

Macroeconomic Policies for AI

Luca Fornaro and Martin Wolf*

April 2026

Abstract

We provide a macroeconomic framework to study monetary and fiscal policies for AI. Advances in AI expand firms' ability to automate production. While higher automation boosts productivity and potential output, it also reduces workers' share of income. Since workers have a high propensity to consume, advances in AI may depress aggregate demand and lead to a slump. Expansionary monetary policy can convert an AI slump into an AI boom, but in doing so it faces two challenges. In the short run, AI worsens the inflation-employment trade off faced by the central bank. In the medium run, monetary policy may be constrained by the zero lower bound, since weak demand lowers the natural rate. Employment subsidies and cuts in labor taxes can usefully complement monetary policy, by reducing firms' cost of labor and inflation, as well as supporting workers' income and aggregate demand.

JEL Codes: E32, E43, E52, O31, O42

Keywords: monetary policy, automation, AI, inflation, liquidity traps, endogenous productivity, wages, artificial intelligence

*Luca Fornaro: CREI, Universitat Pompeu Fabra, Barcelona School of Economics and CEPR; LFornaro@crei.cat. Martin Wolf: University of St. Gallen and CEPR; martin.wolf@unisg.ch. Luca Fornaro acknowledges financial support from the European Research Council through the Consolidator Grant MACROGROWTH-101169982, from the Spanish Ministry of Economy and Competitiveness through the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S), and from the Generalitat de Catalunya, through CERCA and SGR Programme (2021-SGR-01599).

1 Introduction

As we write, artificial intelligence (AI) is transforming our economies and our way of life. Some observers argue that AI may trigger the next industrial revolution, with a similar or even larger potential to improve our standards of living than the industrial revolutions from the past (Smith, 2023). At the same time, there are large concerns about the negative side effects of this new technology. In particular, some worry that AI may lead to technological unemployment and rising inequality by triggering a wave of automation (Acemoglu et al., 2026). It therefore appears crucial to couple advances in AI with an appropriate macroeconomic policy mix, to ensure that AI not only boosts output and productivity, but that it also benefits workers and leads to shared prosperity.

We contribute to this effort by providing a tractable macroeconomic framework to study monetary and fiscal policies for AI. Our model has three key features. First, as in Acemoglu and Restrepo (2018), firms can choose whether to perform some production tasks with labor or automate them with capital. We model advances in AI as a rise in the share of tasks that can be automated.¹ This shock increases productivity and potential output, but also opens up the possibility that capital may displace labor in production. Second, as in Mian et al. (2021), households have non-homothetic preferences, so that the income distribution matters for aggregate demand. In particular, our economy is inhabited by workers and capitalists. As it is the case empirically, workers have a higher propensity to consume compared to capitalists. Third, the presence of nominal rigidities implies that output can deviate from its potential level and monetary policy has real effects.

We start by studying a scenario in which macroeconomic policies do not react to the rise in AI, and show that the result may be a demand-driven slump. The slump materializes because higher automation causes a drop in the labor share, redistributing income from high-spending workers to low-spending capitalists. Absent expansionary macroeconomic policies, the result is a drop in aggregate demand for consumption. Moreover, the prospect of weak demand and low sales reduces firms' investment. Hence, while the rise in automation boosts productivity, its negative impact on aggregate demand ends up depressing output.

What are the other symptoms of this AI-driven slump? First, since productivity rises while output declines, the slump is characterized by high unemployment. Second, the slump affects workers and capitalists very differently. Workers' income takes a double hit, as it falls both because output drops, and the labor share declines. Capitalists' income, instead, is supported by the rise in the capital share. Since capitalists are the richest part of the population, the slump is associated with a rise in inequality.

The root cause of an automation-driven slump is that firms' choice to automate production triggers aggregate demand externalities. That is, firms automate their production to cut costs, so as to increase sales and profits. But each individual firm does not take into account that

¹This is by no means the only channel through which AI may affect our economies. For instance, AI could boost scientific discoveries and technological advances (Davidson et al., 2026). But the automation effect of AI is particularly interesting from a macroeconomic perspective, and that is why we focus on it.

replacing workers with capital depresses aggregate demand, lowering the sales and profits of all other firms. When this general equilibrium effect is sufficiently strong, a paradox of productivity materializes: while automating production to increase productivity is rational from the perspective of an individual firm, the collective adoption of automation technologies decreases firms' sales and profits.

We next consider a central bank that reacts to the rise in AI by setting monetary policy to maintain the economy at full employment. In this case the adoption of AI triggers an output boom. In the short run, the boom is sustained by an increase in firms' investment, which is needed to fully exploit the productivity gains offered by the new automation technologies. Over time, also consumption demand rises. In fact, in the medium run workers' consumption increases because higher production more than compensates the drop in the labor share. There are, however, two challenges that monetary policy faces when trying to sustain full employment during a rise in automation.

First, in the short run higher automation is inflationary. More precisely, the initial phases of the rise in AI look like a cost-push shock, i.e. an adverse shift of the Phillips curve increasing the inflation rate consistent with any level of employment. This finding may come as a surprise, because higher automation boosts productivity, and rising productivity is often seen as a deflationary force. What this logic misses is that higher automation reduces firms' demand for labor. To support full employment, therefore, real wages have to fall. But since nominal wages are rigid, the drop in real wages has to happen through a rise in prices. For this reason, during the first phases of a rise in automation the central bank faces a trade-off between containing inflation and supporting the labor market, and the economy may face an unusual combination of rapid productivity growth, weak labor market and high inflation.

The second challenge for monetary policy comes from the path of the natural interest rate, that is the interest rate consistent with full employment. In the short run, high investment boosts the natural interest rate. In the medium run, however, low consumption demand by workers depresses the natural rate.² If this effect is sufficiently strong, the natural rate goes into negative territory and monetary policy becomes constrained by the zero lower bound. Once this happens, the economy enters an automation-driven liquidity trap, in which weak demand depresses output, employment and investment.

In the last part of the paper, we ask how fiscal policy can help monetary policy. While the array of potential fiscal interventions is large, we focus our attention on employment subsidies, or equivalently on cuts in labor taxes. In our context, employment subsidies are a particularly promising tool because they contain inflation by reducing firms' labor cost, and support demand by boosting workers' income. In fact, we show that, during periods of rapid automation, employment subsidies and cuts in labor taxes can sustain employment and output, while at the same time keeping inflation closer to its target.

²Our model can thus explain the empirical correlation between labor share and the natural interest rate highlighted by [Cho and Williams \(2026\)](#).

Before moving on, let us clarify that we are not predicting that AI will necessarily lead to a slump, or even that this is the most likely outcome. Rather, we argue that macroeconomic policies may determine whether we will experience an AI slump or an AI boom. Moreover, our analysis shows that monetary policy, on its own, may have a hard time converting an AI slump into a boom. Accomplishing this task is likely to require a careful mix of monetary and fiscal policies. In particular, employment subsidies or cuts in labor taxes could be a useful tool to contain inflation and support demand over the coming years.

Related literature. Our paper contributes to the debate on the macroeconomic impact of AI. While we can't do justice to this rapidly-growing literature, some interesting contributions are [Acemoglu \(2025\)](#), [Aghion and Bunel \(2024\)](#), [Acemoglu et al. \(2026\)](#), [Davidson et al. \(2026\)](#) and [Jones and Tonetti \(2026\)](#). These works focus on the impact of AI on the supply-side of the economy. Instead, we argue that the macroeconomic effects of AI will be shaped not only by supply factors, but also by the demand side of the economy. This insight implies that macroeconomic policies are going to be a key determinant of how AI will affect our societies.

