-financial information, and no explanatory power from expectations on the future path of the economy. For this reason, we employ $AIint$ as an internal instrument in our econometric analysis under the assumption that it is contemporaneously exogenous with respect to the endogenous variables of interest. Running the same diagnostic tests including the variables in growth rates yields similar results (Table A1).

Panel (a): Macroeconomic aggregates

	W-stat	P-value	Obs.	Diff. R^2
Long-term Consensus Forecast	0.77	0.38	318	
McCracken and Ng (2016) FRED-MD factors	0.84	0.36	468	
TFP	3.45	0.04	156	<0.001
# patents ($AIint$)	5.18	0.02	468	<0.001

Panel (b): Monthly structural shocks

Shocks	ρ	P-value	Obs.
Baumeister and Hamilton (2019) oil supply	-0.03	0.46	480
Känzig (2021) oil supply news	0.001	0.97	480
Gertler and Karadi (2015) monetary	-0.03	0.60	324
Romer and Romer (2004) monetary	0.05	0.48	204
Baker et al. (2016) EPU	-0.04	0.38	390
Bloom (2009) uncertainty	0.002	0.95	456
Gilchrist and Zakrajšek (2012) EBP	-0.08	0.07	480
Känzig (2022) carbon policy shocks	-0.001	0.99	246

Panel (c): Quarterly structural shocks

Shocks	ρ	P-value	Obs.
Basu et al. (2006) TFP	-0.03	0.76	128
Smets and Wouters (2007) TFP	-0.08	0.44	100
Beaudry and Portier (2014) news	0.02	0.79	131
Barsky and Sims (2011) news	-0.21	0.03	111
Kurmann and Otrok (2013) news	-0.06	0.55	102
Romer and Romer (2010) fiscal	-0.05	0.57	112
Ramey (2011) fiscal	0.006	0.94	124
Fisher and Peters (2010) fiscal	-0.04	0.71	116
Mertens and Ravn (2013) private tax	-0.06	0.51	108
Mertens and Ravn (2013) corporate tax	-0.06	0.56	108

Table 1: Orthogonality of $AIint$

Notes. Panel (a): AIint is regressed on a constant, its own 12 lags, and the explanatory variables of interest. The Wald test statistics correspond to the joint significance test of Long Term Consensus Forecasts for GDP, CPI, and bond yields 1, 3, and 5 years ahead. In the case of FRED-MD factors, the explanatory variable of interest is the first 12 lags of 7 factors extracted from the FRED-MD database. Panel (b)-(c) report the correlation between the AIint residual extracted from an AR(12) process and various structural shocks from the literature. In the case of the quarterly exercise, the residuals are aggregated as quarterly averages. Diff. R^2 is the difference between the R^2 in a regression that includes AIint and the additional explanatory variable and in a regression that includes only the lags of AIint.

4 Empirical Analysis

As explained in the previous section, we employ our measure $AIint$ to proxy for an anticipated (news) shock on future AI-diffusion in technological innovation in the US economy from January 1980 to December 2019. We employ $AIint$ as an internal instrument (Plagborg-Møller and Wolf, 2021) and estimate the dynamic effects by means of local projections (LPs; Jordà, 2005). We choose LPs over VAR methods due to the long and variable lags that are likely to characterize AI technology diffusion (Brynjolfsson et al., 2018), which tilt the bias-variance trade-off across the two methods in favor of LPs (Li et al., 2024). Eq(1) describes the LPs specification that we employ throughout the analysis for each endogenous variable of interest y:

$$y_{t+h} = \alpha_h + \beta_h AIint_t + \delta_h X_{t-1} + \epsilon_t \qquad h = 0, .., 60 \tag{1}$$

where h denotes the horizon of the response, α is a constant, the coefficient β will capture the impulse response functions to an AI innovation shock, as it is associated with $AIint_t$; X is a set of controls that include 12 lags of y, $AIint$, and other variables that are specific to each econometric exercise; ϵ_t is the residual that will display a moving average structure across h, the inference is based on Newey and West (1994) standard errors.

Our analysis proceeds in five steps. First, we focus on the macroeconomic effects on the US economy. Second, we investigate the heterogeneous response of the price of consumption categories. Third, we analyze the impact of AI tech shocks on TFP, highlighting the crucial role played by general equilibrium effects. Fourth, we delve into the labor market by studying occupational flows, the role of education, and sectoral quantities. Finally, we investigate the consequences of AI tech diffusion for inequality.

4.1 Aggregate Effects

Our endogenous variables consist of a mix of macro-financial indicators, with particular attention to the labor market. We include industrial production, the headline Personal Consumption Expenditures (PCE) price index, employment, weekly hours worked, and hourly earnings, the 1- and 10-year Treasury yields, and the S&P500. The estimation is conducted at the monthly frequency and includes 12 lags of each variable as controls.

Figure 3 displays the aggregate responses to a $AIint$ shock. The impact on the US economy is consistent with an expansionary technology news shock as i) all variables respond with delay and ii) output increases while consumer prices fall (Portier, 2014; Miranda-Agrippino et al., 2020). The positive response of the labor market quantities, i.e., employment and hours worked, is consistent with findings in Chahrour et al. (2023). $AIint$ shock is nonetheless distinct from standard TFP shocks because of the response of wages: hourly earnings increase rather than fall. These results suggest that the combination of worker complementary to AI and the aggregate productivity gains (see Section 4.1.2) dominate the displacement effect of this innovation (documented for instance in partial equilibrium by Bonfiglioli et al., 2023).

Another key feature that makes the propagation of an $AIint$ shock different from that of a standard technology shock concerns the speed of the response of the stock market, which is often thought of as discounting relatively quickly the beneficial effects of technological news on the macro economy (Beaudry and Portier, 2014). Similarly to our analysis, Miranda-Agrippino et al. (2020) exploit patent data (as the raw number of patents filing) to identify technology news shocks. Specifically, they proxy TFP news shocks via the number of filed patents and find a rapid response of stock valuations. In our case, stock prices display a delayed response as we use information on the intensive AI margin of the patenting process, whose characteristics and consequences can be harder to gauge for investors than the extensive margin (the number of patents).[8]

Although the response of the headline PCE price index is mild, the fall in core PCE prices is powerful (Figure A6) as food and energy are virtually unaffected by $AIint$ shocks (see Section 4.2).

[8]Results employing $AIshare$ are very similar to our baseline (Figure A1).

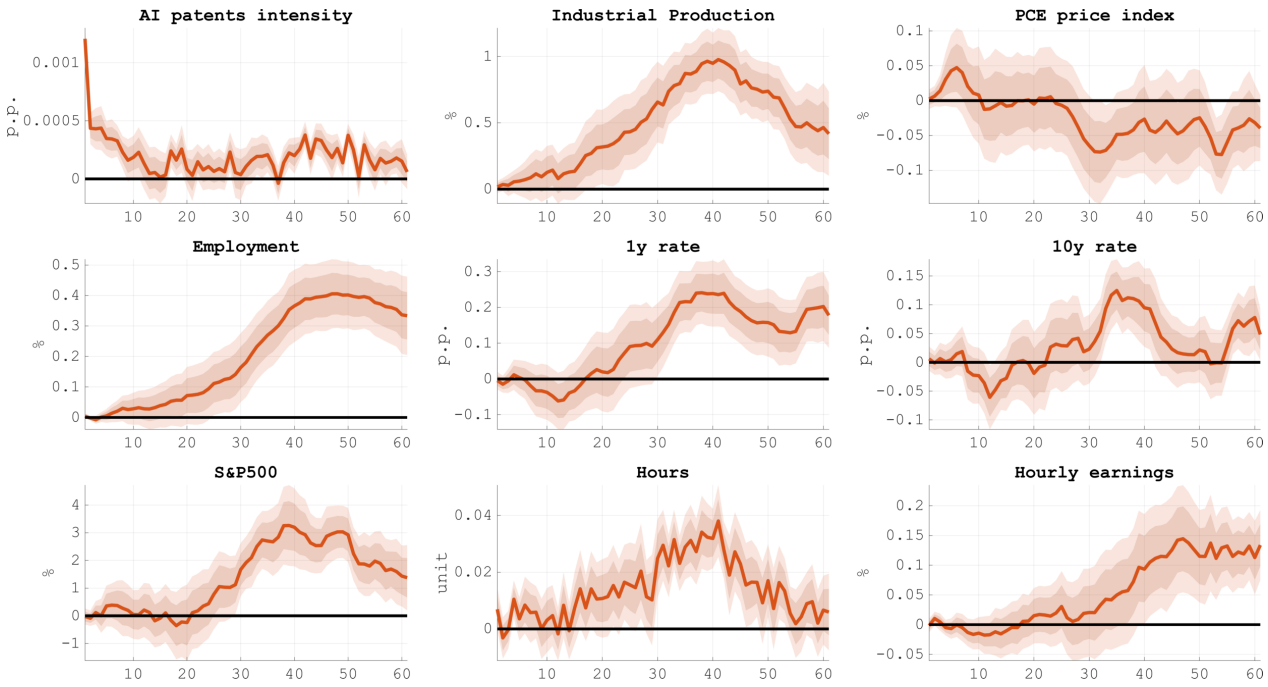


Figure 3: IRFs to an AI technology shock.

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 estimate and 68%-90% confidence bands.

4.1.1 Quantitative implications

The IRFs in Figure 3 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. We follow Gorodnichenko and Lee (2020) and apply their R^2 approach that yields the best performances in finite samples, among other alternatives. The method amounts to estimate the shocks to $AIint$ (the residual from eq. 1 with $y = AIint$; $\widehat{\epsilon}_{AI}$) and employ them in a regression (eq. 2) on the forecast error of the endogenous variable of interest j $\widehat{\epsilon}_j$:

$$\widehat{\epsilon}_{j,t+h|t-1} = \omega_{z,0}\widehat{\epsilon}_{AI,t+h} + \cdots + \omega_{z,h}\widehat{\epsilon}_{AI,t} + \tilde{v}_{t+h|t-1} \qquad \forall j = \text{endog. vars} \tag{2}$$

The R^2 from this regression yields the variance contribution of $AIint$ to y_j. The inference is based on 10000 bootstrap replications.

Results in Figure 4 highlight that overall, $AIint$ shocks have not been a major driver of

the US business cycle. Nonetheless, they explain about 20% of the variance in industrial production, consumer prices, and stock prices at longer horizons. Conversely, their effects on the labor market and monetary policy have been modest in the sample under examination.

