

# The impact of artificial intelligence on output and inflation

I Aldasoro  
BIS

S Doerr  
BIS & CEPR

L Gambacorta  
BIS & CEPR

D Rees  
BIS

July 12, 2024

## Abstract

This paper studies the effects of artificial intelligence (AI) on sectoral and aggregate employment, output and inflation in both the short and long run. We construct an index of industry exposure to AI to calibrate a macroeconomic multi-sector model. Building on studies that find significant increases in workers' output from AI, we model AI as a permanent increase in productivity that differs by sector. We find that AI significantly raises output, consumption and investment in the short and long run. The inflation response depends crucially on households' and firms' anticipation of the impact of AI. If they do not anticipate higher future productivity, AI adoption is initially disinflationary. Over time, general equilibrium forces lead to moderate inflation through demand effects. In contrast, when households and firms anticipate higher future productivity, inflation rises immediately. Inspecting individual sectors and performing counterfactual exercises we find that a sector's initial exposure to AI has little correlation with its long-term increase in output. However, output grows by twice as much for the same increase in aggregate productivity when AI affects sectors producing consumption rather than investment goods, thanks to second round effects through sectoral linkages. We discuss how public policy should foster AI adoption and implications for central banks.

*JEL Codes:* E31, J24, O33, O40.

*Keywords:* artificial intelligence, inflation, output, productivity, monetary policy.

---

We thank seminar participants at the Bank for International Settlements (BIS), Bank of Korea, the Virtual Israel Macro Meeting and the International Monetary Fund for helpful comments and suggestions. Contact: Aldasoro (*corresponding author:* [inaki.aldasoro@bis.org](mailto:inaki.aldasoro@bis.org), Centralbahnplatz 2, 4051, Basel, Switzerland), Doerr ([sebastian.doerr@bis.org](mailto:sebastian.doerr@bis.org)), Gambacorta ([leonardo.gambacorta@bis.org](mailto:leonardo.gambacorta@bis.org)), and Rees ([daniel.rees@bis.org](mailto:daniel.rees@bis.org)). The views expressed here are those of the authors only and not necessarily those of the BIS. Declarations of interest: none.

# 1 Introduction

Recent advances in artificial intelligence (AI) have raised hopes of a boost to economic growth. Many scholars believe that AI has the potential to be “the most important general-purpose technology of our era” (Brynjolfsson et al., 2023). The recent inroads of generative AI in everyday applications in particular promise widespread efficiency gains.<sup>1</sup> Unlike automation through robots, which can accomplish only explicitly understood (i.e., ‘routine’) tasks, AI can infer tacit relationships that are not fully specified by underlying software (Autor, 2022). By transforming occupational tasks, altering corporate strategies, and affecting production efficiency, AI may have significant consequences for labour markets, firms, and whole industries (Agrawal et al., 2019).

A key channel through which AI affects economic growth is through improvements in productivity (Acemoglu and Restrepo, 2018; Aghion et al., 2018). Micro-economic studies find that generative AI can make workers tremendously more productive, especially in occupations that require cognitive abilities (Brynjolfsson et al., 2023; Noy and Zhang, 2023). AI also boosts firm growth and innovation (Babina et al., 2024). At the macro-economic level, analyses suggest that AI could raise annual productivity growth by around 1 percentage point (pp) per annum over the next decade (Baily et al., 2023; Briggs and Kodnani, 2023). The adoption of AI can hence be thought of as an increase in productivity that expands an economy’s output capacity. Compared to information technology (IT), whose impact took years to be reflected in aggregate productivity numbers (Fernald and Wang, 2015), AI is considerably easier to use and implement in processes as it is a general-purpose technology that does not require the deployment of new hardware, deep user know-how, or a substantial reconfiguration of business practices. As a consequence, the impact of AI on productivity will likely be felt in the coming years already (Brynjolfsson et al., 2018; Furman and Seamans, 2019).<sup>2</sup>

In this paper, we investigate the effects of AI on aggregate output and inflation, as

---

<sup>1</sup>Generative AI generally refers to algorithms that can be used to create new content such as text or images, based on patterns detected in large training datasets. The most well-known application is ChatGPT, but there are numerous others.

<sup>2</sup>Think of the advent of IT: firms needed to replace their paper-based systems with computers, familiarise themselves with the new concept of “software”, and train their staff. Meanwhile, public infrastructure, for example in the form of broadband, took years – if not decades – to provide sufficient coverage. AI, on the other hand, can be used with the now near-ubiquitous smartphones and computers. Rather than requiring workers to learn how to use a fundamentally new system (think of an Excel spreadsheet versus a rolodex), AI can be used through an intuitive language-based interface.

well as on output and employment in different sectors. We do so by first constructing a measure of exposure to AI at the industry level. We then embed this exposure measure into a macroeconomic multi-sector model, calibrated to the US economy using input-output tables. We also use the model to perform counterfactual exercises.

We start by constructing an industry-level measure of exposure to AI (AIIE) at the 2-digit NAICS level. Building on the indicator developed in [Felten et al. \(2021\)](#), the measure combines survey evidence on the extent to which AI applications can be used in different workplace abilities with information on the importance of various abilities in different occupations and industries. A higher AIIE score indicates that an industry has many occupations in which AI applications can be used. The industry with the highest AIIE score is finance and insurance, followed by management of companies and enterprises. Those with the lowest score are agriculture, forestry, fishing and hunting as well as transportation and warehousing. Note that our measure does not capture whether AI complements or substitutes for specific occupations, which requires strong assumptions to categorise each specific task ([Pizzinelli et al., 2023](#)).

We then calibrate a macroeconomic model in which AI acts as permanent increase in the level of productivity with a differential impact across sectors. In particular, we assume that AI raises annual productivity growth by 1.5 percentage points for the next decade, in line with plausible estimates in the literature ([Baily et al., 2023](#); [Briggs and Kodnani, 2023](#)).<sup>3</sup> We allocate the aggregate productivity increase across sectors using the AIIE measure. In addition to the usual set of nominal and real rigidities commonly used to account for aggregate economic fluctuations, the model features a detailed industry structure on both the demand and production sides, following the work of [Rees \(2020\)](#). This allows it to account for industry-specific shifts in consumer preferences and work practices, as well as industry-level and aggregate transmission mechanisms of AI. We focus on 20 sectors in our baseline model, which broadly encompass consumption goods sectors closer to final demand, intermediate goods and investment goods.<sup>4</sup>

---

<sup>3</sup>Other assumptions are possible, and our model provides the flexibility to accommodate them. For example, [Acemoglu \(2024\)](#) suggests much more modest estimates of the impact on productivity, building on a conceptual framework featuring a static economy with a unique final good and perfectly competitive markets. As argued in the rejoinder by [Briggs \(2024\)](#), differences in estimates likely arise from the assumption of a lower share of tasks exposed to AI (plus the assumption that only about of fifth of those is cost effective to automate within the next decade) and the assumption that AI automation will not lead to labour reallocation nor the creation of new tasks. Be that as it may, our framework allows for different assumptions about the aggregate impact of AI on productivity.

<sup>4</sup>Whenever convenient to simplify the exposition, we aggregate the 20 sectors into five: primary industries, secondary industries, distribution, professional services, and other services.

Our analysis considers two scenarios. In the first, households and firms observe the impact of AI on productivity at each point in time and form expectations about the future path of consumption, inflation and output based on those developments. However, they do not foresee the boost to productivity from future AI developments. We refer to this as the “unanticipated” case. In the second scenario, households and firms not only observe the effects of AI on productivity that have already occurred, but also anticipate future ones. Accordingly, they adapt more quickly to the changes implied by the AI-induced productivity increase than in the “unanticipated” case. We refer to this as the “anticipated” case. We see these two cases as the extreme scenarios. Actual expectations are likely to fall somewhere in between.<sup>5</sup>

We start with the unanticipated case and investigate the effects of AI on macroeconomic aggregates. Our results show that AI significantly raises output, consumption and investment in the short and long run, reflecting the positive effect of higher productivity growth on economic capacity. For inflation, initially the supply expansion acts as a disinflationary force, as higher TFP increases the economy’s output capacity. Yet after around four years, general equilibrium effects due to rising consumption and investment raise aggregate demand – and hence wages – sufficiently to make AI’s impact inflationary. Responding to inflation, the policy rate first declines but then increases above its initial level to counteract the demand-driven rise in inflation.

We then simulate the model under the anticipated scenario. The long-run response of output, consumption and investment is identical to the unanticipated scenario. However, as households fully anticipate the effects of AI on productivity today and in the future, they increase consumption more forcefully right away and postpone investment. Accordingly, output increases more slowly in this scenario.

The paths of inflation, interest rates and, to a lesser extent, the output gap, are starkly different in this scenario. Because demand increases in anticipation of *future* productivity increases, AI adoption is initially inflationary. Inflation only begins to converge towards the unanticipated scenario as investment becomes positive and the productive capacity of the economy gets replenished. Policy rates, through the policy rules embedded in the

---

<sup>5</sup>It is not obvious which of the two extremes is more likely. On the one hand, results from vector-autoregression models and the history of past general purpose technologies suggest that technology shocks initially have disinflationary effects (Evans and Marshall, 2009), which would be consistent with imperfect anticipation. At the same time, record stock market valuations of companies producing AI or the necessary hardware suggest that at least financial markets anticipate AI to substantially raise growth.

model, rise immediately as inflation increases and decline only after around 10 years.

Next, we inspect individual sectors to understand their relative importance in shaping aggregate dynamics.<sup>6</sup> The first insight we gain is that a sector's initial exposure to AI has little relationship to its ultimate increase in value added output. The reason is that, ultimately, general equilibrium effects arising from higher demand for a sector's output matter much more than the initial increase in productivity as calibrated from our AIIE measure. Moreover, our results show that the ultimate increase in value added output is smaller in more labour-intensive sectors. An increase in productivity implies that more output needs to be produced with a fixed amount of labour, which raises real wages. Sectors in which labour is a more important factor of production hence face higher costs and increase their production by less.

However, which sector is initially most affected by AI matters greatly for the response in aggregate output and inflation. In our baseline calibration based on the AIIE measure, there is limited cross-sectoral variation in exposure to AI. While AIIE provides a solid grounding for studying the impact of AI across sectors, there remains a high degree of uncertainty regarding the effects of AI. Our model allows us to contrast what would happen if AI would, for example, mostly affect the services sector or the manufacturing sector. The counterfactual exercises show that, while output always increases, it can grow by up to twice as much for the same AI-induced increase in productivity when AI affects sectors producing consumption goods rather than investment goods.<sup>7</sup> When AI increases productivity in consumption goods sectors, freed-up labor moves to the investment goods sector, raising production there, too. Since investment goods are used in the production of intermediate and consumption goods, second round effects through sectoral linkages lead to an additional boost in output as the capital stock expands. Because of these second-round effects, the initial decline in inflation is much more pronounced when AI affects consumption-good sectors in the unanticipated case. In contrast, when AI raises productivity only in the investment goods sector, the aggregate responses in output and inflation are weaker.

A second counterfactual exercise illustrates that results are qualitatively unchanged

---

<sup>6</sup>We focus on the unanticipated case for ease of exposition, but the main insights go through more generally as we zoom in on long run effects across industries, which remain largely unchanged across scenarios.

<sup>7</sup>For ease of exposition, our counterfactual exercises feature the productivity shock from AI as a one-off event. Accordingly, the distinction between unanticipated and anticipated is immaterial.

if AI is a factor-specific technology rather than a general purpose technology. Focusing on the long run impact on output and inflation, we explore the consequences of making AI either a labour-augmenting or capital-augmenting technology. The long run effects on both output and inflation are slightly more subdued when AI is factor-specific, especially if it is capital-augmenting. But our baseline result that output and inflation increase in the long run remains qualitatively unchanged.