We focus on the automation effect of AI. In doing so, we build on our previous work [Fornaro and Wolf \(2021\)](#). There we studied monetary policy in an economy in which, in the spirit of [Acemoglu and Restrepo \(2018\)](#), firms' choice to automate production is endogenous. We had already argued in [Fornaro and Wolf \(2021\)](#) that higher automation may lead to technological unemployment due to weak demand. Here we show that AI may lead to a slump by redistributing income from workers to capitalists.³ Variants of this idea have also been explored by [Imas \(2026\)](#) and [Lim \(2026\)](#), while [Moll et al. \(2021\)](#) and [Caballero \(2026\)](#) study other channels through which automation technologies affect the macroeconomy through their impact on the income distribution. Moreover, in this paper we show that the adoption of AI may lead to an unusual combination of high productivity, weak labor market and high inflation. This connects our paper to the literature on automation and inflation ([Basso and Rachedi, 2025](#); [Fueki et al., 2023](#)). Finally, in a recent contribution [Lenzu \(2026\)](#) provides an interesting analysis of some of the implications of AI for monetary policy.

More broadly, this paper is related to an expanding literature studying the medium-run impact of monetary and fiscal policy, in economies in which technology and productivity are endogenous. Some examples of this literature are [Benigno and Fornaro \(2018\)](#), [Moran and Queralto \(2018\)](#), [Anzoategui et al. \(2019\)](#), [Garga and Singh \(2020\)](#), [Schmöller and Spitzer \(2021\)](#), [Fornaro and Wolf \(2023\)](#), [Fornaro \(2025\)](#), [Fornaro and Wolf \(2025\)](#) and [Fornaro et al. \(2025\)](#).⁴ We build on this literature to study macroeconomic policies in times of structural change driven by rapid adoption of AI technologies.

³Our paper is thus connected to the vast literature on monetary policy with heterogeneous agents ([Kaplan et al., 2018](#); [Mian et al., 2021](#); [Debortoli and Galí, 2024](#); [Bilbiie, 2025](#)).

⁴See [Cerra et al. \(2021\)](#) for an excellent review of this literature, while [Fieldhouse and Mertens \(2023\)](#), [Antolin-Diaz and Surico \(2025\)](#), [Jordà et al. \(2024\)](#) and [Ma and Zimmermann \(2023\)](#) provide consistent empirical evidence.

2 Model

This section lays down our baseline model. The economy has three key elements. First, as in [Acemoglu and Restrepo \(2018\)](#), firms can choose whether to automate some production tasks. This decision determines the intensity with which labor and capital are used in production. Second, as in [Mian et al. \(2021\)](#), households have non-homothetic preferences, so that the income distribution matters for aggregate demand. Third, the presence of nominal rigidities implies that output can deviate from its potential level.

Consider an infinite-horizon closed economy. Time is discrete and indexed by $t \in \{0, 1, 2, \dots\}$. The economy is inhabited by households, firms, and by a government that sets macroeconomic policies. For simplicity, we assume perfect foresight.

2.1 Households

There is a continuum of measure one of households deriving utility from consumption of a homogeneous final good. There are two types of households: a mass $1 - \chi$ of workers and a mass χ of capitalists. Within each class, households are identical.

Workers. Workers' utility is given by

$$\sum_{t=0}^{\infty} \beta^t \left(\log c_t^w - \frac{1}{\varphi} \left(\frac{l_t}{\bar{l}} \right)^\varphi \right),$$

where c_t^w and l_t denote respectively consumption and labor supplied for each individual worker, $\beta < 1$ is the discount factor and $\varphi > 1$ determines the labor supply elasticity. Workers do not participate in financial markets, and every period fully consume their labor income

$$P_t c_t^w = W_t l_t,$$

where P_t denotes the nominal price of the consumption good and W_t the nominal wage.

In a frictionless labor market, equilibrium aggregate employment $L_t = (1 - \chi)l_t$ would be determined by households' desired labor supply according to

$$L_t = (1 - \chi) \left(\frac{W_t \bar{l}^\varphi}{P_t c_t^w} \right)^{\frac{1}{\varphi-1}} = \bar{L},$$

where the second equality makes use of workers' budget constraint and we have defined $\bar{L} \equiv (1 - \chi)\bar{l}$. However, due to the presence of nominal wage rigidities to be described below, actual employment L_t may deviate from \bar{L} .

Capitalists. Capitalists own the capital stock and trade in bonds, but have no labor income. More precisely, capitalists can trade in one-period nominal bonds b_t , paying the nominal interest rate i_t . Capitalists can also invest in physical capital k_t , which pays a real return r_t^k and depreciates at rate δ . We denote the real price of capital by q_t .

We think of capitalists as the top percentiles of the income distribution. Following [Mian et al. \(2021\)](#), we assume that these rich households enjoy utility not only from consuming c_t^k , but also from holding wealth $b_{t+1}/P_t + q_t k_{t+1}$. More precisely, the lifetime utility of each capitalist is given by

$$\sum_{t=0}^{\infty} \beta^t \left(\log c_t^k + \frac{\phi}{1-\xi} \left(\frac{b_{t+1}}{P_t} + q_t k_{t+1} \right)^{1-\xi} \right).$$

The parameters $0 \leq \xi < 1$ and $\phi > 0$ determine capitalists' preference for accumulating wealth. We assume that holding wealth is a luxury, which implies that capitalists' marginal propensity to save rises with their income. [Mian et al. \(2020\)](#) show that this class of utility function describes well the saving behavior of the top 1% of the U.S. population.

The problem of the representative capitalist consists in choosing c_t^k , b_{t+1} and k_{t+1} to maximize expected utility, subject to a no-Ponzi constraint and the budget constraint

$$P_t c_t^k + b_{t+1} + P_t q_t k_{t+1} = P_t (r_t^k + (1-\delta)q_t) k_t + b_t (1+i_{t-1}),$$

where P_t is the nominal price of the final good, b_{t+1} is the stock of bonds purchased by each capitalist in period t , and $b_t(1+i_{t-1})$ is the payment received from past investment in bonds. k_{t+1} is the stock of capital held by each capitalist at the end of period t , and used in production in period $t+1$.

Capitalists have two decisions to make. First, their optimal saving decision is described by

$$\frac{1}{c_t^k P_t} = \frac{\beta(1+i_t)}{P_{t+1} c_{t+1}^k} + \frac{\phi}{P_t} \left(\frac{b_{t+1}}{P_t} + q_t k_{t+1} \right)^{-\xi}. \quad (1)$$

This is just a standard Euler equation, except for the additional incentive to save caused by the presence of wealth in utility, captured by the last term on the right-hand side. This term reduces the sensitivity of current consumption to future interest rates, so it introduces additional discounting in the Euler equation ([Michaillat and Saez, 2021](#)). Second, households need to allocate their savings between bonds and capital. No arbitrage between these two assets implies

$$\frac{(1+i_t)P_t}{P_{t+1}} = \frac{r_{t+1}^k + (1-\delta)q_{t+1}}{q_t}. \quad (2)$$

Finally, households' optimal saving behavior obeys the transversality condition⁵

$$\lim_{T \rightarrow \infty} \beta^T \frac{b_{t+T+1}/P_{t+T+1} + q_{t+T+1} k_{t+T+1}}{c_{t+T}^k} = 0.$$

2.2 Final good production

The final good is produced by competitive firms using a continuum of measure one of intermediate inputs, or tasks, $y_{j,t}$, indexed by $j \in [0, 1]$. Denoting by Y_t the output of the final good, the

⁵Since $\beta < 1$, the transversality condition implies that in the long run the wealth to consumption ratio does not tend to infinity. This optimality condition follows from the assumption that wealth is a luxury.

production function is

$$\log Y_t = \int_0^1 \log y_{j,t} dj.$$

Profit maximization implies the demand functions

$$p_{j,t} y_{j,t} = Y_t,$$

where $p_{j,t}$ is the price of intermediate input j in terms of the final good.

2.3 Intermediate inputs production

Intermediate inputs are produced by competitive firms, and are heterogeneous in their production technologies. Following [Acemoglu and Restrepo \(2018\)](#), we model technological constraints on automation by assuming that production tasks indexed by $j > \bar{J}_t$ can be performed only with labor, while the remaining tasks $j \leq \bar{J}_t$ may be performed with capital. An increase in \bar{J}_t thus captures technological progress expanding the potential for automation in the production process.

