BANK OF CANADA
BANQUE DU CANADA

Staff working paper / Document de travail du personnel—2026-27

Last updated: July 24, 2026

Monetary Policy in an AI-Driven Two-Speed Economy

Joshua Brault
Data and Digital Services
Bank of Canada
JBrault@bank-banque-canada.ca

Maryam Haghighi
Data and Digital Services
Bank of Canada
MHaghighi@bank-banque-canada.ca

Jing Yang
Economic and Financial Research Department
Bank of Canada
jyang@bank-banque-canada.ca

DOI: https://doi.org/10.34989/swp-2026-27 | ISSN 1701-9397

Monetary Policy in an AI-Driven Two-Speed Economy

Joshua Brault* — Bank of Canada
Maryam Haghighi† — Bank of Canada
Jing Yang‡ — Bank of Canada

June 2026§

Abstract

We study the monetary policy response to AI adoption in a two-sector New Keynesian model with a task-based microfoundation, sticky prices, and downward nominal wage rigidity. We distinguish between two forms of AI-driven technological change: augmentation, which raises the productivity of labor within existing tasks, and automation, which displaces labor by reallocating tasks from workers to machines, contracting the set of tasks requiring human input. In the short run, both shocks lower labor demand on net, and with downward nominal wage rigidity, unemployment emerges unless monetary policy provides accommodation. But because monetary policy operates through aggregate demand and cannot target sectors differentially, accommodation that reduces unemployment in the AI-affected sector raises inflationary pressure in the unaffected one, opening a sectoral wedge between the policy rates required to clear the two labor markets. Since, for output-equivalent shocks, automation generates a larger decline in labor demand, the associated wedge is wider and the Phillips curve lies above and to the right of the curve for augmentation—restoring full employment comes at a greater cost of inflation. In addition to the nature of the shock, the aggregate inflationary consequences depend on the breadth of AI adoption across the economy. Under augmentation, as the AI-affected sector grows, its falling sectoral price increasingly offsets the inflation generated elsewhere by monetary accommodation—making aggregate inflation an unreliable signal of the underlying trade-off. For automation both sectoral prices rise and no such offset exists. In our framework, sector-specific AI adoption poses an unambiguous short-run labor market stabilization problem, while its implications for aggregate inflation depend on the nature of technological change, the breadth of adoption, and the response of monetary policy.

*Data and Digital Services Department, Bank of Canada, jbrault@bank-banque-canada.ca.
†Data and Digital Services Department, Bank of Canada, mhaghighi@bank-banque-canada.ca.
‡Economic and Financial Research Department, Bank of Canada, jyang@bank-banque-canada.ca.
§The views expressed in this paper are our own and do not necessarily reflect the official views of the Bank of Canada. We thank John Haltiwanger and participants at the Bank of England's *Transformative artificial intelligence (AI): implications for structural change and monetary policy conference* for helpful comments and discussions.

1 Introduction

The rapid advancement of artificial intelligence—especially generative AI—has sparked considerable debate about its macroeconomic consequences, particularly for labor markets and economic growth. While optimists emphasize productivity gains, skeptics point to displacement effects and the uneven distribution of those gains across sectors and workers. Much of this debate centers on long-run outcomes, but the short-run transition may also be disruptive when nominal rigidities prevent prices and wages from adjusting.

AI is likely to affect sectors unevenly, with the largest impacts concentrated in knowledge-intensive occupations and industries. This sectoral asymmetry raises an important question: how should monetary policy respond when technological change alters both productivity and the distribution of labor demand across the economy, and does the answer depend on whether AI primarily augments or automates human labor? Standard New Keynesian models abstract from these asymmetries, yet they may be central to understanding the short-run stabilization challenge AI poses for central banks.

In this paper we study these questions using a two-sector New Keynesian framework with a task-based microfoundation, sticky prices, and downward nominal wage rigidity. We distinguish between two forms of AI-driven technological change. *Augmentation* refers to AI that raises the productivity of labor within existing tasks, complementing workers and expanding what they can produce. *Automation* refers to AI that displaces labor by reallocating tasks from workers to machines, contracting the set of tasks that require human input. In practice both forms occur simultaneously, with their relative magnitudes varying across occupations and sectors, and evolving over time as AI capabilities expand. We study them separately to isolate their distinct mechanisms and implications for monetary policy. Crucially, AI adoption in our setup is confined to a single sector, and this asymmetry lies at the heart of the stabilization challenge—monetary policy cannot differentially target the affected sector.

Our framework builds on Guerrieri et al. (2021), who study monetary policy in the presence of reallocative demand shocks in a two-sector New Keynesian model with nominal rigidities restricted to the impact period. We adopt the same approach but shift focus to permanent technology shocks and enrich the production side with a task-based microfoundation (Acemoglu and Restrepo 2018). Firms produce output from tasks completed by capital and labor, and augmentation and automation enter as distinct parametric shocks with interpretable effects on labor demand and marginal cost.

In this environment, labor demand in the AI-adopting sector is shaped by three forces. Augmentation raises labor productivity, reducing labor demand for any given amount of output. Automation displaces labor directly by reassigning tasks to capital, lowering labor demand for any given factor prices. Lastly, with nominal rigidities, both augmentation and automation produce a common scale effect whereby movements in marginal cost and prices feed back into the sector's real demand, which in turn affects labor demand in the sector.

In the presence of downward nominal wage rigidity, wages cannot adjust to clear the labor market and any decline in labor demand translates into unemployment unless monetary policy provides accommodation. But because monetary policy operates through aggregate demand, accommodation that restores labor demand in the AI-affected sector also raises inflationary pressure in the unaffected sector. Monetary policy therefore faces a sectoral trade-off between unemployment in the AI-affected sector and inflation in the unaffected one, reflected in what we call the *sectoral wedge*: the difference between the policy rates required to clear the two labor markets. Both augmentation and automation produce sectoral wedges, but for output-equivalent shocks, automation produces a larger decline in labor demand and its wedge is systematically wider. Consequently, the rate cut required to restore full employment is larger, and the inflationary cost correspondingly greater.