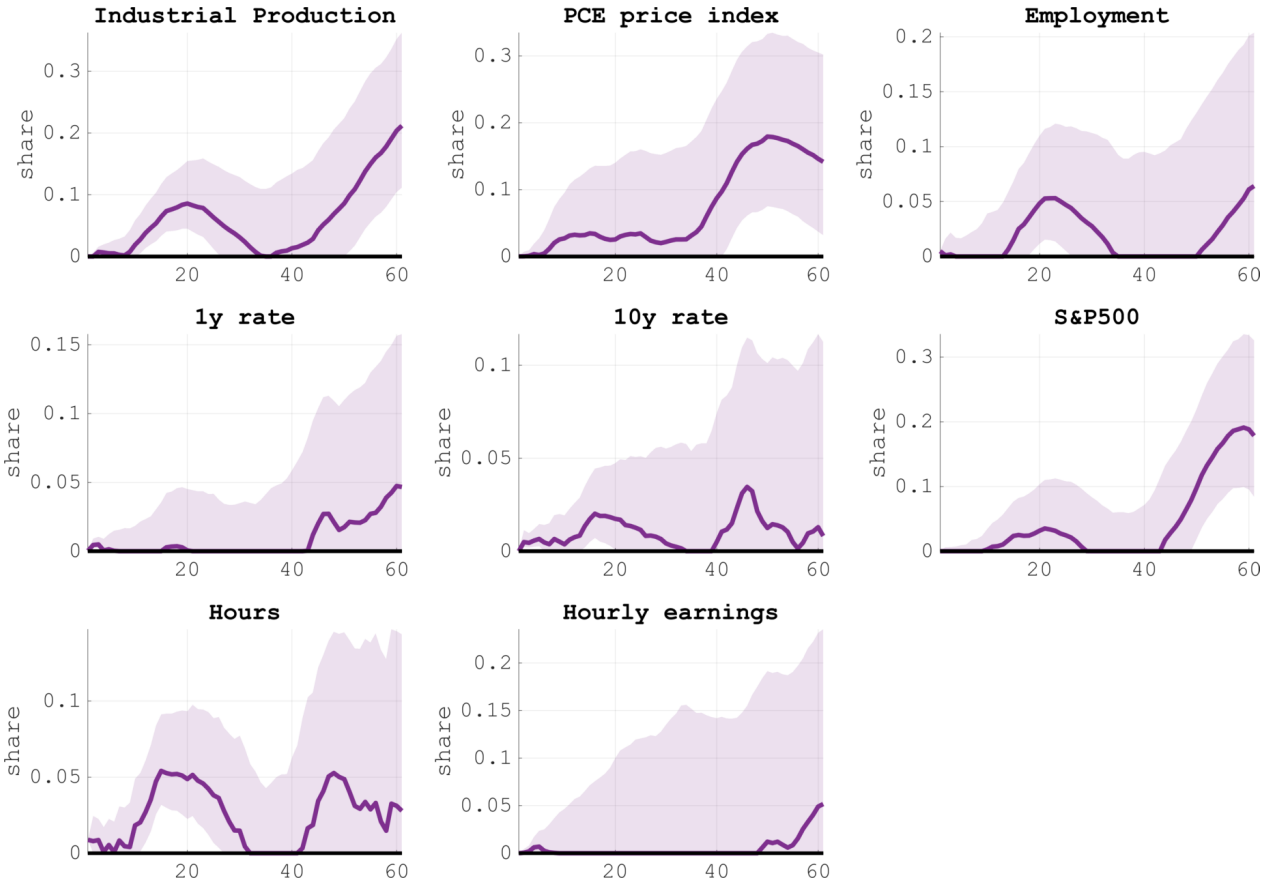


Figure 4: VARIANCE CONTRIBUTION OF AI TECHNOLOGY SHOCKS.

Note. The figure displays the variance contribution of AIint shock. Sample 1980-2019. The estimates are based on local projections with 12 lags. Point estimate and 90% confidence bands from 10000 bootstrap replications.

4.1.2 The role of productivity

The response of TFP rationalizes the aggregate response of the US economy. As TFP is available only at the quarterly frequency, we estimate the effect of $AIint$ shocks on utilization-adjusted TFP by transforming $AIint$ as the average over the quarter. The IRFs in Figure 5 display a delayed and yet very persistent response and contribute to highlight the importance of considering general equilibrium effects when assessing the economic consequence of AI technology diffusion (the "missing intercept" problem, see Wolf, 2023).

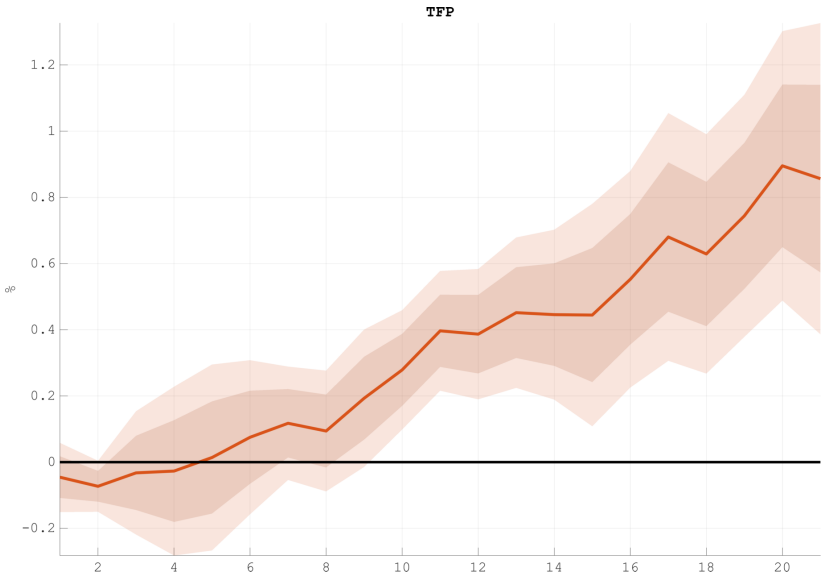


Figure 5: TFP RESPONSE TO AN AI TECHNOLOGY SHOCK.

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 estimate and 68%-90% confidence bands.

4.2 Heterogeneity in consumer prices

Consumer prices tend to fall after an *AIint* shock at the aggregate level. However, by delving into consumption types, we expose a very stark heterogeneity in the response of consumer prices. Figure 6 displays the responses of several consumption categories. The mild aggregate fall in consumer prices is driven by a marked drop in the prices of durable goods and, in particular, by high-tech goods such computers, software, TVs, etc. contained in the category *recreational goods and vehicles*.[9]

[9]This finding is consistent with the companies as IBM, Microsoft, and Google being the major firms filing AI-intensive patents.

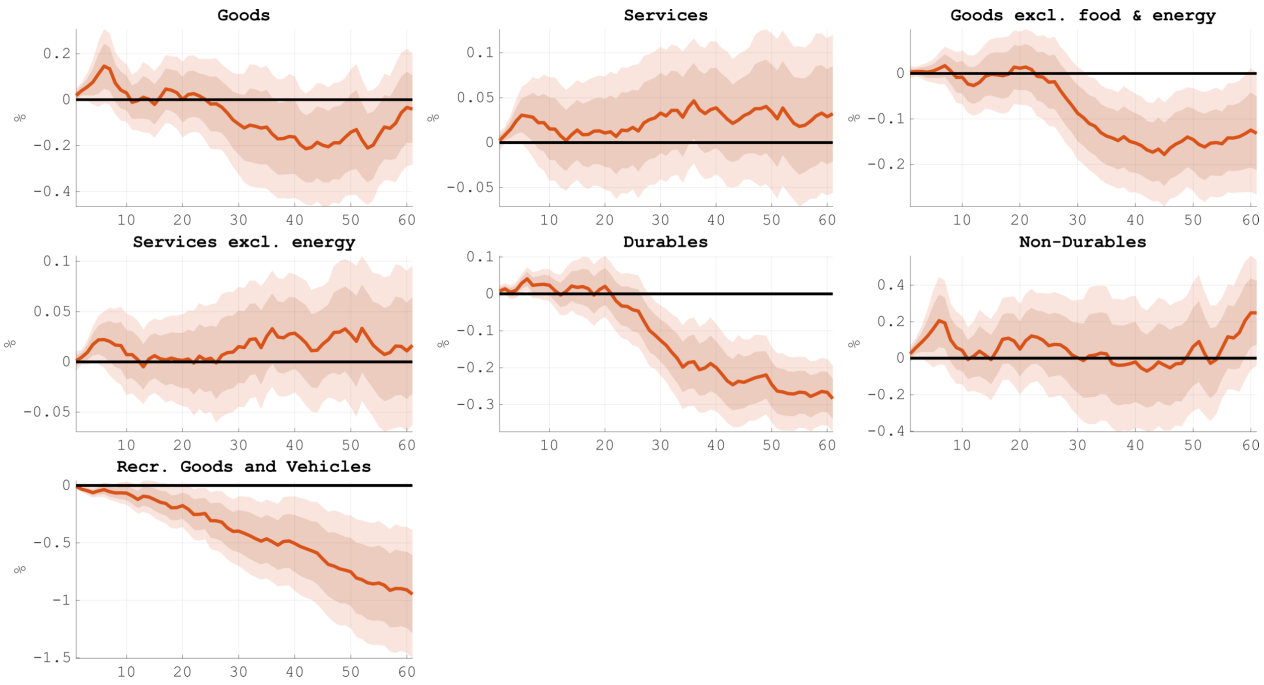


Figure 6: Consumer prices response to an AI technology shock.

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 estimate and 68%-90% confidence bands.

4.3 A zoom on the labor market

Shocks to AI-based technological diffusion transmit crucially through the labor market. We examine in further detail the consequences of *AIint* for occupational flows and sectoral quantities. Figure 7 displays the aggregate responses of occupational flows. Despite improving employment, hours, and wages, *AIint* shocks transform the US labor market by modifying 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 3. By exploring the dynamics of outflows in detail, we find that separations increase due to increased layoffs, whereas the number of quits falls. This pattern is consistent with firms shifting the labor force composition towards workers who are more complementary to AI-based progress to exploit the new technological advancements.

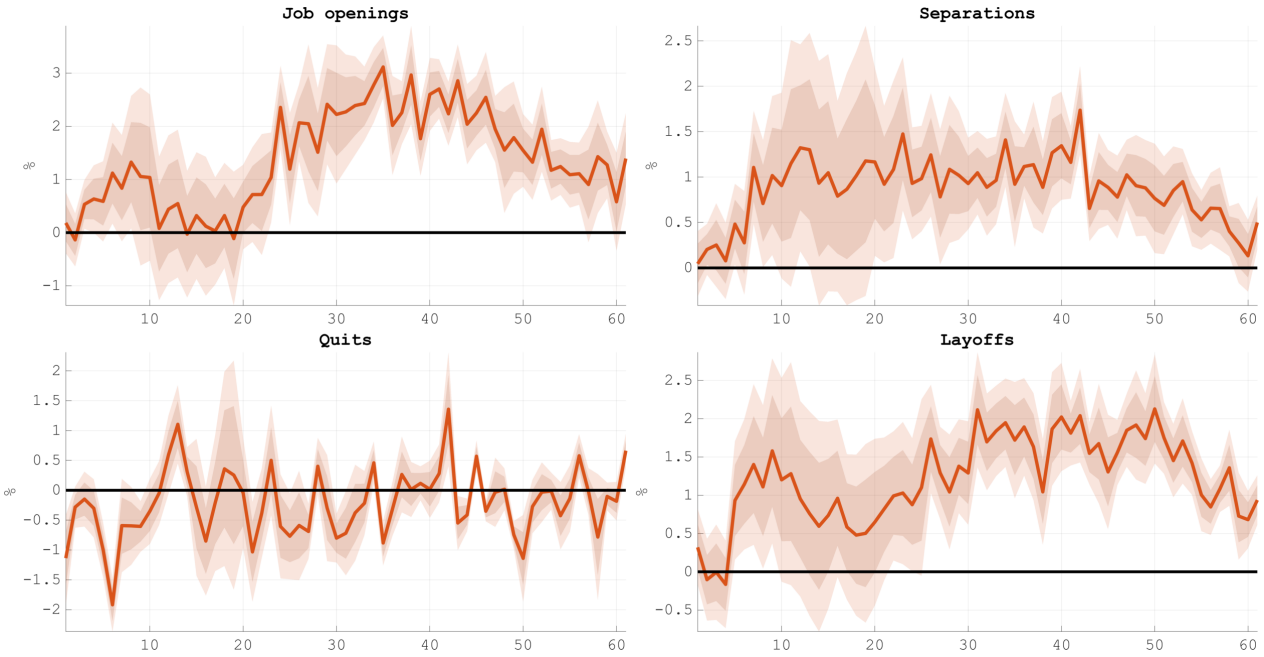


Figure 7: Response of labor market flows to an AI technology shock.

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 estimate and 68%-90% confidence bands.

4.3.1 Sectoral heterogeneity

Next, we investigate the response of some selected labor market variables by sector. Figure 8 displays the LP-based response, cumulated over 60 months, of labor quantities, wages, and flows at the sectoral level. The sectors are ordered according to the AI exposure measure in Acemoglu et al. (2022), which is in turn based on occupation-level exposure reported in Felten et al. (2018).[10] High AI-exposed sectors are those in which the overlap between occupational abilities and the scope of past AI advances is most prominent, so where the effects of AI-based tech shocks on employment and wages might, in principle, also be the largest. In line with our aggregate findings, sectoral evidence is consistent with the presence of sizeable general equilibrium effects, given that we do not observe substantially heterogeneous effects across differently exposed sectors. Wages in sectors less exposed to AI appear to grow more than those in sectors more exposed to AI. These patterns might be consistent with the partial

[10]In Acemoglu et al. (2022), sector-level exposure is the weighted average of occupation-level exposure weighted by the number of vacancies posted by each sector in each occupation. Other papers proposing matching occupations and AI exposure of working tasks are Webb (2019) and Eloundou et al. (2023).

equilibrium displacement effect based on micro-level estimates coupled with strong positive general equilibrium effects.