More precisely, the production function of a generic intermediate input j is

$$y_{j,t} = \gamma_j^k k_{j,t} + \gamma_j^l l_{j,t},$$

where $k_{j,t}$ is the capital used to perform task j , $l_{j,t}$ denotes labor employed in task j , and γ_j^k and γ_j^l denote respectively the productivity of capital and labor in task j .

Intermediates indexed by $j > \bar{J}_t$ can be produced with labor only, and are characterized by $\gamma_j^k = 0$ and $\gamma_j^l = \gamma^l > 0$. These are the production tasks for which automation is not available due to technological constraints. The remaining intermediates $j \leq \bar{J}_t$ can be produced with capital or labor, and they are characterized by $\gamma_j^k = \gamma^k$ and $\gamma_j^l = \gamma^l$.

Perfect competition implies that the price of intermediate inputs is equal to their marginal cost, so that

$$p_{j,t} = \begin{cases} \min\left(\frac{r_t^k}{\gamma^k}, \frac{w_t}{\gamma^l}\right) & \text{if } j \leq \bar{J}_t \\ \frac{w_t}{\gamma^l} & \text{if } j > \bar{J}_t, \end{cases}$$

where $w_t \equiv W_t/P_t$. To produce intermediates $j < \bar{J}_t$ firms employ the cheapest (productivity-adjusted) factor of production.

Now define $J_t^* \leq \bar{J}_t$ such that all the intermediate inputs with $j \leq J_t^*$ are produced with capital, while the rest are produced with labor. To make things simple, throughout the paper we assume that it is always profitable for firms to automate production as much as possible, so that $J_t^* = \bar{J}_t$ for all t . This happens when $r_t^k/\gamma^k < w_t/\gamma^l$, so when capital is cheap compared to labor.

2.4 Aggregate production function

A useful property of this model is that aggregate output can be written as

$$Y_t = \left(\frac{\gamma^k K_t}{\bar{J}_t} \right)^{\bar{J}_t} \left(\frac{\gamma^l L_t}{1 - \bar{J}_t} \right)^{1 - \bar{J}_t}, \quad (3)$$

where $K_t = \chi k_t$ denotes the aggregate capital stock. This is a standard Cobb-Douglas production function in capital and labor, with a twist. The twist is that the intensity with which labor and capital are used in production depends on the technological cap on automation \bar{J}_t . That said, the model preserves all the usual properties of Cobb-Douglas production functions. Hence, equilibrium prices satisfy the condition

$$1 = \left(\frac{r_t^k}{\gamma^k} \right)^{\bar{J}_t} \left(\frac{w_t}{\gamma^l} \right)^{1 - \bar{J}_t}.$$

Moreover, the capital and labor share are respectively given by

$$\frac{r_t^k K_t}{Y_t} = \bar{J}_t \quad (4)$$

$$\frac{w_t L_t}{Y_t} = 1 - \bar{J}_t. \quad (5)$$

Hence, factors' shares depend on the intensity with which automation is used in production. A higher use of automation, that is a higher \bar{J}_t , is associated with a higher capital share.

2.5 Capital production

There is a large number of competitive capital-producing firms. Producing one unit of capital requires

$$1 + \frac{\kappa}{2} \left(\frac{I_t}{I_{t-1}} - 1 \right)^2,$$

units of the final good, where $\kappa \geq 0$ and I_t denotes aggregate investment. When $\kappa > 0$ investment is thus subject to adjustment costs.⁶ As is usual, investment adds to the aggregate capital stock according to

$$K_{t+1} = (1 - \delta)K_t + I_t. \quad (6)$$

Zero profits in the capital-producing sector then implies

$$q_t = 1 + \frac{\kappa}{2} \left(\frac{I_t}{I_{t-1}} - 1 \right)^2. \quad (7)$$

This is a classic condition relating investment to Tobin's Q.

⁶We thus assume that investment adjustment costs are external to the firm, perhaps due to congestion externalities in the capital production process. This assumption simplifies the expressions, but does not affect materially any of the results.

2.6 Wages, prices and monetary policy

We consider an economy with frictions in the adjustment of nominal wages. The presence of nominal wage rigidities plays two roles in the model. First, it creates the possibility that employment may deviate from households' labor supply. Second, it opens the door to real effects of monetary policy interventions. To keep the analysis simple, we assume that wage dynamics are governed by the following reduced-form wage Phillips curve

$$\frac{W_t}{W_{t-1}} = \psi(L_t),$$

where $\psi(\cdot)$ is an increasing function satisfying $\psi(\bar{L}) = 1$.

Monetary policy controls the nominal rate i_t . Because of wage stickiness, movements in the nominal rate affect the real interest rate and other real variables. To see this point, notice that the nominal price of the final good can be written as

$$P_t = \left(\frac{r_t^k}{\gamma^k} \right)^{\frac{\bar{J}_t}{1-\bar{J}_t}} \frac{W_t}{\gamma^l}.$$

This equation implies that the nominal wage rigidity is partly inherited by the price of the final good. Therefore, as it is standard, by setting i_t the central bank can affect the real rate r_t , defined as

$$1 + r_t \equiv (1 + i_t) \frac{P_{t+1}}{P_t}.$$

Throughout the paper, we will consider different assumptions about how monetary policy is set. Moreover, to streamline the analysis, we treat monetary policy as controlling the real interest rate directly, as for example in [Werning \(2015\)](#) and [Mian et al. \(2021\)](#).

2.7 Market clearing and definition of equilibrium

Since all capitalists are identical, equilibrium on the bonds market requires $b_t = 0$. Market clearing for the final good then implies

$$Y_t = (1 - \bar{J}_t)Y_t + \chi c_t^k + I_t \left(1 + \frac{\kappa}{2} \left(\frac{I_t}{I_{t-1}} - 1 \right)^2 \right), \quad (8)$$

The output of the final good is thus either consumed or invested in capital. To derive this expression, we have used the fact that workers fully consume their labor income, which corresponds to a share $1 - \bar{J}_t$ of output.

Turning to the labor market, in absence of nominal rigidities equilibrium employment would satisfy $L_t = \bar{L}$. Hence, we can think of \bar{L} as the natural level of employment. Therefore, when $L_t < \bar{L}$ there is involuntary unemployment and output is below potential. When $L_t > \bar{L}$, instead, the labor market is overheated and output is above potential.

We are now ready to define an equilibrium of the baseline model.

Definition 1 *An equilibrium of the baseline model is a path of real allocations $\{c_t^k, L_t, K_{t+1}, Y_t, I_t\}_t$ and prices $\{w_t, r_t^k, q_t\}_t$, satisfying (1)-(8), given a path for the technological cap on automation $\{\bar{J}_t\}_t$, the - policy determined - real rate $\{r_t\}_t$ and initial values for the capital stock and for investment, $K_0 > 0$ and $I_{-1} > 0$.*

3 The aggregate demand effect of AI

There is substantial uncertainty about how advances in AI technologies will affect our economies. Throughout this paper, we will focus on the automation effect of AI.⁷ We do so because early signals suggest that firms are using AI to automate their production processes, by expanding the number of tasks performed by capital as opposed to labor (Acemoglu et al., 2026). This fact has sparked a lively debate about whether advances in AI will bring faster growth and shared prosperity, or instead displace labor from production and lead to technological unemployment and more inequality (Acemoglu, 2025).

We revisit this debate with the help of our model, which allows for the possibility of involuntary unemployment driven by weak demand. Our first objective is to show that there are plausible scenarios in which advances in automation technologies may lead to an economic slump. The reason is that higher automation expands potential output, but actual output depends on aggregate demand. And rising automation - by redistributing income from high-spending workers to low-spending capitalists - may very well depress aggregate demand. Absent other forces sustaining demand, such as accommodative macroeconomic policies, advances in automation technologies may thus end up reducing economic activity.

