2184

Discussion Papers

Deutsches Institut für Wirtschaftsforschung 2026

It Matters Who Gets Automated: Monetary Policy in an AI Economy

Alina Ermolaeva and Hannah Seidl

Opinions expressed in this paper are those of the author(s) and do not necessarily reflect views of the institute.

IMPRESSUM

© DIW Berlin, 2026

DIW Berlin
German Institute for Economic Research
Anton-Wilhelm-Amo-Str. 58
10117 Berlin

Tel. +49 (30) 897 89-0
Fax +49 (30) 897 89-200
http://www.diw.de

ISSN electronic edition 1619-4535

Papers can be downloaded free of charge from the DIW Berlin website:
http://www.diw.de/discussionpapers

Discussion Papers of DIW Berlin are indexed in RePEc and Econstor:
http://ideas.repec.org/s/diw/diwwpp.html
https://www.econstor.eu/handle/10419/10

It Matters Who Gets Automated: Monetary Policy in an AI Economy

Alina Ermolaeva* Hannah Seidl†

October 12, 2026

Abstract

Does artificial intelligence (AI) change the transmission of monetary policy? We develop a two-sector heterogeneous-agent New Keynesian model with task-based production, in which AI automates part of the work in one exposed sector and leaves the other shielded. In the AI steady state, the economy holds more capital, and labor income shifts from the exposed workers, who lose through hours, to the shielded workers, who gain through higher wages. We then compare the same rate cut across the two economies. It stimulates macroeconomic aggregates more in the AI economy. The deeper capital stock delivers a larger initial stimulus. Investment responds as strongly as without AI, but on a permanently larger base, so the same response means more demand. Machines absorb part of the exposed sector's share of that demand. The shielded sector's share becomes labor income, which is spent in turn. It matters who gets automated: the amplification of monetary policy in an AI economy runs through the workers AI does not replace.

JEL: E21, E24, E32, E52, J23, O33.

Keywords: artificial intelligence, automation, task-based production, HANK, monetary transmission, labor share, sectoral incidence.

*DIW Berlin and Freie Universität Berlin, aermolaeva@diw.de.

†DIW Berlin, hseidl@diw.de. We thank Alex Albright, Rüdiger Bachmann, Wouter den Haan, Alexander Kriwoluzky, Fabian Seyrich, and Mathias Trabandt for their helpful comments.

1 Introduction

Artificial intelligence (AI) is reshaping the economy. Early evidence on AI adoption points to two patterns. First, AI shifts tasks, and with them income, from labor to capital (AEA Committee on Economic Statistics, 2025), which should shrink the labor-income channel of monetary policy, the most important transmission channel in macroeconomic models with realistic household consumption behavior (see, e.g., Kaplan et al., 2018). Second, AI adoption is highly uneven across sectors (U.S. Census Bureau, 2026), which suggests that any change in transmission should originate in the sectors where AI operates (see, e.g., Pastén et al., 2020).

We study how AI changes the transmission of monetary policy. We develop a two-sector heterogeneous-agent New Keynesian (HANK) model with task-based production, in which AI takes over tasks from workers in one sector, the exposed sector, and leaves the other shielded. In the AI steady state, the economy holds more capital, and labor income shifts between the two groups, with exposed workers losing through hours and shielded workers gaining through higher wages. We then compare the effects of the same monetary policy shock across the two steady states. Both expectations fail. Transmission strengthens slightly rather than weakens. And the additional stimulus does not originate where AI operates but reaches households as labor income in the shielded sector.

Our model embeds the task-based production structure of Acemoglu and Restrepo (2018) into a two-sector version of the canonical two-asset HANK model of Auclert et al. (2024). Tasks below an endogenous threshold are performed by capital, tasks above it by labor, and the threshold is determined by the relative cost of the two factors. AI is modeled as higher capital productivity in one of the two sectors, the exposed sector. Cheaper automation advances the threshold, and capital takes over tasks previously performed by workers. Households face uninsurable earnings risk and hold liquid and illiquid assets. They are attached to one of the two sectors, so that each household's labor income comes either from the exposed or from the shielded sector. The two groups differ in patience. The shielded group is less patient, keeps a liquid buffer, and accumulates almost no illiquid wealth. In equilibrium, the exposed group therefore owns almost the entire capital stock, including the capital that automates its own sector. Prices and wages are sticky, with sector-specific reset frequencies. Capital utilization is variable. Monetary policy follows a standard Taylor rule, and the government balances its budget with a labor income tax.

We compare a non-AI steady state, in which both sectors share the same automation threshold, with an AI steady state, in which capital productivity in the exposed sector is higher, by the amount that lowers the aggregate labor share by 5.5 percentage points,

in line with the expert forecasts of Karger et al. (2026). All other differences between the two steady states arise in equilibrium. The capital stock is 30 percent larger, and the capital-output ratio rises from 2.4 to 2.9. Automation spreads to the shielded sector. Output is 11 percent higher in the AI economy, and the additional income is spent on both goods: demand for the shielded good rises, and with it the shielded wage. At the higher wage, automating additional tasks pays: even with unchanged technology, the share of the shielded sector's tasks performed by machines rises by four percentage points. In addition, labor income shifts between the two groups of workers. Exposed workers earn 4 percent less, because their hours fall while their wage is nearly unchanged. Shielded workers earn 5 percent more, because their wage rises by 4.8 percent while their hours are nearly unchanged.

We subject both economies to the same monetary policy shock, an easing that lowers the real rate by 25 basis points on impact, with an identical path thereafter in both economies. The cut is more stimulative in the AI economy: output responds 17 percent more strongly and consumption 7 percent, while the investment response is essentially the same in both economies. We refer to this difference as the amplification. The consumption of exposed households responds almost as if AI had never happened: their response lies on top of its non-AI counterpart. The amplification comes mostly from the shielded households and is driven by their labor income.

These findings follow from the AI steady state, where the capital stock is larger and machines perform a larger share of the tasks. The cut raises the price of capital in both economies. Investment responds by the same percentage in both, but the AI economy invests on a larger base, so the same response means more spending, resulting in a larger initial demand impulse. In the exposed sector, machines produce a larger share of each unit of output than in the non-AI economy, so less labor is needed per unit. However, more units are produced, and the two roughly cancel, such that the exposed labor-income channel responds about as strongly as without AI. In the shielded sector, the labor needed per unit of output is nearly the same as without AI: the extra demand passes through to hours, paid at the permanently higher wage, and labor income rises. That income is spent in turn. The Keynesian-type multiplier is the same as in the non-AI economy, but it operates on more labor income, resulting in the amplification.

The cut also generates capital gains, but these add little to the amplification. They accrue to the illiquid accounts of the exposed households, who hardly spend them, in line with the marginal propensity to consume out of stock-market wealth estimated by Chodorow-Reich et al. (2021). The exposed workers own the capital that displaced them,

but the gains do not restore their share of the transmission.

Two exercises test this mechanism. First, we shut off the investment response to the cut. The consumption amplification disappears, and the responses of the two economies lie on top of each other. The result mirrors Auclert et al. (2020), who show that in HANK models a large part of the consumption response to monetary policy is generated by investment, because investment demand raises labor income and households with high propensities to consume spend it. The AI economy has more of this feedback, as the same investment response operates on a larger base. Second, we solve a representative-agent version of the model. The consumption amplification vanishes as well. The representative agent consumes according to the permanent income hypothesis, so the additional labor income is saved rather than spent. AI shifts income between sectors and factors in that economy too, but with a representative agent, the composition of income is irrelevant for demand.

In the baseline model, workers are attached to one sector. Over time, however, the higher shielded wage may draw workers out of the exposed sector, and the additional labor supply would erode the wage gain. As a robustness exercise, we allow workers to switch sectors at the observed rate of occupational mobility, and the switching decision responds to the wage gap between sectors, as in Beraja and Zorzi (2025). In the AI economy, workers therefore flow into the shielded sector, and the inflow narrows the wage gap. The gap does not close, because switching is costly. In addition, part of the shielded wage gain comes from a rise in wage setting that is common to both sectors, and this part gives workers no reason to move. The shielded wage therefore remains higher than without AI, and the amplification, which runs through the labor income earned at this wage, is essentially unchanged. Mobility changes the composition of workers across sectors, but not the path through which the additional response reaches households. Labor income earned in the shielded sector still accounts for the largest share of the additional consumption response.

Related literature. Our paper builds on the literature on monetary policy transmission in HANK models. This work shows that most of the consumption response to a change in the policy rate comes from labor income rather than from intertemporal substitution (Werning, 2015; McKay et al., 2016; Kaplan et al., 2018; Auclert, 2019; Auclert et al., 2024; Bayer et al., 2024; Broer et al., 2026). In this literature, the labor share is fixed by calibration, so a constant fraction of the income that a rate cut generates arrives as labor income. In our model, automation determines this fraction in equilibrium, and the

sector structure determines whose wages carry it. We find that AI does not weaken the labor-income channel. The exposed channel is essentially unchanged, and the addition comes from the shielded workers.

The task-based theory of automation (Acemoglu and Restrepo, 2018, 2019, 2022; Acemoglu et al., 2025; Moll et al., 2022; Althoff and Reichardt, 2026) shows how automation shifts tasks, and with them income, from labor to capital, so that the labor share falls. We embed task-based production in the HANK framework above, so that the task allocation determines who earns the income that a rate cut generates.

Our calibration draws on the emerging empirical literature on AI. The exposure share follows the occupation-level estimates of Eloundou et al. (2024), and the size of the AI shock follows the expert forecasts of Karger et al. (2026). Evidence on the labor-market effects of AI (Bárány et al., 2026; Freund and Mann, 2026; Lindenlaub et al., 2026) motivates the sectoral structure of the model.

Several papers study automation together with macroeconomic policy. Beraja and Zorzi (2025) study optimal policy toward automation in a heterogeneous-agent economy without nominal rigidities. We add sticky prices and wages and a monetary authority, so that changes in the policy rate have real effects, and ask how automation alters their transmission. We also adopt their specification of sector switching in the mobility exercise of Section 5. Fornaro and Wolf (2021, 2026) study monetary policy with endogenous technology adoption in representative-agent and two-agent economies. In Fornaro and Wolf (2026), automation depresses demand because income shifts to capitalists who spend little of it. Their mechanism concerns the level of demand and the stance of policy. We hold the policy change fixed and ask how its transmission changes, and we find that the income shift does not weaken it. Fueki et al. (2026) embed task-based production in a New Keynesian model with sticky prices and wages and show that monetary policy moves the automation margin itself, flattening the Phillips curve. Their household is representative, so the labor-income channel is absent, and so is the question of who receives the transmission. For a conceptual overview of the channels through which AI may affect monetary policy, see Lenzu (2026).