Our findings inform the debate on the impact of AI on labor markets and output and have implications for public policy. In particular, they suggest that public policy that fosters the adoption of AI could lead to a Goldilocks scenario: in the case of imperfect anticipation, greater use of AI could ease inflationary pressures in the near-term, thereby supporting central banks in their task to bring inflation back to target. In the longer term, as inflation rises because of greater AI-induced demand, central banks' task of controlling inflation would become simpler, as they can dampen demand via monetary tightening. More generally, AI's positive contribution to growth could also offset some of the detrimental secular developments that threaten to depress demand going forward, such as population aging, re-shoring and changes in global supply chains, as well as geopolitical tensions and political fragmentation. These aspects underscore the need for policies that foster the adoption of AI by firms and households. Policy efforts to spur AI adoption should focus on sectors that produce consumption goods, as these promise especially high returns.

Our main contribution is to provide a framework to model the effect of AI on aggregate output and inflation in both the short and long run. We thereby relate to recent studies that investigate the effect of AI on productivity and employment.<sup>8</sup> [Brynjolfsson et al. \(2023\)](#) study the effects of AI on 5,000 customer-support agents working for a large enterprise-software company. The agents were provided with an AI tool, built on the large language models developed by OpenAI, with a staggered timing. Support agents who used the AI tool could handle 13.8% more customer inquiries per hour and work quality, measured by the share of successfully resolved customer problems, improved by 1.3%. [Noy and Zhang \(2023\)](#) asked experienced business professionals from a variety of fields, including marketers, grant writers, data analysts, and human-resource professionals to write two business documents within their field. For the second document, half of the participants were randomly assigned to use ChatGPT. Professionals using ChatGPT were

---

<sup>8</sup>This literature itself builds on previous work on robot adoption and automation – see for example [Faia et al. \(2022\)](#), [Faia et al. \(2023\)](#) and references therein.

almost 60% faster at writing and the rated quality was also higher. Finally, [Peng et al. \(2023\)](#) compare programmers using the GitHub Copilot AI tool to those who did not use AI. They find that those using AI completed programming tasks more than twice as fast. Importantly, less-experienced programmers benefited the most. Early evidence also suggest a positive correlation between AI adoption and firm productivity ([Yang, 2022](#); [Czarnitzki et al., 2023](#)) and innovation ([Babina et al., 2024](#)).

Other papers discuss potential effects of AI on employment. [Acemoglu et al. \(2022\)](#) use establishment-level data on online vacancies in the United States: they find rapid growth in AI-related vacancies over 2010–18 that is driven by establishments whose workers engage in tasks compatible with AI’s capabilities. AI adoption by establishments leads to reduced hiring in non-AI positions and a change in the skill requirements of remaining postings. Yet, there is (so far) no significant effect of AI on aggregate employment and wage growth. [Felten et al. \(2019\)](#) provide evidence that, on average, occupations impacted by AI experience a small but positive change in wages, but no change in employment. [Autor \(2022\)](#) provides an overview of the labour market implications of technological change, with a focus on artificial intelligence. [Lu \(2021\)](#) develops a three-sector endogenous growth model to quantify the implications of AI for growth and welfare.

To the best of our knowledge, our paper is the first to assesses the effect of AI on output and inflation in the short and long run. We do so by constructing industry-level measures of exposure to AI and embedding these into a rich macroeconomic model. Our paper also provides novel insights on how to assess the direct and second-round effect of AI on output, employment and hours worked in different sectors over time.

## 2 Measuring the impact of AI on productivity across occupations and industries

While AI, and in particular generative AI, is a general purpose technology, its impact differs across occupations and industries. Unlike automation through robots, which has predominantly affected jobs and output in occupations that require manual labour, AI is expected to have the largest impact in occupations with more cognitively demanding tasks. The reason is that much like other computer-based technologies, it can substitute for routine cognitive tasks while complementing nonroutine cognitive tasks ([Autor et](#)



al., 2003; Faia et al., 2022). This is best illustrated by contrasting surgeons and meat slaughterers. These two occupations require similar physical abilities (eg dexterity and steadiness), but their cognitive content differs, with various forms of problem solving and logical reasoning being more relevant for surgeons (Felten et al., 2021).

AI could affect productivity through several channels, but two stand out.

The first channel is to directly raise the productivity of (cognitive) workers. For example, Brynjolfsson and McAfee (2017) show that access to a generative AI-based conversational assistant improves customer support agents' productivity by 14%. For college-educated professionals, Noy and Zhang (2023) show that the chatbot ChatGPT substantially raised productivity in solving writing tasks, reducing the time required by 40% and raising output quality by 18%. Meanwhile, Peng et al. (2023) find that software developers that use AI could code more than twice as many projects per week.

The second channel is to spur innovation and thereby future productivity growth (Brynjolfsson et al., 2018; Baily et al., 2023). Most innovation, for example through research and development but also through managerial activities, is generated in occupations that require high cognitive abilities. Improving the efficiency of cognitive work hence holds large potential to generate further innovation that in turn improves efficiency even further.

The macro-economic impact of AI on productivity growth could be sizeable. Assessing this impact requires aggregating industry-specific productivity gains, which can be done by multiplying the size of the productivity increase with the relative size of the sector.<sup>9</sup> Improvements in productivity in an industry can hence have large aggregate effects. Different studies provide estimates for AI's impact on annual labour productivity growth (ie output per employee) over the next decade, with estimates ranging from 1pp to 1.5pp (Baily et al., 2023; Briggs and Kodnani, 2023).

To assess the impact of AI on productivity in different industries, we first construct an industry-level measure of exposure to AI (AIIE) at the 2-digit NAICS level.<sup>10</sup> We build on Felten et al. (2021), who provide an index called AI Occupational Exposure (AIOE)

---

<sup>9</sup>This is well-established in the economics literature in the seminal work of Hulten (1978). See Baqaee and Fahri (2019) for a recent extension to Hulten's theorem.

<sup>10</sup>NAICS refers to the North American Industry Classification System, the standard used by federal statistical agencies in classifying business establishments for the purpose of collecting, analysing, and publishing statistical data. See the dedicated webpage by the US Census Bureau for more details.



that is widely used in the literature (see eg [Acemoglu et al. \(2022\)](#) and [Autor \(2022\)](#)) and constructed as described below. We then attribute the estimated impact of AI on aggregate productivity to each industry, depending on its exposure to AI and its relative size in the economy.

We construct the AIIE measure in four steps. In a first step, ten AI applications covering AI’s most likely use cases are linked to a list of 52 workplace abilities.<sup>11</sup> For each ability, survey respondents need to indicate whether they think the respective AI application can be used. The result is a relatedness measure for each occupation-ability combination that ranges between zero (no relation) and one (high relation). In a second step, each ability’s exposure is constructed as the sum of the relatedness value across all AI applications. It ranges from zero (no exposure) to 10 (high exposure). The third step involves computing each occupation’s exposure to AI (AIOE) by taking the weighted average across the 52 abilities’ exposures to AI, with weights given by abilities’ prevalence in each occupation (provided by O\*Net). We standardise the resulting AIOE variable to range from zero to ten, with higher values indicating a greater importance of AI within an occupation. To construct exposure to AI at the industry level (AIIE), we use data on occupations’ employment shares within each two-digit industry (provided by the Bureau of Labor Statistics) as weights to average across AIOEs. We then standardise AIIE so that the industry with the highest exposure value has a score of one; all other industries’ AIIE are then expressed as a fraction of the AIIE of the highest exposure industry.

[Figure 1](#) plots the five occupations with the highest and lowest AIOE scores. As discussed in [Felten et al. \(2021\)](#), the highest AIOE scores consist mostly of white-collar occupations that require advanced degrees, such as genetic counsellors, financial examiners and actuaries. The lowest-scoring occupations predominately require a high degree of physical effort and include, for example, dancers, fitness trainers or iron and rebar workers.

[Figure 2](#) plots the (standardised) AIIE at the industry level (grey bars) as well as the respective employment shares (red diamonds). The industry with the highest AIIE score is finance and insurance, followed by management of companies and enterprises.

---

<sup>11</sup>The ten applications are: image recognition, visual question answering, image generation, reading comprehension, language modelling, translation, speech recognition, abstract strategy games, real-time video games and instrumental track recognition. A list of 52 cognitive, physical, sensory and psychomotor workplace abilities is provided by the O\*Net database. Cognitive abilities include, for example, deductive reasoning or oral comprehension. Physical abilities include stamina or trunk strength, among others, whereas psychomotor abilities include arm-hand steadiness or finger dexterity.

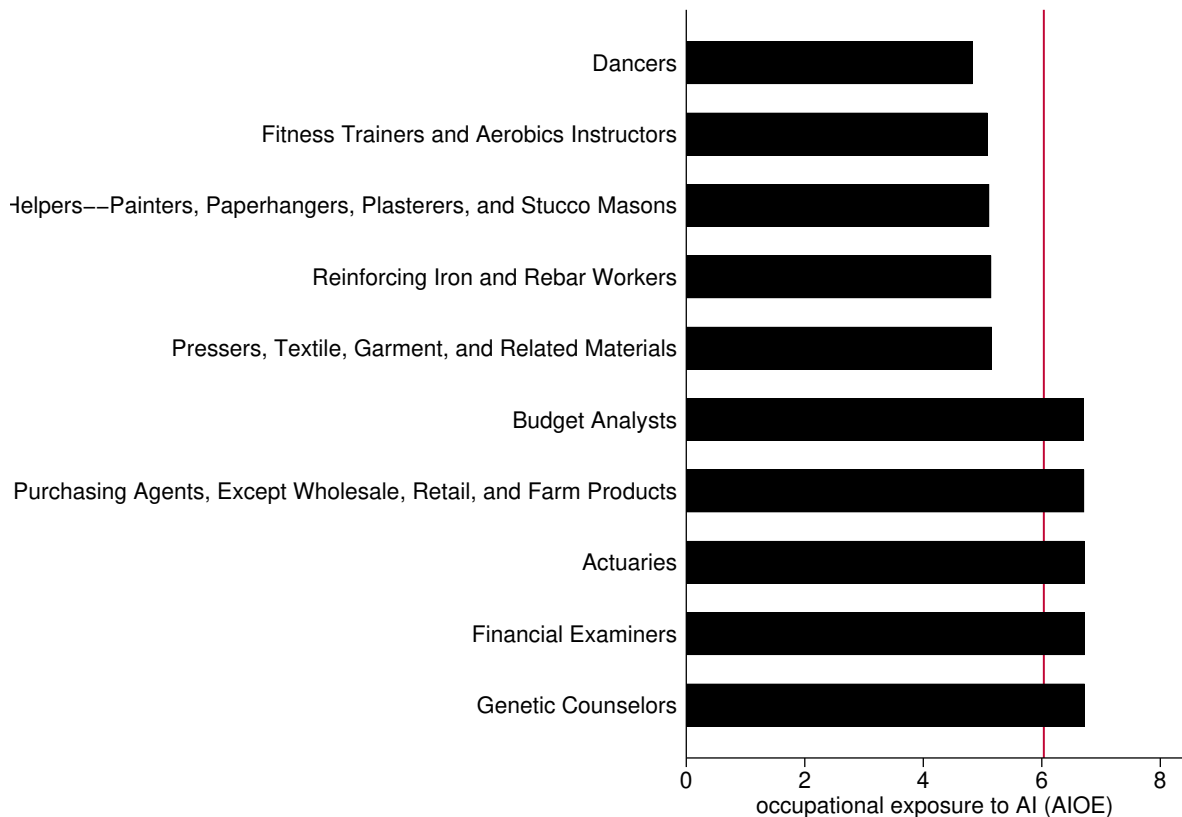


Figure 1: **AI exposure of top/bottom occupations.**

Those with the lowest score are agriculture, forestry, fishing and hunting as well as transportation and warehousing. Overall, there is not very large variation in exposure across industries, the reason being that *within* industries, there are a lot of occupations with various degrees of exposure to artificial intelligence.