3.1 The making of an AI slump

We model the rise in automation triggered by AI as an increase in \bar{J}_t , that is of the number of tasks that capital performs. We consider a simple scenario. Initially, the economy is on a full employment steady state with $\bar{J}_t = \bar{J}_l$. In period $t = 0$, \bar{J}_t permanently increases to a higher value $\bar{J}_h > \bar{J}_l$. This shock is unexpected, but from period 0 on agents perfectly anticipate the path of \bar{J}_t . The outcome of this shock is to move the economy to a new steady state with higher automation.

To close the model, we have to specify a monetary policy. Throughout this section, we assume that the central bank holds the real interest rate fixed and equal to its value in the initial steady state, i.e. $r_t = r$ for all t . This scenario is a useful benchmark, because it illustrates how rising automation may affect the economy absent a macroeconomic policy response. Later on, we will study alternative policies.

A fixed-investment economy. We start by studying a special case of the model, that allows us to illustrate transparently our key results. Assume that investment adjustment costs are extremely large ($\kappa \rightarrow +\infty$) and that $I_{-1} = \delta \bar{K}$. The aggregate capital stock is then constant

⁷To be clear, this is by no means the only channel through which AI can affect labor markets and the macroeconomy. For instance, AI technologies could complement labor, or foster technological progress.

and equal to \bar{K} , and the automation shock induces an immediate jump of the economy from the initial to the final steady state. We will denote with the subscripts l and h the value of a variable respectively in the initial (low-automation $\bar{J}_t = \bar{J}_l$) and in the final (high-automation $\bar{J}_t = \bar{J}_h$) steady states.

Let us first derive the path of potential, or natural, output Y_t^p . This is the value of output consistent with full employment ($L_t = \bar{L}$), and it is given by

$$Y_t^p = \left(\frac{\gamma^k \bar{K}}{\bar{J}_t} \right)^{\bar{J}_t} \left(\frac{\gamma^l \bar{L}}{1 - \bar{J}_t} \right)^{1 - \bar{J}_t}.$$

Since we are assuming that $r_t^k/\gamma^k < w_t/\gamma^l$, advances in automation technologies boost potential output and so $Y_t^p < Y_h^p$.⁸ But what about actual output?

To answer this question, we have to track the impact of higher automation on aggregate demand. With infinitely high capital-adjustment costs, investment is just enough to cover for capital depreciation ($I_t = \delta \bar{K}$). Moreover, we have already seen that workers consume a fraction $1 - \bar{J}_t$ of output. Finally, consumption demand by capitalists can be obtained by rewriting equation (1) as⁹

$$\chi c_t^k = \frac{1 - \beta(1 + r)}{\phi} \chi^{1 - \xi} (q_t \bar{K})^\xi.$$

A lower interest rate thus boosts capitalists' consumption, because it reduces their incentives to save. Moreover, capitalists' demand for consumption is increasing in their wealth, captured by the value of the capital stock $q_t \bar{K}$. Equilibrium on the capital market implies that

$$q_t \bar{K} = \frac{\bar{J}_t Y_t}{r + \delta}.$$

The value of the capital stock is thus given by the fraction of output going to capital $\bar{J}_t Y_t$, discounted by the sum of the interest and the depreciation rate.

Using these results, we can define the level of aggregate demand Y_t^d as

$$Y_t^d = (1 - \bar{J}_t) Y_t + \frac{1 - \beta(1 + r)}{\phi} \chi^{1 - \xi} \left(\frac{\bar{J}_t Y_t}{r + \delta} \right)^\xi + \delta \bar{K}. \quad (\text{AD})$$

In equilibrium $Y_t = Y_t^d$, and hence

$$Y_t = \frac{1}{\bar{J}_t} \left(\left(\frac{1 - \beta(1 + r)}{\phi} \right) \chi^{1 - \xi} \left(\frac{\bar{J}_t Y_t}{r + \delta} \right)^\xi + \delta \bar{K} \right). \quad (9)$$

This expression shows that the increase in \bar{J}_t triggers two contrasting effects on aggregate demand and output. On the one hand, a higher capital share reduces the fraction of national income that accrues to workers. Since workers have a high propensity to spend, this amounts to a negative

⁸Consider that $r_t^k/\gamma^k < w_t/\gamma^l$ implies that $r_t^k < \gamma^k$. In turn, this condition implies that $\bar{J}_t \gamma^l \bar{L} < (1 - \bar{J}_t) \gamma^k \bar{K}$. This condition is necessary and sufficient to ensure that Y_t^p is increasing in \bar{J}_t .

⁹Recall that in this special case the economy is always on a steady state.




Figure 1: Rise in automation and the Keynesian cross.

demand shock. On the other hand, advances in automation technologies increase the value of the capital stock and capitalists' wealth, boosting their demand for consumption. The first effect, however, always dominates the second. Therefore, higher automation lowers aggregate demand and equilibrium output. In fact, it is easy to see that $Y_h/Y_l = \bar{J}_l/\bar{J}_h < 1$.¹⁰

Figure 1 illustrates this argument using a Keynesian cross diagram, which plots the AD schedule against the market clearing condition $Y = Y^d$. In the initial low-automation steady state aggregate demand is just enough to maintain the economy at potential and so $Y_l = Y_l^p$. The rise in automation lowers aggregate demand by redistributing income from workers, which have a high propensity to spend, to capitalists. Graphically, this is captured by a downward shift of the AD schedule. As a result, even though potential output rises from Y_l^p to Y_h^p , actual output drops below its initial value ($Y_h < Y_l^p$).

This insight is simple, but powerful. While the automation effect of AI is likely to boost potential output, it is also likely to depress aggregate demand by reducing the share of income going to high-spending workers. Absent other contrasting forces, such as expansionary macroeconomic policies, advances in AI may thus lead to an economic slump.

What are other symptoms of this AI-driven slump? The first one is involuntary unemployment. In fact, equilibrium employment is given by

$$L_t = \bar{L} \left(\frac{Y_t}{Y_t^p} \right)^{\frac{1}{1-\bar{J}_t}}.$$

The shortfall of actual output below its potential value, therefore, implies that employment falls below workers' desired labor supply. This result is intuitive, given that higher automation reduces firms' demand for labor, and captures the notion that the spread of AI may lead to technological unemployment.

¹⁰Equation (9) implies that $\bar{J}_t Y_t$ is independent of \bar{J}_t , and hence that $\bar{J}_l Y_l = \bar{J}_h Y_h$. Rearranging and using that $\bar{J}_l < \bar{J}_h$ yields the result.

Moreover, a distinctive feature of an automation-driven slump is that the burden of adjustment falls on workers. Workers' consumption takes a double hit, as it drops both because output is depressed, and because the labor share of income has declined. Capitalists' consumption, instead, is less affected. This happens because the negative impact of lower output is mitigated by the rise in the capital share.¹¹ Since capitalists are the richest part of the population, an automation-driven slump is therefore associated with an increase in inequality.

Investment and capital accumulation. We now bring investment and capital accumulation into the picture. To make things simple, suppose that there are no investment adjustment costs ($\kappa = 0$). In this case, under a constant interest rate policy, the capital stock evolves according to

$$K_{t+1} = \frac{\bar{J}_{t+1} Y_{t+1}}{r + \delta}.$$

Higher automation, i.e. an increase in \bar{J}_{t+1} , naturally induces firms to boost investment and the capital stock. But investment also depends on expectations of future output, because the return to investment is increasing in future demand Y_{t+1} . It turns out that under the simple policy that we study investment does not react to the automation shock, because $\bar{J}_{t+1} Y_{t+1}$ remains constant and so the two effects exactly balance out.¹² Besides depressing employment, weak aggregate demand thus ends up also depressing firms' investment.

