# The macroeconomic effects of AI innovation

Andrea Gazzani* Filippo Natoli†

PRELIMINARY DRAFT

July 15, 2024

**Abstract**

What are the macroeconomic consequences of AI innovation? We answer this question by constructing a measure of *AI intensity of US innovation* based on a scoring of the AI content of each patent. An increase in the AI intensity of US innovation leads to a delayed surge in industrial production and a slight decline in consumer prices, in line with the transmission of a positive supply shock. Such positive effects descend from a positive, albeit lagged, response of total factor productivity. Our estimates also show that employment, hours, and wages *increase* following a positive shock to AI innovation, even in high AI-exposed sectors, underscoring the crucial role played by general equilibrium effects when studying the aggregate implications of AI technology. The expansionary effect of an AI-driven technological development comes at the cost of an increase in inequality across the income and wealth distributions.

*JEL:* E32, O33, O34, C36

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

*Bank of Italy, Economics and Statistics Department.

†Bank of Italy, Economics and Statistics Department.

Contacts: andreagiovanni.gazzani@bancaditalia.it, filippo.natoli@bancaditalia.it. 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) have argued that for its effectiveness to spur innovation across different economic sectors, such technology might potentially be seen as a GPT (Cockburn et al., 2019; Goldfarb et al., 2023). Such initial 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 pertaining 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] AI-based technology might have repercussions on multiple dimensions such as productivity, economic growth, inflation, labor market and inequality both in the long-run and at the business cycle frequency. As such, these innovations may have profound repercussions also for the conduct of monetary and fiscal policies.[2]

We take up this issue by investigating the macroeconomic implications of AI innovation in the United States. For this purpose, we rely on granular data on US patents to construct an aggregate measure of *AI intensity in innovation*. We construct this measure by averaging, at the monthly frequency, patent-level AI scores – ranging from 0 to 1 – proxying the artificial intelligence content in the underlying technology, which are available for all patents filed at the US patent office since late 1970s. Scores are averaged based on the *filing* month of each patent – when the information content of a new technology is disclosed to the public for the first time, as in Miranda-Agrippino et al. (2020) – and used to proxy *news* of a future AI-driven technology development. In other words, fluctuations in the AI content of innovation constitute an anticipated shock to the AI *intensity* of future technology.

[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.

We estimate the dynamic effects of anticipated AI technology shocks on the US macroeconomy between January 1980 and December 2019 by means of local projections (Jordà, 2005). We find that an increase in the AI intensity of US innovation leads to a delayed surge in industrial production and a slight 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 expansionary effect of the shock induces a shift in the yield curve, with a stronger impact on its short end, suggesting an endogenous contractionary response of the monetary authority.

The overall expansionary macroeconomic effect can be linked to the implications of AI innovation to the labor market, which is the linchpin of its aggregate transmission. Indeed, while an amplification of 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 contemporaneous increase in layoffs. Overall, positive shocks lead to an increase in the number of employed people, in the number of hours worked and in wages (measured by hourly earnings). These results underscore the importance of considering the general equilibrium effects of 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 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, do not vary significantly between highly AI-exposed *vis-à-vis* lowly AI-exposed sectors. Such homogeneity underscores the the power of the overall (general equilibrium) expansionary push of the AI technology in driving the results and, potentially, a certain degree of complementarity (rather than simply substutability) between AI and human work.

Our aggregate and sector-level evidence neglect potential distributional consequences.