In contemporaneous work, Lim (2026) finds in a one-sector two-asset HANK model with endogenous automation that automation amplifies the response to a rate cut. Our contribution relative to this work is twofold. First, the production side. Lim closes the task model with Cobb-Douglas aggregation, which ties the labor share to the automation level. We keep tasks as complements, in line with the empirical evidence (Oberfield and Raval, 2021), so the labor share responds to prices. Second, we add a second sector, with

households attached to it, which makes the incidence of transmission across automated and non-automated workers an outcome of the model. In Lim (2026), the amplification comes from the composition of output. In our model, it is labor income. The sector split turns out to be the crucial one. Although AI operates only in the exposed sector, the amplification runs through the shielded workers.

The rest of the paper is organized as follows. Section 2 presents the model. Section 3 describes the calibration and the construction of the two economies. Section 4 reports the responses to the monetary shock and their decomposition. Section 5 shows that the results are robust to sector labor mobility. Section 6 concludes.

2 Model

We develop a two-asset HANK model with two sectors and combine it with the task-based production technology of Acemoglu and Restrepo (2018). Sector 1 is exposed to AI and Sector 2 is shielded. Households face uninsurable earnings risk and save in a liquid and an illiquid asset, as in Auclert et al. (2024). Each household is attached to one of the two sectors. Unions set sticky wages in each sector. In production, an endogenous automation threshold divides the tasks of each sector between capital and labor. This threshold is where AI enters the model. A final good aggregates the two sector outputs, whose prices are sticky. Financial intermediaries channel household wealth into capital and bonds, and the government conducts monetary and fiscal policy. We solve the model with the sequence-space Jacobian method of Auclert et al. (2021).

2.1 Households

There are two household groups $g \in \{1, 2\}$ of masses ϕ and $1 - \phi$, attached to sectors 1 and 2, respectively. Hours and wage income per worker are $1/\phi$ and $1/(1 - \phi)$ times the sector values. A household in group g chooses consumption c_t to maximize

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta_g^t \log c_t,$$

so the elasticity of intertemporal substitution is one. The group discount factors are $\beta_1 = \beta$ and $\beta_2 = \beta - \Delta_\beta$, so the shielded group is the less patient one (Krusell and Smith, 1998; Carroll et al., 2017). Both discount factors are the same in the two economies.

Budget constraint and portfolio friction. Households hold two accounts, a liquid account with balance $b \geq 0$ and an illiquid account with balance $a \geq 0$. Moving resources into or out of the illiquid account incurs a convex cost, as in Auclert et al. (2024),

$$\Psi(a', a) = \frac{\chi_1}{\chi_2} \frac{|a' - (1 + r^a)a|^{\chi_2}}{[(1 + r^a)a + \chi_0]^{\chi_2 - 1}}.$$

The budget constraint of a group-g household is:

$$c_t + b_{t+1} + a_{t+1} + \Psi(a_{t+1}, a_t) \;=\; (1 + r_t^b)\, b_t + (1 + r_t^a)\, a_t + z_{g,t}\, e_t + Tr,$$

where Tr is a flat lump-sum transfer and $z_{g,t}$ is the per-capita after-tax income of group g,

$$z_{1,t} = (1 - \tau_t) \left(\frac{w_{1,t} N_{1,t}}{\phi} + PH_t \right), \qquad z_{2,t} = (1 - \tau_t) \left(\frac{w_{2,t} N_{2,t}}{1 - \phi} + PH_t \right).$$

Here $PH_t = (1 - \varpi)(\Pi_{1,t} + \Pi_{2,t})$ is the share of pure profits distributed to households, paid in proportion to productivity e as in Kaplan et al. (2018). The remaining share ϖ is retained by the equity fund and shows up in the return on illiquid wealth. Hours are set by the unions, so households take $N_{i,t}$ as given. Individual productivity e_t follows a Markov chain with normalization $\mathbb{E}[e] = 1$.

Returns. Both accounts earn the return on a common portfolio of government bonds and firm equity, r^{post}, defined in Section 2.5. The liquid account carries a constant wedge,

$$r_t^a = r_t^{post}, \qquad r_t^b = r_t^{post} - \zeta\,(1 + r_t^{post}),$$

where ζ is the liquid-illiquid spread as in Kaplan et al. (2018).

Group aggregates such as C_g are means over the households of group g. Economy-wide aggregates weight them by population shares, $C = \phi C_1 + (1 - \phi) C_2$, and likewise for illiquid wealth A, liquid wealth A^{liq}, and the portfolio adjustment costs X^{Ψ}.

2.2 Wage setting

A union in each sector sets the nominal wage for its attached workers. Sector wage Phillips curves take the standard form

$$\kappa_{w,i} \left(\frac{\mu_w\, \varphi_i\, N_{i,t}^{1/\nu}}{U_{i,t}^{CE}\,(1 - \tau_t)\, w_{i,t}} - 1 \right) + \beta\, \pi_{i,t+1}^w (1 + \pi_{i,t+1}^w) - \pi_{i,t}^w (1 + \pi_{i,t}^w) = 0,$$

where ν is the Frisch elasticity, U_i^{CE} the group's average marginal utility of consumption, φ_i the disutility level, and wage inflation is $1+\pi_{i,t}^w = (1+\pi_t)\, w_{i,t}/w_{i,t-1}$. Slopes come from sector-specific Calvo reset frequencies, given by:

$$\kappa_{w,i} = \frac{1}{1+\Gamma}\left(1-\beta_i\left(1-f_{w,i}\right)\right)\frac{f_{w,i}}{1-f_{w,i}}, \qquad \Gamma = \varepsilon_w/\nu.$$

2.3 Production: tasks, automation and factor prices

Production follows the task-based technology of Acemoglu and Restrepo (2018). Each sector produces with a continuum of tasks $j \in [0,1]$, aggregated with elasticity σ,

$$Y_i = \left(\int_0^1 y_i(j)^{\frac{\sigma-1}{\sigma}}\, dj\right)^{\frac{\sigma}{\sigma-1}}, \qquad y_i(j) = \begin{cases} \gamma_i^k\, k_i(j), & j \le I_i^*, \\ \gamma_i^l\, e^{\vartheta_i j}\, n_i(j), & j > I_i^*. \end{cases}$$

Tasks below the automation frontier I_i^* are performed by capital services with productivity γ_i^k. Tasks above it require labor, whose comparative advantage rises with the task index at rate ϑ_i. Integrating over tasks gives the sector production function,

$$Y_{i,t} = \Theta_i \left[\underbrace{I_i^{*1/\sigma}}_{\text{capital weight}} \left(\gamma_i^k K_{i,t}^s\right)^{\frac{\sigma-1}{\sigma}} + \underbrace{\Lambda_i(I_i^*)^{1/\sigma}}_{\text{labor weight}} \left(\gamma_i^l N_{i,t}\right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}, \tag{1}$$

with

$$\Lambda_i(I^*) = \frac{e^{(\sigma-1)\vartheta_i} - e^{(\sigma-1)\vartheta_i I^*}}{(\sigma-1)\vartheta_i},$$

where $K_{i,t}^s = u_{i,t}\bar{K}_{i,t-1}$ are capital services, utilization times the predetermined stock. Equation (1) is a CES in capital and labor whose weights are functions of the automation frontier. Advancing I_i^* raises the capital weight and lowers the labor weight. With $\sigma < 1$, the labor share can fall in response to automation. Under factor-augmenting technical change with the same elasticity, higher capital productivity would raise the labor share instead (Acemoglu and Restrepo, 2019). The task structure is what allows AI to lower it. Cost minimization prices both factors at marginal cost,

$$w_{i,t} = mc_{i,t}\frac{\partial Y_{i,t}}{\partial N_{i,t}}, \qquad r_{i,t}^k = mc_{i,t}\frac{\partial Y_{i,t}}{\partial K_{i,t}^s}.$$

The automation frontier. The marginal task is assigned to whichever factor performs it more cheaply, so $I^*_{i,t}$ satisfies:

$$\frac{r^k_{i,t}}{\gamma^k_i} \;=\; \frac{w_{i,t}\, e^{-\vartheta_i I^*_{i,t}}}{\gamma^l_i}. \tag{2}$$

The automation margin is a factor-price margin with elasticity $1/\vartheta_i$. A rate cut that raises wages relative to the rental rate therefore induces further automation within the period. In our experiments this margin is active but quantitatively small. Because the rental rate equals $\bar{r} + \delta$ in both sectors in steady state, condition (2) ties the shielded frontier to the shielded wage, so a change in the exposed sector's technology that raises w_2 moves I^*_2 as well.

Capital utilization and investment. Utilization carries a convex resource cost per unit of stock, $\Psi_u(u) = \sigma_b[\frac{\sigma_a}{2}u^2 + (1-\sigma_a)u + \frac{\sigma_a}{2} - 1]$ with $\Psi_u(1) = 0$. The first-order condition is $r^k_{i,t} = \sigma_b\big(1 + \sigma_a(u_{i,t} - 1)\big)$ with $\sigma_b = \bar{r} + \delta$, so that $u_i = 1$ in steady state and $1/\sigma_a$ is the elasticity of utilization to the rental rate. A sector capital firm owns the stock, accumulates $\bar{K}_{i,t} = (1-\delta)\bar{K}_{i,t-1} + I_{i,t}$ subject to a quadratic adjustment cost with elasticity ϵ_I, and satisfies the usual Tobin-Q conditions:

$$Q_{i,t} = 1 + \frac{1}{\delta\epsilon_I}\left(\frac{\bar{K}_{i,t}}{\bar{K}_{i,t-1}} - 1\right),$$

$$(1+r_t)\, Q_{i,t} = \big[u_{i,t+1} r^k_{i,t+1} - \Psi_u(u_{i,t+1})\big] - \left[\frac{\bar{K}_{i,t+1}}{\bar{K}_{i,t}} - (1-\delta) + \frac{(\bar{K}_{i,t+1}/\bar{K}_{i,t} - 1)^2}{2\delta\epsilon_I}\right] + \frac{\bar{K}_{i,t+1}}{\bar{K}_{i,t}}\, Q_{i,t+1}.$$

Here $Q_{i,t}$ is the price of installed capital. The first condition makes it rise with the pace of accumulation, because adjusting the stock quickly is costly. The second is a no-arbitrage condition. Holding a unit of capital must earn the real rate, through the utilization rents it generates net of costs and through its resale value. A lower path of the real rate therefore raises $Q_{i,t}$, which stimulates investment and revalues the existing stock.