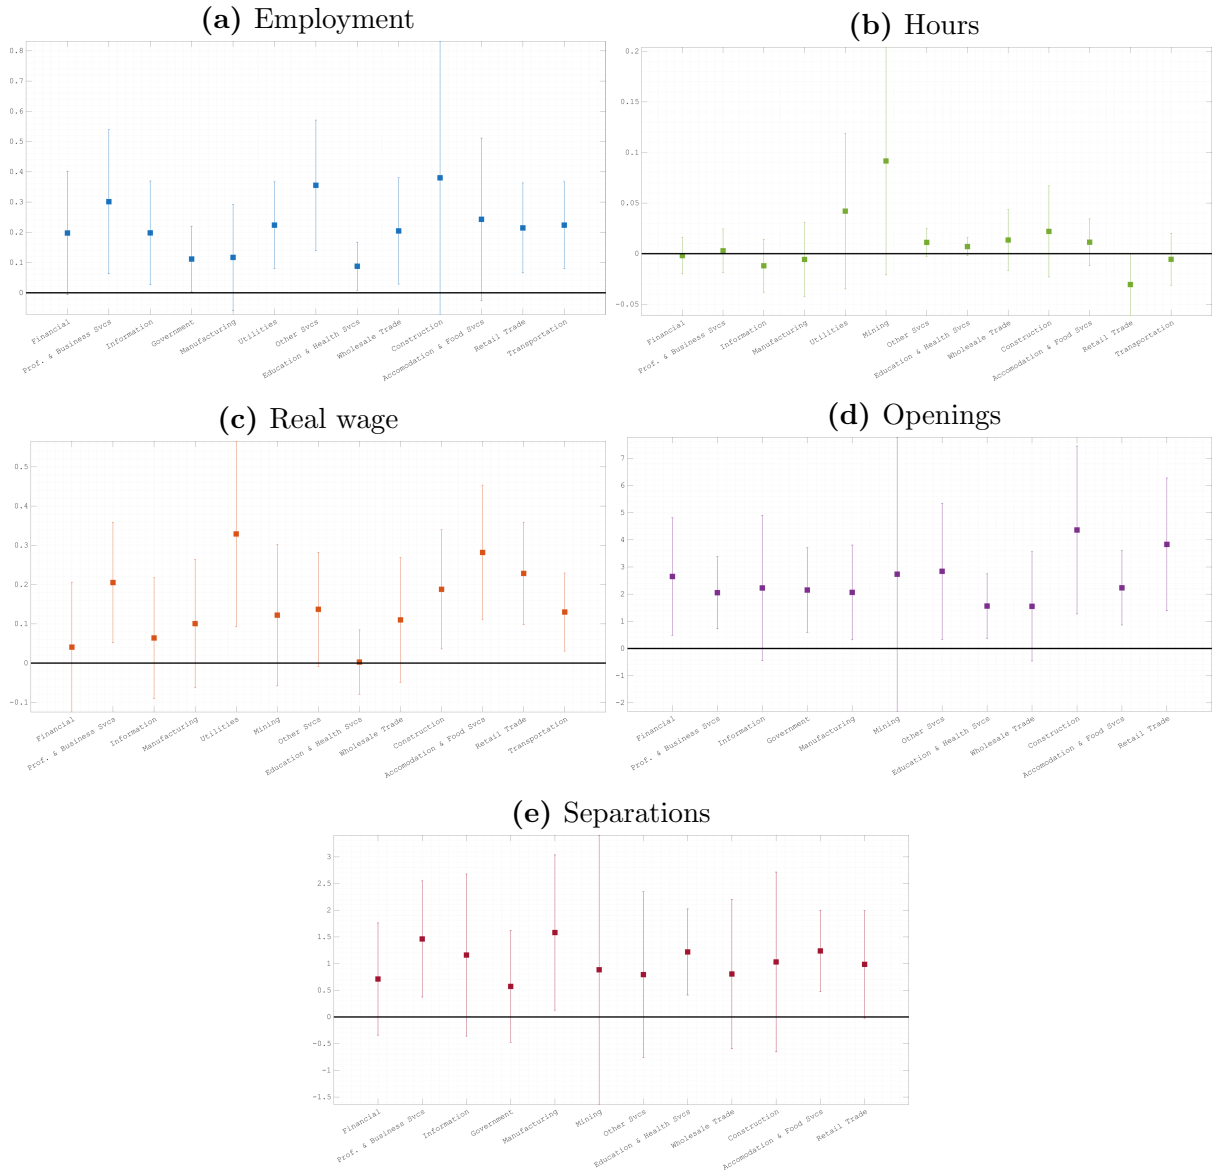


Figure 8: Response of labor market flows to an AI technology shock.

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 estimate and 90% confidence bands.

4.3.2 The Role of Education

Education is another potential source of heterogeneity for the transmission of *AIint* shocks in terms of employment and wages. We investigate the response of these two variables across different degrees of educational attainment, spanning from lower than high school to higher than a bachelor's degree. Figures 9 show that, in terms of employment, the two intermediate groups are the ones that benefit the most, where those who did not attend high school appear to be worst off, whereas the group with top education benefit only marginally. Consequently, AI-driven technological advances may have contributed to the secular decline in the demand for low-skilled workers in developed economies Berman et al. (1998). The picture is different for earnings (Figure 10) as, in this case, all groups appear to benefit more homogeneously.

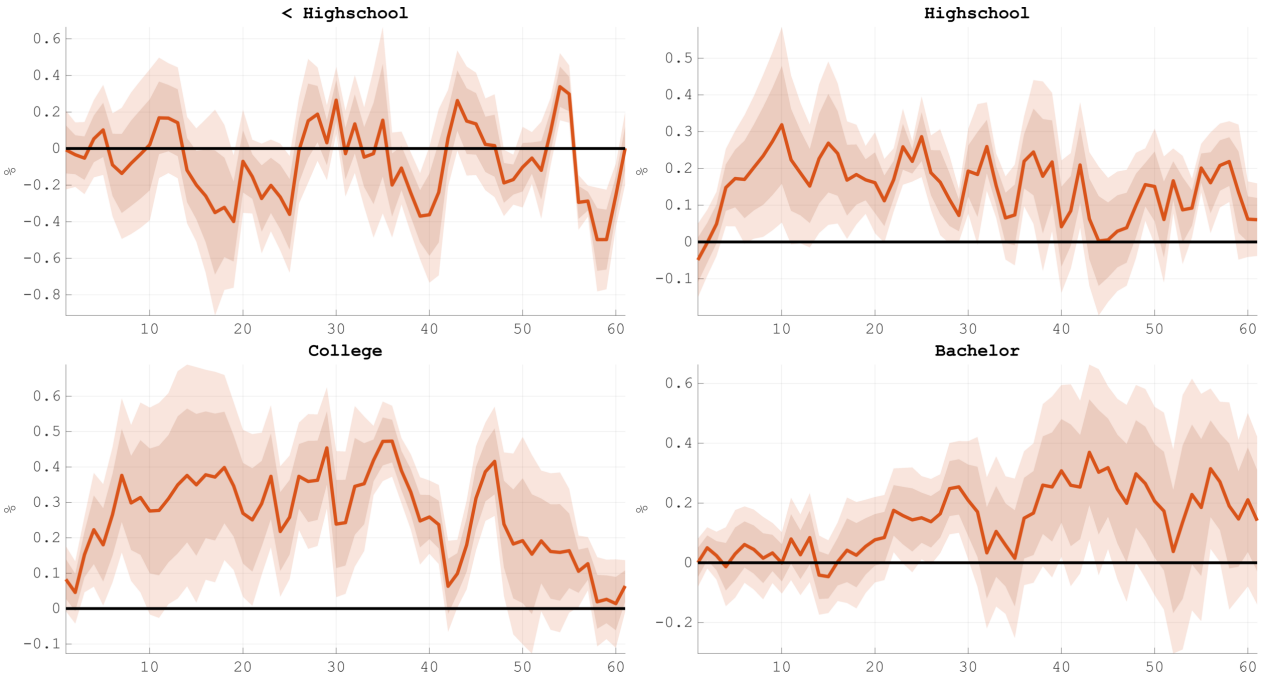


Figure 9: Response of employment by educational attainment to an AI technology shock.

Note. The figure displays the IRFs to a AIint shock. Sample 1992-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.

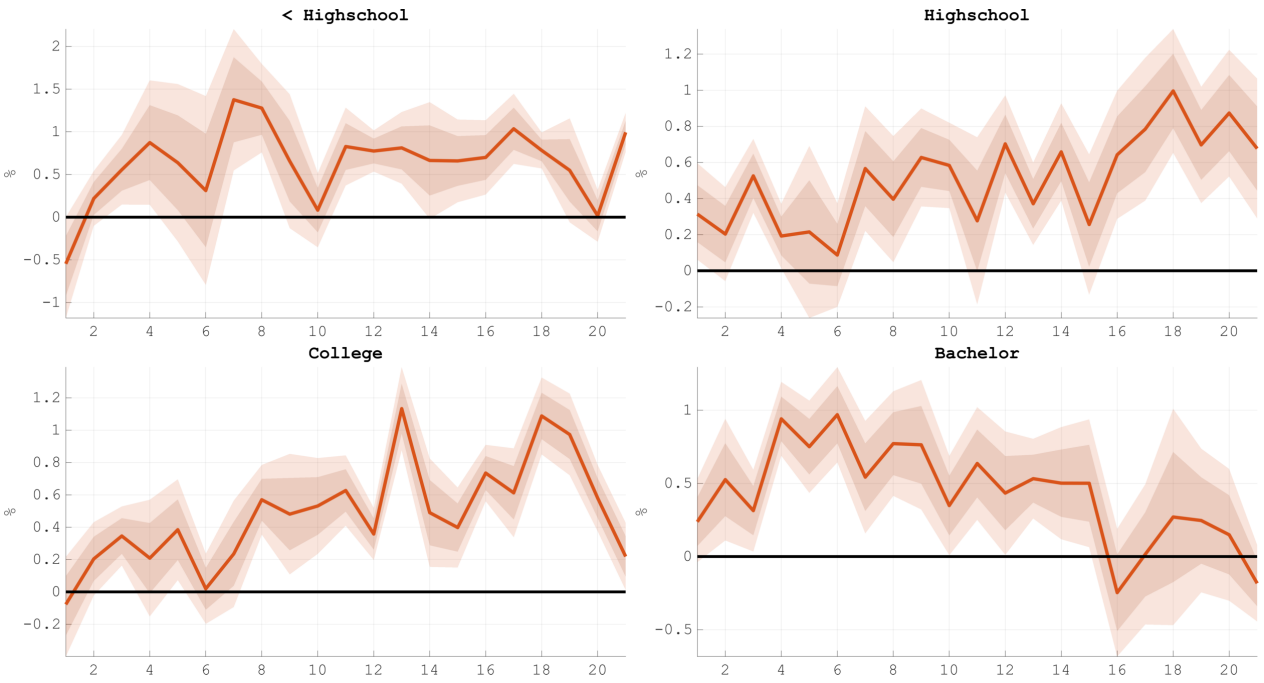


Figure 10: Response of earnings by educational attainment to an AI technology shock.

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 estimate and 68%-90% confidence bands.