Thus, we investigate the heterogeneous impact of our shock on total income, wealth, and labor income for specific percentiles of the respective distributions (data come from Blanchet et al., 2022), showing that an increase in the AI intensity of US innovation leads to a rise in labor income and wealth in the top 10% of the distribution (more than the top 1%) and a fall of the same variables in the bottom 50th percentile. These results inform the literature on the inequality effects of automation and technology more generally (Acemoglu and Restrepo, 2020; Prettner and Strulik, 2020). By driving substitution within the workforce, AI-driven technological advances may have contributed to the secular decline in the demand of low-skilled workers in developed economies highlighted by 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 on the relative importance of either partial equilibrium effects, on which the literature has mostly focused, and 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, explained in a general framework by Wolf (2023).[3] To the best of our knowledge, this work constitutes the first empirical analysis of the aggregate economic implications of AI technology that, as such, takes into account general equilibrium effects. As, according to our findings, AI diffusion acts as an expansionary technology shock, the aggregate productivity gains dominate the potential displacement effects coming from such type of innovation, although AI technology diffusion leads to a transformation of the skills demanded by firms. The findings contained in this paper also have fiscal policy implications: while they highlight the potential of promoting a type of private R&D tilted toward the diffusion and adoption of artificial intelligence in business (that pays off by expanding total factor productivity in the economy), they also underscore the potential problems related to income and wealth inequality.

[3] Results in this literature are mixed. For instance, Bonfiglioli et al. (2023) find a negative displacement effect negative 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 free-lancers and financial analysts, respectively. Abis and Veldkamp (2024) suggest that higher data intensity led a decline in the labor share in the investment management industry. Conversely, Brynjolfsson et al. (2023) highlight instead the complementary of AI with respect to workers by showing that the productivity of customer support agents has surged.

## 2 Data

While AI advances are significantly 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). 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 filed 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 development 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.

We use this granular patent-level database to construct a macroeconomic shock to AI-driven technological progress. Specifically, we build a measure of 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 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; Ferriani et al., 2024). The intuition about the news effect of patenting for future technology is that the filing of a patent discloses information to the public –

granting can come years after filing – on future technology developments, inducing significant macroeconomic effects.

Figure 1 displays the dynamics of our measure of $AIint$ together with an alternative measure ($AIshare$) constructed as the share of the number of "AI patents" (defined by the AIPD as those with score exceeding the 0.5 threshold) over the total number of patents filed in month t.[4] The two series contain very similar information so using one or the other in the empirical analysis does not materially change the results. They both show an upward trend since the beginning of our sample 1980, which was reinforced 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 in each month were AI patents. In our baseline estimation, we employ $AIint$ as its continuous nature is likely to capture more precisely the AI-technology diffusion.

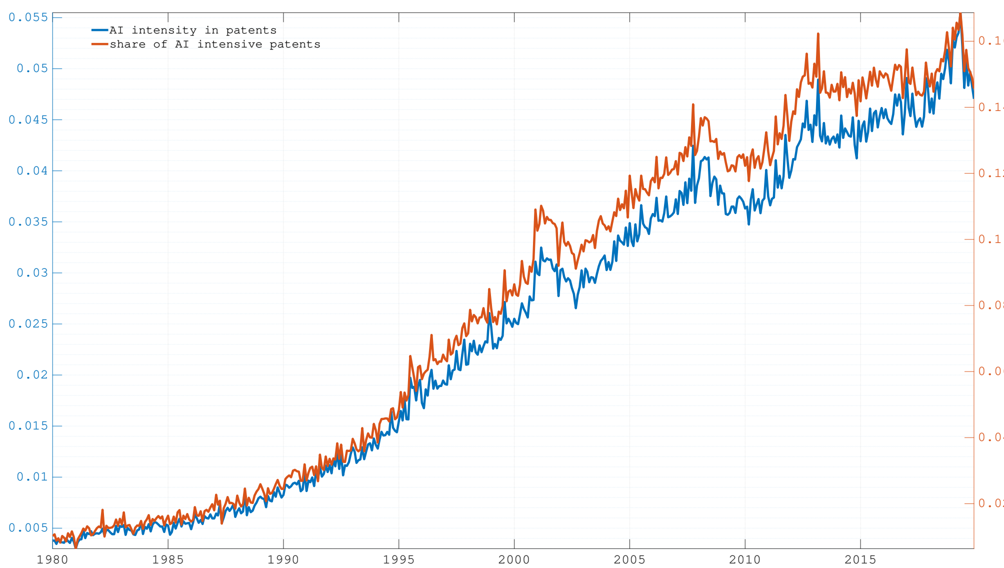


**Figure 1:** AI intensity in US patenting, based on the AIPD dataset.

*Note. The figure displays $AIint$ and $AIshare$ at the monthly frequency between 1980 and 2019.*

[4]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).