Price setting. Sector price Phillips curves are given by:

$$\kappa_{p,i}\left(\mu_p \frac{mc_{i,t}}{p_{i,t}} - 1\right) + \frac{1}{1+r_t}\,\frac{Y_{i,t+1}}{Y_{i,t}}\,\pi_{i,t+1}(1+\pi_{i,t+1}) - \pi_{i,t}(1+\pi_{i,t}) = 0,$$

with a common markup μ_p and slopes from sector Calvo frequencies. The sector dividend paid to the equity fund and the sector's pure profits are

$$div_{i,t} = p_{i,t}Y_{i,t} - w_{i,t}N_{i,t} - I_{i,t} - \Psi_u(u_{i,t})\,\bar{K}_{i,t-1} - \Phi^K_{i,t},$$
$$\Pi_{i,t} = p_{i,t}Y_{i,t} - w_{i,t}N_{i,t} - r^k_{i,t}K^s_{i,t} - \Phi^P_{i,t}.$$

Here $\Phi^K_{i,t}$ is the quadratic capital-adjustment cost and $\Phi^P_{i,t}$ the resource cost of price adjustment. Both are part of the aggregate resource costs Ξ_t in the market-clearing condition below.

2.4 Final demand

A final-good aggregator combines sector outputs with weight ω_1 and elasticity η_c,

$$Y_t = \left[\omega_1^{1/\eta_c} Y_{1,t}^{\frac{\eta_c-1}{\eta_c}} + (1-\omega_1)^{1/\eta_c} Y_{2,t}^{\frac{\eta_c-1}{\eta_c}}\right]^{\frac{\eta_c}{\eta_c-1}},$$

with demand curves $Y_{i,t} = \omega_i\, p_{i,t}^{-\eta_c}\, Y_t$, where $p_{i,t} \equiv P_{i,t}/P_t$ is the price of sector i relative to the final-good price index P_t, which we normalize to one. The relative price $q_t = p_{1,t}/p_{2,t}$ evolves with the sector inflation rates, and aggregate inflation is the CES-weighted index of the sector inflation rates.

2.5 Financial intermediation and the return on wealth

An equity fund holds the sector capital firms and the retained share of pure profits,

$$V_t = div^{agg}_t + \frac{V_{t+1}}{1+r_t}, \qquad div^{agg}_t = div_{1,t} + div_{2,t} - (1-\varpi)(\Pi_{1,t} + \Pi_{2,t}),$$

and the household portfolio blends government bonds and fund equity with the last-period weight $\theta^B = B/(B + V^{end})$, giving the ex-post return

$$1 + r^{post}_t = \theta^B\,(1 + r_{t-1}) + (1-\theta^B)\,\frac{V_t}{V^{end}_{t-1}}, \qquad V^{end}_t = \frac{V_{t+1}}{1+r_t}.$$

The second term is where a rate cut revalues wealth. A lower path of r raises V_0 relative to V^{end}_{-1}, and the capital gain accrues to every household in proportion to its holdings. The Fisher equation, $1 + i_t = (1 + r_t)(1 + \pi_{t+1})$, links the policy rate to the real rate.

2.6 Government

Fiscal policy. Debt B, spending G and the transfer Tr are constant within an economy, and a flat tax rate on labor income and distributed profits adjusts to balance the budget period by period:

$$\tau_t = \frac{(1 + r_{t-1})B + G + Tr - B}{w_{1,t}N_{1,t} + w_{2,t}N_{2,t} + PH_t}.$$

This fiscal rule matters for the decomposition below. When a rate cut lowers debt service and raises the tax base, the tax rate falls, and the resulting increase in after-tax income is an automatic rebate that is part of the transmission.

Monetary policy. The central bank sets the nominal interest rate according to a Taylor rule,

$$i_t = \bar{r} + \phi_\pi \pi_t + \varepsilon_t^m,$$

where $\bar{r}$ is the steady-state real rate, ϕ_π governs the response to inflation, and ε_t^m is the monetary policy shock of our experiments. In Section 4, we choose the path of ε_t^m so that the real rate follows the identical path in both economies, which makes their responses directly comparable.

2.7 Market clearing

The goods market clears when final output covers consumption, investment, government spending and all real resource costs,

$$Y_t = C_t + I_t + G + \Xi_t + X_t^\Psi,$$

where $I_t = I_{1,t} + I_{2,t}$ is aggregate investment and Ξ_t collects the real resource costs of the economy: price-adjustment, capital-adjustment and utilization costs in the two sectors, and X_t^Ψ are households' portfolio-adjustment costs. The asset market clears when household wealth equals government debt plus the ex-dividend value of the equity fund,

$$A_t = B + V_t^{end}.$$

3 Constructing the non-AI and the AI economy

This section builds the two economies that we compare. Section 3.1 calibrates the non-AI economy to US data at quarterly frequency. Section 3.2 constructs the AI economy. Section 3.3 compares the two steady states.

3.1 Calibration of the non-AI economy

Table 1 summarizes the calibration. On the household side, idiosyncratic labor productivity follows the estimated earnings process of Kaplan et al. (2018), discretized on their 33-state grid at quarterly frequency. The discount factor β of the exposed group is calibrated in the non-AI economy so that the asset market clears at the Kaplan et al. (2018) real rate of 5.7 percent a year. In the AI economy, β stays fixed and the real rate adjusts to clear the asset market. The shielded group discounts at $\beta_2 = \beta - \Delta_\beta$, with $\Delta_\beta = 0.0075$ per quarter, which corresponds to the spread between the impatient and patient types in Krusell and Smith (1998) at quarterly frequency. The portfolio adjustment cost has the shape parameters (χ_0, χ_2) of Auclert et al. (2024) and the level $\chi_1 = 36$. The liquid account earns 2 percent a year, so the wedge ζ equals the spread to the 5.7 percent illiquid return. The exposed employment share is $\phi = 0.40$, the share of employment in occupations with high measured AI exposure (Eloundou et al., 2024; Felten et al., 2023).

Price and wage reset frequencies are sector-specific. Prices reset every 10 months in the exposed sector ($f_{p,1} = 0.31$) and every 20 months in the shielded one ($f_{p,2} = 0.15$), following Nakamura and Steinsson (2008) and the sectoral evidence in Pastén et al. (2020). Wages reset every 9 months in the exposed sector, where pay has a large bonus component ($f_{w,1} = 0.35$), and every 20 months in the shielded one ($f_{w,2} = 0.15$) (Grigsby et al., 2021). The final-good weights are symmetric, $\omega_1 = 0.5$, and the two outputs are complements, $\eta_c = 0.5$. The two sectors are therefore symmetric in technology and demand in the non-AI economy. They differ in patience and in nominal rigidity, as the exposed sector adjusts both prices and wages more often.

On the policy side, government debt is 26 percent of annual GDP (Kaplan et al., 2018), government spending is 13 percent of GDP, and transfers are 5 percent, which implies a tax rate of 0.30. The Taylor coefficient takes the standard value $\phi_\pi = 1.5$. The monetary shock ε_t^m follows an AR(1) with quarterly persistence $\rho_m = 0.85$, and we scale it so that the real rate falls by 25 basis points (annualized) on impact.

Both sectors have the same automation threshold, $I_1^* = I_2^* = 0.33$, and the technology levels (Θ_i, γ_i^l) are chosen so that in steady state $N_i = Y_i = 1$, $q = 1$, $u_i = 1$, $mc_i = p_i/\mu_p$,

and the arbitrage condition (2) holds at $r^k = \bar{r} + \delta$. The disutility levels φ_i are set so that both wage Phillips curves hold at $N_i = 1$. This results in a labor share of 0.59.

The calibration does not target household moments beyond the real rate. In the non-AI steady state, the population-weighted quarterly MPC is 0.22, 14 percent of households are poor and 25 percent wealthy hand-to-mouth, and the top 10 percent hold 68 percent of wealth (Table 2).

3.2 Constructing the AI economy

AI enters the model as higher capital productivity in the exposed sector. In the AI economy, γ_1^k rises from 1 to 1.088. This results in a fall in the aggregate labor share of 5.5 percentage points, which targets the expert forecast of Karger et al. (2026). All other technology parameters $(\Theta_i, \gamma_i^l, \vartheta_i, \gamma_2^k)$, the preference and transaction parameters (β_g, χ_0, χ_1, χ_2), and the fiscal parameters (B, G, Tr) are held at their non-AI values. The tax rate balances the government budget and the real rate clears the asset market. The automation frontiers, factor prices, employment and output are equilibrium outcomes.

Figure 1 shows how the automation frontiers respond to AI. Each task is performed by the cheaper factor. The unit cost of performing a task with capital is the same for all tasks, while the unit cost with labor falls with the task index, so the frontier lies where the two costs are equal (Section 2.3). Higher capital productivity lowers the cost of performing tasks with capital. The exposed frontier therefore advances from 0.33 to 0.41. The tasks in between move from workers to capital.

In the shielded sector, technology does not change, yet the frontier also moves, from 0.33 to 0.37. AI makes the exposed good cheaper, and the two goods are complements, so demand for the shielded good rises. Meeting this demand requires more shielded labor, and the shielded wage rises by 4.8 percent.[1] The cost of capital hardly changes, because the arbitrage condition (2) ties the rental rate to the real rate, which rises by only four basis points. At a higher wage and an unchanged cost of capital, additional tasks are cheaper to perform with capital, and the frontier advances.

3.3 The two steady states

Table 2 compares the non-AI and the AI steady state. All differences between them are equilibrium consequences of the higher capital productivity in the exposed sector.

[1] Workers are attached to their sector, so the higher shielded wage cannot be arbitraged away by workers switching into the shielded sector. Section 5 relaxes this assumption. The wage gap between the sectors narrows without closing, and the results are essentially unchanged.