4.4 Implications for Inequality

Several commentators suggest that one of the major impacts of AI tech diffusion concerns inequality. We explore this mechanism by exploiting the data compiled by Blanchet et al. (2022).[11] Figure 11 reports the effect of *AIint* shocks on overall income, wealth, and labor income shares. The total income distribution (top raw) is the least affected; broadly speaking, the top shares go up, although the growth appears weak in terms of economic magnitudes and statistical precision. A different pattern emerges for the wealth distribution (second row): the top 1% and especially the top 10% of the wealth share increases while the AI tech diffusion is detrimental for the bottom 50%. Notably, the responses of the latter two groups are very persistent. The negative wealth effect on the bottom 50% is not just a matter of shares: additional estimates on the *level* of wealth across different percentiles of the distribution, displayed in Figure A20, suggest that an increase in AI intensity also anticipates a wealth

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

loss for the poor. These results echo the findings in Moll et al. (2022) on the impact of automation on inequality via asset returns and, thus, wealth. A similar pattern emerges for the labor income share. An important difference is that the share for the bottom quantile of the distribution falls but, in the case of labor income, the effects appear to be transitory.[12] These results deal exclusively with redistributive consequences while not being informative of net income or wealth gains for each group. In the Appendix, we show that all groups tend to benefit in terms of real absolute income but not in terms of real absolute wealth.

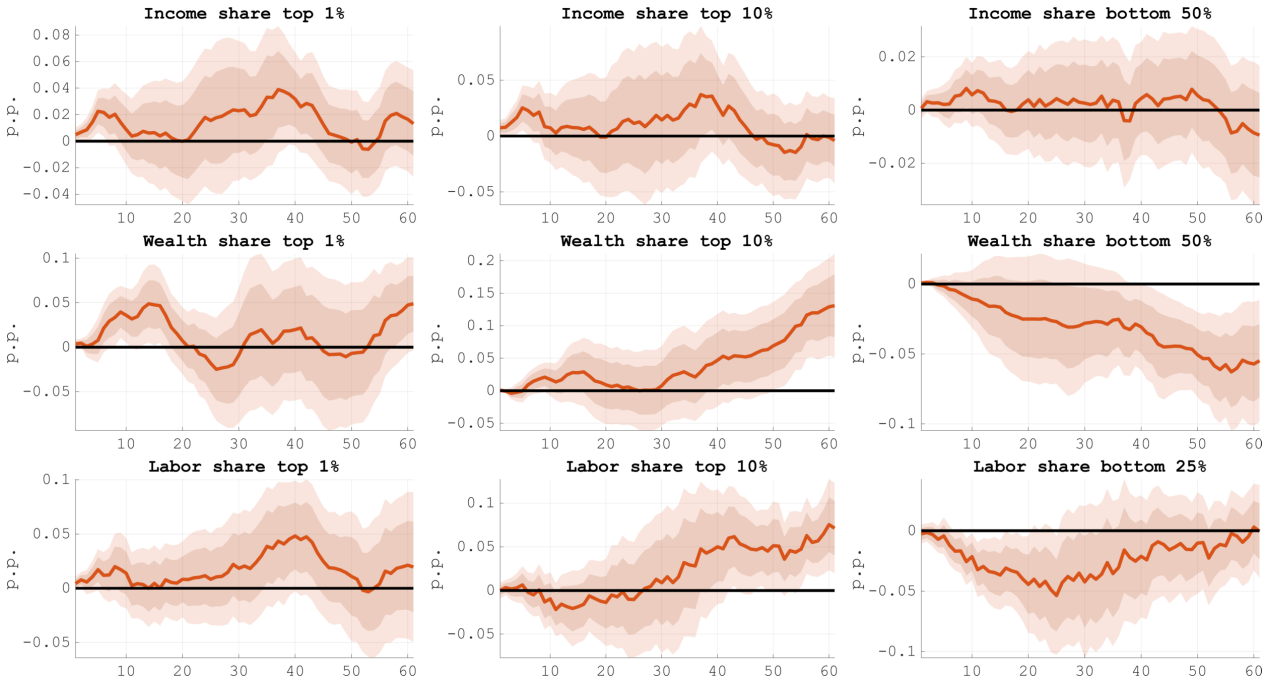


Figure 11: Response of inequality to an AI technology shock.

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 estimate and 68%-90% confidence bands.

5 Robustness Exercises

The economic insights from our analysis are robust to various perturbations of the baseline econometric specification. Results are reported in Appendix.

AI share rather than intensity. We consider an alternative measure ($AIshare$) con-

[12] The income and wealth breakdown provided in https://realtimeinequality.org/ is different from the one for labor income.

structed as the share of the number of "AI patents" (defined by the AIPD as those with a score exceeding the 0.5 threshold) over the total number of (granted) patents filed in month t.[13] This alternative measure is important in light of a potential concern: if AI developments are correlated with general technological progress, our analysis based on $AIint$ would spuriously impute general technological innovation to AI advancement. Nonetheless, the two indicators contain a very similar picture of the AI content in US innovation and yield similar economic insights in our empirical analysis. Results employing $AIshare$ rather than $AIint$, as defined in Section 3, are consistent with the baseline (Figure A1).

Raw number of AI patents. Another possibility is that, independently of the general technological process, what matters is the raw number of AI-intensive patents. Figure A4 shows that an increase in the number of AI-intensive patents produces similar macroeconomic implications to our baseline estimation. These results also hold when controlling for the lagged (Fig) or the contemporaneous level (Figure A3) of total patents (Figure A2).

Using $AIint$ controlling for the number of patents. Another econometric approach consists of employing $AIint$ but controlling for general technological trends. $AIint$ displays indeed a positive correlation with the number of patents filed each month (Table 1). Although we have shown that the explanatory power of the number of patents for $AIint$ is negligible, we have performed the LP estimation controlling for the number of patents (even the contemporaneous values). Results are not materially affected (Figure A5), and the response of the number of patents is substantially null.

Alternative LPs specification. The choice of variables does not significantly affect our results. In the Appendix, we display alternative specifications where we (i) substitute the headline PCE index with the Consumer Price Index (CPI) index or their core definition; ii) and the equity index (Nasdaq or Dow Jones rather than S&P500, or a relative measure of Nasdaq out of Dow Jones); (iii) include other variables such as real private consumption expenditures.

Stationarity. In our baseline, $AIint$ enters in levels that are trending. Similar results hold if we remove the linear or quadric trend from $AIint$ or if we use the variable in growth rates.

[13] In formulas $AIshare = \frac{\#\text{ AI patents}_t}{\#\text{ total patents}_t}$. A measure of a news-based shock of this kind, constructed using the number of climate change mitigation patents over the total number of patents to get an anticipated shock to a greener technology mix, is proposed in Ferriani et al. (2024).

Additionally, results are robust if we include a trend directly in the LP estimation.

Estimation with iid shocks. An alternative modeling strategy is to extract the residuals from an AR(12) process fitted to $AIint$ and use them in our LP estimation. Results are very similar to the baseline.

Excluding overlap with robotics innovation. AI innovation may be potentially correlated with innovation in robotics and automation. If this was the case, we would erroneously combine the aggregate effects of AI and automation. To exclude this contamination, we perform two exercises. First, we check the AI intensity in robotic patents: this is similar to the AI intensity in non-robotic patents. Second, we include the share or number of robotic patents in our LP estimation and find results that are unaffected. Moreover, the response of robotic patents is basically null.

Controlling for financial conditions. Our economic conclusions are unaffected if we control for financial condition by including the excess bond premium (Gilchrist and Zakrajšek, 2012) or economic policy uncertainty (Baker et al., 2016) in our estimation.

6 Conclusions

AI-based innovations such as *ChatGPT* have brought AI to the center of the academic and policy debate. The future socio-economic implications comprise a broad spectrum of views from idyllic to apocalyptic scenarios. Although the future direction and speed of AI technology are tough to anticipate, in this paper, we have exploited data on the AI intensity of US patents between 1980 and 2019 to study the macroeconomic consequences of AI-based innovation. Our findings can help to discipline the theoretical modeling of the economic effects of AI. The key insight is that the general equilibrium effects of AI innovation can be rather sizeable and should be appropriately taken into account when drawing policy implications. In particular, partial equilibrium analyses miss the contribution of AI to 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 in terms of employment, hours worked, and salaries. However, the more intensive use of AI technology in innovation produces ad-

verse distributional effects. The share of wealth owned by the bottom 50% distribution of the population falls, potentially calling for a policy response that should be explored in future research.

References

Abis, S. and L. Veldkamp (2024): "The changing economics of knowledge production," *The Review of Financial Studies*, 37, 89–118.

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

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

Albanesi, S., A. D. Silva, J. F. Jimeno, A. Lamo, and A. Wabitsch (2023): "New Technologies and Jobs in Europe," IZA Discussion Papers 16227, Institute of Labor Economics (IZA).

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.

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 (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. (2009): "The impact of uncertainty shocks," *econometrica*, 77, 623–685.

Bonfiglioli, A., G. Gancia, I. Papadakis, and R. Crinò (2023): "Artificial Intelligence and Jobs: Evidence from US Commuting Zones," Tech. rep., CESifo Working Paper.

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. R. Raymond (2023): "Generative AI at work," Tech. rep., National Bureau of Economic Research.

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.

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

Chahrour, R., S. K. Chugh, and T. Potter (2023): "Anticipated productivity and the labor market," *Quantitative Economics*, 14, 897–934.

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," .

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

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.

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. and R. Peters (2010): "Using Stock Returns to Identify Government Spending Shocks*," *The Economic Journal*, 120, 414–436.

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.

Gordon, R. J. (2018): "Why Has Economic Growth Slowed When Innovation Appears To Be Accelerating?" CEPR Discussion Papers 13039, C.E.P.R. Discussion Papers.

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

Grennan, J. and R. Michaely (2020): "Artificial intelligence and high-skilled work: Evidence from analysts," *Swiss Finance Institute Research Paper*.

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

Hampole, M., D. Papanikolaou, L. D. Schmidt, and B. Seegmiller (2025): "Artificial Intelligence and the Labor Market," Working Paper 33509, National Bureau of Economic Research.

Hui, X., O. Reshef, and L. Zhou (2023): "The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market," *Available at SSRN 4527336*.

Ilzetzki, E. and S. Jain (2023): "The impact of artificial intelligence on growth and employment," Tech. rep., VoxEU.

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.

——— (2022): "The unequal economic consequences of carbon pricing," *Available at SSRN 3786030*.

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.

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

McCarthy, J. et al. (2007): "What is artificial intelligence," .

McCracken, M. W. and S. Ng (2016): "FRED-MD: A monthly database for macroeconomic research," *Journal of Business & Economic Statistics*, 34, 574–589.

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. H. Hoke, and K. Bluwstein (2020): "Patents, News, and Business Cycles," CEPR Discussion Papers 15062, C.E.P.R. Discussion Papers.

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

Morikawa, M. (2016): "The Effects of AI and Robotics on Business and Employment: Evidence from a Survey on Japanese Firms," Tech. Rep. 16-E-066, RIETI Discussion Paper.

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

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.

Panel (a): Macroeconomic aggregates

	W-stat	P-value	Obs.	Diff. R^2
Long-term Consensus Forecast	0.31	0.57	317	
McCracken and Ng (2016) FRED-MD factors	0.32	0.56	467	
TFP	2.24	0.14	155	<0.001
# patents (*AIint*)	9.95	0.001	467	<0.01

Table A1: Orthogonality of *AIint* - growth rates

Notes. Panel (a): AIint, in growth rates, is regressed on a constant, its own 12 lags, and the explanatory variables of interest (in growth rates if appropriate - e.g. TFP). The Wald test statistics correspond to the joint significance test of Long Term Consensus Forecasts for GDP, CPI, and bond yields 1, 3, and 5 years ahead. In the case of FRED-MD factors, the explanatory variable of interest is the first 12 lags of 7 factors extracted from the FRED-MD database. Diff. R^2 is the difference between the R^2 in a regression that includes AIint and the additional explanatory variable and in a regression that includes only the lags of AIint.