The aggregate counterpart to the sectoral wedge is a piecewise-linear Phillips curve, describing the continuum of inflation–unemployment equilibria available to the central bank. Both augmentation and automation produce Phillips curves with three distinct regions: a flat region, where both sectors are demand-constrained and additional accommodation reduces unemployment with negligible inflationary pressure; a binding trade-off region steepening the Phillips curve, where additional accommodation raises prices in the capacity-constrained unaffected sector while reducing unemployment in the AI-adopting sector; and a vertical region, where both sectors are supply-constrained and additional accommodation is purely inflationary. Because automation produces the larger sectoral wedge, its Phillips curve lies above and to the right of augmentation's—at any given unemployment rate inflation is higher under automation than under augmentation.

Our main contribution is to characterize the monetary policy trade-off that arises when AI adoption is sector-specific and nominal rigidities restrict the adjustment of prices and wages. By distinguishing between augmentation and automation, we show that the nature of technological change determines both the severity of the trade-off and whether aggregate inflation is a reliable signal of the underlying sectoral disruption. We also emphasize that the aggregate inflationary consequences of AI are inherently ambiguous, depending on the size of the AI-adopting sector and the policy response rather than on technology alone.

Based on our framework, we report six main findings:

1. **Sector-specific AI creates a two-speed economy that poses a challenge for monetary policy.** Because monetary policy operates through aggregate demand and cannot direct spending across sectors, accommodation that reduces unemployment in the AI-adopting sector raises inflationary pressure in the unaffected one. This opens a *sectoral wedge*—the gap between the policy rates required to clear the two labor markets, $\Delta_i = i_A^{\text{full}} - i_B^{\text{full}} > 0$—and it appears under both forms of AI technological change, reflecting the limits of a single aggregate instrument.

2. **The "benign" case of augmentation still poses a short-run trade-off.** The sectoral wedge is not special to disruptive automation—augmentation, typically viewed as the optimistic case for workers, opens it too. With sticky prices, the fall in marginal cost passes only partly into prices, so the resulting expansion in real demand (and thus labor demand) does not fully offset the decline in labor demand from higher productivity. Labor demand falls in sector B, the wage floor binds, and unemployment arises from a shock that is purely labor-enhancing. This echoes the contractionary employment effect of a positive productivity shock in the single-sector New Keynesian model (Galí 1999), but the two-sector structure denies any costless escape—the accommodation that restores employment in sector B produces inflation in sector A.

3. **Automation poses a more demanding stabilization problem than augmentation.** For output-equivalent shocks, the decline in labor demand is substantially larger under automation. The displacement effect reduces labor demand directly, even at unchanged factor prices. This is compounded by a *capital-substitution externality* in our model: as all firms simultaneously substitute toward capital, the cost of capital rises, pushing up marginal cost and prices and contracting real demand, further reducing labor demand. The sectoral wedge is therefore wider under automation, its Phillips curve lies above and to the right, and restoring full employment requires significantly more accommodation (a cut of 1.48pp for augmentation vs. 3.34pp for automation).

4. **The scope of AI changes how the trade-off appears, not whether it exists.** In our framework, the size of the AI-adopting sector indexes the economy-wide use of AI. A wider footprint does *not* shrink the underlying problem—the sectoral wedge is invariant to sector size. What changes is the wedge's visibility in aggregate inflation. Under augmentation, the adopting sector's price falls, and as that decline carries progressively more weight in the aggregate price index, a sufficiently large adopting sector can allow the central bank to achieve full employment together with zero or even negative aggregate inflation—even as the sectoral wedge, rising prices in sector A, and the underlying sectoral tension persist. Under automation, both sectoral prices rise and no such offset exists, so the aggregate trade-off remains strictly positive at any sector size.

5. **Aggregate inflation is a shock-dependent signal of the trade-off.** Under automation, prices rise in both sectors with accommodation, so aggregate inflation co-moves with the trade-off and reliably tracks its severity regardless of sector size. Under augmentation, sectoral prices move in opposite directions, and as the size of the adopting sector grows, declining P_B increasingly offsets rising P_A in the aggregate index. A central bank watching only headline inflation will therefore under-diagnose the disruption under broad augmentation. The apparent stability is cancellation, not balance.

6. **Stabilizing employment in the short run can delay the long-run structural change required by AI.** Although accommodation can reduce unemployment in the short run, it may delay the necessary structural adjustment that AI-driven technological change requires. Automation, for example, implies a long-run decline in the labor share of the AI-affected sector, and sustained accommodation that supports employment at prevailing wages may postpone rather than facilitate that transition. A central bank that fully offsets the employment consequences of automation risks prolonging the adjustment, substituting monetary stimulus for the structural change that would otherwise occur. This underscores the importance of diagnosing the nature of the shock, since labor market weakness may reflect insufficient demand warranting accommodation or long-run structural change warranting more restraint.

Related Literature. Our paper connects to the rapidly expanding literature on the macroeconomics of AI and automation. Methodologically it rests on the task-based production framework (Acemoglu and Restrepo 2018, 2019) and on two-sector New Keynesian models with nominal rigidities (Guerrieri et al. 2021, 2022), and draws on a large literature on the long-run growth and labor-market consequences of AI (Aghion et al. 2019, Acemoglu 2024, Acemoglu and Restrepo 2019, Aghion and Bunel 2024). We concentrate on the work closest to ours, at the intersection of AI and monetary policy, where the literature remains comparatively sparse. While much of the broader literature emphasizes long-run outcomes, we focus on short-run stabilization and show that the inflationary consequences of AI are not structurally determined but depend on the nature of the shock, sectoral composition, and the policy response.

Within this work, a first strand studies how automation can flatten the Phillips curve. Basso and Rachedi (2025) show that automation erodes workers' bargaining power by threatening displacement, weakening the response of wages to labor-market tightness. Empirically, robot adoption reduces the sensitivity of non-tradable inflation to unemployment by roughly fifteen to seventeen percent. Fueki et al. (2023) emphasize a complementary supply-side channel in a medium-scale NK DSGE framework. Monetary easing raises wages and thereby induces firms to automate, which lifts labor productivity, lowers real marginal costs through a weight effect, and restrains inflation.