# 3 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 across 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 local projections (Li et al., 2024). Our analysis proceeds in four steps. First, we focus on the macroeconomic effects on the US economy. Second, we analyze the effect of AI tech shocks on TFP, highlighting the crucial role played by general equilibrium effects. Third, we delve into the labor market by studying occupational flows and sectoral quantities. Finally, we investigate the consequences of AI tech diffusion for inequality.

## 3.1 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 2 displays the aggregate responses to a $AIint$ shock. The impact on the US economy is overall 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 3.1.1) 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 compared to the extensive margin (the number of patents).[5]

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

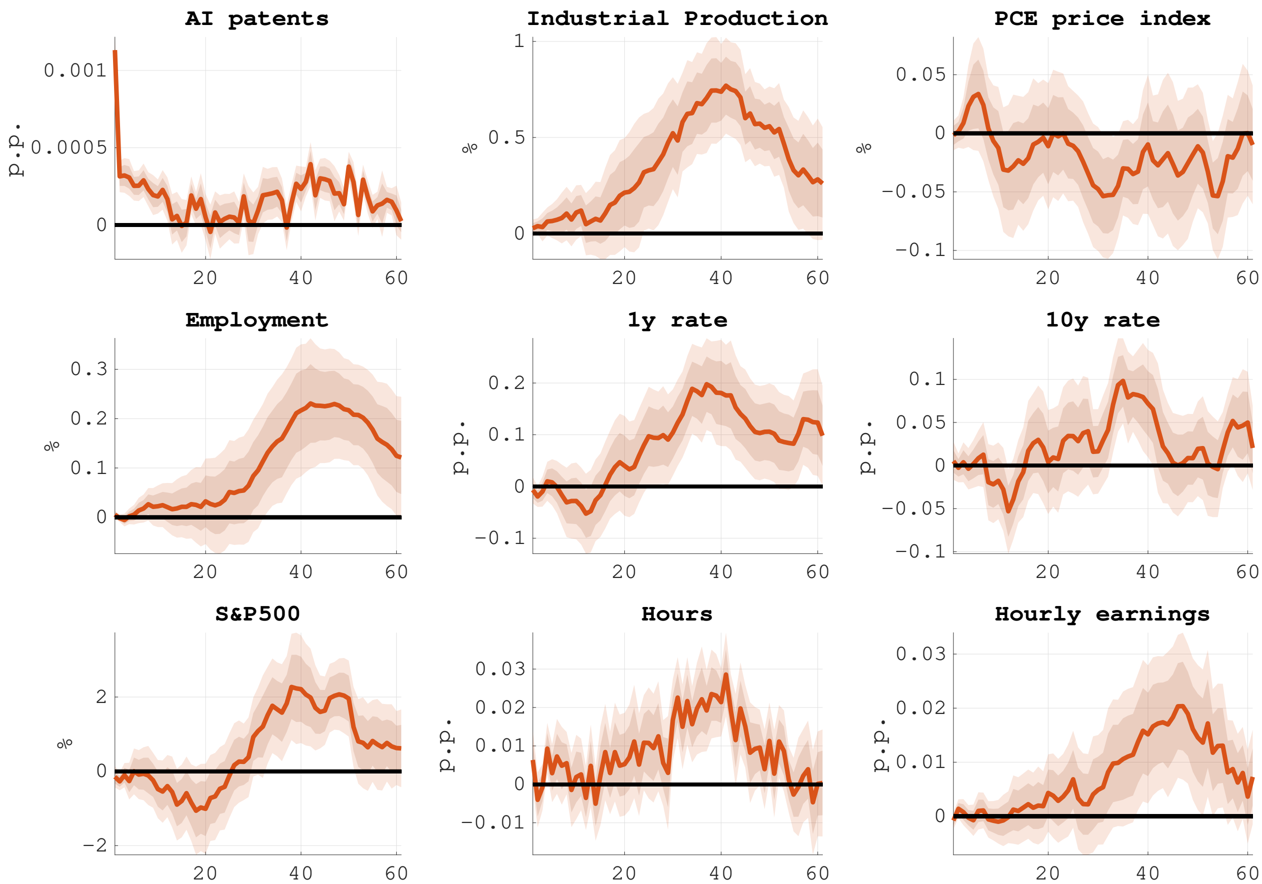


**Figure 2:** IRFs to a AI technology shock.

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

### 3.1.1 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 3 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).