This also matters for potential output. In the economy with variable investment, potential output is endogenous as it depends on the capital stock and therefore on past investment decisions. In fact, in any period $t > 0$, the potential level of output is given by

$$Y_t^p = \left(\frac{\gamma^k K_t}{\bar{J}_t} \right)^{\bar{J}_t} \left(\frac{\gamma^l \bar{L}}{1 - \bar{J}_t} \right)^{1 - \bar{J}_t} = \left(\frac{\gamma^k Y_t}{r + \delta} \right)^{\bar{J}_t} \left(\frac{\gamma^l \bar{L}}{1 - \bar{J}_t} \right)^{1 - \bar{J}_t}.$$

Through its negative impact on investment, low aggregate demand reduces potential output. The implication is that a narrow focus on unemployment may underestimate substantially the output losses due to an automation-driven slump.

A paradox of productivity. The reason why higher automation may lead to a slump is that firms' technological decisions trigger aggregate demand externalities. That is, firms automate their production processes to increase sales and profits. But each individual firm does not take into account that replacing workers with capital depresses aggregate demand, lowering the sales

¹¹In the simple scenario that we just studied the two effects exactly balance, so that capitalists' consumption is insulated from the automation shock.

¹²To derive this result, consider that with endogenous investment steady state output is given by

$$Y = \frac{1}{\bar{J}} \left(\left(\frac{1 - \beta(1+r)}{\phi} \right) \chi^{1-\xi} \left(\frac{\bar{J}Y}{r+\delta} \right)^\xi + \frac{\delta}{r+\delta} \bar{J}Y \right).$$

Hence, under a constant interest rate policy, $\bar{J}Y$ is unaffected by the automation shock and so is investment.

and profits of all other firms. This can be seen by noticing that

$$\frac{y_{j,h}}{y_{j,l}} = \underbrace{\frac{p_{j,l}}{p_{j,h}}}_{\text{cost-saving effect}} \underbrace{\frac{Y_h}{Y_l}}_{\text{aggregate demand effect}}.$$

The sales of a generic firm j thus depend on the price of its product p_j , which is determined by its marginal cost, and aggregate demand Y . Since $Y_h < Y_l$, sales are lower in the high-automation steady state for all the firms that do not change technology (i.e. for $j < \bar{J}_l$ and $j > \bar{J}_h$).

Perhaps surprisingly, sales may drop even for the firms that exploit the advances in automation to replace workers with capital (i.e. for $\bar{J}_l \leq j \leq \bar{J}_h$).¹³ Effectively, these firms experience a paradox of productivity. Individually, replacing workers with capital increases profits, by reducing production costs and boosting sales. Collectively, however, the decision to automate may lower aggregate demand so much so that sales drop even for automating firms.

An automation-driven slump can thus be seen as the result of a coordination failure. While from an individual perspective adopting the new automation technology to increase productivity makes sense, collectively this decision may lead to adverse outcomes. This paradox of productivity implies that firms' attempts to save on production costs and expand sales may - by depressing aggregate demand - be self-defeating in general equilibrium.

3.2 Numerical example

In this section we illustrate our results by resorting to some simple numerical simulations. To be clear, the objective of this exercise is not to provide a careful quantitative evaluation of the framework, a task that would require a much richer model. Rather, our aim is to show that, for reasonable values of the parameters, the magnitudes implied by the model are quantitatively relevant.

One period of the model corresponds to one year. Since the bulk of the capital stock in the United States is owned by the richest 10% of the population, we set $\chi = 0.1$. Following Mian et al. (2021), we calibrate the wealth-in-utility curvature ξ so that in the initial steady state capitalists have a marginal propensity to consume out of wealth of 1%, which yields $\xi = 0.1818$. We set $\phi = 0.6376$ to hit a real return on wealth in the initial steady state of $r = 5.5\%$, again in line with Mian et al. (2021). For the subjective discount factor, we follow the evidence in Michailat and Saez (2021) and set $\beta = 0.57$. We normalize $\bar{l} = 1/(1 - \chi)$, so that $\bar{L} = 1$.¹⁴

¹³We can easily derive conditions under which this is the case. Consider a generic firm $\bar{J}_l < \tilde{j} < \bar{J}_h$. In the initial steady state, this firm produces using labor and its sales are given by $y_{\tilde{j},l} = Y_l/p_{\tilde{j},l} = \gamma^l Y_l/w_l$, while in the high-automation steady state this firm produces using capital and its sales are given by $y_{\tilde{j},h} = Y_h/p_{\tilde{j},h} = \gamma^k Y_h/r_h^k$. Taking the ratio gives that

$$\frac{y_{\tilde{j},h}}{y_{\tilde{j},l}} = \frac{\gamma^k w_l Y_h}{\gamma^l r_h^k Y_l} = \frac{\gamma^k w_l \bar{J}_l}{\gamma^l r_l^k \bar{J}_h},$$

where the second equality makes use of the fact that, with a constant real rate, $r_h^k = r_l^k$ and $Y_h/Y_l = \bar{J}_l/\bar{J}_h$. Intuitively, automating firms experience a drop in profits if the cost-saving effect of automation $\gamma^k w_l > \gamma^l r_l^k$ is small relative to the aggregate demand effect $\bar{J}_l < \bar{J}_h$.

¹⁴The other parameter describing workers' utility is the Frisch elasticity of labor supply φ . However, this parameter

Turning to technology, we set $\gamma^l = 0.5405$ to normalize output to 1 in the initial steady state. We set $\bar{J}_l = 0.4$ so that the capital share of income in the initial steady state is in line with its empirical counterpart in the United States. In our numerical experiments we set $\bar{J}_t - \bar{J}_h = \rho(\bar{J}_{t-1} - \bar{J}_h)$, where $\bar{J}_{-1} = \bar{J}_l$, with the adjustment parameter $\rho = 0.5$, and with the terminal condition $\bar{J}_h = 0.43$. We thus assume that automation occurs gradually with a half-life of one year, and that the capital share rises by 3 percentage points over time. We also set $\gamma^k = 0.2398$ so that the rise in automation boosts potential output in steady state by 10%. This calibration strikes a balance between the pessimistic estimates by [Acemoglu \(2025\)](#), which suggests that AI will boost output by 1% over the next decade, and much more optimistic estimates such as those provided by [Aghion and Bunel \(2024\)](#). Moreover, it is compatible with the analysis in [Acemoglu \(2025\)](#), who estimates that an AI boom which raises the capital share by 0.3 percentage points leads to long-run output gains in the order of 1%. Of course, this calibration is highly speculative, given the large uncertainty about the future impact of AI technologies.

Finally, we set $\delta = 0.15$, corresponding to an annual depreciation rate of capital of 15%, and abstract from capital adjustment costs in our baseline calibration ($\kappa = 0$). We do so because in the U.S. AI-related investments are rising very rapidly, pointing to a low degree of investment adjustment costs. Moreover, at the yearly frequency that we study in our experiment, adjustment costs are less likely to be empirically relevant compared to quarterly calibrations. That said varying κ over reasonable ranges does not alter much our results.

Figure 2 shows the results. The left column refers to a full-employment economy, in which the central bank keeps output always at its potential value. Output increases gradually over time, because the economy needs to build a higher capital stock to fully exploit the productivity gains from higher automation. In fact, the output boom is initially driven by investment, while consumption drops on impact. The initial drop in consumption is particularly severe for workers, due to the fact that on impact higher automation leads to a large drop in the labor share, but only to small productivity gains.¹⁵ However, as capital accumulates and output rises workers' consumption eventually increases above its value in the initial steady state.¹⁶

The right column, instead, refers to the demand-deficient economy that we studied in the previous section, in which the central bank holds the real interest rate fixed at its value in the initial steady state. The picture painted is very different. The drop in workers' income caused by higher automation depresses aggregate consumption. Moreover, the prospect of weak future aggregate demand discourages firms' investment. As a result, rather than a boom, the rise in automation generates a slump. This slump hurts workers particularly hard, since their consumption remains permanently depressed below its value in the initial steady state.

does not affect the numerical simulations.