A second strand places AI within the broader conduct of stabilization policy. In Fornaro and Wolf (2026), AI automation redistributes income toward capitalists with lower marginal propensities to consume, potentially generating demand-driven slumps that force the central bank to accept cost-push inflation or lean on complementary policies at the ZLB. Lenzu (2026) offers a broad framework in which AI reshapes monetary policy through cyclical transmission, structural transition, and financial stability, altering how shocks map into inflation and shifting the natural benchmarks around which policy is calibrated.

We differ by embedding the augmentation–automation distinction in a two-sector economy with downward nominal wage rigidity, motivated by the empirical concentration of AI exposure in white-collar, cognitively intensive industries such as finance, legal, and professional services (Felten et al. 2021, Eloundou et al. 2023). This generates a sectoral inflation–employment trade-off that monetary policy cannot address in isolation, and implies that the inflationary consequences of AI depend on sectoral composition and the policy response rather than on technology alone.

2 Model

This section describes the setup of the model. Two features are central. First, the economy has two sectors, each populated by a continuum of monopolistically competitive firms producing differentiated varieties with a task-based production structure. A competitive final goods firm bundles sectoral output. Second, there are two nominal rigidities—firms have sticky prices via Rotemberg adjustment costs and nominal wages are downward-rigid—which creates a role for monetary policy. As in Guerrieri et al. (2021), we restrict nominal rigidities to the impact period, and assume wages and prices are fully flexible thereafter. We analyze the model through its *impact-period* response to permanent technology shocks rather than the demand shock in their paper. We consider two types of technological change: *Augmentation*, a labor-augmenting improvement, and *automation*, an expansion in the set of tasks that can be performed by capital.

2.1 Households

A representative household has log preferences and consumes a Cobb–Douglas bundle of sectoral goods (the index omits a multiplicative constant, which leaves all real allocations unchanged):

$$C = C_A^{1-\omega} C_B^{\omega} \qquad \omega \in (0,1), \tag{1}$$

subject to a budget constraint. Optimization implies constant nominal expenditure shares: $P_A C_A = (1-\omega)X$ and $P_B C_B = \omega X$, where X is total nominal expenditure. The consumer price index is given by

$$P = P_A^{1-\omega}\, P_B^{\omega}. \tag{2}$$

Households have access to one-period nominal bonds in zero net supply, which pins down the standard Euler equation $1/X_t = \beta \mathbf{E}_t[(1+i_t)/X_{t+1}]$ under log preferences. To keep the impact period tractable and to compare augmentation and automation, we anchor future expenditure at its pre-shock steady state, $X_{t+1} = \bar{X}$, collapsing the Euler equation so that the current nominal interest i_t uniquely pins down current aggregate expenditure X_t.[1]

[1]We also explored anchoring future expenditure at the value that ensured aggregate price stability after the shock, but the wedges changed by such a small amount ($< 0.01pp$) that we opted to keep it fixed for simplicity.

$$1 + i_t = \frac{1}{\beta}\frac{\bar{X}}{X_t} \tag{3}$$

2.2 Production

2.2.1 Final-goods producers

Each sector j contains a competitive final-goods firm that combines the continuum of varieties $\{y_j(i)\}_{i\in[0,1]}$ into sectoral aggregates using the following CES technology:

$$Y_j \;=\; \left(\int_0^1 y_j(i)^{(\varepsilon-1)/\varepsilon}\,di\right)^{\varepsilon/(\varepsilon-1)}. \tag{4}$$

where $\varepsilon > 1$ is the elasticity of substitution across varieties. Cost minimization yields the standard downward-sloping demand function for each variety in sector j and the sectoral price index:

$$y_j(i) \;=\; \left(\frac{p_j(i)}{P_j}\right)^{-\varepsilon} Y_j, \qquad P_j \;=\; \left(\int_0^1 p_j(i)^{1-\varepsilon}\,di\right)^{1/(1-\varepsilon)}. \tag{5}$$

2.2.2 Differentiated varieties with task-based production

We adopt a task-based microstructure that, under particular conditions, permits an aggregation of a continuum of task-level production technologies into a tractable sector-level production function. Our setup draws on results from earlier contributions (Zeira 1998, Acemoglu and Restrepo 2018, Fornaro and Wolf 2022). Each firm in sector j produces its variety $y_j(i)$ from a continuum of internal production tasks $z \in [0,1]$, aggregated through a Cobb–Douglas technology over tasks.[2] Firm i allocates its own capital $K_j(i)$ and labor $L_j(i)$ across the task continuum subject to the resource constraints

$$\int_0^{\alpha_j} k_j(i,z)\,dz \;=\; K_j(i), \qquad \int_{\alpha_j}^1 \ell_j(i,z)\,dz \;=\; L_j(i). \tag{6}$$

A technological constraint is imposed on automation: tasks in the interval $[0,\alpha_j]$ can be performed by either capital or labor, but tasks in $(\alpha_j,1]$ can be performed only by labor. An increase in α_j therefore represents automation, expanding the set of tasks that can be performed by capital.

We adopt a standard *cost-ordering* assumption. For any task $z \in [0,\alpha_j]$ that can be performed by capital, the effective unit cost of capital is no greater than that of labor, $R_j/Z_j \le W_j/A_j$. That is, capital delivers more output per unit of expenditure than labor on these tasks, and the split of tasks into $[0,\alpha_j]$ and $(\alpha_j,1]$ coincides with the cost-minimizing assignment of factors to tasks. Under our calibrations this ordering holds throughout, and is preserved under shocks that move equilibrium factor prices. Thus, an increase in α_j shifts the task boundary outward, reassigning tasks from labor to capital and displacing workers at prevailing prices.

[2] Fueki et al. (2023) adopt a similar task-based microfoundation, but split the within-firm task block into two competitive layers. Task producers complete individual tasks using capital or labor, and task aggregators combine these tasks via a CES aggregator and sell the bundle to intermediate goods firms. We instead let each intermediate firm aggregate its tasks internally via a Cobb–Douglas aggregator.