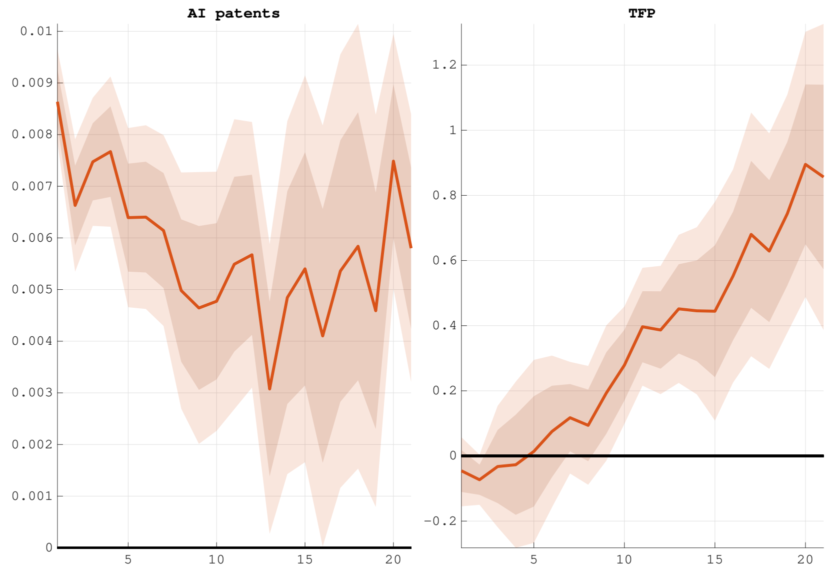


**Figure 3:** TFP response to a AI technology shock.

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

## 3.2 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 4 displays the aggregate responses of occupational flows. Despite the improvement in 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 2. Digging further into the dynamics of outflows, separations increase due to an increase in layoffs, whereas the number of quits falls. This pattern is consistent with firms shifting the labor force composition towards workers that are more complementary with the AI-based progress to exploit the new technological advancements.