Our measure reflects exposure to AI and does not capture whether AI will substitute or complement any particular occupation. While there is general agreement that AI is positive for productivity, there is no consensus on the scope for AI complementing or substituting tasks. One option is to use judgement to categorise each specific occupation as more/less at risk of displacement (Pizzinelli et al., 2023; Cazzaniga et al., 2024). The drawback is that this approach requires strong assumptions that are difficult to verify, which risks assuming the final impact across occupations from inception. Instead, we calibrate exposure to AI and let the model decide which industries witness increased or decreased employment. In this way, the final impact across industries is disciplined by

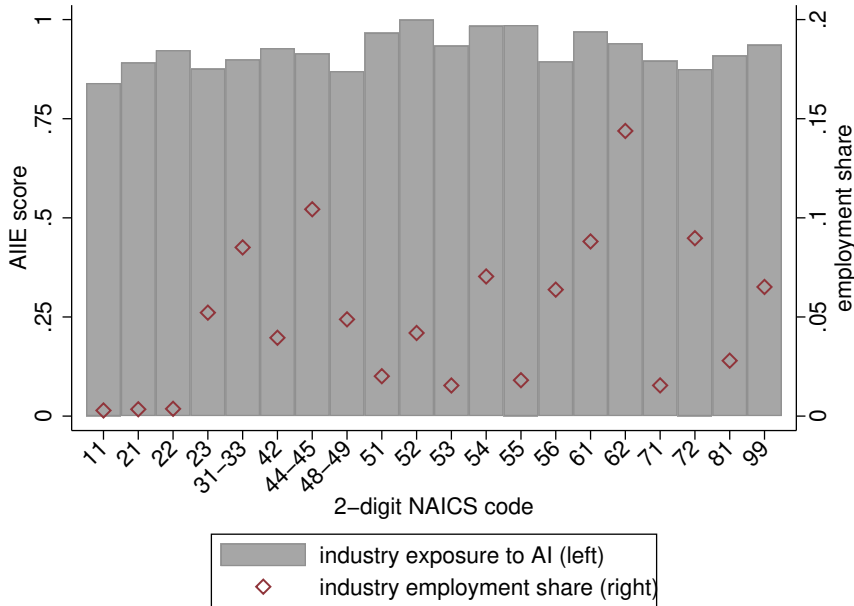


Figure 2: **The direct impact of AI across industries**

general equilibrium effects and is not a function of our priors.

### 3 A multi-sector model

This section describes the building blocks of our multi-sector model and explains how we use it to chart the economic effects of AI. We then provide intuition for our core results using simplified versions of the full model.

#### 3.1 The model

We work with the multi-industry New-Keynesian Dynamic Stochastic General Equilibrium model presented in [Rees \(2020\)](#). The model features a detailed industry structure in both its demand and supply sides. This allows it to capture the key industry-level and aggregate transmission mechanisms of AI. We use the model to account for the direct effects of AI on industry-level productivity and then to trace through the effects of these changes across the economy to assess the implications for aggregate outcomes.

The model consists of a closed economy featuring households, firms, the government and the central bank.<sup>12</sup>

Households make consumption, work, investment and saving decisions to maximise their lifetime utility, subject to an intertemporal budget constraint. Their utility function is given by:

$$\sum_{t=0}^{\infty} \beta^t \left[ \log(C_t - hC_{t-1}) - \frac{A_N}{1+\nu} N_t^{1+\nu} \right] \quad (1)$$

where  $C_t$  and  $N_t$  are household consumption and labour supply. The parameters  $\beta$ ,  $h$  and  $\nu$  are the household's intertemporal discount rate, its habits parameter and its Frisch labour supply elasticity.

The intertemporal budget constraint in turn is given by:

$$P_{C,t}C_t + P_{I,t}I_t + \frac{B_{t+1}}{R_t} \leq B_t + \sum_{j=1}^{\mathcal{F}} \left( P_{C,t}r_{j,t}^K k_{j,t} + w_{j,t}n_{j,t} \right) \quad (2)$$

where  $I_t$  is the household's total investment in physical capital,  $P_{C,t}$  and  $P_{I,t}$  are the prices of the consumption and investment goods,  $B_{t+1}$  is a risk free nominal bond that pays one unit of the consumption good in period  $t+1$  and  $R_t$  is the interest rate of that bond.

The economy features  $\mathcal{F}$  industries. The variables  $k_{j,t}$  and  $n_{j,t}$  represent the total supply of capital and labour from the household to industry  $j$ . In turn,  $r_{j,t}^K$  and  $w_{j,t}$  are the return on capital and nominal wage paid by that industry.

The aggregate consumption and investment goods in Equation (2) consist of bundles of products from individual industries. For example, the aggregate consumption bundle is:

$$C_t = \left[ \sum_{j=1}^{\mathcal{F}} \omega_{c_j}^{\frac{1}{\eta}} c_{j,t}^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}} \quad (3)$$

---

<sup>12</sup>For the sake of expositional clarity, we focus on the aspects of the model that are most relevant for our application. We provide the full model as a separate file.

where  $c_{j,t}$  is the amount of output of industry  $j$  used to produce consumption goods at time  $t$ . Conditional on the relative prices of industries' output, the parameter  $\omega_{cj}$  determines the weight of industry  $j$  in the aggregate consumption bundle. The parameter  $\eta$  determines the degree of substitutability between the output of different industries in consumption: a higher (lower) value means that the output of different industries is more (less) substitutable in consumption.

The aggregate labour supply that appears in the household utility function is also a weighted sum of labour supply to individual industries:

$$N_t = \left[ \sum_{j=1}^{\mathcal{F}} \omega_{nj}^{-\frac{1}{\xi}} n_{j,t}^{\frac{\xi+1}{\xi}} \right]^{\frac{\xi}{\xi+1}} \quad (4)$$

where  $\omega_{nj}$  captures the relative disutility the household receives from supplying labour to industry  $j$  and  $\xi$  controls the substitutability of work across industries. In the limit, if  $\xi = \infty$ , workers are indifferent between working in different industries.

On the production side of the model, each industry consists of many firms producing differentiated product varieties under monopolistic competition. Individual firms produce output using a multi-stage production process. The first stage combines labour and capital according to the following production function:

$$f_{j,t}(i) = \left[ \omega_{f,j}^{\frac{1}{\zeta}} n_{j,t}(i)^{\frac{\zeta-1}{\zeta}} + (1 - \omega_{f,j})^{\frac{1}{\zeta}} k_{j,t}^s(i)^{\frac{\zeta-1}{\zeta}} \right]^{\frac{\zeta}{\zeta-1}} \quad (5)$$

where  $f_{j,t}$  is an aggregate of labour and capital used by firm  $i$  in industry  $j$  and  $n_{j,t}(i)$  and  $k_{j,t}^s(i)$  are the amount of labour and capital services employed by the firm. The parameter  $\zeta$  is the elasticity of substitution between capital and labour.

In a second stage, firms combine the labour and capital bundle,  $f_{j,t}$ , with intermediate inputs sourced from other industries:

$$y_{j,t}(i) = a_{j,t} \left[ \omega_{y,j}^{\frac{1}{\varphi}} f_{j,t}(i)^{\frac{\varphi-1}{\varphi}} + (1 - \omega_{y,j}) x_{j,t}(i)^{\frac{\varphi-1}{\varphi}} \right]^{\frac{\varphi}{\varphi-1}} \quad (6)$$

where  $y_{j,t}(i)$  is the gross output of firm  $i$  in industry  $j$  and  $x_{j,t}(i)$  is the amount of intermediate inputs used by the firm. The term  $a_{j,t}$  represents total factor productivity,

which is industry-specific and common to all firms in that industry. The parameter  $\varphi$  is the elasticity of substitution between intermediate inputs and the aggregate of labour and capital.

The intermediate input is itself a bundle of intermediate goods from the other industries, with the elasticity of substitution between different varieties of intermediate goods given by  $\psi$ .

Value-added output is equal to a firm's gross output minus its intermediate inputs. In the case where all final goods prices are:

$$y_{j,t}^{va}(i) = y_{j,t}(i) - x_{j,t}(i) \quad (7)$$

Market clearing requires that the gross output of industry  $j$  ( $y_{j,t}$ ) equals the sum of demand for the good as a consumption good ( $c_{j,t}$ ), investment good ( $i_{j,t}$ ) and public demand good ( $g_{j,t}$ ), or as an intermediate input (where  $x_{j,k,t}$  is the output of industry  $j$  used as an intermediate input in industry  $k$ ):

$$y_{j,t} = c_{j,t} + i_{j,t} + g_{j,t} + \sum_{k=1}^{\mathcal{F}} x_{j,k,t} \quad (8)$$

The presence of intermediate inputs creates a rich network of linkages across its industries. For instance, higher productivity in the manufacturing industry lowers costs for firms operating in industries that use manufacturing goods as an input. Two factors dictate the importance of inter-industry linkages in the model: the weight of intermediate inputs in industry production functions and the substitutability between intermediate inputs and other factors.

The model includes a number of nominal and real rigidities. Firms face Calvo-style price rigidities, with the degree of price stickiness varying across industries. Households are assumed to unionise, giving them a degree of monopoly power in the labour market. There too, Calvo wage rigidities exist. The model's real rigidities include habits, investment adjustment costs and capital utilisation costs. While these rigidities do not alter the long run consequences of AI adoption, they materially influence the model's short-run dynamics.

The two remaining agents are the government and the central bank. We assume that

government expenditure as a share of nominal GDP follows an autoregressive process, funded by lump-sum taxation. The central bank adjusts its policy interest rate in response to deviations of inflation from target and the output gap, defined as the deviation of real GDP from the model’s flexible price benchmark.

The model includes 20 industries, corresponding to the two-digit NAICS classification described in Section 2. We calibrate the model to match key features of the US economy, using input-output tables to pin down the weights of capital, labour and industry-specific intermediate inputs in the industry production functions, the weights of consumption, investment and government spending in domestic demand and the weights of each industry in the consumption, investment and government spending bundles. Most of the other parameters controlling the dynamics of the model, such as the habits and investment adjustment cost parameters, or the aggregate labour supply elasticity, are taken from the literature. Table 1 provides a summary of these – for details we refer the reader to Rees (2020).

Table 1: Calibration of key parameters

Parameter	Description	Value
$\beta$	Discount rate	0.99
$h$	Habits	0.70
$Spp$	Investment adjustment cost	3.00
$\eta$	Elasticity of substitution in demand CES	0.90
$\zeta$	Elasticity of substitution between capital and labour	0.95
$\varphi$	Elasticity of substitution between intermediates and capital / labour	0.60
$\psi$	Elasticity of substitution between intermediates	0.40
$\nu$	Frisch labour supply elasticity	2.00
$\varepsilon_w$	Labour supply elasticity across industries	5.00
$\delta$	Depreciation rate	0.02
$\rho_r$	Taylor rule - autoregressive parameter	0.80
$\phi_\pi$	Taylor rule - response to inflation	1.50
$\phi_{gap}$	Taylor rule - response to output gap	0.25
$\theta_{sticky}$	Calvo - sticky price sectors	0.80
$\theta_{semi-flex}$	Calvo - semi-flexible price sectors	0.50
$\theta_{flex}$	Calvo - flexible price sectors	0.25
$\chi_p$	Price indexation	0.20
$\theta_w$	Calvo - wages	0.75

Note: Sticky price sectors are Agriculture and Mining; Semi-flexible price sectors are Utilities, Manufacturing, Retail trade, Wholesale trade and Transport.

Finally, we use the AIIE measure constructed above to simulate the impact of AI adoption by industries. Guided by the estimates in the literature, we do this in a way such that the annual impact on total TFP growth aggregates to a value of 1.5% for a



decade.

### 3.2 Modelling the effects of AI

We simulate the effects of AI as follows.

We first linearise the model along its balanced growth path. This provides us with a baseline case in which there is no AI-induced boost to productivity. We present all of our results as deviations from this baseline. The model's linearised structural equations take the form:

$$\mathbf{A}y_t = \mathbf{C} + \mathbf{B}y_{t-1} + \mathbf{D}\mathbb{E}y_{t+1} + \mathbf{F}\varepsilon_t \quad (9)$$

where  $y_t$  is the vector of the model's endogenous variables and  $\varepsilon_t$  is the vector of structural shocks, which we take to be i.i.d. without loss of generality. Note that in solving the model we keep track of the steady state of the linearised variables, resulting in the inclusion of the vector  $\mathbf{C}$  in Equation (9).