Under this assumption, task production takes the piecewise form

$$y_j(i,z) \;=\; \begin{cases} Z_j\, k_j(i,z) & \text{if } z \in [0,\, \alpha_j], \\ A_j\, \ell_j(i,z) & \text{if } z \in (\alpha_j,\, 1], \end{cases} \tag{7}$$

where $y_j(i,z)$ is firm i's output on task z and $k_j(i,z)$, $\ell_j(i,z)$ are its capital and labor allocations to that task.

We also impose a symmetry assumption. Within each factor region, productivity is constant—Z_j for capital on $[0,\alpha_j]$, A_j for labor on $(\alpha_j,1]$—and the Cobb–Douglas aggregator assigns equal weight to every task. Together these imply a uniform allocation of each factor across regions at firm i's cost-minimizing solution,

$$k_j(i,z) \;=\; \frac{K_j(i)}{\alpha_j} \text{ for } z \in [0,\alpha_j], \qquad \ell_j(i,z) \;=\; \frac{L_j(i)}{1-\alpha_j} \text{ for } z \in (\alpha_j,1]. \tag{8}$$

2.2.3 Cobb–Douglas reduced form

It is straightforward to show that by substituting (7) and (8) into the Cobb–Douglas task aggregator $\log y_j(i) = \int_0^1 \log y_j(i,z)\, dz$, a production function for firm i in sector j can be written as

$$y_j(i) \;=\; \left(\frac{Z_j K_j(i)}{\alpha_j}\right)^{\alpha_j} \left(\frac{A_j L_j(i)}{1-\alpha_j}\right)^{1-\alpha_j}. \tag{9}$$

From firms to sector. In equilibrium, all firms producing differentiated varieties within sector j are identical and face symmetric demand from (5), so they choose the same factor inputs $K_j(i) = K_j$, $L_j(i) = L_j$ and produce the same variety output $y_j(i) = y_j$. The Dixit–Stiglitz aggregator (4) then delivers

$$Y_j = y_j = \left(\frac{Z_j K_j}{\alpha_j}\right)^{\alpha_j} \left(\frac{A_j L_j}{1-\alpha_j}\right)^{1-\alpha_j}, \tag{10}$$

where K_j and L_j are per-firm factor inputs, equal to sectoral aggregates in symmetric equilibrium with a continuum of firms.

Cost minimization yields the nominal marginal cost function

$$MC_j = \left(\frac{R_j}{Z_j}\right)^{\alpha_j} \left(\frac{W_j}{A_j}\right)^{1-\alpha_j}, \tag{11}$$

where R_j is the rental rate and W_j is the nominal wage in sector j. The rental rate is determined by the relative factor-demand equation

$$R_j = W_j\, \frac{\alpha_j}{1-\alpha_j}\, \frac{L_j}{K_j}, \tag{12}$$

which, in conjunction with the sectoral capital market clearing condition $K_j = \bar{K}_j$, pins down the equilibrium rental rate. Note that capital is fixed at the sectoral level ($K_j = \bar{K}_j$), but firms treat it as a choice variable in their cost minimization. The rental rate adjusts so that firms' capital demand equals the fixed supply, clearing the market.

2.2.4 Augmentation and automation shocks

We study two technology shocks affecting sector B, holding sector A technology fixed. Each acts on a distinct primitive:

- *Augmentation* ($A_B \uparrow$): raises labor productivity within the labor-task region $(\alpha_B, 1]$ and leaves the task split between capital and labor unchanged.
- *Automation* ($\alpha_B \uparrow$): shifts the task split toward capital, converting marginal labor tasks into capital tasks and displacing labor at prevailing prices.

2.3 Price-setting and Rotemberg adjustment costs

Monopolistically competitive firms set prices subject to a quadratic Rotemberg (Rotemberg 1982) adjustment cost,

$$AC_j = \frac{\varphi}{2}\left(\frac{P_j}{P_{j,-1}} - 1\right)^2 Y_j, \tag{13}$$

where the cost is incurred in units of the final consumption bundle and is pecuniary,[3] and $\varphi \geq 0$ is the price-adjustment cost parameter. Firms within a sector are identical and choose the same price in equilibrium, so we suppress the variety index and write the optimal-pricing problem in terms of the sectoral price P_j. Since nominal rigidities are restricted to the impact period, the pricing rule can be reduced to a closed-form Phillips curve, as in Guerrieri et al. (2021).[4] This is given by

$$P_j(P_j - 1) = \kappa\left(\tfrac{P_j^*}{P} - \tfrac{P_j}{P}\right), \qquad P_j^* \equiv \mu\, MC_j, \qquad \kappa \equiv \frac{\varepsilon - 1}{\varphi}, \tag{14}$$

where φ governs the degree of price stickiness, $P_j^* \equiv \mu \cdot MC_j$ is the *desired price*, the frictionless optimum at which firms would set prices absent adjustment costs, with $\mu \equiv \varepsilon/(\varepsilon - 1)$ the desired markup over marginal cost MC_j, MC_j/P is real marginal cost, and P_j/P is the relative price. When P_j^* deviates from the current price, firms adjust partially rather than jumping to P_j^*. The forward-looking term in the Rotemberg pricing equation vanishes with the assumption that prices are flexible in the period following the shock, yielding the static pricing rule above.

2.4 Labor market and downward nominal wage rigidity

Households supply labor inelastically at the sector level. Each sector j has a fixed labor endowment N_j, with $N_A + N_B = 1$, and employment satisfies $L_j \in [0, N_j]$. Our benchmark setup does not permit labor mobility across sectors. Nominal wages in each sector are subject to downward nominal wage rigidity (DNWR), represented by the constraint

$$W_j \geq W_{j,-1}, \qquad j \in \{A, B\}, \tag{15}$$

where $W_{j,-1} = W_j^{ss} = 1$ under our steady-state normalization. The constraint binds whenever sectoral labor demand falls below labor supply. In that case $L_j < N_j$, workers are rationed, and each available

[3] We assume that each firm pays $P\,AC_j$, which is rebated lump-sum to the representative household and therefore nets out of aggregate profit income, so the resource constraint is $C_j = Y_j$ in every sector. Because the rebate is lump-sum, firms do not internalize it when setting prices and perceive the full marginal cost of adjustment, leaving the Phillips curve (14) unaffected.