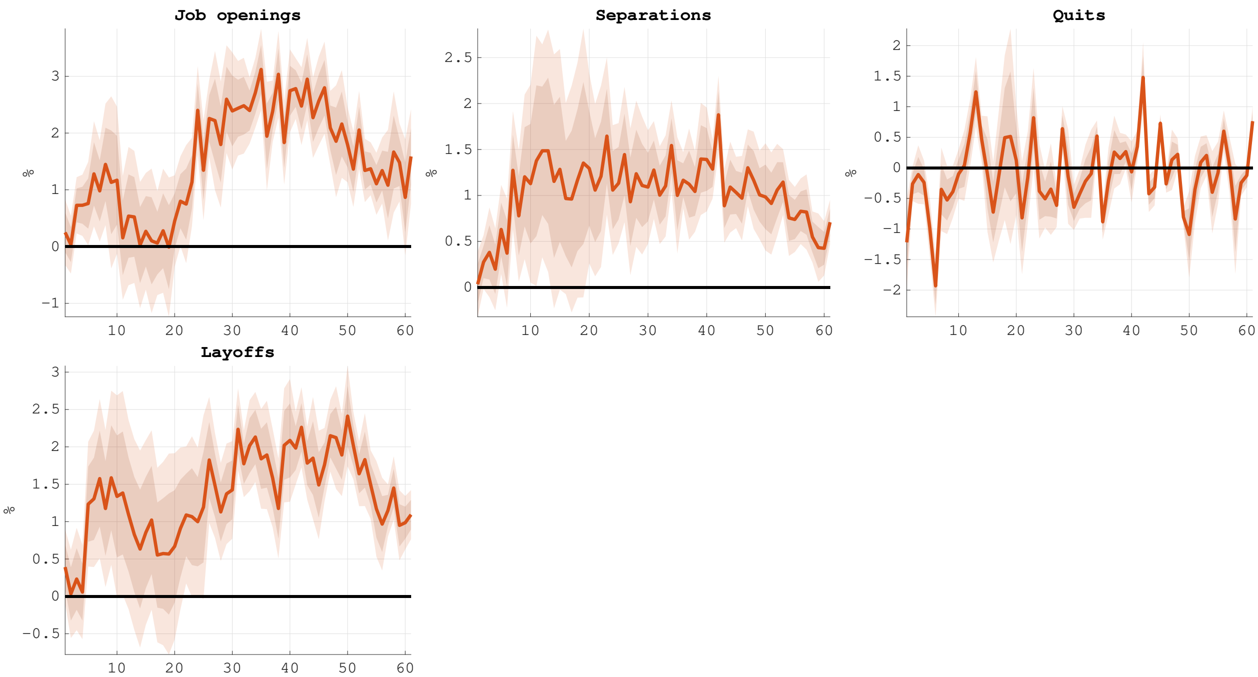


**Figure 4:** Response of labor market flows to a AI technology shock.

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

Next, we investigate the response of some selected labor market variables by sector. Figure 5 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).[6] High AI-exposed sectors are those in which the overlap between occupational abilities and the scope of past AI advances is largest, so where the effects of AI-based tech shocks on employment and wages might, in principle, also be largest. In line with our aggregate findings, sectoral evidence is consistent with the presence of large general equilibrium effects given that we do not observe monotonous effects based on sectoral AI exposure. Wages in sectors less exposed to AI appear to grow more than those in sectors more exposed to AI. This patterns might be consistent with the partial equilibrium displacement effect based on micro-level estimates, coupled with strong positive general equilibrium effects.

[6] 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).