Parameter	Description	Value	Source / target
Households			
β	discount factor, exposed group	0.9777	real rate $r = 5.7\%$ p.a. (Kaplan et al., 2018)
ν	Frisch elasticity of labor supply	1	Auclert et al. (2024)
Δ_β	discount gap, $\beta_2 = \beta - \Delta_\beta$	0.0075	Krusell and Smith (1998)
EIS	intertemporal elasticity of substitution	1	log utility
e, sd$(\log e)$	earnings process, states and dispersion	33, 0.858	Kaplan et al. (2018)
χ_0, χ_2	adjustment-cost shape	0.25, 2	Auclert et al. (2024)
χ_1	adjustment-cost level	36	
ϕ	exposed employment share	0.40	Eloundou et al. (2024); Felten et al. (2023)
Returns, fiscal and monetary policy			
ζ	liquid wedge	–	liquid return 2% p.a.
B	government debt / annual GDP	0.26	Kaplan et al. (2018)
G	government spending / GDP	0.13	
Tr	transfers / GDP	0.05	tax rate $\tau = 0.30$
ϖ	profit share retained by fund	0.33	Kaplan et al. (2018)
ϕ_π	Taylor-rule coefficient	1.5	standard
ρ_m	monetary shock persistence	0.85	calibration choice
Production			
σ	task elasticity	0.8	Acemoglu and Restrepo (2019)
ϑ_i	comparative-advantage slope	1	normalization
I_i^*	automation frontier, non-AI	0.33	labor share 0.59
δ	depreciation, annual	0.07	Kaplan et al. (2018)
ϵ_I	investment adjustment	4	Auclert et al. (2024)
σ_a	utilization-cost slope	0.01	near-costless utilization
μ_p	price markup	10/9	Kaplan et al. (2018)
μ_w	wage markup	1.001	normalization
Nominal rigidities and demand			
$f_{p,1}, f_{p,2}$	price reset frequency, exposed / shielded	0.31, 0.15	Nakamura and Steinsson (2008); Pastén et al. (2020)
$f_{w,1}, f_{w,2}$	wage reset frequency, exposed / shielded	0.35, 0.15	Grigsby et al. (2021)
Γ	ε_w/ν	5	Galí (2011)
ω_1	final-good weight, exposed	0.5	symmetric sectors
η_c	cross-sector elasticity	0.5	complements
AI economy			
γ_1^k	capital productivity, exposed sector	1.088	labor share -5.5 pp (Karger et al., 2026)

Table 1: Calibration, quarterly unless stated otherwise. All other parameters are identical in the non-AI and the AI economy; in the non-AI economy $\gamma_1^k = 1$.

More capital and a larger investment base. The AI economy holds more capital. Advancing the exposed frontier from 0.33 to 0.41 hands eight percent of that sector's

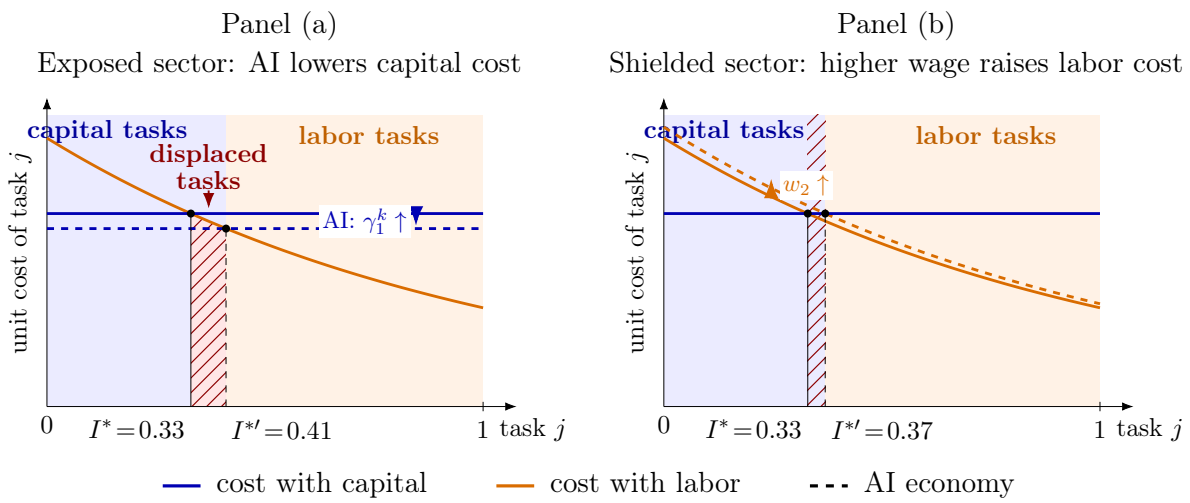

Figure 1: The automation frontiers in the non-AI and the AI economy. Each task is performed by the cheaper factor. The cost of performing task j with capital is r^k/γ_i^k, the cost with labor is $w_i\, e^{-\vartheta j}/\gamma_i^l$ (Section 2.3). In the exposed sector, AI lowers the capital cost and the frontier advances from 0.33 to 0.41. In the shielded sector, technology is unchanged and the higher wage raises the labor cost, advancing the frontier from 0.33 to 0.37.

tasks from workers to machines, and the spillover moves the shielded frontier from 0.33 to 0.37 (Figure 1). Both shifts raise the capital weight in equation (1). With the rental rate pinned at $\bar{r} + \delta$, firms therefore want to hold more capital (Figure 2, bottom row). The stock is 30 percent larger, the capital-output ratio rises from 2.43 to 2.85, and output is 11 percent higher. The labor share falls from 0.59 to 0.53, and the capital share rises from 0.31 to 0.37. More capital also means more investment. In steady state, investment replaces depreciation, and a larger stock requires more of it. Investment rises from 17.5 to 20.5 percent of output. The AI economy thus invests on a permanently larger base. Because households own the capital through the equity fund, a rate cut also has more wealth to revalue.

Labor income shifts between the groups. Labor income per worker falls by 4.4 percent in the exposed sector and rises by 5.0 percent in the shielded one (Figure 2). The two groups adjust on different margins. Exposed workers lose through hours, which fall by 4.3 percent at a nearly unchanged wage. Shielded workers gain through the wage, which rises by 4.8 percent at nearly unchanged hours.

These margins follow from how labor demand and wage setting move (Figure 2, top row). Labor demand moves in opposite directions in the two sectors. In the exposed sector, capital takes over tasks, so firms need fewer hours per unit of output. In the

	Steady state	
	Non-AI	AI
Technology and factor prices		
Capital productivity, exposed sector γ_1^k	1.000	1.088
Automation frontier exposed / shielded	0.33 / 0.33	0.41 / 0.37
Labor share	0.59	0.53
Capital income share	0.31	0.37
Capital / annual GDP	2.43	2.85
Output (non-AI = 100)	100.0	110.7
Investment / GDP	0.17	0.21
Real interest rate (%, annualized)	5.70	5.74
Relative price $q = p_1/p_2$	1.00	0.94
Hours, exposed / shielded	1.00 / 1.00	0.96 / 1.00
Labor income per worker, exposed / shielded	1.47 / 0.98	1.41 / 1.03
Tax rate τ	0.30	0.30
Households		
Discount factor β (exposed)	0.9777	0.9777
MPC, exposed / shielded	0.26 / 0.19	0.29 / 0.19
MPC, population-weighted	0.22	0.23
Top-10% / top-1% wealth share	0.68 / 0.18	0.66 / 0.17
Hand-to-mouth, poor / wealthy	0.14 / 0.25	0.12 / 0.26
Wealthy hand-to-mouth, exposed / shielded	0.47 / 0.10	0.54 / 0.08

Table 2: Comparison of the non-AI and the AI steady state. Output is indexed to 100 in the non-AI economy, and hours are normalized to one in the non-AI economy. MPCs are quarterly. The two economies differ only in γ_1^k (Table 1). All other differences are equilibrium outcomes.

shielded sector, labor demand rises with the demand for the shielded good (Section 3.2). Wage setting shifts in the same direction in both sectors. Households consume more in the AI economy, so their marginal utility is lower, and at given hours the unions' wage-setting condition requires a wage about 4.5 percent higher. In the exposed sector, the lower labor demand meets the higher wage-setting curve at fewer hours and a nearly unchanged wage. In the shielded sector, the higher labor demand meets it at a higher wage and nearly unchanged hours.

Because the exposed sector employs the smaller group, its workers remain the higher paid. Labor income per worker is 1.47 against 0.98 before AI and 1.41 against 1.03 after.

Propensities to consume. AI barely moves the propensities to consume. The quarterly impact MPCs are 0.26 in the exposed and 0.19 in the shielded group of the non-AI economy, and 0.29 and 0.19 in the AI economy. The population-weighted averages are 0.22 and 0.23. The exposed group has the higher MPC in both economies although it is the higher-paid one. It holds the illiquid capital and little liquid wealth, and 47 percent

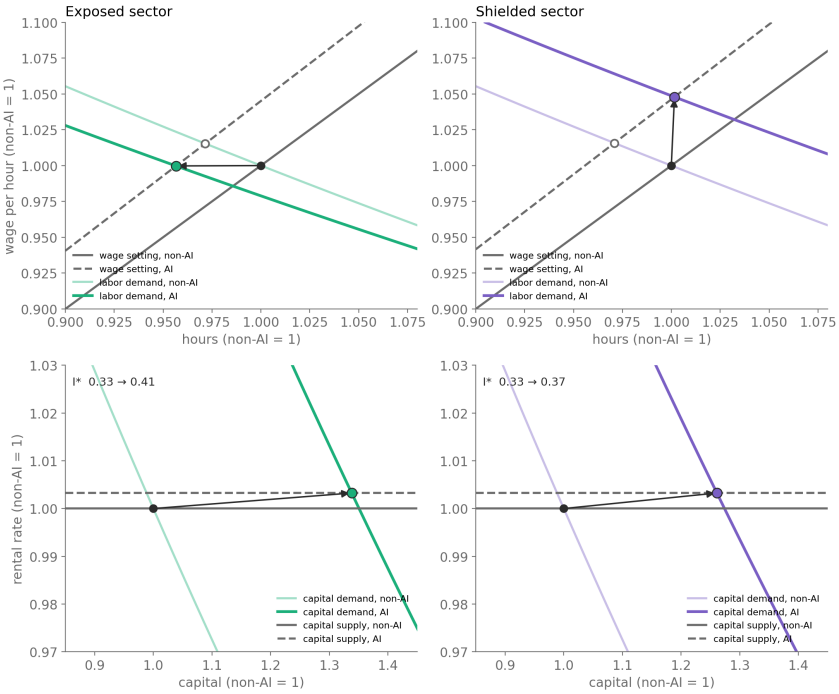

Figure 2: Factor markets in the non-AI and the AI steady state (non-AI = 1). The top row shows the labor market of each sector, with labor demand and the wage-setting curve. The bottom row shows the capital market, with capital demand and the capital supply pinned by the real rate. The arrows mark the move from the non-AI to the AI equilibrium. In the exposed sector, the labor market adjusts in hours at a nearly unchanged wage. In the shielded sector, it adjusts in the wage at nearly unchanged hours. In both sectors, capital demand shifts out at a nearly unchanged rental rate.

of its members are wealthy hand-to-mouth in the non-AI economy and 54 percent in the AI economy. The shielded group, being the less patient, holds almost no illiquid wealth but keeps a liquid buffer, which lowers its propensity to consume. When the responses to the monetary shock differ between the two economies in Section 4, the difference therefore comes from the income that reaches households, not from a change in their spending behavior.