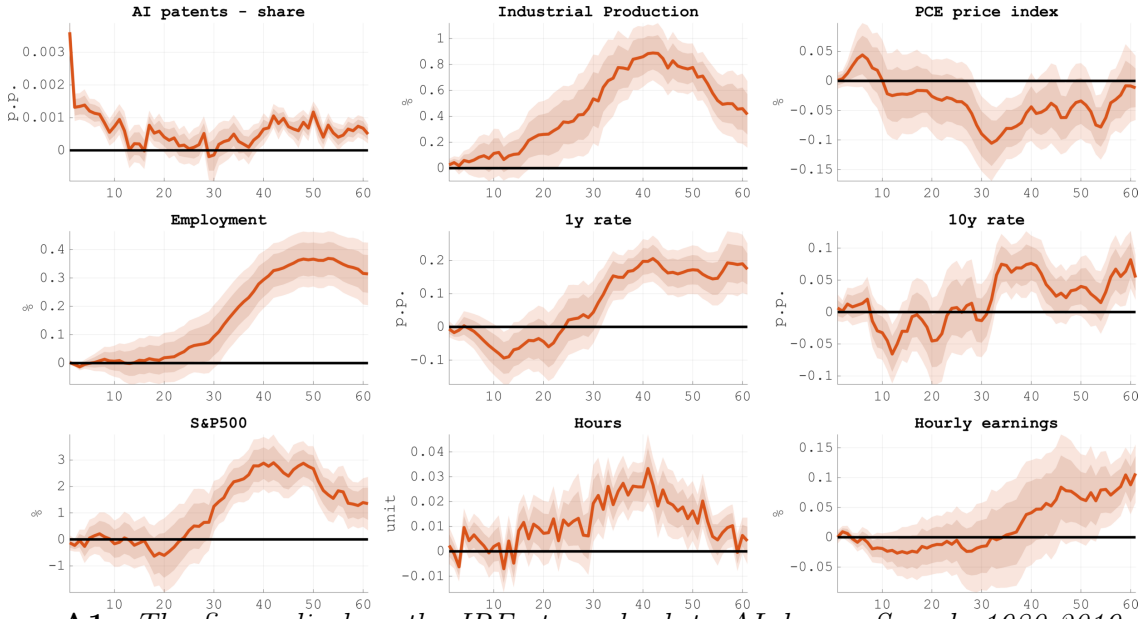


Figure A1: *The figure displays the IRFs to a shock to AIshare. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.*

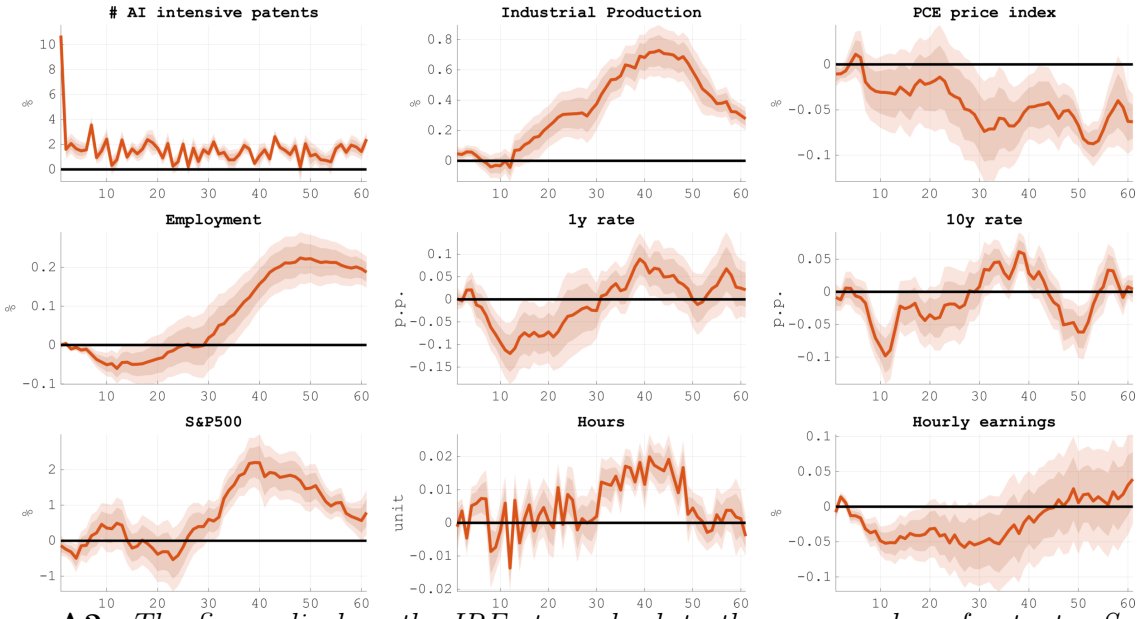


Figure A2: *The figure displays the IRFs to a shock to the raw number of patents. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.*

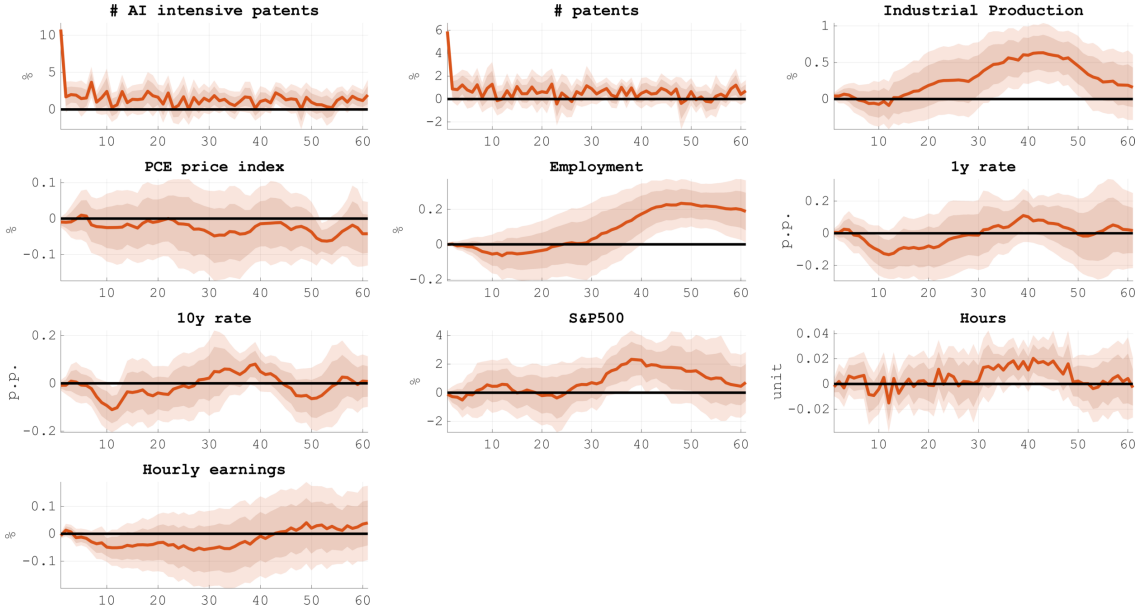


Figure A3: *The figure displays the IRFs to a shock to the raw number of patents, controlling for lags of total patents. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.*

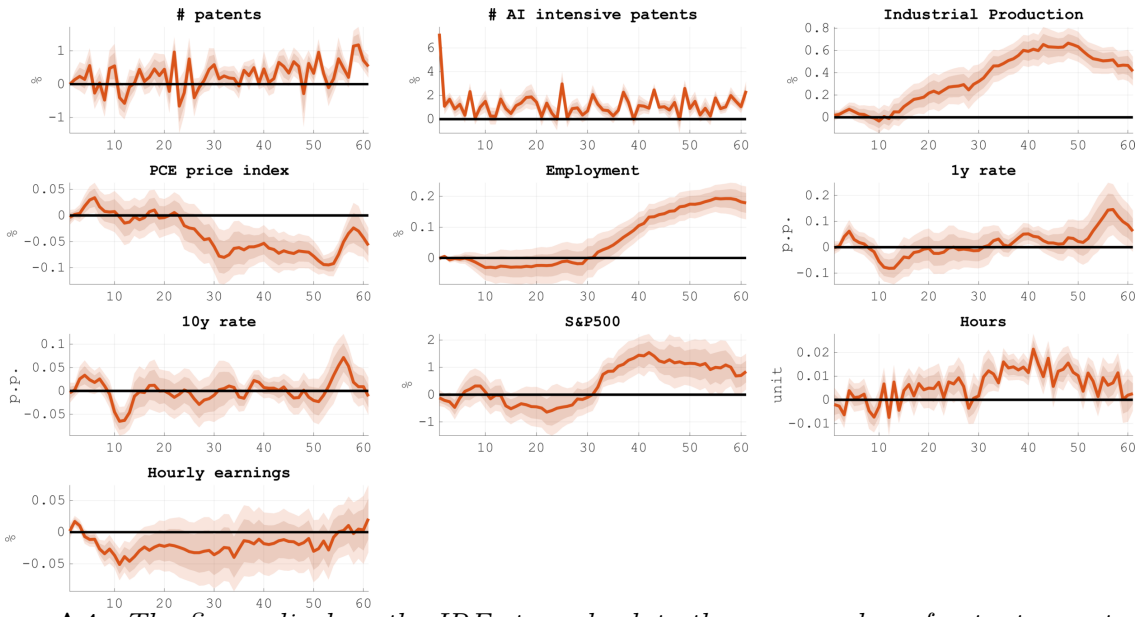


Figure A4: *The figure displays the IRFs to a shock to the raw number of patents, controlling for contemporaneous values of total patents. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.*

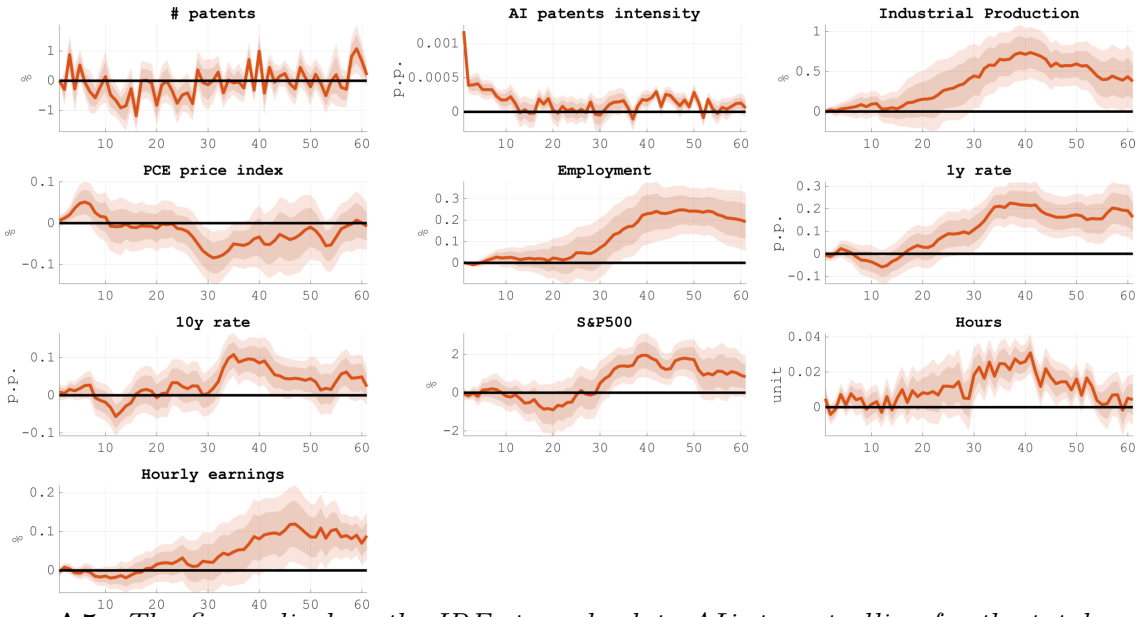


Figure A5: *The figure displays the IRFs to a shock to AIint controlling for the total number of patents. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.*