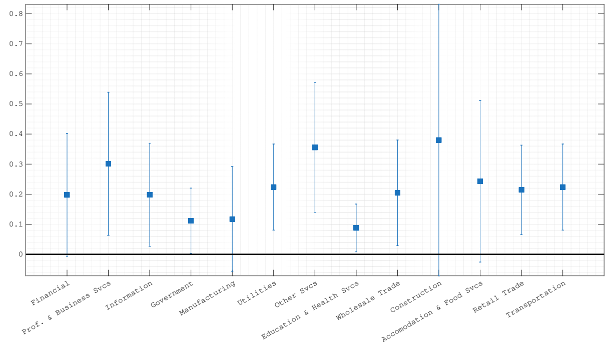


**(a)** Employment

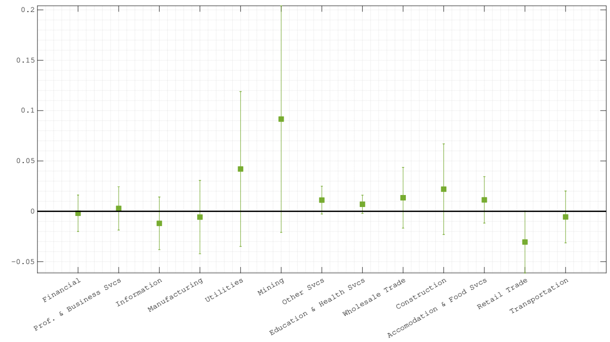


**(b)** Hours

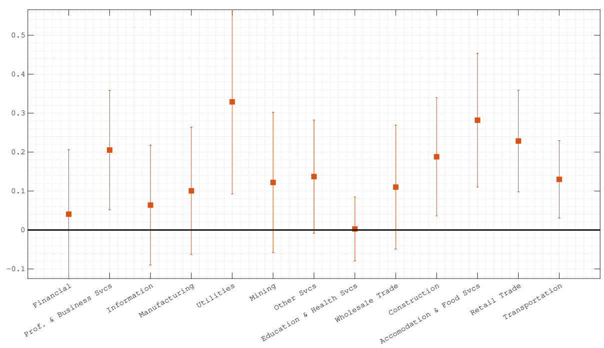


**(c)** Real wage

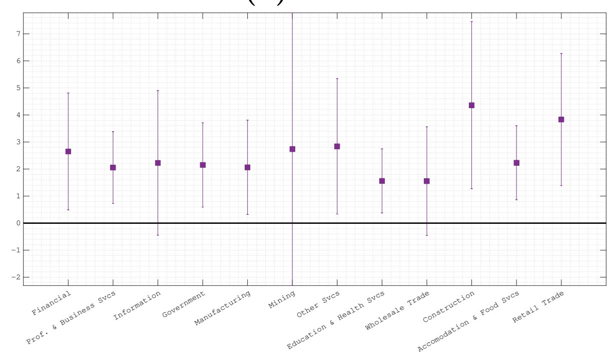


**(d)** Openings

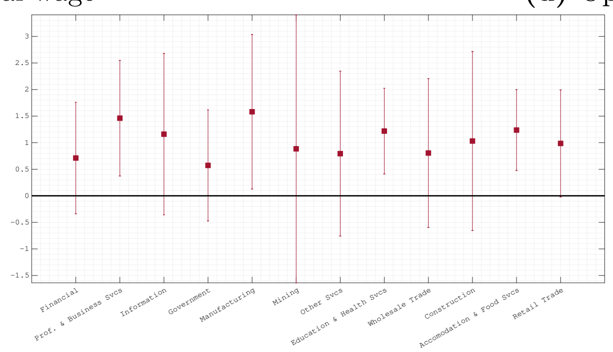


**(e)** Separations

**Figure 5:** Response of labor market flows to a 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.*

## 3.3 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).[7] Figure 6 reports the effect of *AIint* shocks on overall income, wealth, and labor

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

income shares. These results deal exclusively with redistributive consequences while not being informative of the net income or wealth gains for each group. 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 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. 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.[8]