[4] In Guerrieri et al. (2021), production is linear in labor ($Y_j = L_j$), so nominal marginal cost equals the wage rate. Our task-based reduced form involves both capital and labor, so our Phillips curve instead contains the more general marginal cost term.

worker is employed with probability $\rho_j = L_j/N_j$, the sectoral job-finding rate. Sectoral unemployment is $u_j = 1 - L_j/N_j = 1 - \rho_j$. When sectoral demand is high enough to require full employment ($L_j = N_j$), the wage rises above the floor and (15) is slack. Aggregate unemployment is the labor-force-weighted average $u = N_A u_A + N_B u_B$.

2.5 Monetary policy

The central bank conducts monetary policy by setting the nominal interest rate i. The Euler equation (3) makes aggregate nominal expenditure X a strictly decreasing function of the policy rate i, so lower i means more accommodation. Monetary policy therefore selects among the feasible allocations indexed by the policy rate. A critical constraint, however, is that while the central bank can control the total level of nominal spending via the policy rate, its allocation across sectors is determined by the fixed expenditure shares ω and $1 - \omega$ arising from household preferences. Thus, monetary policy cannot target demand toward a specific sector.

2.6 Steady state, market clearing, and equilibrium

The economy is closed and capital is predetermined within the period at $K_j = \bar{K}_j$, with the levels $\bar{K}_j$ pinned down by the steady-state calibration. In the steady state we normalize all productivities, prices, and wages to one, which determines the rental rate through the condition that price equals the markup over marginal cost. We then calibrate each sector's capital so the capital market clears at full employment given this rental rate, and pre-shock nominal expenditure so that demand absorbs each sector's full-employment output at these unit prices. An equilibrium definition is as follows:

Definition 1. *Given technology*

$$(\alpha_j, A_j, Z_j)_{j\in\{A,B\}},$$

factor endowments

$$(\bar{K}_j, N_j)_{j\in\{A,B\}},$$

the pre-shock steady-state expenditure $\bar{X}$, and the central bank's policy rate i, an equilibrium is a collection

$$\mathcal{E} = \{(L_j, Y_j, C_j, W_j, R_j, P_j, MC_j)\}_{j\in\{A,B\}},$$

aggregate nominal expenditure X, and the aggregate price index P satisfying the Euler equation (3), the sectoral demand conditions

$$P_A C_A = (1-\omega)X, \qquad P_B C_B = \omega X,$$

the production and cost relations (10)–(12), the Phillips curve (14), the aggregate price index (2), the wage-rigidity constraint (15), labor-market complementary slackness $(N_j - L_j)(W_j - W_{j,-1}) = 0$ with $L_j \leq N_j$, and market clearing in goods and capital.

2.7 Calibration

Our benchmark calibration is reported in Table 1 in the Appendix. Briefly, the fraction of tasks performed by capital is $\alpha_{\text{baseline}} = 0.30$, the Rotemberg adjustment cost is $\varphi = 11$, and the sectoral expenditure share is $\omega = 0.3$. The remaining parameters are: $\varepsilon = 6$ (markup $\mu = 1.2$), $\beta = 0.96$, $N_A = 0.70$, and $N_B = 0.30$. The augmentation (A_B from 1.00 to 1.03) and automation (α_B from 0.30 to 0.34) shocks are sized to deliver the same sector-B output gain at full employment.

3 How Augmentation and Automation Reshape Production

Before turning to the full two-sector model with nominal rigidities, we first illustrate the production-side mechanics of the two technology shocks in a single-sector flexible-price environment. Stripping out nominal frictions and sector A, we compute the response of sector-B full-employment output and marginal cost to each shock. Labor demand is fixed at $L_B = N_B$ since the demand-side mechanisms through which L_B falls below N_B require nominal frictions.

3.1 Augmentation

The augmentation case is straightforward and intuitive, so we only provide a brief discussion. A labor-augmenting shock $d\log A_B > 0$ raises the productivity of labor within the labor-task region. Totally differentiating the production function (10) at $L_B = N_B$ and $K_B = \bar{K}_B$ yields

$$\frac{d\log Y_B^{\text{full}}}{d\log A_B} \;=\; 1-\alpha_B \;>\; 0. \tag{16}$$

Full-employment output rises in proportion to the labor share. Output depends only on quantities and technology, and both factors are at their fixed sectoral supplies, so the response is the same in partial and general equilibrium.

Marginal costs, however, depend on quantities/technology and equilibrium factor prices. The augmentation shock then affects MC_B in two ways: directly, via A_B in (11), and indirectly, through equilibrium prices R_B and W_B. Holding factor prices fixed, only the first channel is active:

$$\frac{d\log MC_B}{d\log A_B} \;=\; -(1-\alpha_B) \;<\; 0, \tag{17}$$

However, in general equilibrium, both factors remain at their fixed sectoral supplies and the MRTS (12) pins the factor-price ratio R_B/W_B, and under Cobb–Douglas constant factor shares the levels of W_B and R_B are pinned by sector-B's nominal demand ωX. Since this section holds X at its pre-shock value, neither R_B nor W_B moves with A_B, and the general equilibrium response coincides with the partial equilibrium response.

3.2 Automation

The automation case is less straightforward. A task-displacement shock $d\alpha_B > 0$ redistributes tasks from labor to capital. From the production function (10) at $L_B = N_B$ and $K_B = \bar{K}_B$,

$$\frac{d\log Y_B^{\text{full}}}{d\alpha_B} \;=\; \log\!\left(\frac{Z_B\bar{K}_B/\alpha_B}{A_B N_B/(1-\alpha_B)}\right) \;>\; 0. \tag{18}$$

The sign follows directly from the cost-ordering assumption introduced earlier: $R_B/Z_B \leq W_B/A_B$, meaning that capital delivers more output per unit of expenditure than labor on capital-feasible tasks. This translates into $Z_B\bar{K}_B/\alpha_B > A_B N_B/(1-\alpha_B)$, and so shifting a marginal task from labor to capital moves it to its more productive application and full-employment output rises.