If it exists and is unique, the standard rational expectations solution to Equation (9) is:

$$y_t = \mathbf{J} + \mathbf{Q}y_{t-1} + \mathbf{G}\varepsilon_t \quad (10)$$

We model the effects of AI as a sequence of *permanent* increases in the level of industry-specific productivity (i.e. the steady state of the terms  $a_{j,t}$  in Equation 6). As well as their direct effects on industry output, these changes have general equilibrium effects that alter the economy's entire balanced growth path. We account for these changes using the approach outlined in [Kulish and Pagan \(2017\)](#). Specifically, for each period in our simulations, we calculate a new system of linearised equations:

$$\mathbf{A}_t y_t = \mathbf{C}_t + \mathbf{B}_t y_{t-1} + \mathbf{D}_t \mathbb{E} y_{t+1} + \mathbf{F}_t \varepsilon_t \quad (11)$$

where  $\mathbf{A}_t$ ,  $\mathbf{B}_t$ ,  $\mathbf{C}_t$ ,  $\mathbf{D}_t$  and  $\mathbf{F}_t$  represent the linearised structural equations for the industry-specific productivity levels in period  $t$ .

The presence of forward-looking expectations in Equation (11) requires us to take a stance on economic agents' beliefs about the future path of productivity to solve for the reduced form of this system.

Various plausible alternatives suggest themselves. Advances in AI technology have received considerable public attention, particularly after the release of ChatGPT 3.0 in November 2022. Rapid increases in the share prices of AI-related firms suggest that financial market participants, at least, have started to factor the transformative possibilities of these technologies into their investment decisions. That said, considerable uncertainty exists about the real-world application of existing AI models, not to mention their future evolution. As such, it would not be surprising if households and firms responded cautiously to AI developments, adjusting their economic decisions as the technology evolves and its real-world applications become more apparent.

Acknowledging this uncertainty, we consider two cases. In the first, which we refer to as the “unanticipated” case, model’s agents observe AI-induced productivity increases when they occur, but do not anticipate further gains in the future. In this case, the solution to the reduced form of the model is a time varying VAR of the form:

$$y_t = \mathbf{J}_t + \mathbf{Q}_t y_{t-1} + \mathbf{G}_t \varepsilon_t \quad (12)$$

where  $\mathbf{J}_t$ ,  $\mathbf{Q}_t$  and  $\mathbf{G}_t$  are the standard rational expectations solutions to Equation (11) in each period.

In the second case, agents observe AI-induced productivity increases as they occur *and* correctly anticipate their further evolution. In this “anticipated” case we solve the model recursively. We start in period  $T$ , at which point we assume that no further AI-induced productivity increases will occur. After solving for the reduced form solution matrices,  $\tilde{\mathbf{J}}_T$ ,  $\tilde{\mathbf{Q}}_T$  and  $\tilde{\mathbf{G}}_T$  in that period, we can then use the formulas provided in [Kulish and Pagan \(2017\)](#) to determine all previous reduced form matrices.

### 3.3 Building intuition through parsimonious models

The effects of AI on macroeconomic aggregates in our full model reflect a rich set of dynamics and mechanisms. To build intuition for our results, we first show how AI-induced increases in industry productivity propagate in two simpler versions of the model. Each highlights a specific mechanism at work in the full model.

The first scaled-down model illustrates the role of cross-industry differences in production technologies. The model has three industries: the first is labour intensive, the second

is capital intensive and a third is ‘balanced’ between those two extremes. We calibrate the model so that each industry accounts for the same share of aggregate consumption, investment and intermediate inputs before the introduction of AI. We vary these weights in the second scaled-down model below.

We consider four scenarios. In the first three scenarios, AI adoption leads to an immediate, permanent and one-time 10% increase in the TFP of a *single* industry, and has no effect on productivity in the other industries. In the fourth, AI adoption raises productivity in *each* industry by  $3\frac{1}{3}\%$ .<sup>13</sup>

The path of GDP is similar across the scenarios (Figure 3, left-hand panel). That is, the response of GDP to AI does not depend on whether its effects are concentrated in labour- or capital-intensive industries. It rises immediately by around 3%.<sup>14</sup> It then increases further over time, even though the level of productivity in each industry remains constant after the first period. The additional output increase, beyond that implied directly by the increase in productivity, occurs because higher productivity enables an expansion in the economy’s capital stock and the production of more intermediate inputs, which both deliver an additional boost to the economy’s productive capacity. The final increase in GDP of around 12% is hence much larger than the initial rise.

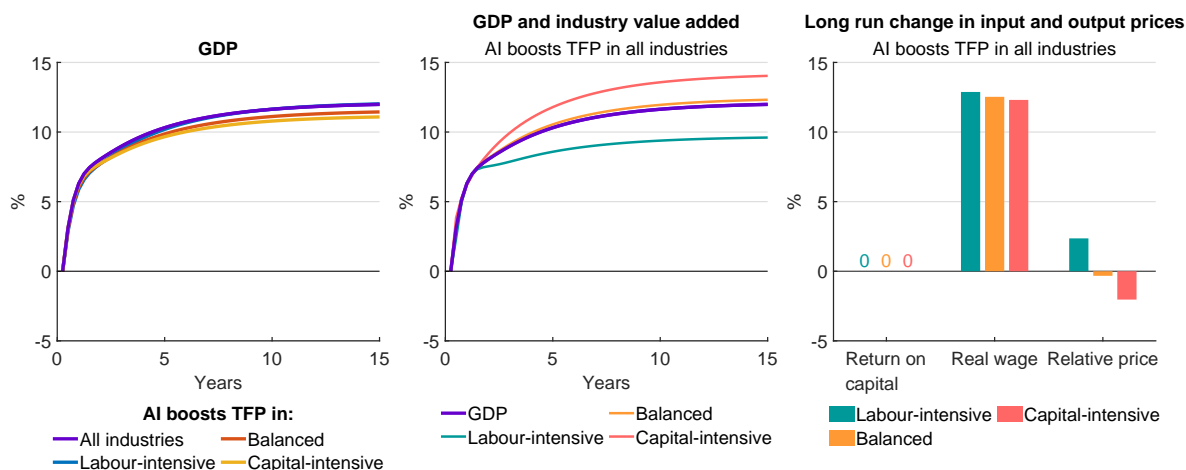


Figure 3: Simple model 1 – Cross-industry differences in production technologies

At the industry level, output trajectories differ substantially. In Scenarios 1-3, the

<sup>13</sup>Because the industries are initially equally-sized, the weighted average increase in industry productivity in Scenario 4 is the same as that in Scenarios 1–3.

<sup>14</sup>This is roughly the increase in GDP that would result if the introduction of AI had no effect on the supply of other factors of production and influenced GDP only through its effect on TFP.

industry that experiences the productivity boost naturally experiences the largest increase in output. But even in Scenario 4, where AI delivers the same productivity boost to each industry, cross-industry output differences arise (Figure 3, centre panel). In particular, value added output rises more in capital-intensive industries than labour-intensive ones.

The behaviour of input and output prices provides intuition for this result.

In the long-run, the size of each industry's capital stock adjusts to equalise the rate of return on capital across industries. AI adoption that raises the return on capital in a given industry encourages investment in that industry. In time, higher investment increases the size of that industry's capital stock and drives down its return.<sup>15</sup> Hence, although capital returns vary along the *transition* to a new balanced growth path, as long as the effects of AI on productivity *growth* are temporary, and do not affect households' discount rates or capital depreciation rates, the return on capital will ultimately converge to its original level (Figure 3, right-hand panel, first three bars).

The scope to increase the labour force is much more limited. Hence, for workers, AI-induced productivity gains lead to permanently higher real wages. And, when AI raises productivity by the same amount in each industry, the increases in real wages – deflated by aggregate consumer prices – are also similar (Figure 3, right-hand panel, centre bars).<sup>16</sup>

The behaviour of rates of return to capital (flat) and real wages (higher) implies that a proportional increase in TFP will raise input costs more in labour intensive industries. As firms price output as a markup over their marginal costs, the relative output prices of labour-intensive industries rise and those of capital-intensive industries fall (Figure 3, right-hand panel, right-hand bars). This, in turn, induces households and firms to substitute towards more capital-intensive goods and services, explaining the differences in value-added across industries in Figure 3, centre panel.

The second scaled-down model illustrates the implications of differences in the use of industry output. As before, we build a three-sector version of the model, calibrated so that each sector accounts for the same share of value added output before the adoption of AI. In this case, industries share same production technology, i.e. their use of capital,

---

<sup>15</sup>Investment adjustment costs prevent the adjustment in the size of the capital stock from taking place immediately.

<sup>16</sup>Because we assume that labour is imperfectly substitutable across industries, wage rates do not equalise.

labour and intermediate inputs are initially identical. However, they differ in how their products are used. We calibrate the weights of the three industries in the consumption, investment and intermediate-input CES bundles so that each industry specialises in the production of goods for one particular use. We simulate the same four scenarios as for the previous model.

In contrast to the first scaled-down model, the trajectory of aggregate GDP differs markedly across scenarios (Figure 4, left-hand panel). When the effects of AI are concentrated in the industry specialising in the production of consumer goods, GDP rises by more than 20% in the long run (blue line). When instead they are concentrated in the industry producing investment goods, the long-run rise in GDP is around 5% (orange line). The other two scenarios lie between these two extremes.

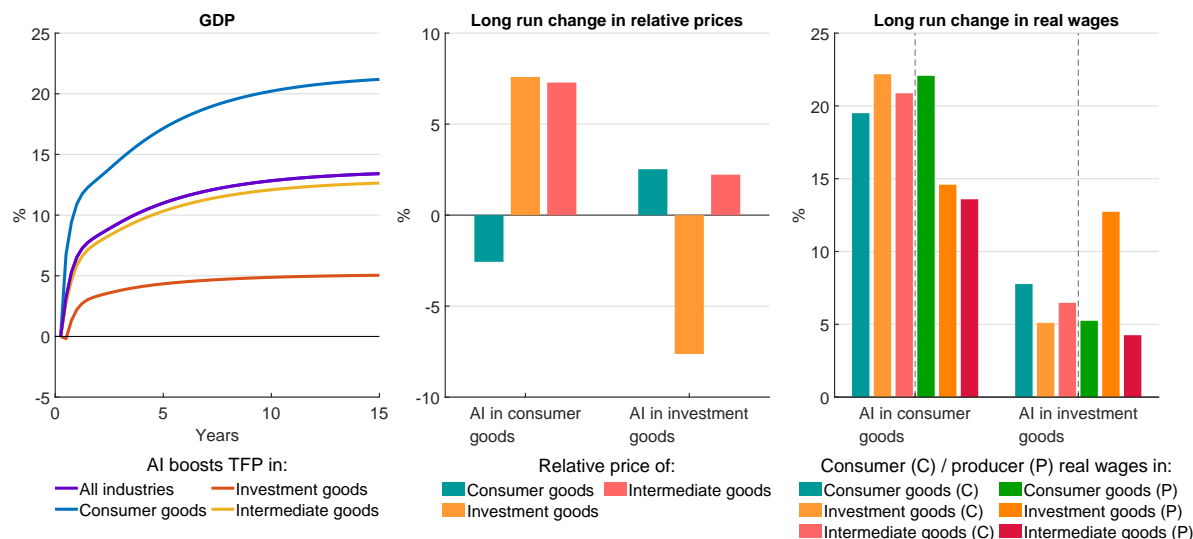


Figure 4: Simple model 2 – Differences in the use of industry output

Once again, the intuition for this result starts with an examination of the behaviour of prices and input costs.

A disproportionately large increase in productivity due to AI adoption within a particular industry lowers the relative price of that industry’s output. For example, if higher productivity occurs in the consumer goods sector, the relative price of consumption goods falls (Figure 4, centre panel, left bars). If instead it is concentrated in the investment goods sector, the relative price of investment goods falls (right bars).