Distribution. Figure 3 shows the wealth distribution in the two economies. The AI economy's larger capital stock is held by the same households under the same preferences, so the distribution of wealth compresses only slightly. The top-10 percent share falls from 0.68 to 0.66. What changes is the distribution of income, which shifts from the exposed sector's wage bill toward capital income. Ownership is concentrated by group. The shielded households hold almost no illiquid wealth, so the capital stock, including the

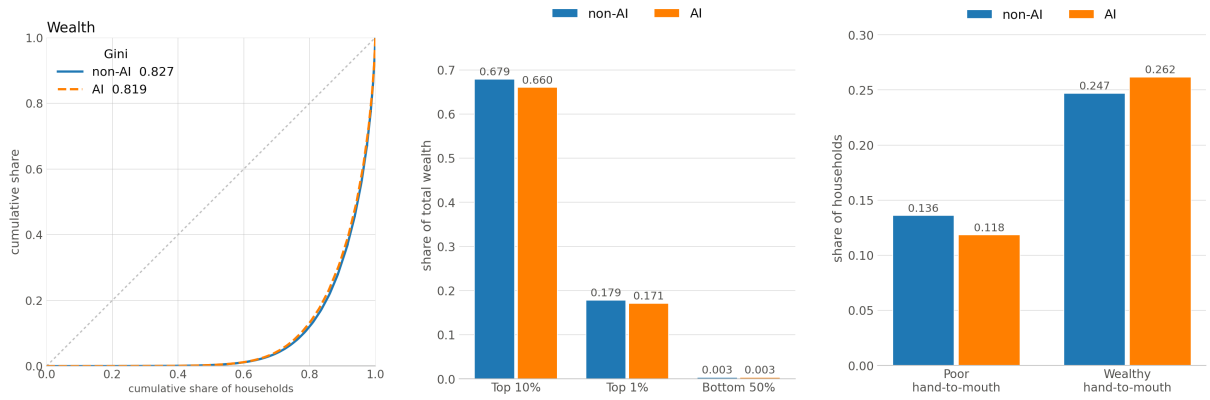

Figure 3: Steady-state wealth distributions in the two economies.

capital that automates the exposed sector, is owned by the exposed group. We assume only the difference in patience, and both groups can hold both assets. The concentration of ownership is an equilibrium outcome, and it matters for the capital gains in Section 4.

Implications for transmission. The differences documented in this section shape the response to the monetary policy shock in Section 4. Investment operates on a permanently larger base, so the same percentage response is more spending. The exposed sector needs less labor per unit of output, so machines serve part of any additional demand. The shielded wage is permanently higher, so the hours that additional demand creates there earn more income. The propensities to consume are nearly unchanged, so these income differences drive the response, not spending behavior.

4 Monetary transmission in an AI economy

In this section, we analyze the transmission of monetary policy in the AI economy relative to the non-AI economy. Section 4.1 presents the aggregate responses and shows that the same monetary policy shock stimulates the AI economy more. Section 4.2 decomposes the consumption response, the largest component of GDP, by household group and shows that the amplification comes mostly from the shielded households. Section 4.3 decomposes it by transmission channel and shows that the channels themselves are unchanged, while the labor income of the shielded sector carries most of the difference. Section 4.4 explains these differences as a chain that runs from the AI steady state back into demand. The same investment response operates on a larger capital base, machines absorb part of the exposed sector's share of the additional demand, and the shielded sector's share becomes

labor income that is spent in turn. Section 4.5 tests the mechanism. The consumption amplification disappears when the investment response is shut off, and it vanishes in a representative-agent version of the model.

4.1 Aggregate responses

We now study the aggregate effects of monetary policy in the AI economy relative to the non-AI economy. Figure 4 shows the responses of the main aggregates in the two economies to the same monetary policy shock, an easing that lowers the real rate by 25 basis points annualized on impact and phases out gradually. The real-rate path is identical across the two economies by construction, as the first panel confirms, and responses are expressed in percent of the respective economy's own steady state.

Output, employment, and consumption respond more strongly in the AI economy. On impact, output rises by 1.0 percent against 0.86 in the non-AI economy, 17 percent more. Hours rise by 0.93 against 0.80, 16 percent more. Consumption rises by 0.62 against 0.58 percent, 7 percent more. Investment and inflation are essentially the same in both economies. Investment rises by 1.05 percent in both and follows the same path in every period. Inflation rises by 0.159 percentage points against 0.152.

The same cut thus stimulates the AI economy more. We refer to this difference as the amplification. Section 4.2 shows who carries it.

4.2 Who carries the amplification

Figure 5 splits the consumption response by household group. The two groups are the steady-state groups of Section 3.3, the exposed workers who lost income through hours and the shielded workers who gained through the wage.

The consumption of the exposed households responds almost as if AI had never happened. Their response in the AI economy lies essentially on top of its non-AI counterpart at every horizon. The consumption of the shielded households responds more strongly, and the gap between their two responses accounts for most of the difference in aggregate consumption.

Monetary policy still works through both groups, in both economies. What AI adds is the amplification. The next subsection decomposes the consumption response by transmission channel to study where this amplification stems from.

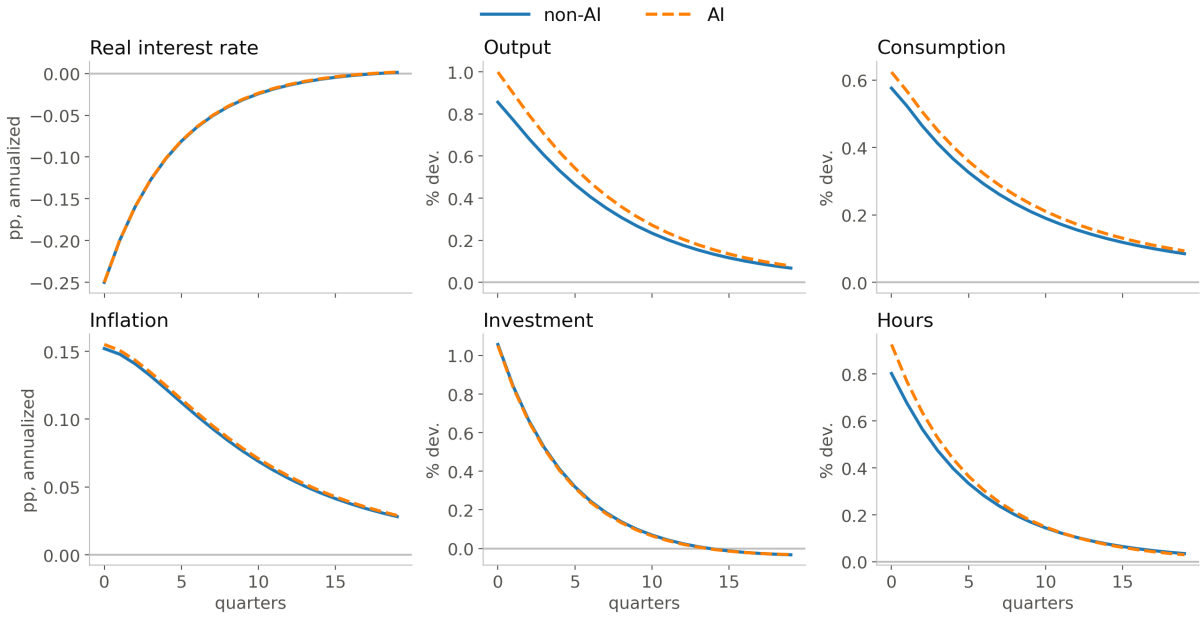

Figure 4: Impulse responses to a −25bp annualized easing in the non-AI economy (solid) and the AI economy (dashed). Real rate and inflation in annualized percentage points; other variables in percent deviation from the respective steady state.

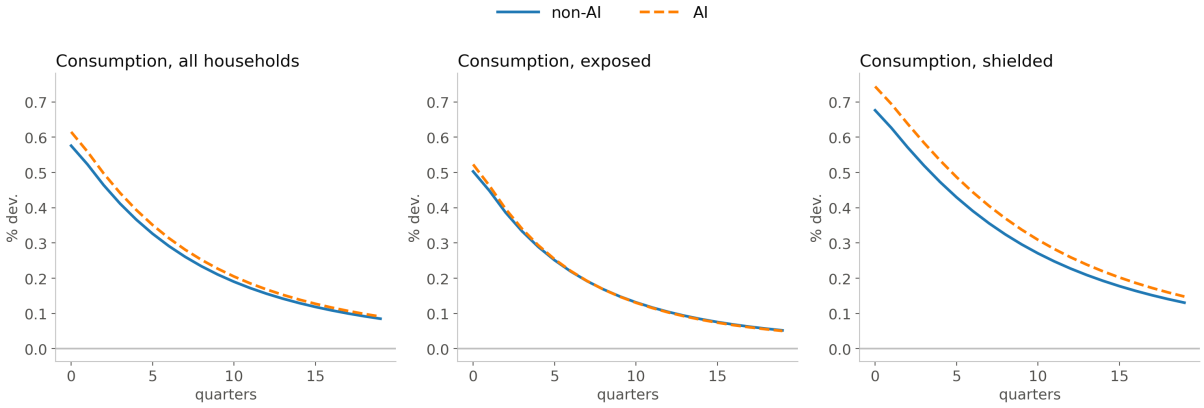

Figure 5: Consumption response by household group, in percent of the group's steady-state consumption. Non-AI economy solid, AI economy dashed.