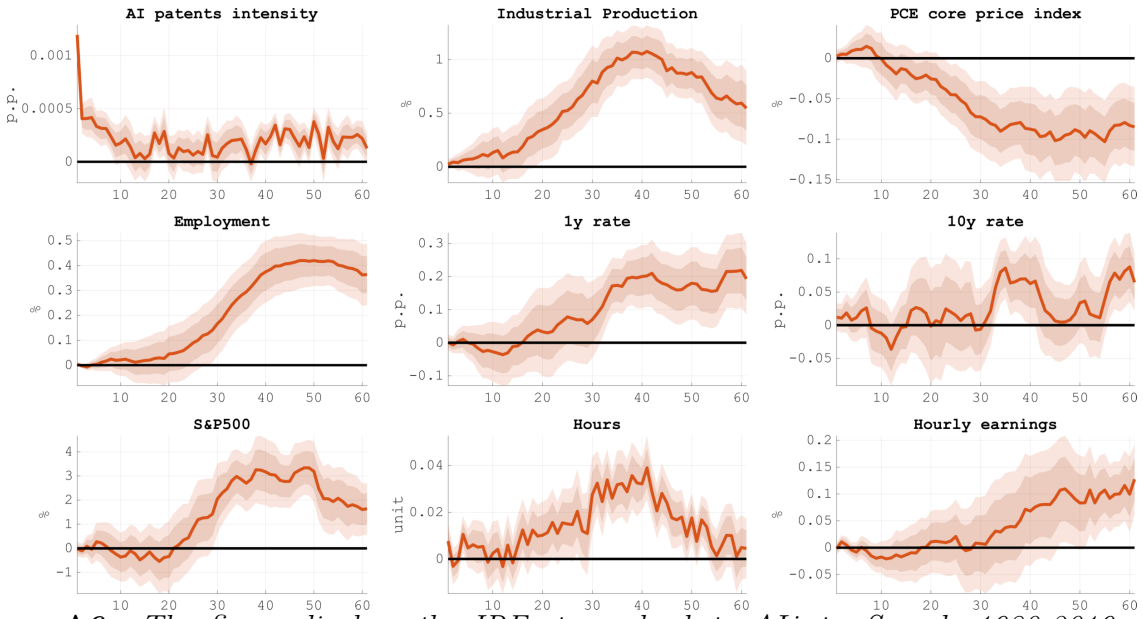


Figure A6: *The figure displays the IRFs to a shock to AIint. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.*

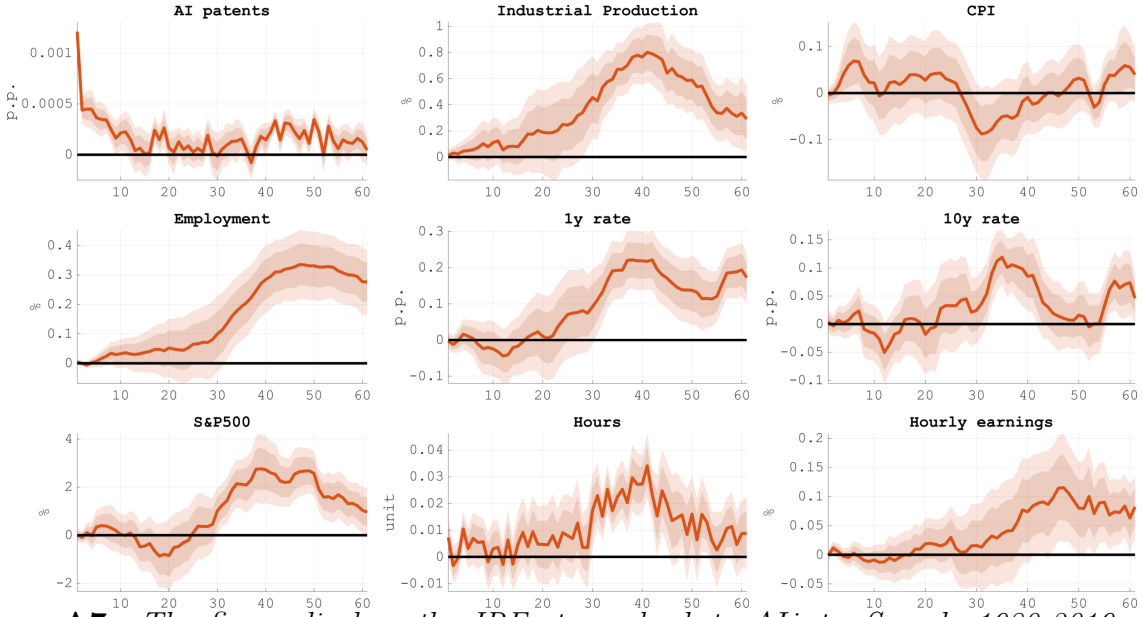


Figure A7: *The figure displays the IRFs to a shock to AIint. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.*

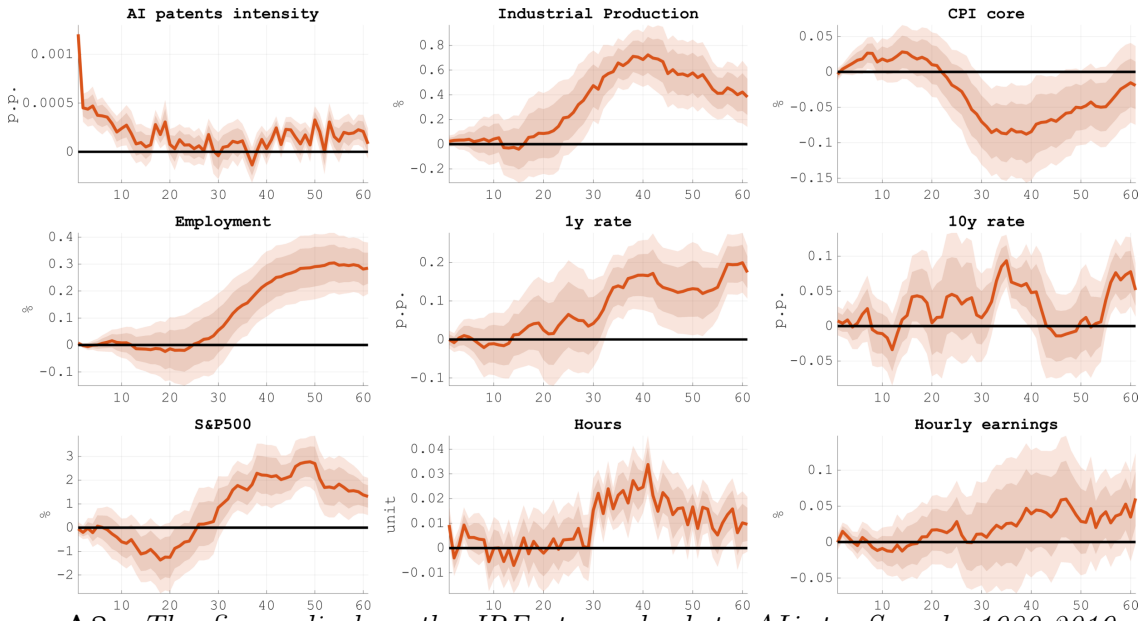


Figure A8: *The figure displays the IRFs to a shock to AIint. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.*

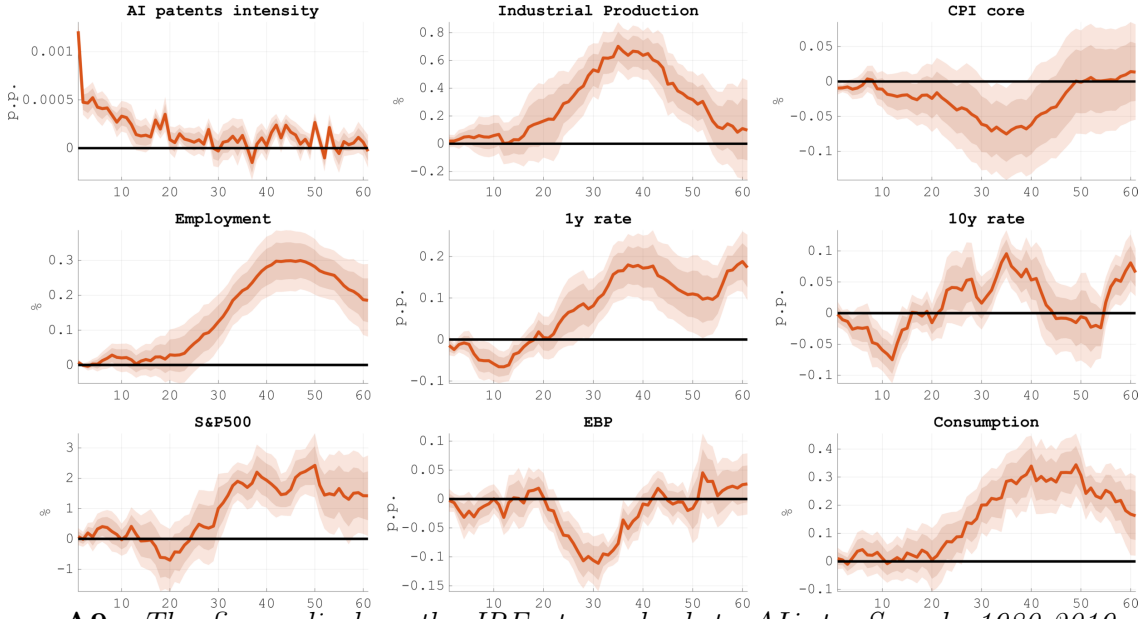


Figure A9: *The figure displays the IRFs to a shock to AIint. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.*

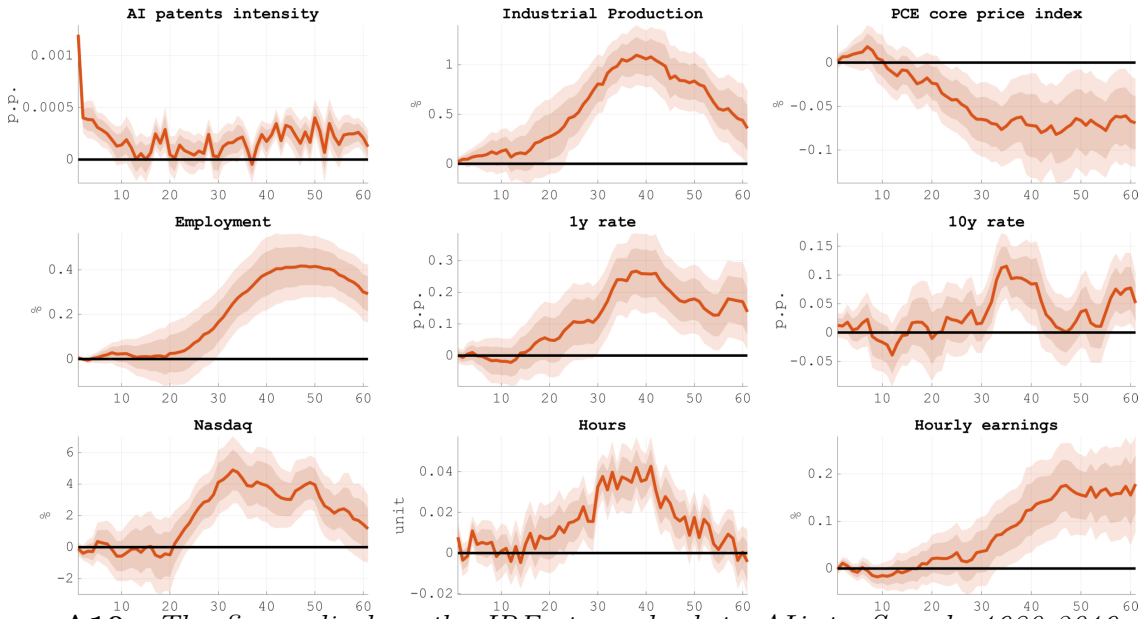


Figure A10: *The figure displays the IRFs to a shock to AIint. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.*

	Robotic patent?	
	0 (no)	1 (yes)
# observations	13,675,265 (99.8%)	30,129 (0.2%)
AI score	0.034 (0.099)	0.035 (0.104)
AI intensive patent	0.115 (0.319)	0.134 (0.341)
AI prediction score from ML model	0.018 (0.114)	0.028 (0.147)
AI prediction score from evolutionary model	0.009 (0.053)	0.010 (0.053)
AI prediction score from NLP model	0.014 (0.094)	0.007 (0.058)
AI prediction score from speech model	0.009 (0.077)	0.007 (0.063)
AI prediction score from vision model	0.036 (0.151)	0.069 (0.210)
AI prediction score from knowledge model	0.068 (0.229)	0.085 (0.256)
AI prediction score from planning/control model	0.075 (0.233)	0.076 (0.228)
AI prediction score from AI hardware model	0.048 (0.161)	0.050 (0.171)