As factor mobility limits divergences in wages and capital returns, AI adoption leads

to a similar rise in real wages – deflated by consumer prices – in all industries. And it does so even when the effects of AI are concentrated in only one industry (Figure 4, right-hand panel).

But what matters for firms are factor costs deflated by their own output prices. And movements in relative prices mean that real producer wages (i.e., nominal wages deflated by a firm’s own prices) can vary substantially, even when real consumer wages (i.e., nominal wages deflated by the overall consumer price index) move by a similar amount. For example, when the effects of AI are concentrated in the consumer goods sector, real producer wages for firms producing consumer goods rise much more than those in other sectors because of the relative decline in consumer goods prices. When the effects of AI concentrate in the investment goods sector, the difference in real producer wages is even larger – again a function of the large relative price swings for investment goods in that scenario.

Profit maximisation by firms creates a link between a firm’s costs and its production. At the margin, firms that face higher costs – evaluated in producer prices – will cut back on production. When the effects of AI are concentrated in consumer goods industries, which lie at the end of the economy’s production chain, this constraint on production has positive spillovers. It frees up resources to increase the production of investment and intermediate goods, both of which further expand the economy’s productive capacity. In contrast, if lower relative prices of investment goods decrease that industry’s production, the spillovers are negative as the economy’s ability to produce other goods and services is also diminished.

The scaled-down models also provide some intuition for the possible implications of AI adoption for inflation and interest rates (Figures 13 and 14 in Appendix A).

In the first model, inflation trajectories are similar across the four scenarios. For the specific calibration we consider, the introduction of AI initially lowers inflation (Figure 13, left-hand panel). Inflation is only temporarily lower, however. It soon turns positive, alongside a positive output gap and a rise in the policy rate (centre and right-hand panels).

The responses of inflation, the output gap and policy rates for the scenarios constructed using the second scaled-down model are qualitatively similar, albeit with greater variation in outcomes across the four scenarios (Figure 14). In one extreme, when the

effect of AI concentrates on the consumer goods industry, the effects on inflation, output gap and policy rates substantially more pronounced. In the other extreme, effects are weakest when AI boosts TFP in the investment goods sector.

Why is the introduction of AI – which expands aggregate supply – inflationary in these scenarios? The answer is two-fold. First, the introduction of AI raises incomes, which increases aggregate demand. Second, it creates the need for additional investment to expand the capital stock and the production of intermediate inputs. This additional investment increases aggregate demand immediately, and aggregate supply only with a lag, and hence is inflationary.

Taken together, the results from our scaled-down model variations deliver useful insights to assess the results from our main model, which we present in the next section. Differences in production technologies as reflected in the relative use of factors of production are not the main driver of the evolution of key macroeconomic outcome variables.<sup>17</sup> Differences in the use of industry output, in contrast, can lead to quite different dynamics in response to an AI-induced increase in productivity. When the effects of AI concentrate in industries that are closer to final demand (i.e., “downstream”), general equilibrium price effects shift resources to upstream industries, which in turn reinforces the positive direct effects on output through their high production linkages.<sup>18</sup>

## 4 The macroeconomic impact of artificial intelligence

Having built intuition for the key mechanisms driving our model dynamics, we now present results from the full model. Recall that AI is represented as an increase in productivity growth of 1.5 ppt annually over the next decade, allocated to sectors based on the AIIE measure constructed in Section 2. In what follows we first discuss how AI adoption affects key macroeconomic aggregates, before analysing its impact across industries.

---

<sup>17</sup>To be sure, there are cross-industry differences even when AI delivers the same productivity boost to all industries. As discussed above, this can be rationalised by assessing general equilibrium effects through prices. But they will not be a driver of large differences in the aggregate.

<sup>18</sup>That said, we emphasise that the patterns of inflation, the output gap and policy rates obtain from these models are specific to the calibration used in this exercise.



## 4.1 The macroeconomic impact of AI

**The unanticipated case.** We first discuss results when households and firms do not anticipate future AI-induced productivity increases.

Productivity improvements from AI adoption lead to a significant increase in GDP. Growth is fastest in the first 10 years – i.e. the period in which AI directly raises industry-level TFP – at which point GDP is almost 30% higher than it would have been without AI adoption (Figure 5, green line in panel (a)). GDP continues to increase even after the direct productivity gains from AI adoption are exhausted, albeit at a slower pace, as it takes time for firms to adjust their capital stock and use of intermediate inputs to take full advantage of AI. The level of GDP ultimately stabilises around 35% above the no-AI baseline.

The paths of aggregate consumption and investment broadly resemble that of GDP (Figure 5, green lines in panels (b) and (c)). The level of investment overshoots, thereby delivering the required increase in the economy’s capital stock, before converging to its long-run level.

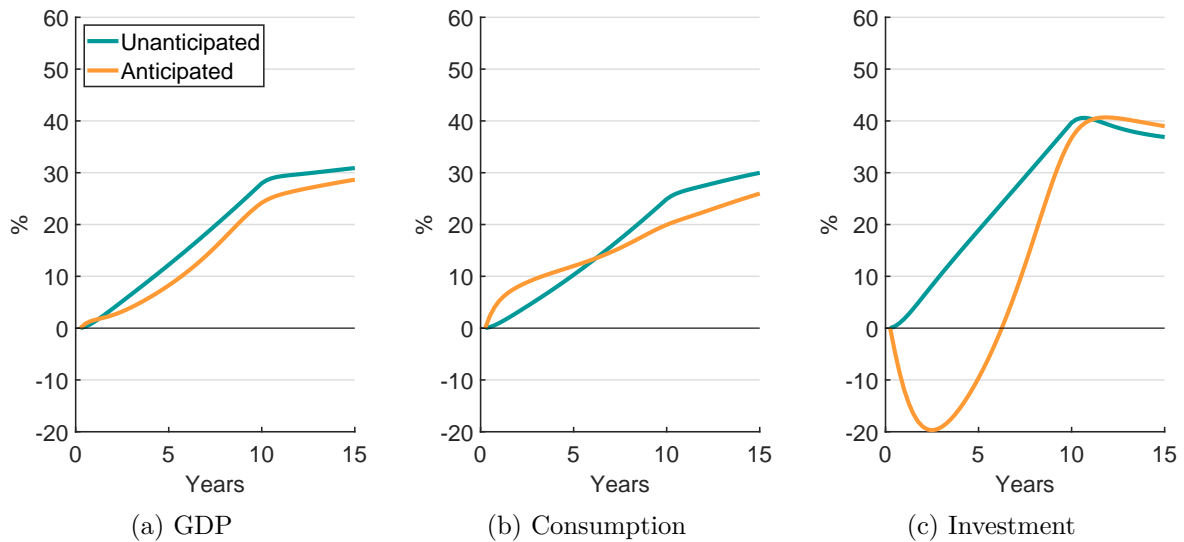


Figure 5: **Output, consumption, and investment**

In this scenario, AI adoption initially lowers inflation (Figure 6, green line in panel (a)). Faster productivity growth increases the economy’s productive capacity, lowering firm costs and relaxing supply constraints, thus acting as a disinflationary force. The

demand response, on the other hand, takes time to materialise, in part because of frictions such as habits in consumption and investment adjustment costs.<sup>19</sup> Policy rates respond accordingly in this initial phase and decline with inflation.<sup>20</sup> Over time, however, the effects of AI adoption on aggregate demand start to dominate its effects on aggregate supply. Consequently, AI adoption ultimately raises the inflation rate after around four years. Inflation peaks about 0.75 percentage points above its level without AI adoption. In response to higher demand-driven inflation, the central bank steadily raises its policy rate (Figure 6, green line in panel (c)).

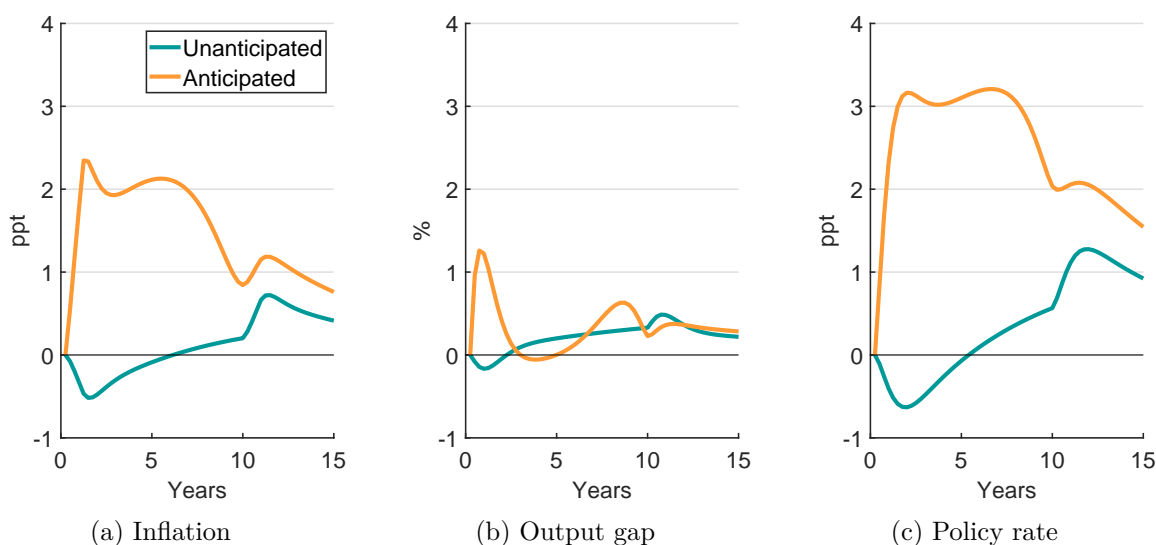


Figure 6: **Inflation, output gap and policy rate**

**The anticipated case.** We now discuss results when agents correctly anticipate the entire future path of AI-induced productivity increases.

GDP increases more slowly in the anticipated case than the unanticipated one (Figure 5, orange line in panel (a)). This is because households, who correctly foresee that AI will raise productivity in the future, bring forward consumption in order to smooth its

<sup>19</sup>The timing of inflation is also a function of, among other things, the price stickiness of the industries affected and where they sit in the production chain (i.e. whether they supply consumption or investment goods), since the inflation measure we consider in the Figure is for consumption goods. This is discussed in some more detail in the context of our scaled-down models in Section 3.3.

<sup>20</sup>The lagged term in the Taylor rule that determines the policy rate response prevents the central bank from lowering rates quickly enough to prevent disinflation (or raising them quickly enough to prevent inflation later).

trajectory over time (panel (b)). To accommodate higher consumption, which runs ahead of actual AI-induced productivity improvements, investment falls (panel (c)). As a result, the economy’s capital stock grows more slowly than in the unanticipated case, resulting in a lower level of GDP compared to the unanticipated case.

The paths of inflation, interest rates and the output gap differ markedly in the anticipated scenario from their trajectories in the unanticipated one (Figure 6, orange lines). Although the level of output is lower in the anticipated case than the unanticipated one, inflation is significantly higher, peaking at more than 2 ppt higher than it would have been without AI adoption. It then remains above its baseline level throughout the scenario. This, in turn, induces a large rise in policy rates. The output gap is modestly positive throughout the scenario.

Taken together, these results illustrate that while different expectation formation mechanisms change the transitory dynamics of macro variables, they do not affect the long-term impact of AI on the economy. Our results can be seen as supporting the productivity effect in the task-based framework of Acemoglu and Restrepo (2018). An important caveat to bear in mind when interpreting our results, however, is that they cannot account for future changes to industrial structure nor the disappearance of existing occupations or the appearance of new ones. For example, Autor et al. (2024) documents that around 60% of contemporary work is done in occupations that did not exist 80 years ago. If recent advances in AI prove to be disruptive, we can also expect new tasks and occupations to appear, which could provide further support to our results.