¹⁵On impact, capital is predetermined. Moreover, in the natural allocation L_t is fixed at \bar{L} . Thus, GDP rises only through the rise in TFP associated with higher automation. In our calibration those TFP gains are small, and not enough to compensate for workers' income loss due to the lower labor income share.

¹⁶[Jones and Tonetti \(2026\)](#) emphasize how the presence of "weak links" may tame the growth explosion due to AI. In our framework, the slow accumulation of capital can be seen as such a weak link, by taming the pace with which AI advances translate into productivity gains.




Figure 2: Impact of automation on output and consumption. Notes: Left column (a): Natural allocation. Right column (b): Constant real interest rate policy. The upper row decomposes output into its spending components: (aggregate) consumption and investment. Both are expressed in percentage deviation from the initial steady state, weighted by the initial steady state share in GDP. The lower row shows the components of aggregate consumption, namely worker consumption and capitalist consumption.

Before moving on, let us stress again that we are not arguing that advances in AI technologies will necessarily lead to a deep and prolonged slump.¹⁷ Rather, our message is that under plausible scenarios the automation effect of AI may depress aggregate demand substantially, so that accommodative macroeconomic policies may be needed to exploit the potential benefits of AI.

4 Two challenges for monetary policy

If advances in AI start putting downward pressure on employment and economic activity, governments will likely rely on a mix of expansionary monetary and fiscal policy to sustain demand and restore full employment. In this section, we focus on monetary policy and point toward two limits of using monetary policy to implement full employment. First, maintaining full employment may come at the expense of a temporary rise in inflation. Second, if equilibrium interest rates are low, the rise in automation may push the economy into a liquidity trap.

¹⁷In fact, one could say that we are taking a pessimistic perspective on AI. For instance, some commentators argue that AI will boost the trend growth rate of our economies, by fostering scientific discoveries and technological development. In that case, AI may even boost aggregate demand, because high growth prospects lift investment and capitalists' consumption.




Figure 3: Monetary policy responses to advances in AI. Notes: Blue solid lines: Targeting full employment (natural allocation). Red dashed lines: Targeting constant inflation.

4.1 Will AI be inflationary?

What happens if the central bank sets monetary policy so as to ensure that the economy always operates at full employment? As we have seen in Section 3.2, the rise in automation then causes a sustained output boom. The blue lines in Figure 3 point out two additional interesting effects.¹⁸ First, the real interest rate needed to maintain full employment - i.e. the natural interest rate - rises on impact, but then gradually falls below its value in the initial steady state. We will go back to this point in Section 4.2. Second, higher automation is associated with an initial spike in inflation. The first phases of the rise in automation thus look like a cost-push shock, increasing the inflation rate consistent with full employment.

This cost-push shock aspect of automation may come as a surprise, given that higher automation boosts productivity. Standard logic, in fact, would suggest that higher productivity reduces inflation by lowering firms' production costs. Consistent with this view, Kevin Warsh, who may soon become the Chairman of the Federal Reserve, has recently argued that AI acts as a powerful deflationary force through its positive impact on productivity (Markman, 2026).

To resolve this puzzle notice that higher automation, in spite of its positive impact on productivity, may depress firms' labor demand. This happens because higher automation increases the set of tasks performed by capital as opposed to labor. When this effect is sufficiently strong, maintaining full employment requires a drop in the cost of labor, i.e. lower real wages. Since nominal wages are sticky, the drop in real wages has to happen through price inflation.

¹⁸To pin down the inflation dynamics, we have to take a stance on the shape of the wage Phillips curve. In the data, wages are downwardly rigid (Hazell and Taska, 2025). To capture this fact, we assume that the wage Phillips curve is nonlinear and such that $\psi(L_t) = (L_t/\bar{L})^{\tilde{\psi}(L_t)}$, where $\tilde{\psi}(L_t) = .1$ if $L_t \leq \bar{L}$ and $\tilde{\psi}(L_t) = .5$ if $L_t > \bar{L}$.

Automation and the Phillips curve. To isolate the forces that shape inflation, it is useful to go back to the fixed-investment economy ($\kappa = +\infty$) and to the case where in period 0, \bar{J}_t jumps immediately to its new steady state value. Let us also assume that nominal wages are constant, so that $W_t = W$ for all t . Under these assumptions, the price level evolves according to

$$P_t = \frac{W}{1 - \bar{J}_t} \frac{\bar{L}}{Y_t^p} \left(\frac{Y_t}{Y_t^p} \right)^{\frac{1}{1 - \bar{J}_t}}. \quad (\text{PC})$$

As shown in Figure 4, this expression implies a positive relationship between the price level and economic activity, captured by the output gap Y_t/Y_t^p . Intuitively, a higher output gap puts upward pressure on firms' marginal costs, causing an increase in prices.

How does the rise in \bar{J}_t affect this Phillips curve? There are three effects to consider. First, the rise in \bar{J}_t boosts potential output and labor productivity Y_t^p/\bar{L} - the productivity effect of automation. This effect decreases production costs and pushes prices down. Second, higher automation lowers the number of tasks performed by labor $1 - \bar{J}_t$ - the displacement effect of automation. This effect reduces firms' demand for labor and the real wage consistent with any level of employment, and so pushes prices up. Third, rises in \bar{J}_t also steepen the Phillips curve, i.e. the rate at which the output gap translates into changes in inflation. Intuitively, this happens because the production factor with the flexible price (capital) becomes more important in production.

Now assume that, as in Figure 3, the central bank keeps the economy at full employment when the rise in automation happens. The inflation response is then ambiguous, as it is driven by the productivity effect, pushing prices down, versus the displacement effect, pushing prices up. Figure 4 shows a case in which the inflationary effect dominates, so that the rise in \bar{J}_t causes an upward shift of the Phillips curve. This is exactly what happens in the simulation underlying Figure 3, in which inflation rises initially because the negative impact of higher automation on real wages quantitatively dominates the productivity effect. Over time, as capital accumulation increases labor productivity, the productivity effect of automation becomes more important, and the initial inflation spike turns into a period of inflation mildly below target.

Automation advances which raise productivity only mildly, i.e. which have a small productivity effect, have been dubbed “so-so technologies” by Acemoglu and Restrepo (2019). Our framework then suggests that the adoption of so-so automation technologies is likely to be accompanied by a bout of inflation.

These results suggest that the automation effect of AI may create a dilemma between sustaining employment and maintaining inflation equal to target for central banks. To illustrate this trade-off, the dashed lines in Figure 3 show what happens if the central bank adopts a policy of strict inflation targeting, such that $P_t = P_{t-1}$ for all t . To fully insulate inflation from the rise in automation, the central bank has to implement a tight monetary policy at the start of the transition, leading to a drop in employment and to a negative output gap.

The automation effect of AI may thus contemporaneously lead to fast productivity growth, inflation above target and a weak labor market, which is an unusual macroeconomic combination.




Figure 4: Rise in automation and the Phillips curve.

If this scenario materializes, it will be a difficult one to navigate for central banks.

4.2 The natural rate and the liquidity trap

Let us now go back to the path of the real interest rate consistent with full employment, i.e. of the natural rate, shown in Figure 3. In the short run, the natural rate rises, because the high investments needed to automate production push up aggregate demand.¹⁹ Once the investment boom subsides, however, the natural rate falls below its initial value. This is a reflection of the drop in consumption demand due to the automation effect of AI, as we have shown in Section 3. To lift demand back up to the level consistent with full employment, the interest rate then needs to fall. Our model is thus consistent with the empirical correlation between the labor share and the natural rate highlighted by [Cho and Williams \(2026\)](#).

We will next argue that the zero lower bound on the policy rate may prevent the central bank from sustaining full employment. Looking at the simulation in Figure 3, the zero lower bound may seem like a remote issue, since the natural rate remains well above zero. But the interest rate in our baseline model should be interpreted as the return on investment in illiquid capital, which in the data is substantially higher than the return on liquid bonds.