4.3 Which channel carries it

Figure 6 decomposes the aggregate consumption response into the channels through which the cut reaches households, following Kaplan et al. (2018). Each channel measures the consumption response that one equilibrium path generates on its own, the path of labor income, of taxes, of the liquid return, or of illiquid returns and profits, while all other paths are held at their steady-state values. Mechanically, each path is fed through the

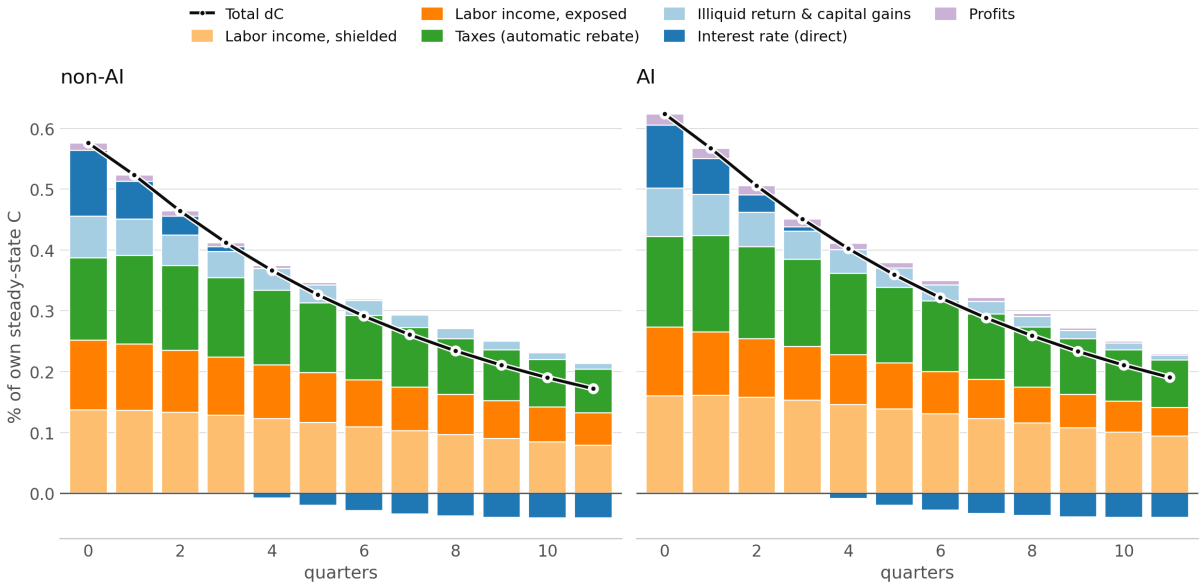

Figure 6: Decomposition of the consumption response, in percent of each economy's steady-state consumption. Each bar feeds one equilibrium price path through the household Jacobians holding the others at steady state; the labor-income term is split by sector. "Taxes" is the rebate implied by the balanced-budget rule.

household Jacobians, and the labor-income term is split by the sector that pays the wage. The sum of the channels equals the total response in Figure 4.

The figure shows that the non-AI economy reproduces the transmission anatomy of Kaplan et al. (2018). The labor-income channels dominate the response, the tax rebate comes next, and the direct effect of the liquid return is small. The same pattern holds in the AI economy, with the same channels in the same order. What changes is how much each channel carries.

Figure 7 shows that shift directly, as the difference between the two decompositions, channel by channel. The black line is the difference between the two consumption responses. The capital-gains channel contributes early and fades. The direct effect of the liquid return contributes negatively; it is essentially unchanged in absolute terms and simply counts for less relative to the AI economy's larger response. The tax rebate contributes throughout; with a balanced budget, it scales with the labor income that the other channels create. The channel that stays is the labor income of the shielded sector, two-fifths of the difference on impact and one half cumulated over the first year. The exposed sector's labor income contributes a little at first and then turns negative. The additional transmission of the AI economy reaches households as labor income in the sector whose technology has not changed; Section 4.4 explains why.

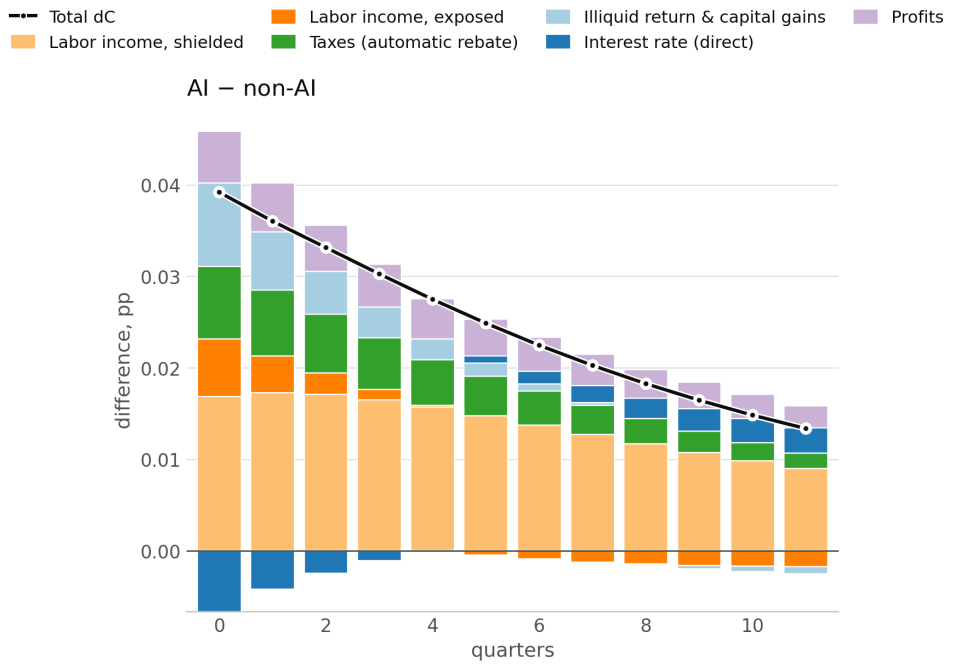

Figure 7: Difference between the two decompositions, AI minus non-AI, channel by channel. The black line is the difference between the two consumption responses.

4.4 The mechanism

The amplification starts from the capital stock that AI has accumulated and reaches households as labor income. The cut raises the price of capital in both economies and triggers the same investment response. However, in the AI economy, that response operates on a larger base, so the initial demand impulse is larger. The exposed sector serves its part of the additional demand partly with machines, the shielded sector with hours, and the resulting labor income is spent in turn. The capital gains from the cut, in contrast, accrue to the illiquid account and are mostly saved. The following paragraphs take these steps in order.

The impulse is larger. The investment response is nearly the same in the two economies in percentage terms (Figure 4), but the AI economy invests on a permanently larger base. Steady-state investment is 20.5 percent of output in the AI economy against 17.5 percent in the non-AI economy (Section 3.3), so the same percentage response is more investment spending relative to the size of the economy. This amplifies the consumption response. The additional investment spending has to be produced, and the production generates labor income that households spend. This feedback from investment to consumption is

a general property of HANK models. Auclert et al. (2020) show that about half of the consumption response to a monetary policy shock disappears when the investment response is shut off, because investment demand raises labor income and households with high propensities to consume spend it. In the AI economy, the same feedback operates on more investment demand.

The exposed channel responds about as much as without AI. The expectation that automation weakens the labor-income channel rests on partial-equilibrium reasoning. It holds the response of aggregate demand fixed and asks how a given increase in demand is divided between labor and capital. Under this reasoning, the channel shrinks with the steady-state labor income that it transmits, which is about four percent lower in the exposed sector (Table 2). However, in general equilibrium, the demand response is not fixed. AI automates part of the work, so each unit of exposed output generates less labor income than without AI. But the impulse of the previous paragraph raises the number of units, and the additional units require labor as well. Less labor income per unit meets more units, and the two roughly offset. Exposed labor income responds 1.15 times as strongly as in the non-AI economy. Spending behavior barely changes either. The quarterly MPC of the exposed group rises by three cents, as more of its members become constrained (Section 3.3). Taken together, the lower income base, the slightly stronger response, and the slightly higher propensity to consume leave the exposed channel about five percent larger than its non-AI value.

Shielded labor income carries more. The larger impulse raises spending on the shielded good as well. The shielded sector serves this additional demand with workers. The labor needed per unit of output is nearly the same as without AI, so the demand passes through to hours. Wages are sticky in the short run, and more so in the shielded sector (Table 1), so the adjustment falls on quantities, and each hour is paid at the permanently higher steady-state wage. Shielded labor income therefore responds 1.16 times as strongly as in the non-AI economy, and it is spent in turn. The multiplier is the same as in the non-AI economy, but it operates on more labor income. The impact propensities in Table 2 understate this spending. They are responses to a one-off transfer, while a monetary easing raises income for many quarters, and the shielded households, who hold a liquid buffer, spend considerably more out of a persistent path than out of a windfall. The shielded channel ends up twelve percent larger.

Capital gains add little to consumption. The cut also revalues the existing capital stock, and the AI economy has more capital to revalue. It holds 17 percent more capital relative to GDP, 2.85 against 2.43 times annual GDP, and 22 percent more relative to consumption. The capital-gains channel is accordingly 1.13 times larger in the AI economy. On impact this is the second-largest piece of the difference between the two consumption responses, about a quarter of it, but it decays quickly. Cumulated over the first year it accounts for a sixth. The profit channel is smaller still.

Three features of the model keep the channel small. The revaluation itself is modest, because the shock is transitory. The gains accrue to the illiquid account, out of which households consume little at the margin. And they accrue to the exposed group. The shielded households, being the less patient, hold almost no illiquid wealth, so the capital stock is owned by the households whose labor income AI has lowered. Of these, 47 percent are wealthy hand-to-mouth in the non-AI economy and 54 percent in the AI economy, and the remainder are unconstrained savers. Neither type spends much of a gain that arrives in the illiquid account. Chodorow-Reich et al. (2021) estimate a marginal propensity to consume out of stock-market wealth of about three cents per dollar, operating mainly through local labor demand rather than through shareholders' own spending. The channel here is of the same order.

This has two implications. The larger capital stock of the AI economy does not make monetary policy a wealth-effect instrument. The transmission remains a labor-market phenomenon in both economies. And the shift of income from labor to capital does not, by itself, change the response of demand to a rate cut. The capital owners spend little of their capital income in either economy, so what they gain in the AI economy they did not spend in the non-AI economy either. The two-agent mechanism of Fornaro and Wolf (2026), in which automation depresses demand by shifting income to low-spending capitalists, therefore operates on the level of demand in our model as well, but it does not weaken the response of demand to policy.

4.5 Testing the mechanism

The chain of Section 4.4 claims that the source of the additional impulse is the investment response, and the reason the resulting income becomes consumption is household heterogeneity. We remove each in turn.

Shutting off the investment response. We repeat the experiment with the investment response shut off, holding the same real-rate path. Figure 8 reports the result.

The consumption amplification disappears, meaning that the consumption responses of the two economies lie on top of each other. The small remaining difference is carried by capital gains and profits, while labor income on net contributes negatively. The shielded channel stays slightly positive, the exposed channel turns negative, and the two roughly cancel.