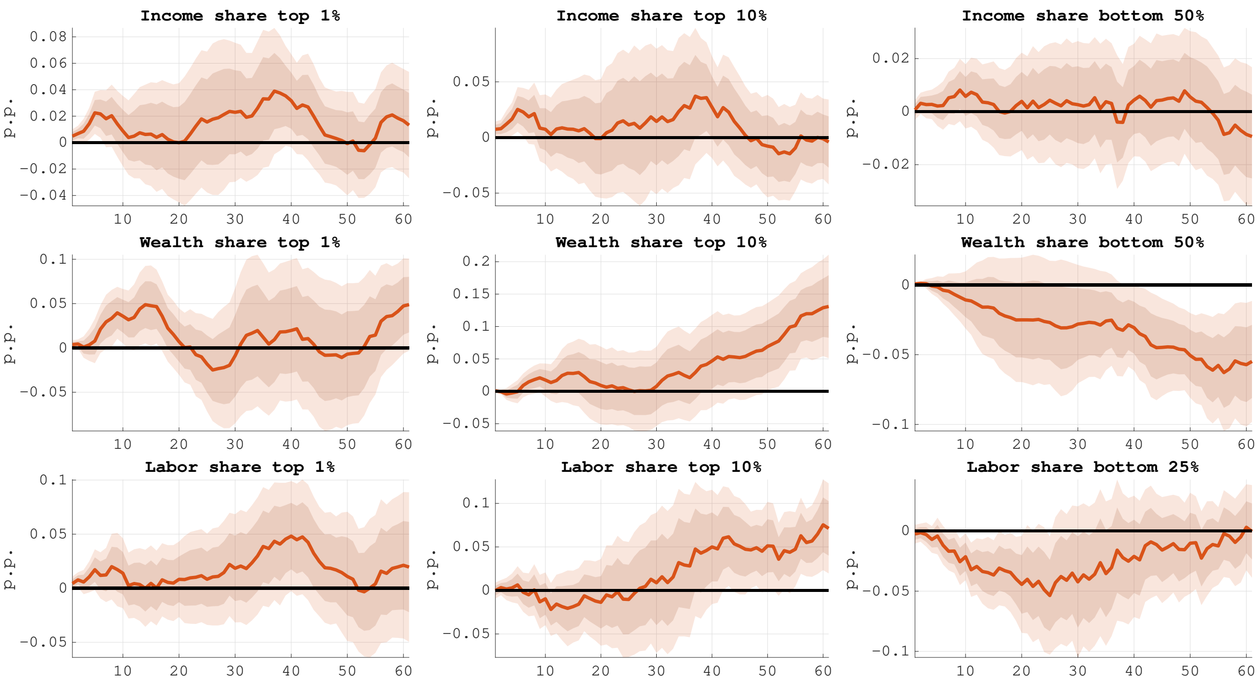


**Figure 6:** RESPONSE OF INEQUALITY TO A AI TECHNOLOGY SHOCK.

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

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

## 4 Conclusions

AI-based innovations such as *ChatGPT* have brought AI to the center of the academic and policy debate. By exploiting data on the AI intensity of US patents between 1980 and 2019 we have investigated the macroeconomic consequences of AI-based innovation. An increase in the AI intensity of US innovation leads to a delayed surge in industrial production and a slight decline in consumer prices, in line with the transmission of a positive supply shock. Such positive effects descend from the positive response of total factor productivity. However, this expansionary effect comes at the cost of an increase in inequality: the more intensive use of AI technology in innovation also leads to increased inequality as the wealth and labor income share of the top 10% workers benefits from AI whereas the bottom 50%-25% see their shares falling. These results might have important implications for the conduct of fiscal policy.

# References


Abis, S. and Veldkamp, L. (2024). The changing economics of knowledge production. *The Review of Financial Studies*, 37(1):89–118.

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

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

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

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

Berman, E., Bound, J., and Machin, S. (1998). Implications of Skill-Biased Technological Change: International Evidence*. *The Quarterly Journal of Economics*, 113(4):1245–1279.

Blanchet, T., Saez, E., and Zucman, G. (2022). Real-time inequality. Technical report, National Bureau of Economic Research.

Bonfiglioli, A., Gancia, G., Papadakis, I., and Crinò, R. (2023). Artificial intelligence and jobs: Evidence from us commuting zones. Technical report, CESifo Working Paper.

Bresnahan, T. F. and Trajtenberg, M. (1995). General purpose technologies 'engines of growth'? *Journal of Econometrics*, 65(1):83–108.

Brynjolfsson, E., Li, D., and Raymond, L. R. (2023). Generative ai at work. Technical report, National Bureau of Economic Research.

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

Chahrour, R., Chugh, S. K., and Potter, T. (2023). Anticipated productivity and the labor market. *Quantitative Economics*, 14(3):897–934.

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

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

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

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

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (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., Raj, M., and Seamans, R. (2018). A method to link advances in artificial intelligence to occupational abilities. *AEA Papers and Proceedings*, 108:54–57.

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

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

Goldfarb, A., Taska, B., and Teodoridis, F. (2023). Could machine learning be a general purpose technology? a comparison of emerging technologies using data from online job postings. *Research Policy*, 52(1):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.

Grennan, J. and Michaely, R. (2020). Artificial intelligence and high-skilled work: Evidence from analysts. *Swiss Finance Institute Research Paper*, (20-84).

Grimm, F. and Gathmann, C. (2022). The Diffusion of Digital Technologies and its Consequences in the Labor Market. Technical report.

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

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American economic review*, 95(1):161–182.

Li, D., Plagborg-Møller, M., and Wolf, C. K. (2024). Local projections vs. vars: Lessons from thousands of dgps. *Journal of Econometrics*, page 105722.

Lu, Y. and Zhou, Y. (2021). A review on the economics of artificial intelligence. *Journal of Economic Surveys*, 35(4):1045–1072.