## 4.2 The impact of artificial intelligence across sectors

We now turn to the effects of AI across sectors. We focus on long-run outcomes, which are the same in both the anticipated and unanticipated cases.

Figure 7 shows the long run increase in value-added output by industry. For exposition, we color-code the industries into five groups – primary industries, secondary industries, distribution, professional services and other services. As a reference point, the horizontal line shows the economy-wide long-run increase in aggregate GDP due to the introduction of AI.

Three observations stand out. First, value-added output rises in all industries, re-

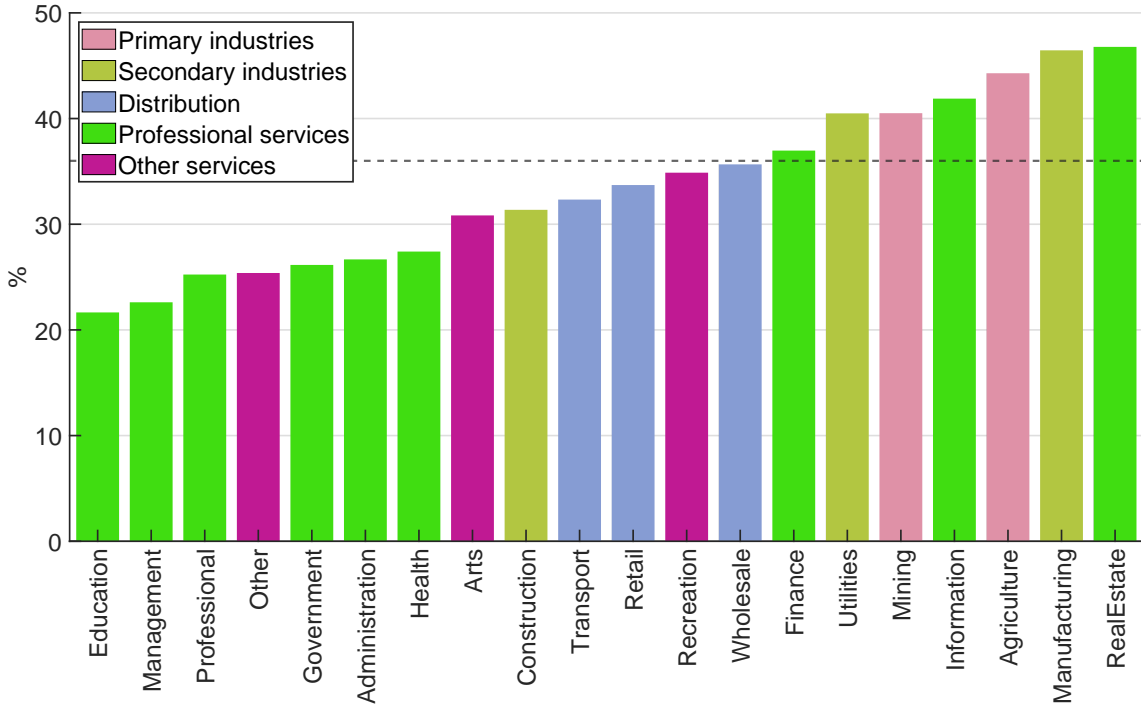


Figure 7: Long-run increase in industry value added

flecting the nature of AI as a general-purpose technology. Second, the impact varies significantly across industries, ranging from a nearly 50% increase in value-added output in manufacturing and real estate services to about 20% in education and management services. Third, there is no direct mapping between an industry’s initial exposure to AI and the long run increase in value added output (see Figure 8a). Some sectors that are expected to receive the smallest initial productivity boost from AI, such as arts & recreation services, record a relatively large output increase. Others that are expected to receive a particularly large direct boost from AI, such as management services, see smaller output increases. In general, primary and secondary industries display the largest increases in value added output, while professional services are in the bottom half of the distribution, with a couple of notable exceptions such as Information & Communications and Real Estate Services.

What explains the differences in value added growth across industries? It turns out that the mechanism in the first simple model presented in Section 3 – differences in the capital and labour intensity of production – is crucially important. The relationship between an industry’s labour-intensity, as proxied by the income share of labour in that industry’s gross output, and its increase in value-added output as a result of AI adoption

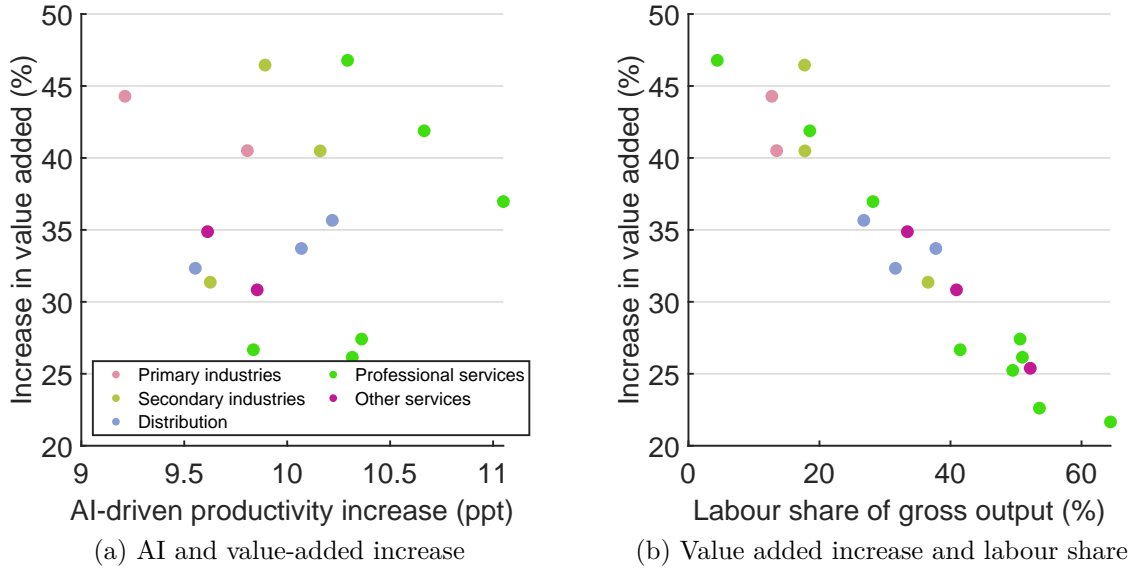


Figure 8: **Value-added, AI and labour shares**

is close. More labour-intensive industries record the smallest increases in value added output (Figure 8, panel (b)). As in the simple model, differences in labour intensity translate into movements in relative prices, with the relative prices of labour-intensive industries rising relative to more capital intensive ones (Figure 9a).

AI adoption also leads to a reallocation of labour across industries (Figure 9b). In general, employment increases in services industries that experience higher relative prices and the smallest increases in value added output. Higher selling prices allow firms in these industries to raise wages more than firms in industries whose relative prices decline. This induces workers to adjust their labour supply towards these industries. In contrast, in the capital-intensive industries that record the largest increases in value-added, and where relative prices decline, hours worked fall.

Why do the industries that record the largest increases in hours worked not record larger increases in output than those where hours worked declines? The answer lies in the behaviour of capital. The productivity boost from AI adoption allows for a material increase in the economy's capital stock. Because labour and capital are imperfect substitutes, a given proportional increase in the capital stock delivers a larger output increase in capital intensive industries (i.e. those where  $\omega_{y,j}$  is lower) than labour intensive ones.

The above discussion describes changes in *value added* output at the industry level. In

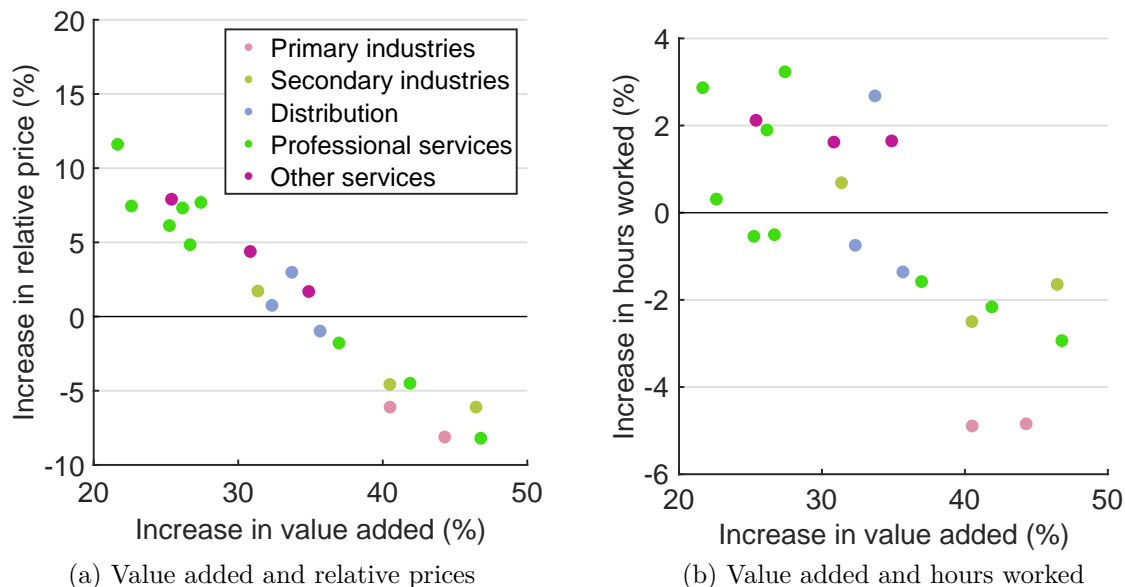


Figure 9: **Value-added, hours worked and relative prices**

practice, much of the increase in the output of primary and secondary industries will be used in downstream industries as intermediate inputs in the production of final goods and services for households and businesses. To illustrate these effects, Figure 10 decomposes the increase in value added by industry into the increase in its value added that shows up as final demand (i.e. consumption, investment or government expenditure) and the part that is used as intermediate inputs by other industries. It further breaks down these two categories into the change due to “income” effects (i.e. higher aggregate demand) and “substitution” effects (i.e. changes in relative prices). The overall picture that emerges is that income effects matter much more than substitution effects.

## 5 Counterfactual scenarios

Our setting has two distinctive features that matter for the interpretation of our main results. First, the effect of AI as calibrated from the AIIE measure shows relatively little variation across sectors. The measure is both theoretically and empirically grounded. However, substantial uncertainty remains as to how AI will ultimately affect individual industries. Second, AI is neither capital- nor labour-specific, but is rather a general purpose technology affecting overall TFP. This is in line with the common understanding

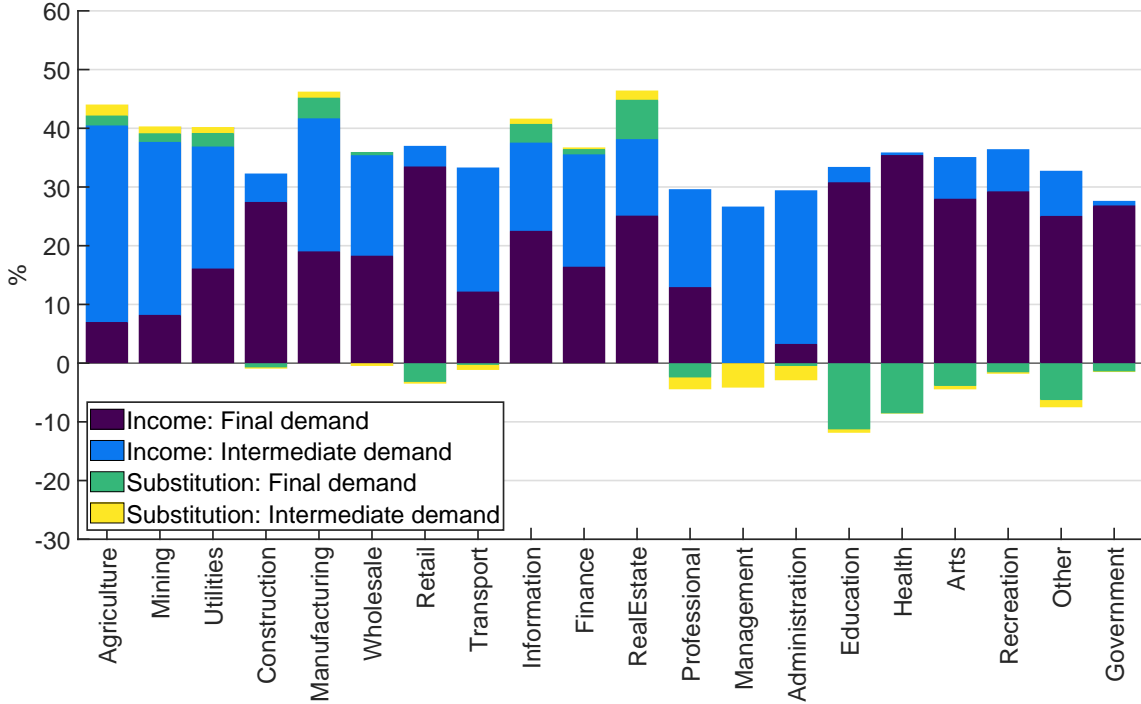


Figure 10: **Decomposition of industry value-added**

in the literature. But again, uncertainty remains as to whether AI affects more capital or labour.