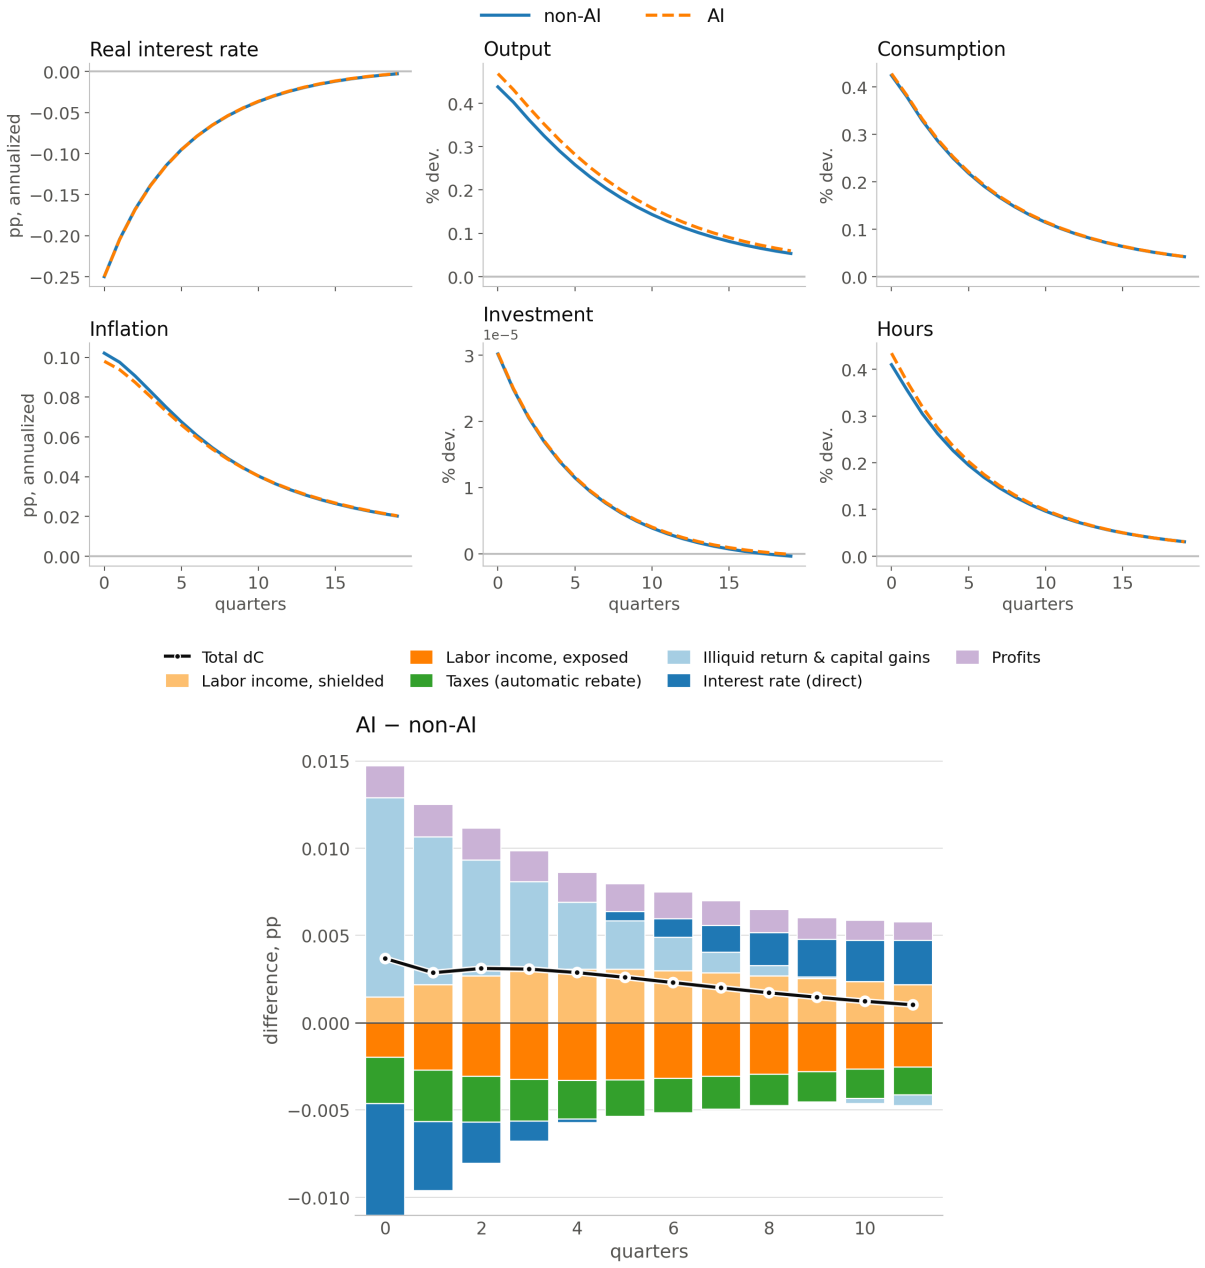

Figure 8: The experiment with the investment response shut off. Top: aggregate responses, non-AI (solid) and AI (dashed). Bottom: difference between the two consumption decompositions, AI minus non-AI.

The experiment also shows what the partial-equilibrium expectation of Section 4.4 got right. Without the quantity response of investment, less of the spending in the exposed sector becomes labor income, and the exposed channel contributes negatively, exactly as the expectation predicts. The prediction fails in the full model because quantities respond. The additional investment demand raises output in both sectors, and the exposed sector's lower labor content per unit meets more units.

A representative-agent economy. We solve a representative-agent version of the model with the same production side and the same two steady states. Figure 9 shows the responses to the same shock. The consumption amplification vanishes. The representative agent consumes according to the permanent income hypothesis, so the additional labor income that the AI economy generates is saved rather than spent, and the investment feedback of Section 4.4 has nothing to work with. AI shifts income between sectors and factors in that economy too, but with a representative agent, the composition of income is irrelevant for demand.

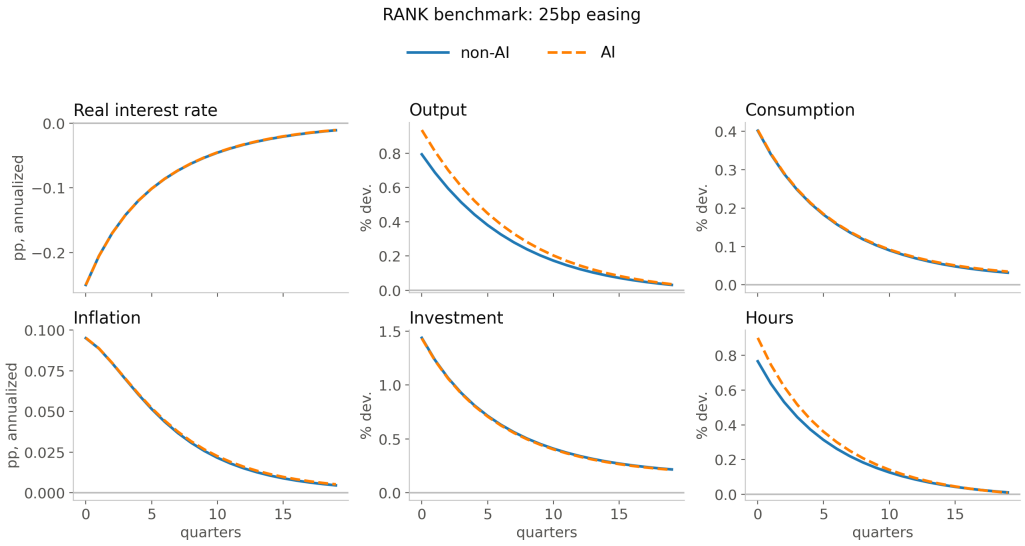

Figure 9: Aggregate responses in the representative-agent version, non-AI (solid) and AI (dashed).

Together, the two tests give the result its anatomy. The larger capital stock supplies the additional demand, through an investment response that operates on a larger base. The task allocation decides who serves that demand, machines in the exposed sector and workers in the shielded one. And household heterogeneity decides whether the resulting income becomes consumption. By removing the investment response, the amplification loses its source. By removing the heterogeneity, it loses its carrier.

5 Robustness: sector labor mobility

The baseline assumes that households are attached to one sector. Over the horizon of a monetary policy shock this is the empirically relevant case, because occupational switching is slow and costly. Over the long run, however, workers can move toward the sector where wages have risen. This could erode the higher shielded wage in the AI economy, and with it the amplification on which our results rest. In this section we therefore relax the assumption and allow workers to move between sectors.

The switching technology follows Beraja and Zorzi (2025). Workers switch sectors at an annual rate of 10 percent, and switching responds to the wage gap. When the shielded sector pays 1 percent more relative to the exposed one, the number of shielded relative to exposed workers rises by 1 percent in the long run. Those who switch spend on average about four months (1.25 quarters) without labor income before starting in the new sector. Each household keeps a home sector in which it spends 90 percent of its working time, so switching reallocates workers at the margin rather than dissolving the sector attachment.

Each household keeps its discount factor when it switches sectors, so both sectors employ patient and impatient households.[2] The AI economy is constructed as in Section 3.2, with the same fall in the labor share of 5.5 percentage points.

Figure 10 shows that the amplification is of the same order as in the baseline. The consumption response is still 6 percent larger in the AI economy than in the non-AI economy, against 7 percent before, and the output response is 16 percent larger, against 17 percent. In the AI economy, workers move toward the higher shielded wage, so the shielded sector employs more workers and the exposed sector fewer than in the non-AI economy. The inflow narrows the wage gap between the two sectors, but only partly. The shielded wage still rises clearly, and so does labor income per shielded worker. The labor income of the shielded sector remains the largest component of the additional consumption response. Switching itself barely responds to the monetary shock, so mobility matters through the steady state rather than through the dynamics.

The shielded wage gain survives because only part of it is open to arbitrage. The shielded wage rises in the AI economy for two reasons. The AI economy is richer, so lower marginal utility pushes wage setting up in both sectors alike. Because this affects both sectors equally, it gives workers no reason to move. In addition, the exposed good becomes cheaper and, the two goods being complements, demand for the shielded good rises. Only this second reason opens a wage gap between the sectors, and only the gap draws workers

[2] Unions price hours per member, so that entry into a sector lowers its wage. With constant sector sizes, this wage-setting condition coincides with the baseline one.