To capture the liquidity premium on bonds, imagine that capitalists' utility is given by

$$\sum_{t=0}^{\infty} \beta^t \left(\log c_t^k + \frac{\phi}{1-\xi} \left(e^{\zeta} \frac{b_{t+1}}{P_t} + q_t k_{t+1} \right)^{1-\xi} \right), \quad (10)$$

where $\zeta \geq 0$ determines households' preference for investing in liquid bonds. The no arbitrage

¹⁹With high enough investment adjustment costs, the short-run rise in the natural interest rate can be overturned, and the natural rate declines throughout the transition. Intuitively, this happens because investment adjustment costs dampen the boom in aggregate demand due to higher investment.




Figure 5: Rise in automation and the liquidity trap.

condition between bonds and capital then implies²⁰

$$1 + r_t = \frac{r_{t+1}^k + (1 - \delta)q_{t+1}}{q_t} - \left(e^\zeta - 1 \right) \phi c_{t+1}^k (q_t k_{t+1})^{-\xi}.$$

When $\zeta > 0$, the interest rate on liquid bonds lies below the return on illiquid capital.

Let us also assume that the central bank faces a lower bound r_{lb} on its policy rate. For simplicity, we frame the lower bound directly in terms of the real interest rate, so that $r_t \geq r_{lb}$. Monetary policy is now described by

$$(Y_t - Y_t^p)(r_t - r_{lb}) = 0, \quad (\text{MP})$$

with $r_t \geq r_{lb}$. In words, the central bank seeks to maintain the economy at full employment and output at potential, but may not reach this target if the lower bound on the interest rate binds.

Automation and the liquidity trap. Now consider again our fixed-investment economy, that jumps immediately to its new steady state in $t = 0$. The aggregate demand equation is now

$$Y_t^d = (1 - \bar{J}_t) Y_t + \frac{1 - \beta(1 + r_t)}{\phi} \chi^{1-\xi} \left(\frac{\bar{J}_t Y_t}{r_t + \delta + (1 - e^{-\zeta})(1 - \beta(1 + r_t))} \right)^\xi + \delta \bar{K}. \quad (\text{AD})$$

Hence, it is still the case that aggregate demand is a decreasing function of the policy rate r_t . Moreover, holding constant the real rate, a rise in automation depresses aggregate demand.

Figure 5 shows the AD and MP schedules. In the initial steady state with $\bar{J}_t = \bar{J}_l$, aggregate

²⁰To obtain this expression, consider that the optimality conditions for bonds and capital are respectively

$$\frac{1}{c_t^k} = \frac{\beta(1 + r_t)}{c_{t+1}^k} + \phi e^\zeta \left(\frac{b_{t+1}}{P_t} + q_t k_{t+1} \right)^{-\xi}$$

$$\frac{1}{c_t^k} = \beta \frac{r_{t+1}^k + (1 - \delta)q_{t+1}}{q_t} \frac{1}{c_{t+1}^k} + \phi \left(\frac{b_{t+1}}{P_t} + q_t k_{t+1} \right)^{-\xi}.$$

Combining these expressions and imposing the equilibrium condition $b_{t+1} = 0$ gives the expression in the main text.




Figure 6: AI boom and liquidity trap. Notes: Blue solid lines: Targeting full employment as long as $r_t > 0$, set $r_t = 0$ else. Red-dashed lines: Targeting full employment, ignoring the zero lower bound.

demand is high (AD_l) and the central bank manages to keep output at potential ($Y_l = Y_l^p$) by setting $r = r_l > r_{lb}$. The rise in automation triggers a drop in aggregate demand, corresponding to a downward shift of the demand schedule from AD_l to AD_h . If the shock is large enough, as shown in the diagram, maintaining output at potential would require setting an interest rate lower than r_{lb} . Since this is not attainable, the central bank sets $r = r_{lb}$ and the economy operates below potential $Y_h < Y_h^p$. Higher automation thus plunges the economy in a liquidity trap, in which the economy operates below potential.

Let us now go back to our numerical example, in which investment is endogenous and the transition to the final steady state takes some time. We set $\zeta = 0.069$ so that in the initial steady state the return on bonds is 5% lower than the one on illiquid capital, and assume that $r_{lb} = 0$. The rest of the parameters are set as in Section 3.2.

Figure 6 illustrates the results, by comparing an economy constrained by the lower bound on the policy rate (solid lines) to a counterfactual scenario in which the lower bound is absent (dashed lines). As expected, over the medium run monetary policy runs against the zero lower bound constraint, and the economy enters a liquidity trap characterized by unemployment and a negative output gap. Moreover, weak demand also causes an investment slump, which depresses potential output. These two forces imply that the liquidity trap imposes severe output losses.

5 Employment subsidies as macroeconomic policy

As we have seen, the automation effect of AI may challenge central banks' ability to maintain low inflation and support demand. We now discuss how fiscal policy can help. While the array




Figure 7: Macroeconomic impact of employment subsidies.

of potential fiscal interventions is large, we focus our attention on employment subsidies. In our context, employment subsidies are a particularly promising tool because they contain inflation by reducing firms' labor cost, and support demand by boosting workers' income.

Assume that the government subsidizes firms' wage bill at rate s_t , financed with a lump-sum tax paid by capitalists. This subsidy reduces firms' cost of labor, since now hiring a unit of labor costs $(1 - s_t)w_t$. Suppose that the subsidy is small enough such that it does not affect firms' decision to automate their production.²¹ The subsidy then lifts firms' labor demand only at the intensive margin (final good firms rely more on those intermediate goods that are produced with labor), but not at the extensive margin (the share of intermediate goods produced with labor is unaffected). Notice that an increase in this employment subsidy could also be interpreted as a cut in labor taxes.

Now consider one last time our fixed-investment economy, in which investment adjustment costs are high and \bar{J}_t jumps immediately to its value in the final steady state. Using the fact that the share of income going to labor is now $w_t L_t / Y_t = (1 - \bar{J}_t) / (1 - s_t)$, the aggregate demand equation now becomes

$$Y_t^d = \frac{1 - \bar{J}_t}{1 - s_t} Y_t + \frac{1 - \beta(1 + r_t)}{\phi} \chi^{1-\xi} \left(\frac{\bar{J}_t Y_t}{r_t + \delta + (1 - e^{-\xi})(1 - \beta(1 + r_t))} \right)^\xi + \delta \bar{K}. \quad (\text{AD})$$

This expression shows that subsidizing employment lifts aggregate demand, because it redistributes income towards high-spending workers. As shown in the left panel of Figure 7, this implies that increasing s_t during a liquidity trap boosts output.

Turning to inflation, assuming for simplicity a constant nominal wage, we can now write the Phillips curve as

$$P_t = \frac{(1 - s_t)W}{1 - \bar{J}_t} \frac{\bar{L}}{Y_t^p} \left(\frac{Y_t}{Y_t^p} \right)^{\frac{1}{1 - \bar{J}_t}}. \quad (\text{PC})$$

²¹That is, we assume that $r_t^k / \gamma^k < (1 - s_t)w_t / \gamma^l$.




Figure 8: Employment subsidies. Notes: Blue solid lines: no subsidies. Red dashed lines: constant employment subsidy.

This expression shows that a rise in s_t acts as a deflationary force, i.e. it induces a downward shift of the Phillips curve. This is intuitive, since subsidizing employment lowers firms' production costs. As illustrated by the right panel of Figure 7, the implication is that subsidizing employment reduces the inflation rate consistent with maintaining output equal to its potential value.

Figure 8 shows the impact of subsidizing employment in our running numerical example. The figure compares a baseline economy without fiscal interventions (solid lines), to one in which the government subsidizes employment at a constant rate (dashed lines). The subsidy serves a dual purpose. In the first phases of the transition, the subsidy helps to contain inflation. Over the medium run, the subsidy mitigates the output losses due to weak aggregate demand. As a result, subsidizing employment sustains employment and output, while at the same time keeping inflation closer to its target.