In this section we leverage the flexibility of our model to perform counterfactual analyses that explore how the results change when we relax each of these features. To simplify the exposition, we show results from only the case where future TFP adoption is unanticipated. However, the qualitative message from the anticipated case is similar.

**Heterogeneity in initial impact across industries.** We first consider how our results would change if the productivity boost from AI adoption was concentrated in specific industries. We focus on the impact on output and inflation and consider four scenarios marked by which sectors are affected. Concretely, we assess the effects of AI raising productivity only in professional services, construction, healthcare or manufacturing. We choose the first industries because their output is used primarily as intermediate inputs, investment goods and consumption services, respectively. We choose the manufacturing industry because of the diversity of its final output destinations. To make the exercises comparable, for each industry we assume that the increase in productivity growth from



AI adoption is equal to 1.5% divided by that industry’s share of total value added. For comparison, we also show the results from the baseline model.<sup>21</sup> All other aspects of the model remain as in the baseline case.

Figure 11 presents the results of these exercises. Output grows irrespective of which industries the effects of AI concentrate on, although there is meaningful cross-industry variation (panel (a)). For example, when the effects of AI affect only professional services or construction, the long run effect on output is smaller than when all industries are equally affected. Conversely, if the effect of AI concentrates on healthcare or manufacturing, the long run effect on output can over one third larger than in the equally-affected case. The result that the increase in GDP is larger when the productivity boost from AI is concentrated in consumer-facing industries mirrors that in the second simple model described in Section 3.3.

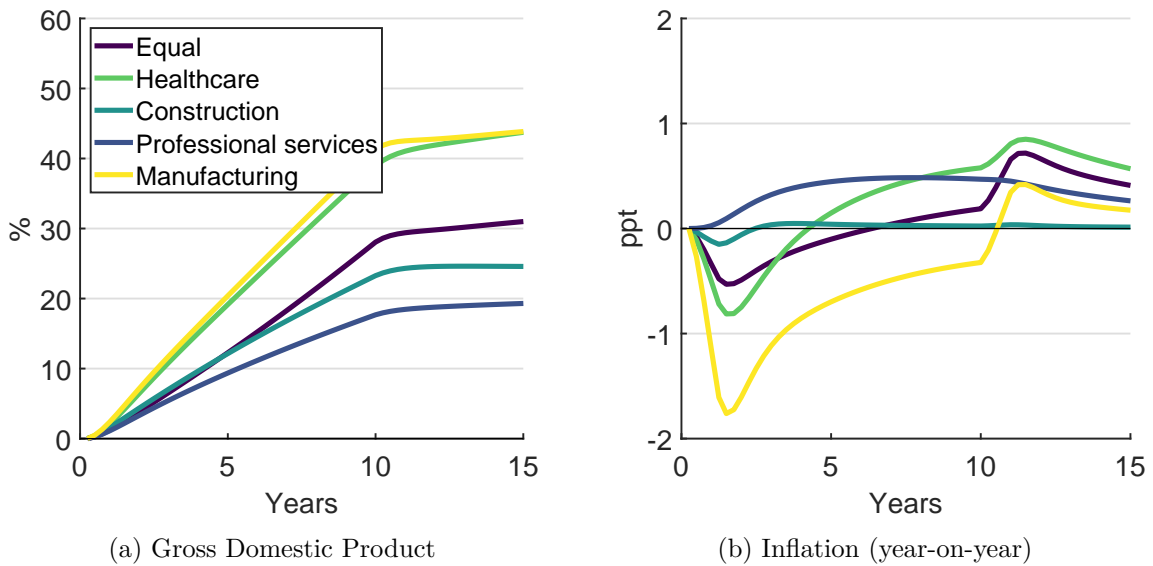


Figure 11: **GDP and Inflation**

Inflation displays considerable heterogeneity across the scenarios. When the effects of AI are concentrated in the professional services industry, there is no initial disinflation at all, even though that scenario features the smallest increase in GDP. When only construction is affected, the long run effect on inflation is small despite an initial mild disinflation. In contrast, when the effect concentrates on manufacturing, there is substantial

<sup>21</sup>The results when we assume that the effects of AI are felt equally in all industries are similar to the baseline.

disinflation early on, and it takes at least a decade for inflation to become positive.

The different inflation responses reflect, in part, different required paths of *relative* prices. Higher productivity in a single industry lowers its relative price. When the productivity boost occurs in an industry that accounts for a large share of consumption output – like healthcare – the relative price adjustment shows up largely in lower consumer price inflation and higher inflation in the investment goods deflator. In contrast, when the productivity boost occurs in an upstream industry, like professional services or construction, inflation in the investment goods deflator tends to fall significantly, while the consumer price inflation decreases only slightly, or even rises.

**General purpose versus factor-specific technology.** We now use our model to study the evolution of output and inflation under two counter-factual scenarios: 1) AI as a labour-specific productivity increase and 2) AI as a capital-specific productivity increase.

Figure 12 presents the results of this exercise, where for reference we also present the baseline scenario of a general purpose technology. The differences between the alternatives considered are smaller here than in the previous counterfactual exercise. To be sure, there are some qualitative differences in terms of the effect on output: the long run impact is smaller when AI is a factor-specific technology (especially so for capital-augmenting TFP). But regardless of how AI enters the production function, the long run impact on output is positive and considerable. For inflation there are even less differences across the alternatives, which all show very similar dynamics. That said, AI as capital-augmenting TFP leads to a smaller long run impact in terms of inflation.

## 6 Conclusion

Considering AI as a general purpose technology that improves productivity growth, our analysis has shown that AI adoption raises aggregate output, consumption and investment. The impact on inflation depends on households' and firms' anticipation of future income growth from AI. In the case of imperfect anticipation, inflation declines in the short run but eventually increases relative to its initial level. In contrast, when income increases due to the AI-induced rise in productivity are fully anticipated today, inflation

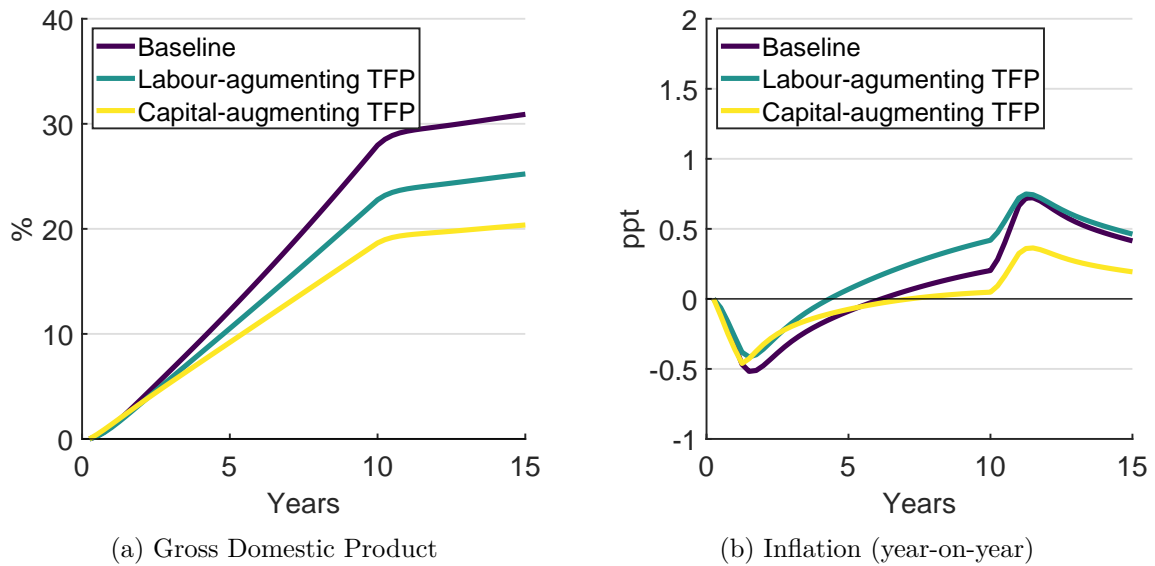


Figure 12: **GDP and Inflation**

increases already in the short run and remains elevated in the longer run.

Our analysis also delivers important insights on the impact of AI on different industries. For one, a sector’s initial exposure to AI is uncorrelated with the ultimate impact of AI on that sector’s output, as indirect effects – through spillovers and linkages – are what matters. However, for aggregate output, which sectors are initially most affected by AI is important. When the effect of AI concentrates on sectors producing consumption rather than investment goods, output can grow by twice as much. Policies that foster the adoption of AI by firms should thus focus on sectors producing consumption goods.

In addition, the adoption of AI could lead to a Goldilocks scenario for monetary policy: greater use of AI could ease inflationary pressures in the near-term, thereby supporting central banks in their task to bring inflation back to target. In the longer term, as inflation rises because of greater AI-induced demand, central banks’ task of controlling inflation would become easier, as they can dampen demand via monetary tightening. However, if households and firms anticipate the boost to growth from AI, policy rates need to rise rapidly already today. More research is hence needed to understand households’ and firms’ expectations about gen AI, and how they differ across subgroups of the population (Aldasoro et al., 2024).

## References

- Acemoglu, Daron**, “The simple macroeconomics of AI,” *Economic Policy*, 2024, *forthcoming*.
- **and Pascual Restrepo**, “Artificial Intelligence, Automation and Work,” *NBER Working Papers*, 2018, (24196).
- , **David Autor, Jon Hazell, and Pascual Restrepo**, “Artificial Intelligence and Jobs: Evidence from Online Vacancies,” *Journal of Labor Economics*, 2022, *40* (S1), S293–S340.
- Aghion, Philippe, Benjamin F Jones, and Charles I Jones**, “Artificial intelligence and economic growth,” in “The economics of artificial intelligence: An agenda,” University of Chicago Press, 2018, pp. 237–282.
- Agrawal, Ajay, Joshua Gans, and Avi Goldfarb**, *The economics of artificial intelligence: an agenda*, University of Chicago Press, 2019.
- Aldasoro, Iñaki, Olivier Armantier, Sebastian Doerr, Leonardo Gambacorta, and Tommaso Oliviero**, “Survey evidence on gen AI and households: job prospects amid trust concerns,” *BIS Bulletin*, 2024, *86*.
- Autor, David**, “The Labor Market Impacts of Technological Change: From Unbridled Enthusiasm to Qualified Optimism to Vast Uncertainty,” *NBER Working Paper*, 2022, (w30074).
- , **Caroline Chin, Anna Salomons, and Bryan Seegmiller**, “New Frontiers: The Origins and Content of New Work, 1940–2018\*,” *The Quarterly Journal of Economics*, 03 2024, *139* (3), 1399–1465.
- , **Frank Levy, and Richard J. Murnane**, “The Skill Content of Recent Technological Change: An Empirical Exploration,” *The Quarterly Journal of Economics*, 2003, *118* (4), 1279–1333.
- Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson**, “Artificial intelligence, firm growth, and product innovation,” *Journal of Financial Economics*, 2024, *151*, 103745.