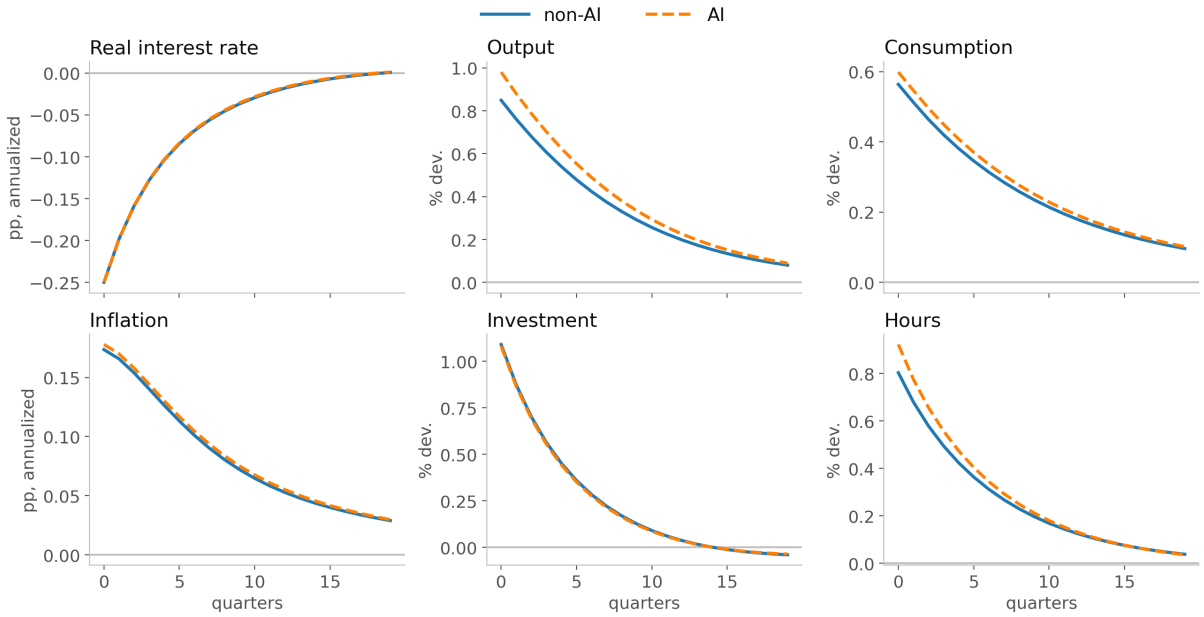

Figure 10: Impulse responses to the same −25bp easing in the economies with sector labor mobility, non-AI (solid) and AI (dashed).

into the shielded sector. Their arrival lowers the shielded wage, and the net flows stop once the remaining gap no longer compensates for the four months without income that a switch costs. The gap therefore narrows without closing. Mobility compresses the gap between the sectors and leaves the common rise in wage setting untouched.

Mobility does change the level of transmission. Households insure against the spells without a job with liquid buffers, which lowers the share of hand-to-mouth households (from 14 to 5 percent poor and from 25 to 14 percent wealthy) and the propensities to consume. The consumption response to the cut is therefore smaller in both economies, but the ratio between them is not.

Mobility thus changes who earns the shielded wage and how much households save, but not the anatomy of the result. The larger capital stock still supplies the additional demand, the task allocation still decides who serves it, and household heterogeneity still decides whether the resulting income becomes consumption. The result does not hinge on households being attached to a particular sector.

6 Conclusion

We build a two-sector HANK model with task-based production to study whether AI changes the transmission of monetary policy. We find that the same rate cut moves

output, employment, and consumption more strongly in the AI economy than in the non-AI economy. The additional impulse comes from investment. Its response is the same in both economies, but the AI economy invests on a permanently larger capital base, and when we shut off the investment response, the consumption amplification disappears. The additional transmission runs through the shielded, not the automated sector. Automation lowers exposed workers' steady-state labor income but not its response to the shock, so the exposed labor-income channel is essentially unchanged, while the extra demand reaching the shielded sector becomes hours, paid at the higher wage, and income that is spent in turn. In a representative-agent version of the model the consumption amplification vanishes, because the composition of income is irrelevant for demand there. Automation thus shifts the incidence of monetary policy toward the workers it does not replace.

We leave two extensions for future work. The first is sector-specific displacement risk. The baseline holds exposed workers' earnings risk fixed across the two economies. However, the rise in permanent income risk at the top of the distribution documented by Braxton et al. (2026) suggests that AI may increase their income risk. Higher income risk would induce precautionary saving and lower exposed workers' propensity to consume, thereby weakening the amplification through the household consumption channel. Following Broer et al. (2026) and Patterson (2023), an employment state with sector-specific, cyclically varying separation and job-finding rates would allow the model to discipline the risk gap between exposed and shielded workers using evidence on AI exposure.

The second extension is the ownership of the automating capital. In the present model, the capital stock is held through a single fund that blends both sectors, so the model cannot distinguish a worker who owns the machines replacing her from one who owns the machines in the other sector. Sector-specific claims would instead tie the revaluation of automating capital to the workers it displaces, allowing us to study how ownership of automating capital affects the capital-gains channel of monetary policy.

References

Acemoglu, D., F. Kong, and P. Restrepo (2025): "Tasks at Work: Comparative Advantage, Technology and Labor Demand," in *Handbook of Labor Economics*, Elsevier, vol. 6.

Acemoglu, D. and P. Restrepo (2018): "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment," *American Economic Review*, 108, 1488–1542.

——— (2019): "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, 33, 3–30.

——— (2022): "Tasks, Automation, and the Rise in U.S. Wage Inequality," *Econometrica*, 90, 1973–2016.

AEA Committee on Economic Statistics (2025): "Measuring the Economic Effects of AI," Report.

Althoff, L. and H. Reichardt (2026): "Task-Specific Technical Change and Comparative Advantage," NBER Working Paper 35353, National Bureau of Economic Research.

Auclert, A. (2019): "Monetary Policy and the Redistribution Channel," *American Economic Review*, 109, 2333–2367.

Auclert, A., B. Bardóczy, M. Rognlie, and L. Straub (2021): "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 89, 2375–2408.

Auclert, A., M. Rognlie, and L. Straub (2020): "Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model," Tech. rep., National Bureau of Economic Research.

——— (2024): "The intertemporal keynesian cross," *Journal of Political Economy*, 132, 4068–4121.

Bárány, Z. L., A. Patel, and C. Siegel (2026): "Does Automation Lower the Labor Share?" CEPR Discussion Paper 21700.

Bayer, C., B. Born, and R. Luetticke (2024): "Shocks, Frictions, and Inequality in US Business Cycles," *American Economic Review*, 114, 1211–1247.

Beraja, M. and N. Zorzi (2025): "Inefficient automation," *Review of Economic Studies*, 92, 69–96.

Braxton, J. C., K. F. Herkenhoff, C. Huang, M. Nattinger, J. L. Rothbaum, and L. D. W. Schmidt (2026): "Rising Income Risk at the Top," NBER Working Paper 35664, National Bureau of Economic Research.

Broer, T., J. Kramer, and K. Mitman (2026): "The Curious Incidence of Monetary Policy Across the Income Distribution," CEPR working paper.

Carroll, C. D., J. Slacalek, K. Tokuoka, and M. N. White (2017): "The Distribution of Wealth and the Marginal Propensity to Consume," *Quantitative Economics*, 8, 977–1020.

Chodorow-Reich, G., P. T. Nenov, and A. Simsek (2021): "Stock Market Wealth and the Real Economy: A Local Labor Market Approach," *American Economic Review*, 111, 1613–1657.

Eloundou, T., S. Manning, P. Mishkin, and D. Rock (2024): "GPTs are GPTs:

Labor Market Impact Potential of LLMs," *Science*, 384, 1306–1308.

Felten, E. W., M. Raj, and R. Seamans (2023): "How will Language Modelers like ChatGPT Affect Occupations and Industries?" Working paper.

Fornaro, L. and M. Wolf (2021): "Monetary Policy in the Age of Automation," CEPR Discussion Paper 16416.

——— (2026): "Macroeconomic Policies for AI," Mimeo.

Freund, L. B. and L. F. Mann (2026): "Job Transformation, Specialization, and the Labor Market Effects of AI," Working paper; first version August 2025.

Fueki, T., S. Katsuki, I. Muto, and Y. Sugisaki (2026): "Automation and Nominal Rigidities," Mimeo.

Galí, J. (2011): "The Return of the Wage Phillips Curve," *Journal of the European Economic Association*, 9, 436–461.

Grigsby, J., E. Hurst, and A. Yildirmaz (2021): "Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data," *American Economic Review*, 111, 428–471.

Kaplan, G., B. Moll, and G. L. Violante (2018): "Monetary Policy According to HANK," *American Economic Review*, 108, 697–743.

Karger, E., O. Kuusela, J. Abaluck, K. A. Bryan, B. Halperin, T. R. Jones, C. Murphy, P. Trammell, M. Reynolds, D. Mayland, R. Viswanathan, A. Mittal, R. Ceppas de Castro, and J. Rosenberg (2026): "Forecasting the Economic Effects of AI," NBER Working Paper 35046, National Bureau of Economic Research.

Krusell, P. and A. A. Smith, Jr. (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106, 867–896.

Lenzu, S. (2026): "Artificial Intelligence and Monetary Policy: A Framework and Perspective on Cyclical Transmission, Structural Transition, and Financial Stability," Staff Report 1192, Federal Reserve Bank of New York.

Lim, H. (2026): "Automation, Income Distribution, and Aggregate Demand in a Two-Asset HANK Model," Working paper.

Lindenlaub, I., R. Oh, M. A. Rodríguez, and L. Veldkamp (2026): "Beyond Exposure: Predicting AI Adoption Based on Comparative Advantage," Working paper.

McKay, A., E. Nakamura, and J. Steinsson (2016): "The Power of Forward Guidance Revisited," *American Economic Review*, 106, 3133–3158.

Moll, B., L. Rachel, and P. Restrepo (2022): "Uneven Growth: Automation's Impact on Income and Wealth Inequality," *Econometrica*, 90, 2645–2683.

Nakamura, E. and J. Steinsson (2008): "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, 123, 1415–1464.

Oberfield, E. and D. Raval (2021): "Micro Data and Macro Technology," *Econometrica*, 89, 703–732.

Pastén, E., R. Schoenle, and M. Weber (2020): "The Propagation of Monetary Policy Shocks in a Heterogeneous Production Economy," *Journal of Monetary Economics*, 116, 1–22.

Patterson, C. (2023): "The Matching Multiplier and the Amplification of Recessions," *American Economic Review*, 113, 982–1012.

U.S. Census Bureau (2026): "Business Trends and Outlook Survey," BTOS AI Supplement.

Werning, I. (2015): "Incomplete Markets and Aggregate Demand," NBER Working Paper 21448, National Bureau of Economic Research.