Table A2: *The table presents summary statistics for AI intensity in robotic patents.*

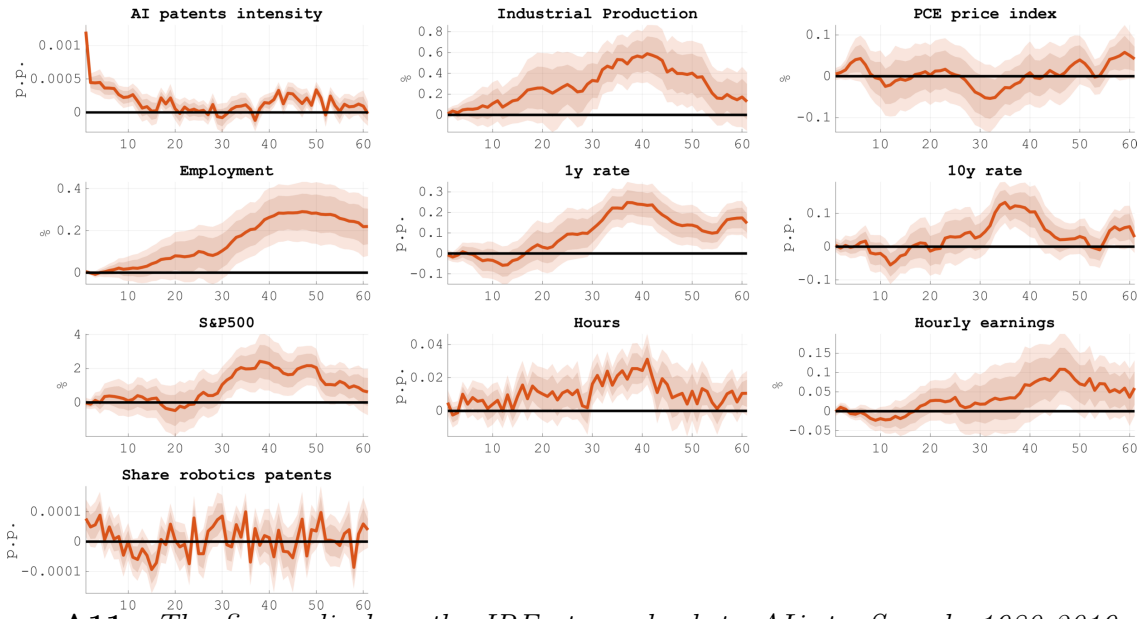


Figure A11: *The figure displays the IRFs to a shock to AIint. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.*

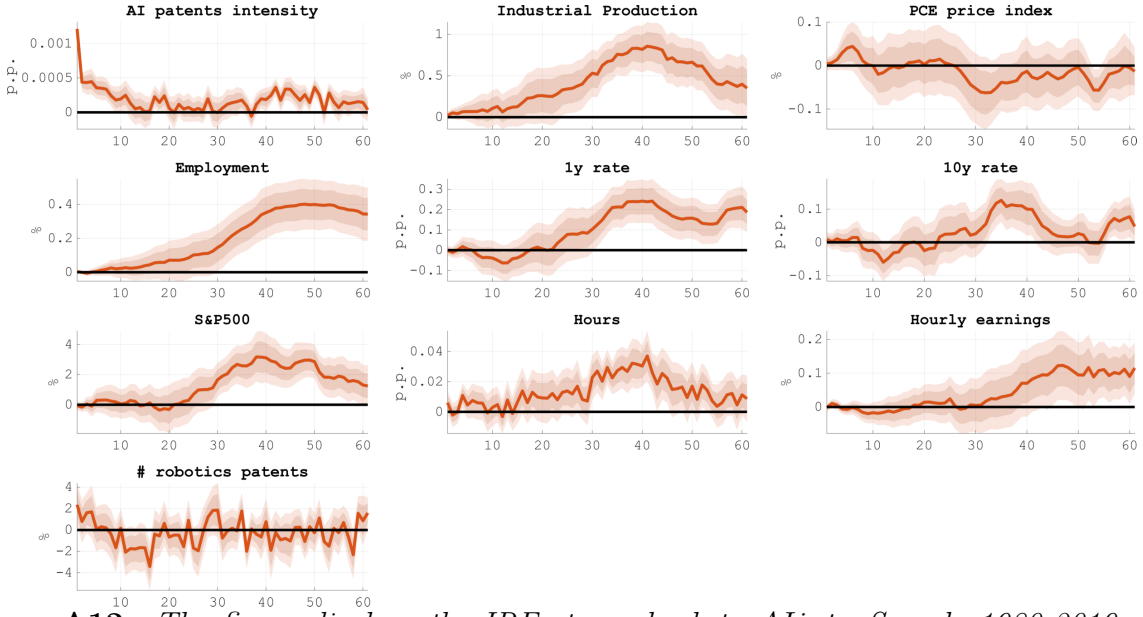


Figure A12: *The figure displays the IRFs to a shock to AIint. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.*

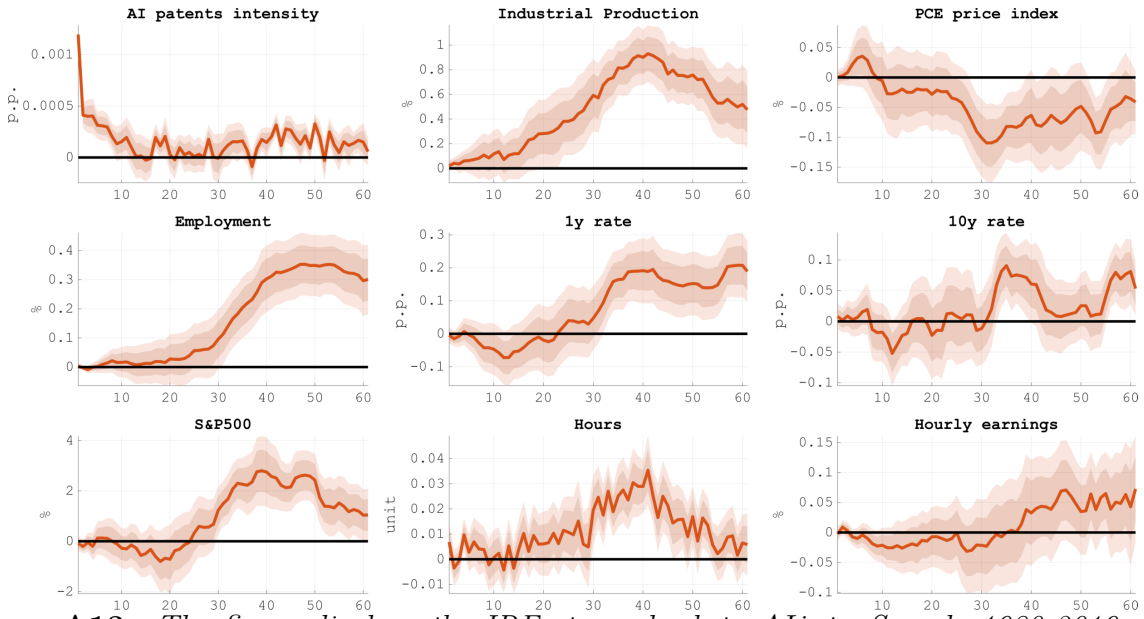


Figure A13: *The figure displays the IRFs to a shock to AIint. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.*

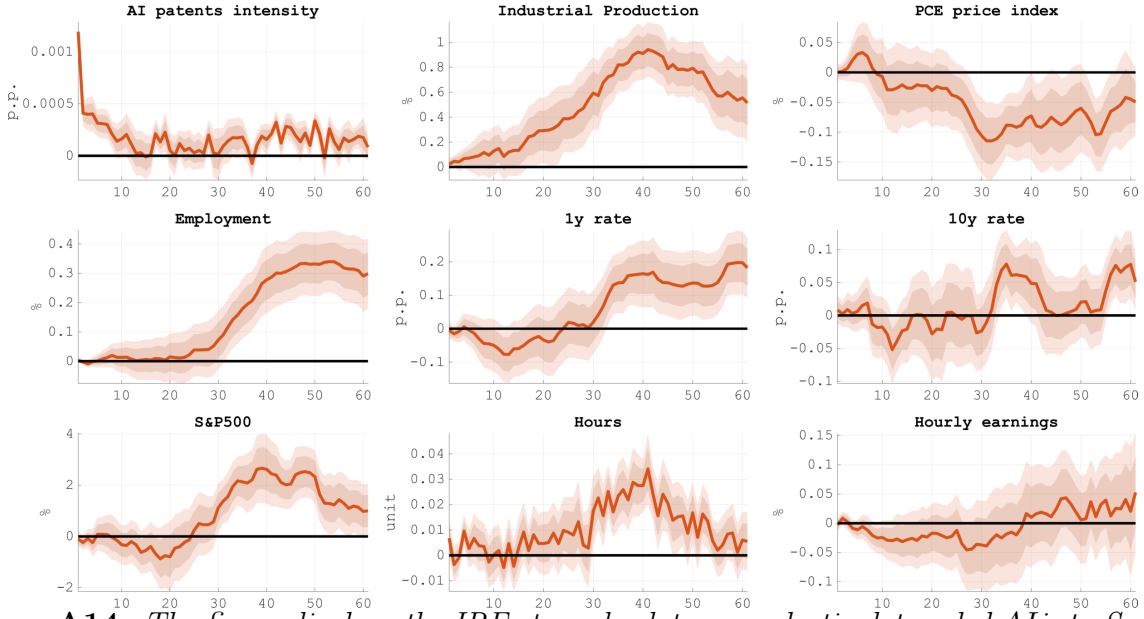


Figure A14: *The figure displays the IRFs to a shock to a quadratic detrended AIint. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.*

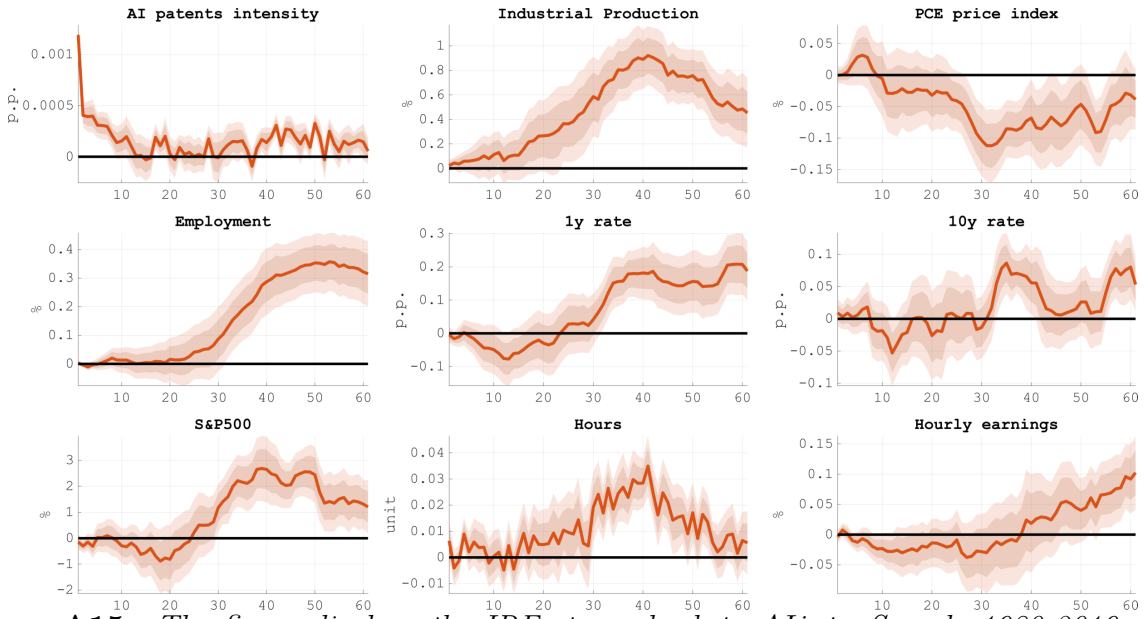


Figure A15: *The figure displays the IRFs to a shock to AIint. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.*