- Baily, Martin N., Erik Brynjolfsson, and Anton Korinek**, “Machines of Mind: The Case for an AI-Powered Productivity Boom,” *Brookings*, May 2023.
- Baqaee, David and Emmanuel Fahri**, “The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten’s Theorem,” *Econometrica*, 2019, 87 (4), 1155–1203.
- Briggs, Joseph**, “Reconciling Estimates of the Growth Impact of Generative AI,” *Goldman Sachs: Global Economics Comment*, 2024.
- **and Devesh Kodnani**, “Generative AI: Hype, or Truly Transformative?,” *Goldman Sachs: Global Macro Research Report*, 2023, (120).
- Brynjolfsson, Erik and Andrew McAfee**, “The Business of Artificial Intelligence,” *Harvard Business Review*, 2017.
- , **Daniel Li, and Luis R. Raymond**, “Generative AI at Work,” *NBER Working Paper*, 2023, (w31161).
- , **Daniel Rock, and Chad Syverson**, “Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics,” in “The economics of artificial intelligence: An agenda,” University of Chicago Press, 2018, pp. 23–57.
- Cazzaniga, Mauro, Florence Jaumotte, Longji Li, Giovanni Melina, Augustus Panton, Carlo Pizzinelli, Emma Rockall, and Marina M Tavares**, “Gen-AI: Artificial Intelligence and the Future of Work,” *IMF Staff Discussion Notes*, 2024, (2024/001).
- Czarnitzki, Dirk, Gonzalo P. Fernández, and Christian Rammer**, “Artificial Intelligence and Firm-Level Productivity,” *Journal of Economic Behavior & Organization*, 2023, (211), 188–205.
- Evans, Charles L and David A Marshall**, “Fundamental economic shocks and the macroeconomy,” *Journal of Money, Credit and Banking*, 2009, 41 (8), 1515–1555.
- Faia, Ester, Gianmarco Ottaviano, and Saverio Spinella**, “Robot adoption, worker-firm sorting and wage inequality: evidence from administrative panel data,” CEP Discussion Papers dp1902, Centre for Economic Performance, LSE February 2023.
- , **Sebastien Lafitte, Maximilian Mayer, and Gianmarco Ottaviano**, “On the employment consequences of automation and offshoring: A Labor Market Sorting View,” in “Robots and AI: A new economic era,” Routledge, 2022.

- Felten, Edward, Meera Raj, and Robert Seamans**, “The Effect of Artificial Intelligence on Human Labor: An Ability-Based Approach,” 2019, *2019* (1).
- , – , and – , “Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Use,” *Strategic Management Journal*, 2021, *42* (12), 2195–2217.
- Fernald, John and Bing Wang**, “The Recent Rise and Fall of Rapid Productivity Growth,” *FRBSF Economic Letter*, 2015, (4).
- Furman, Jason and Robert Seamans**, “AI and the Economy,” *Innovation policy and the economy*, 2019, *19* (1), 161–191.
- Hulten, Charles R.**, “Growth Accounting With Intermediate Inputs,” *The Review of Economic Studies*, 1978, pp. 511–518.
- Kulish, Mariano and Adrian Pagan**, “Estimation and Solution of Models with Expectations and Structural Changes,” *Journal of Applied Econometrics*, 2017, *32* (2), 255–274.
- Lu, Chia-Hui**, “The impact of artificial intelligence on economic growth and welfare,” *Journal of Macroeconomics*, 2021, *69*, 103342.
- Noy, Shiri and Wei Zhang**, “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence,” *Science*, 2023, *381* (6654), 187–192.
- Peng, Siyuan, Eirini Kalliamvakou, Peter Cihon, and Mehmet Demirer**, “The Impact of AI on Developer Productivity: Evidence from Github Copilot,” *arXiv preprint*, 2023.
- Pizzinelli, Carlo, Augustus Panton, Marina Mendes Tavares, Mauro Cazzaniga, and Longji Li**, “Labor Market Exposure to AI: Cross-country Differences and Distributional Implications,” *IMF Working Papers*, 2023, (2023/2016).
- Rees, David**, “What Comes Next?,” *BIS Working Paper*, 2020, (898).
- Yang, Chun-Hung**, “How Artificial Intelligence Technology Affects Productivity and Employment: Firm-Level Evidence from Taiwan,” *Research Policy*, 2022, (51.6).

# Appendix

# A Additional figures

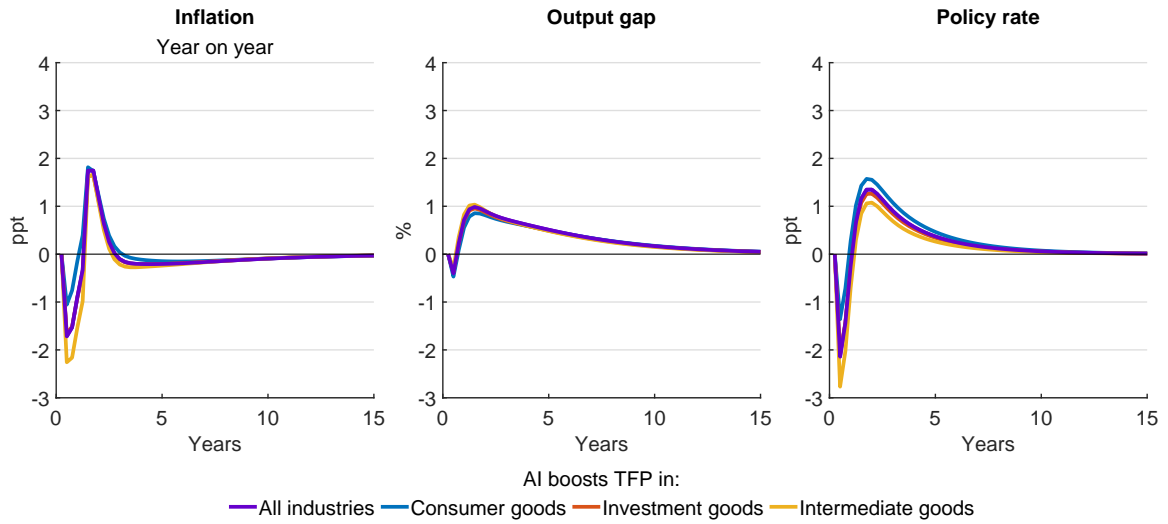


Figure 13: Simple model 1

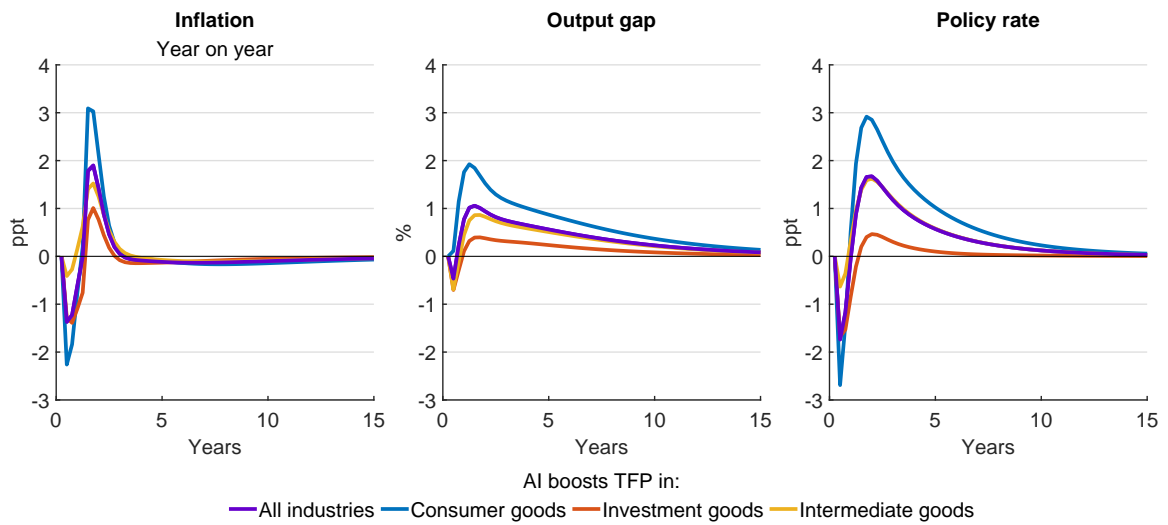


Figure 14: Simple model 2



## B Decomposition of industry value added

This appendix describes the decomposition of industry value added into the contributions of final and intermediate demand and income and substitution effects described in Section 4.2.

An industry's gross output can be expressed as:

$$\underbrace{y_{j,t}}_{\text{Gross output}} = \underbrace{c_{j,t} + i_{j,t} + g_{j,t}}_{\text{Final demand}} + \underbrace{\sum_{k=1}^{\mathcal{F}} x_{k,j,t}}_{\text{Intermediate demand}} \quad (\text{B.1})$$

Demand for an industry's output as a final or intermediate good are given by:

$$c_{j,t} = \omega_{c,j} (\gamma_{j,t})^{-\eta} C_t \quad (\text{B.2})$$

$$i_{j,t} = \omega_{i,j} \left( \frac{\gamma_{j,t}}{\gamma_{I,t}} \right)^{-\eta} I_t \quad (\text{B.3})$$

$$g_{j,t} = \omega_{g,j} \left( \frac{\gamma_{j,t}}{\gamma_{G,t}} \right)^{-\eta} G_t \quad (\text{B.4})$$

$$x_{k,j,t} = \omega_{x,k,j} \left( \frac{\gamma_{j,t}}{\gamma_{X,k,t}} \right)^{-\psi} X_{k,t} \quad (\text{B.5})$$

In the model's initial steady state all relative prices are equal to 1. We define final demand income effects,  $FDIE_t$  as any change in the final demand for an industry's output from its initial steady state that is unrelated to changes in relative prices. This is given by:

$$FDIE_t = \omega_{c,j}(C_t - C_0) + \omega_{i,j}(I_t - I_0) + \omega_{g,j}(G_t - G_0) \quad (\text{B.6})$$

where  $C_0$ ,  $I_0$  and  $G_0$  are aggregate consumption, investment and government expenditure in the model's initial steady state.

We define intermediate demand income effects,  $IDIE_t$ , equivalently as:

$$IDIE_t = \sum_{k=1}^{\mathcal{F}} \omega_{k,j} (X_{k,t} - X_{k,0}) \quad (\text{B.7})$$

where  $X_{k,0}$  is the total demand for intermediates in industry  $k$  in the model's initial steady state. We define final demand substitution effects,  $FDSE_t$ , as any change in final

demand due to changes in relative prices. Because any change in industry-level demand that is not due to changes in aggregate demand must be due to price changes, this is given by:

$$FDSE_t = c_{j,t} + i_{j,t} + g_{j,t} - FDIE_t \quad (\text{B.8})$$

and equivalently, intermediate demand substitution effects,  $IDSE_t$  are given by:

$$IDSE_t = \sum_{k=1}^{\mathcal{F}} x_{j,t} - IDIE_t \quad (\text{B.9})$$

The above calculations decompose gross output. To decompose value-added output, we first calculate the total value of intermediate inputs used in an industry:

$$x_{j,t} = \sum_{k=1}^{\mathcal{F}} x_{j,k,t} \quad (\text{B.10})$$

Denoting the value added decompositions by the superscript  $VA$ , the equivalent calculations are:

$$FDIE_t^{VA} = FDIE_t (1 - x_{j,t}/y_{j,t}) \quad (\text{B.11})$$

$$IDIE_t^{VA} = IDIE_t (1 - x_{j,t}/y_{j,t}) \quad (\text{B.12})$$

$$FDSE_t^{VA} = FDSE_t (1 - x_{j,t}/y_{j,t}) \quad (\text{B.13})$$

$$IDSE_t^{VA} = IDSE_t (1 - x_{j,t}/y_{j,t}) \quad (\text{B.14})$$