The response of marginal cost is also more subtle. Holding factor prices fixed, the partial equilibrium response to a task-displacement shock is

$$\left.\frac{d\log MC_B}{d\alpha_B}\right|_{\text{PE}} \;=\; \log\!\left(\frac{R_B/Z_B}{W_B/A_B}\right) \;<\; 0, \tag{19}$$

which follows directly from the cost-ordering assumption that capital is cheaper than labor in effective units, with the magnitude determined by how much cheaper.

In general equilibrium, however, factor prices adjust. Total-differentiating $\log MC_B = \alpha_B \log(R_B/Z_B) + (1-\alpha_B)\log(W_B/A_B)$ with respect to α_B,

$$\left.\frac{d\log MC_B}{d\alpha_B}\right|_{\text{GE}} \;=\; \underbrace{\log\!\left(\frac{R_B/Z_B}{W_B/A_B}\right)}_{\text{PE cost saving}} \;+\; \underbrace{\alpha_B\,\frac{d\log R_B}{d\alpha_B}}_{\text{rental-rate response}} \;+\; \underbrace{(1-\alpha_B)\,\frac{d\log W_B}{d\alpha_B}}_{\text{wage response}}. \tag{20}$$

The first term captures the direct cost saving. The second and third terms capture the general-equilibrium adjustment of factor prices as all firms simultaneously demand more capital, typically resulting in a higher rental rate and a lower wage. The overall sign depends on whether the factor-price adjustment outweighs the partial equilibrium cost saving.

The full model: binding wage floor. This section abstracts from nominal rigidities, but it is worth previewing a key mechanism that operates in the full model under automation. Downward nominal wage rigidity prevents wages from falling in sector B, pinning $W_B = W_B^{ss} = 1$ at the floor and forcing the entire general-equilibrium adjustment onto the rental rate. With W_B fixed and factor quantities at $K_B = \bar{K}_B$ and $L_B = N_B$, the MRTS (12) gives

$$\left.\frac{d\log R_B}{d\alpha_B}\right|_{W_B\text{ fixed}} \;=\; \frac{1}{\alpha_B(1-\alpha_B)}, \tag{21}$$

and (20) reduces to

$$\left.\frac{d\log MC_B}{d\alpha_B}\right|_{\text{GE, DNWR binds}} \;=\; \log\!\left(\frac{R_B/Z_B}{W_B/A_B}\right) + \frac{1}{1-\alpha_B}. \tag{22}$$

The amplification term $1/(1-\alpha_B)$ exceeds one for any $\alpha_B \in (0,1)$, but whether it dominates the partial-equilibrium saving depends on the size of that saving, which is pinned by the calibration rather than free. At the pre-shock steady state, with $W_B = A_B = Z_B = 1$ and $R_B = R_B^{ss} = \mu^{-1/\alpha_B}$, the saving equals the cost-ordering slack $\log(W_B/A_B) - \log(R_B/Z_B) = \frac{1}{\alpha_B}\log\mu$. The marginal-cost response therefore flips sign whenever $\frac{1}{1-\alpha_B} > \frac{1}{\alpha_B}\log\mu$, equivalently $\mu < e^{\alpha_B/(1-\alpha_B)}$. At our calibration ($\mu = 1.2$, $\alpha_B = 0.30$) this requires $\mu < 1.54$ and holds with room to spare, so automation raises marginal cost in general equilibrium.[5]

The capital-substitution externality. The mechanism behind the PE-to-GE divergence, which we refer to as the *capital-substitution externality*, is as follows: Each firm, taking the rental rate as given, sees automation as a cost-saving substitution. Collectively, every firm wants more capital after the automation shock expands the set of capital-feasible tasks. But since capital supply is inelastic, the rental rate rises to clear the capital market, and consequently raises the effective price of capital across all automated tasks,

[5]The amplification $1/[\alpha_B(1-\alpha_B)]$ in (21) is evaluated at full employment ($L_B = N_B$). Once the floor binds and $L_B < N_B$, the marginal cost response carries an additional term $\alpha_B d\log L_B/d\alpha_B < 0$, so the realized amplification is smaller. We therefore confirm the sign in the binding-floor equilibrium itself. At the do-nothing rate the realized response is $d\log MC_B/d\alpha_B \approx 0.49$ and $d\log P_B/d\alpha_B \approx 0.15$, both positive, so P_B rises under automation.

not just the marginal one. Individual firms fail to internalize the equilibrium rental-rate effect of their capital-substitution decisions, and the resulting cost-push is borne collectively.[6]

The qualitative sign reversal does not require our assumption that capital is fixed in each sector. It requires only that capital demand outpaces supply by enough to dominate the partial-equilibrium cost saving. Fixed sectoral capital makes the mechanism quantitatively sharp by forcing the entire equilibrium adjustment onto the rental rate, but the qualitative direction survives more generally. If capital is elastic within the period, through sectoral reallocation or new investment, the rental-rate response attenuates and may no longer be large enough to flip the sign. Whether the sign reversal survives is then a quantitative question that depends on the elasticity of capital supply and the size of the partial-equilibrium cost saving.

4 Nominal Frictions and the Inflation–Unemployment Trade-off

This section reintroduces both nominal frictions and sector A to characterize the propagation of augmentation and automation shocks through the two-sector economy. Under sticky prices and downward nominal wage rigidity, output is demand-determined and employment can deviate from full employment—creating a role for monetary policy to shape labor-market outcomes. We begin by tracing the sectoral responses to each shock. We then map the aggregate inflation–unemployment menu available to the central bank across policy rates. Finally, we examine how the trade-off varies with the relative size of the AI-adopting sector.

Experimental design. Both sectors begin in a symmetric steady state in which $A_A = A_B = Z_A = Z_B = 1$, prices and wages are normalized to one, and both sectors are at full employment. Aggregate nominal spending is set to $X = \bar{X}$ to yield this steady state. We then consider two separate experiments, each starting from this initial position. In the *augmentation* experiment, sector B receives a labor-augmenting productivity shock. In the *automation* experiment, sector B's capital task share expands. Sector A's technology is unchanged in both experiments. The two cases are never combined; each is analyzed independently as a deviation from the same pre-shock steady state.