Miranda-Agrippino, S., Hacioglu Hoke, S., and Bluwstein, K. (2020). Patents, News, and Business Cycles. CEPR Discussion Papers 15062, C.E.P.R. Discussion Papers.

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

Portier, F. (2014). Comment on" whither news shocks?". Technical report, National Bureau of Economic Research.

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

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(8):2232–2269.

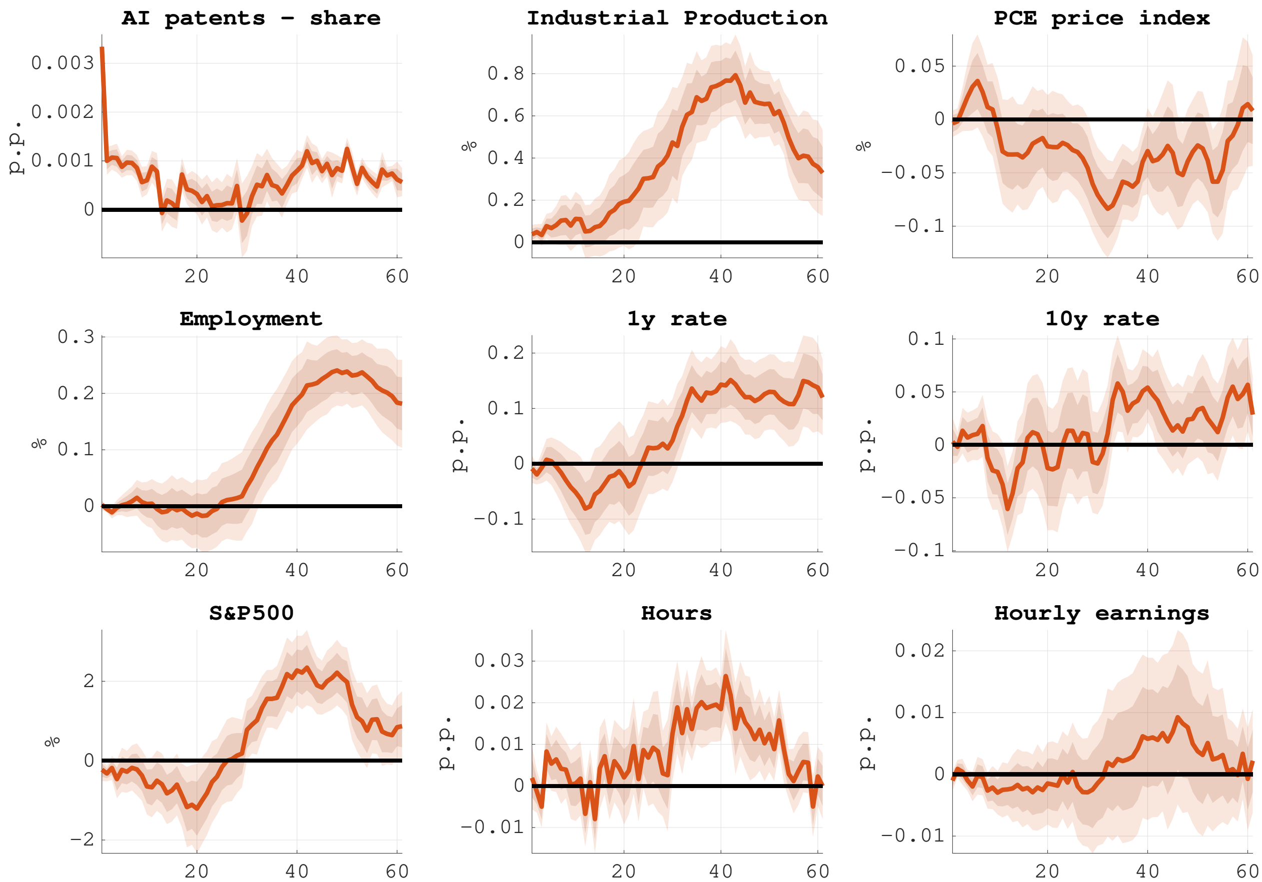


**Figure A1:** IRFs to a AI technology shock.

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