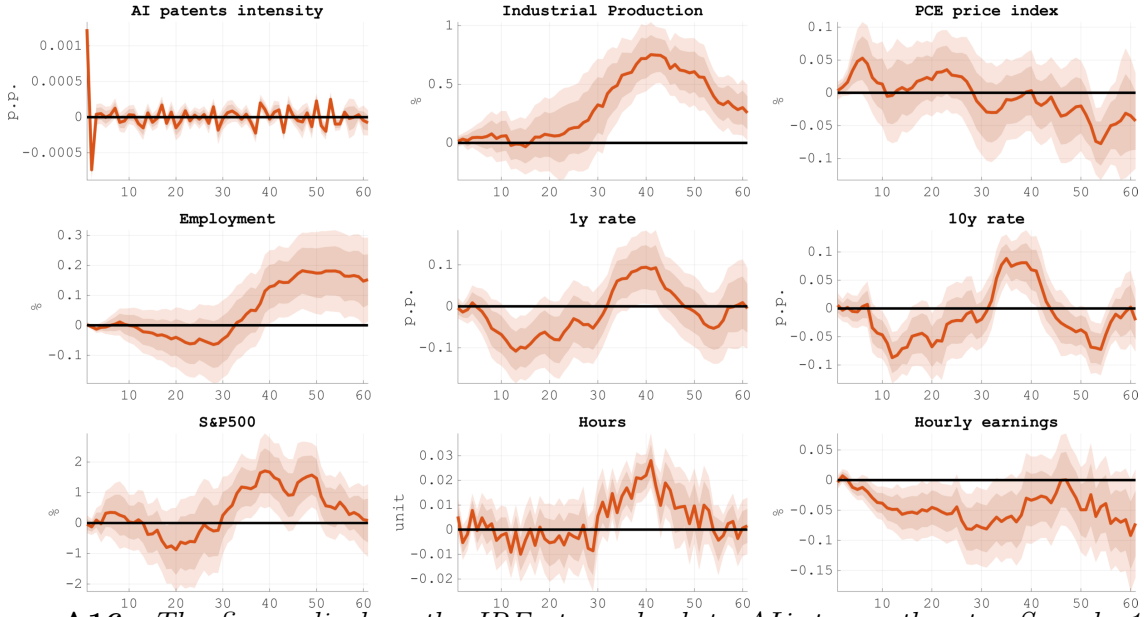


Figure A16: *The figure displays the IRFs to a shock to AIint growth rate. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.*

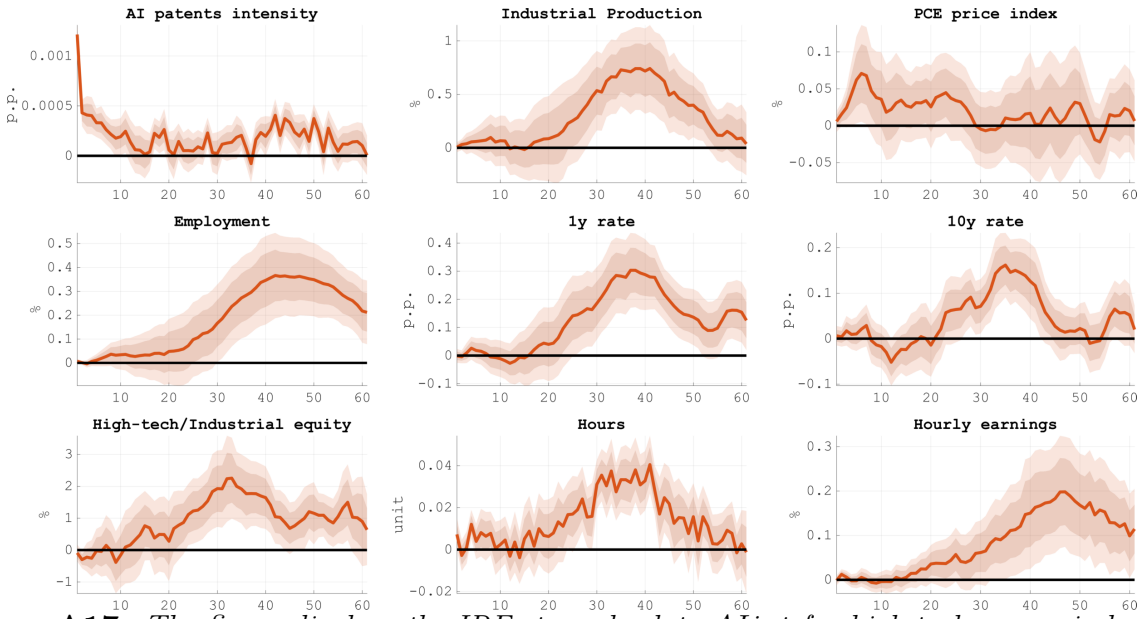


Figure A17: *The figure displays the IRFs to a shock to AIint for high tech versus industrial stocks. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.*

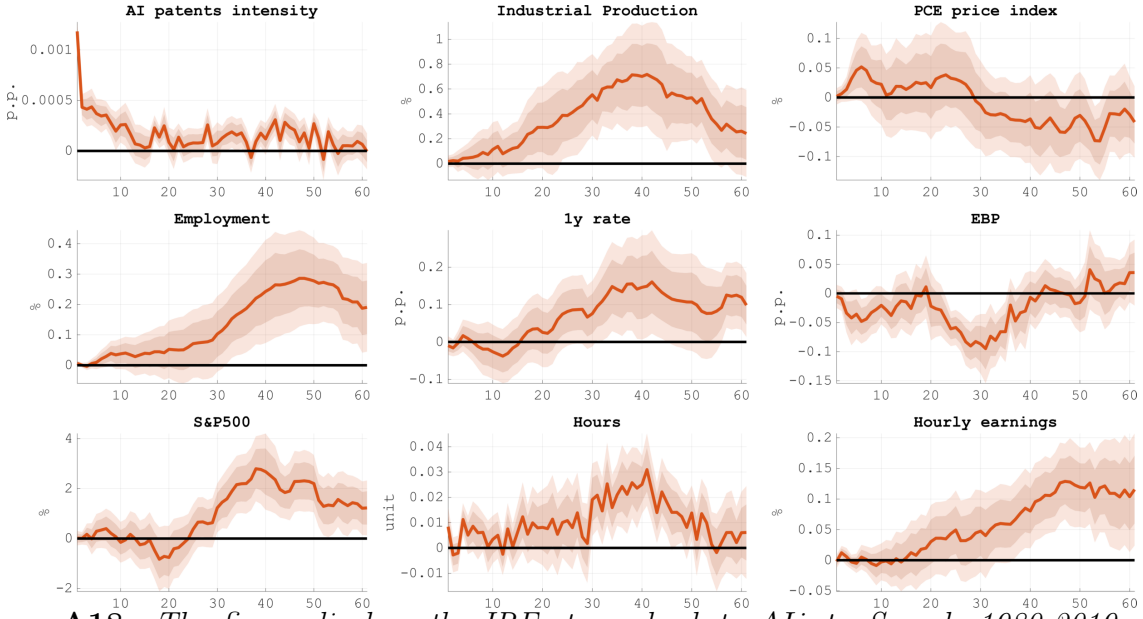


Figure A18: *The figure displays the IRFs to a shock to AIint. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.*

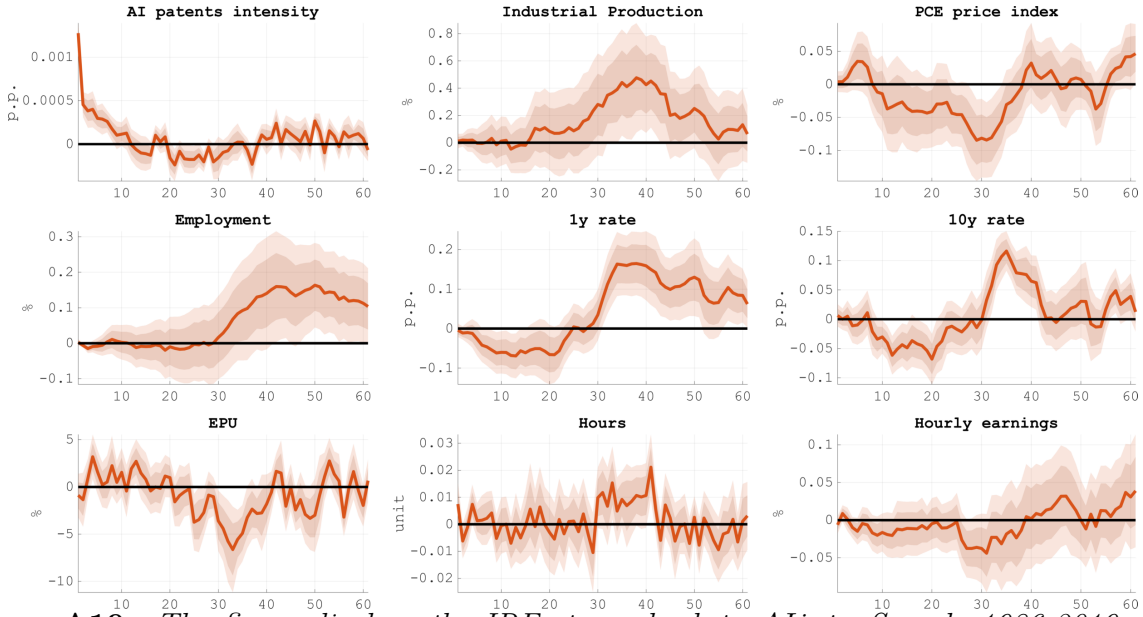


Figure A19: *The figure displays the IRFs to a shock to AIint. Sample 1986-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.*

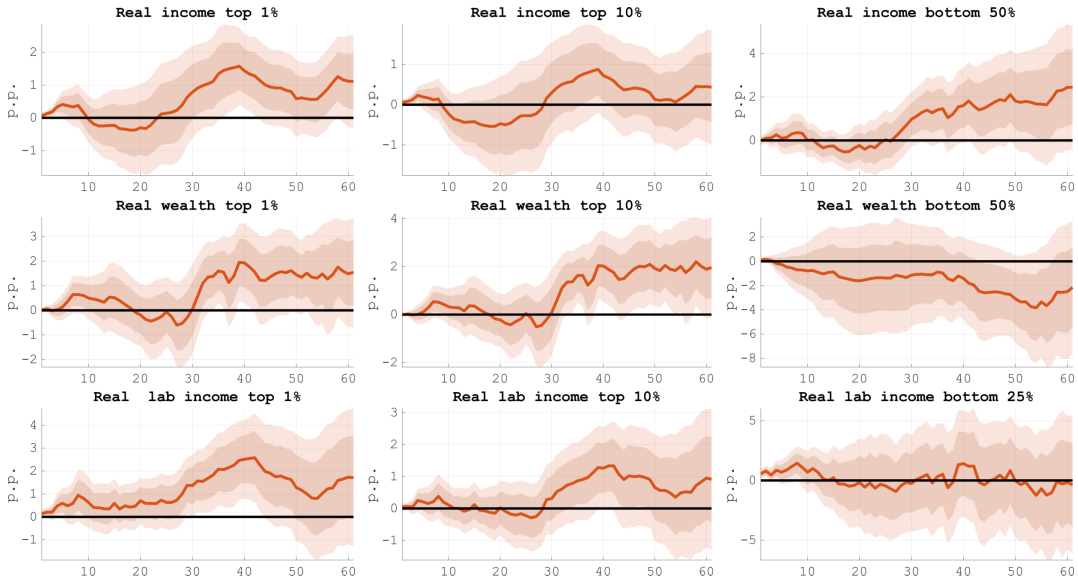


Figure A20: *The figure displays the IRFs to a AIint shock. Sample 1980-2019. The estimates are based on local projections with Newey-West standard errors. Point estimate and 68%-90% confidence bands.*