4.1 Augmentation, Automation, and the Sectoral Wedge

With nominal rigidities, augmentation and automation affect labor demand in sector B through distinct primary channels. Augmentation raises labor productivity directly through A_B, reducing the labor input required to produce any given output—a *productivity effect*. Automation instead shifts the task partition, lowering the labor share $1 - \alpha_B$ and contracting the labor-demand schedule at any output level and factor prices—a *displacement effect*. The cost-ordering assumption ensures the reassigned tasks are cheaper to perform with capital, so the shift is efficient, but the reduction in labor requirements reflects task displacement rather than a productivity improvement.

Both shocks also operate through a common *scale effect*. Because output is demand-determined, movements in marginal cost pass through to prices and shift real demand $\omega X/P_B$, expanding or contracting labor demand at any given productivity level.

Absent DNWR, wages would fall to clear the labor market in sector B under either shock. With DNWR, the wage floor binds ($W_B = 1$), labor demand falls short of supply N_B, and unemployment emerges. Workers are

[6]This is similar to the labor-mobility externality in Guerrieri et al. (2021), in which workers fail to internalize the equilibrium price effects of their reallocation across sectors.

rationed at job-finding rate $\rho_B = L_B/N_B < 1$.

Monetary policy can reduce (or eliminate) unemployment in Sector B by cutting the policy rate i and expanding aggregate nominal expenditure. Since Cobb–Douglas preferences yield fixed expenditure shares, a fraction ω of any increase in X flows to sector B and a fraction $1-\omega$ flows to sector A. Sector A is already at capacity—there is full employment for any $i \le i^{ss}$ (the steady-state policy rate, $i^{ss} \equiv 1/\beta - 1$)—so additional spending materializes entirely as increases in P_A. Sector B is below capacity and absorbs additional nominal expenditure primarily through higher output, raising labor demand until the market clears at $L_B = N_B$. Beyond this point sector B is itself capacity-constrained, and further accommodation only raises P_B, exactly as in sector A.

The sectoral wedge Δ_i. To quantify the monetary policy trade-off, we define the critical interest rate at which sector j reaches full employment:

$$i_j^{\text{full}} \equiv \max\{i : L_j(i) = N_j\}, \qquad j \in \{A, B\}. \tag{23}$$

Since sector A is unaffected by either shock, $i_A^{\text{full}} = i^{ss}$. The *sectoral wedge*

$$\Delta_i \equiv i_A^{\text{full}} - i_B^{\text{full}} = i^{ss} - i_B^{\text{full}} > 0 \tag{24}$$

measures how far the central bank must cut rates beyond i^{ss} to clear sector B's labor market, and indexes the range of rates over which it faces a binding trade-off between unemployment in sector B and inflation in sector A. Because labor demand in sector B falls by more under automation than under augmentation, Δ_i is larger and the trade-off is more severe.

To illustrate this, Figure 1 plots sectoral employment, output, prices, and nominal wages as functions of the policy rate i, outlining the continuum of equilibria available to the central bank under each shock. Sector A behaves identically under both shocks. For sector B, augmentation is plotted in orange and automation in red. The grey shaded region indicates the binding trade-off region under augmentation, and the combined grey and red region indicates the wider trade-off region under automation.

Panel (a) shows sectoral employment as a function of the policy rate i. Sector A, whose technology is unchanged, is at full employment for any $i \le i^{ss}$. Above i^{ss}, tighter policy reduces aggregate expenditure below its steady-state level, and with sticky prices and downward rigid wages, the fall in nominal spending passes into lower real output and unemployment emerges. The threshold i^{ss} thus marks sector A's transition from a capacity-constrained regime, in which additional spending is pure inflation, to a demand-constrained one, in which insufficient spending generates unemployment. Because sector A's technology and demand are unaffected by the sector-B shock, this cutoff is identical across the two experiments.

Unemployment emerges in Sector B under both shocks when accommodation is held fixed, but rises more under automation. As the central bank cuts i, unemployment in sector B declines, reaching zero at $i_B^{\text{full}} = 2.68\%$ under augmentation and $i_B^{\text{full}} = 0.83\%$ under automation—a rate cut more than twice as deep.

Panel (b) shows output. At $i = i^{ss}$, sector A is at full capacity while sector B operates below potential under both shocks. Relative to steady state, and holding monetary policy unchanged, output rises under augmentation but falls under automation. This reflects the differential price responses: augmentation lowers marginal cost and prices, expanding real demand, while automation raises marginal cost and prices through