The figure also highlights how subsidizing employment can have a nuanced impact on investment and labor productivity. If the economy operates at full employment, which happens in the first phases of the transition, subsidizing employment crowds out investment. This is due to the fact that the employment subsidy increases workers' consumption, forcing the central bank to increase the interest rate which in turn lowers investment. However, if output is depressed below its potential value - as it is the case in the final steady state - subsidizing employment encourages investment. The reason is that in a depressed economy higher aggregate demand increases the return to investment. This result suggests that employment subsidies should be carefully calibrated, taking into account their impact on investment and future productivity.

Before concluding, let us spend a few words on how these employment subsidies could be

implemented in practice. In most advanced economies, labor taxes represent a major source of fiscal revenue. These taxes increase firms' cost of hiring workers, and lower the post-tax share of income accruing to workers. Our results thus suggest that, during periods of rapid adoption of new automation technologies, lowering labor taxes may contribute to macroeconomic stabilization.

6 Conclusion

In this paper, we have provided a tractable framework to study monetary and fiscal policies for AI. Our key message is that the macroeconomic effects of AI will be shaped not only by supply factors, but also by the demand side of the economy. And AI, by triggering a wave of automation and redistributing income from high-spending workers to low-spending capitalists, may have a negative impact on aggregate demand. If this effect is sufficiently strong, we might experience an AI slump in the future.

Macroeconomic policies are a key determinant of whether advances in AI will trigger a slump or a boom. In this respect, our key insight is that monetary policy, if left on its own, may have a hard time supporting an AI boom. Accomplishing this task is likely to require a careful mix of monetary and fiscal policies. In particular, our model suggests that employment subsidies or cuts in labor taxes may usefully complement monetary policy, and ensure that advances in AI technologies lead to output gains and shared prosperity.

References

- Acemoglu, Daron (2025) "The simple macroeconomics of AI," *Economic Policy*, Vol. 40, No. 121, pp. 13–58.
- Acemoglu, Daron, David Autor, and Simon Johnson (2026) "Building pro-worker artificial intelligence," Technical report, National Bureau of Economic Research.
- Acemoglu, Daron and Pascual Restrepo (2018) "The race between man and machine: Implications of technology for growth, factor shares, and employment," *American Economic Review*, Vol. 108, No. 6, pp. 1488–1542.
- (2019) "Automation and new tasks: How technology displaces and reinstates labor," *Journal of Economic Perspectives*, Vol. 33, No. 2, pp. 3–30.
- Aghion, Philippe and Simon Bunel (2024) "AI and growth: Where do we stand?" *policy note*.
- Antolin-Diaz, Juan and Paolo Surico (2025) "The long-run effects of government spending," *American Economic Review*, Vol. 115, No. 7, pp. 2376–2413.
- Anzoategui, Diego, Diego Comin, Mark Gertler, and Joseba Martinez (2019) "Endogenous technology adoption and R&D as sources of business cycle persistence," *American Economic Journal: Macroeconomics*, Vol. 11, No. 3, pp. 67–110.

- Basso, Henrique S and Omar Rachedi (2025) “Robot adoption and inflation dynamics,” *Journal of Monetary Economics*, p. 103856.
- Benigno, Gianluca and Luca Fornaro (2018) “Stagnation traps,” *Review of Economic Studies*, Vol. 85, No. 3, pp. 1425–1470.
- Bilbiie, Florin O (2025) “Monetary policy and heterogeneity: An analytical framework,” *Review of Economic Studies*, Vol. 92, No. 4, pp. 2398–2436.
- Caballero, Ricardo J (2026) “Speculative Growth and the AI” Bubble”, Technical report, National Bureau of Economic Research.
- Cerra, V., A. Fatas, and S. Saxena (2021) “Hysteresis and Business Cycles,” *Journal of Economic Literature*.
- Cho, Sophia and John C Williams (2026) “The R*–Labor Share Nexus,” Technical report, Federal Reserve Bank of New York.
- Davidson, Tom, Basil Halperin, Thomas Houlden, and Anton Korinek (2026) “When Does Automating AI Research Produce Explosive Growth?”.
- Debortoli, Davide and Jordi Galí (2024) “Idiosyncratic income risk and aggregate fluctuations,” *American Economic Journal: Macroeconomics*, Vol. 16, No. 4, pp. 279–310.
- Fieldhouse, Andrew J and Karel Mertens (2023) *The returns to government R&D: Evidence from US appropriations shocks*: Federal Reserve Bank of Dallas, Research Department.
- Fornaro, L. and M. Wolf (2025) “Fiscal stagnation,” CREI working paper.
- Fornaro, Luca (2025) “Fiscal Stimulus with Supply Constraints,” *IMF Economic Review*, pp. 1–25.
- Fornaro, Luca, Veronica Guerrieri, and Lucrezia Reichlin (2025) “Monetary policy for the green transition,” report prepared for the BIS.
- Fornaro, Luca and Martin Wolf (2021) “Monetary policy in the age of automation,” CREI working paper.
- (2023) “The scars of supply shocks: Implications for monetary policy,” *Journal of Monetary Economics*, Vol. 140, pp. S18–S36.
- Fueki, Takuji, Shinnosuke Katsuki, Ichiro Muto, and Yu Sugisaki (2023) “Automation and nominal rigidities,” Technical report, Institute for Monetary and Economic Studies, Bank of Japan.
- Garga, Vaishali and Sanjay R. Singh (2020) “Output hysteresis and optimal monetary policy,” *Journal of Monetary Economics*.
- Hazell, Jonathon and Bledi Taska (2025) “Downward rigidity in the wage for new hires,” *American Economic Review*, Vol. 115, No. 12, pp. 4183–4217.

- Imas, A. (2026) “Can advanced AI lead to negative economic growth?,” Substack post.
- Jones, Charles I. and Christopher Tonetti (2026) “Past Automation and Future AI: How Weak Links Tame Technological Progress,” Working paper.
- Jordà, Òscar, Sanjay R Singh, and Alan M Taylor (2024) “The long-run effects of monetary policy,” *Review of Economics and Statistics*, pp. 1–49.
- Kaplan, Greg, Benjamin Moll, and Giovanni L Violante (2018) “Monetary policy according to HANK,” *American Economic Review*, Vol. 108, No. 3, pp. 697–743.
- Lenzu, Simone (2026) “Artificial Intelligence and Monetary Policy: A Framework and Perspective on Cyclical Transmission, Structural Transition, and Financial Stability,” *FRB of New York Staff Report*, No. 1192.
- Lim, Hyunjoon (2026) “Automation, Income Distribution, and Aggregate Demand in a Two-Asset HANK Model.”
- Ma, Yueran and Kaspar Zimmermann (2023) “Monetary policy and innovation,” Technical report, National Bureau of Economic Research.
- Markman, Jon (2026) “Kevin Warsh’s New Playbook: AI, Productivity And A Deflation Bet,” February.
- Mian, Atif R, Ludwig Straub, and Amir Sufi (2020) “The saving glut of the rich,” Technical report, National Bureau of Economic Research.
- (2021) “Indebted demand,” *Quarterly Journal of Economics*.
- Michaillat, Pascal and Emmanuel Saez (2021) “Resolving New Keynesian anomalies with wealth in the utility function,” *Review of Economics and Statistics*, Vol. 103, No. 2, pp. 197–215.
- Moll, Benjamin, Lukasz Rachel, and Pascual Restrepo (2021) “Uneven Growth: Automation’s Impact on Income and Wealth Inequality,” LSE working paper.
- Moran, Patrick and Albert Queralto (2018) “Innovation, productivity, and monetary policy,” *Journal of Monetary Economics*, Vol. 93, pp. 24 – 41.
- Schmöller, Michaela Elfsbacka and Martin Spitzer (2021) “Deep recessions, slowing productivity and missing (dis-) inflation in the euro area,” *European Economic Review*, Vol. 134, p. 103708.
- Smith, Noah (2023) “Will AI take away the coding jobs?,” February, Noahpinion (Substack blog).
- Werning, Iván (2015) “Incomplete markets and aggregate demand,” NBER Working Paper No. 21448.