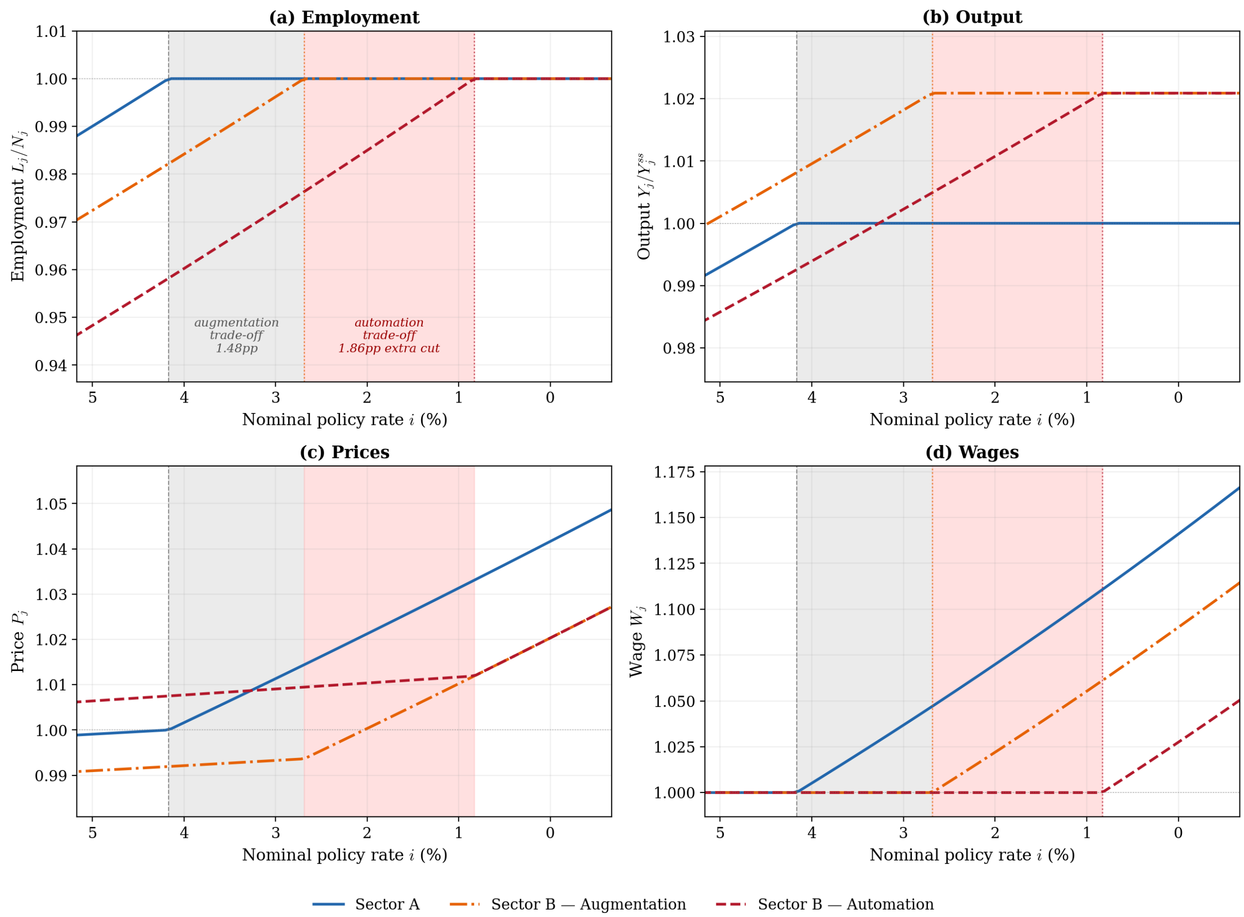

Figure 1: Sectoral employment, output, prices, and nominal wages as functions of the nominal interest rate i. Employment and output are normalized by their respective steady-state values to facilitate comparison across sectors of different sizes. Blue lines: sector A; orange lines: sector B under augmentation; red lines: sector B under automation.

the capital-substitution externality, depressing real output. As i falls, rising ωX expands sector-B output until the labor market clears.

Panel (c) shows prices. Once i falls below i^{ss}, sector A is supply-constrained and absorbs any additional monetary accommodation entirely through P_A. Sector B, still below capacity, absorbs its share primarily through higher output, so P_B rises less than P_A for any given rate cut. The relative price P_A/P_B therefore rises monotonically as i falls, reflecting the differing capacity positions of the two sectors rather than any direct interaction between them.[7]

Panel (d) shows nominal wages. Under both shocks, the wage floor binds in sector B throughout the trade-off region—$W_B = 1$ while unemployment persists—and rises above one only once full employment is restored at $i \leq i_B^{\text{full}}$. Sector A wages rise as soon as i falls below i^{ss}. With sector A already at full employment, additional accommodation cannot raise output and instead feeds into wages, marginal costs, and subsequently prices.

[7]Sector B absorbs additional accommodation primarily through higher output, but with capital fixed at $\bar{K}_B$, the resulting fall in the capital–labor ratio raises the rental rate R_B, feeding into marginal cost and P_B. Prices in sector B thus rise even before full employment is restored.

4.2 Aggregate Trade-offs: The Phillips Curve and the Role of Monetary Policy

Figure 2 translates the sectoral dynamics of the previous section into aggregate trade-offs conditional on the type of AI shock, plotting aggregate inflation against aggregate unemployment for augmentation and automation. While the Phillips curves trace the continuum of equilibria available to the central bank, we indicate for reference the equilibrium selected by a standard Taylor rule, $i = i^{ss} + \phi_\pi \pi + \phi_u u$, with $\phi_\pi = 1.5$ and $\phi_u = -1$.

Under both shocks, no combination of policy delivers both zero inflation and zero unemployment. Automation, however, generates a markedly worse trade-off. The Phillips curve is piecewise linear with three regions. In the flat region, both sectors are demand-constrained and additional accommodation reduces unemployment without generating inflation. In the binding trade-off region, sector A is capacity-constrained and additional accommodation raises P_A while reducing unemployment in sector B. In the vertical region, both sectors are supply-constrained and further accommodation is purely inflationary. Within the binding trade-off region, the automation Phillips curve lies above and to the right of the augmentation curve—for any given unemployment rate inflation is higher, and for any given inflation rate unemployment is higher. This reflects the larger sectoral wedge, which requires greater accommodation to restore full employment and generates correspondingly more inflation in the unaffected sector.

The aggregate Phillips curve also crosses zero inflation at a strictly positive unemployment rate under both shocks. At the policy rate consistent with aggregate price stability, some sector-B unemployment persists—not because it is technically impossible to reduce, but because eliminating it requires accommodation that generates inflation in sector A. This unemployment rate is markedly higher under automation than under augmentation for two reasons. First, automation produces a larger decline in labor demand in sector B, requiring greater accommodation to restore full employment. Second, automation is inflationary on impact, so restoring aggregate price stability requires raising rates above i^{ss}, pushing sector A into the demand-constrained region and generating unemployment there as well.

The marker in Figure 2 shows the equilibrium allocation selected under a standard Taylor rule for each shock. Under augmentation, the rule lands near the origin—low unemployment with near-zero aggregate inflation—because modest deflation in sector B partially offsets inflation in sector A. Under automation, the same rule yields both higher unemployment and higher aggregate inflation. The wider wedge prevents the central bank from achieving either objective, and the offset disappears because sector B prices rise rather than fall. A Taylor rule calibrated to standard macroeconomic conditions therefore produces qualitatively different outcomes across the two shocks, with automation exposing the trade-off that augmentation partially conceals.

The sacrifice ratio. A natural question for policymakers is whether the worse trade-off under automation reflects a steeper Phillips curve or an outward shift. The sacrifice ratio over the binding trade-off region—the percentage-point increase in inflation required to reduce unemployment by one percentage point—measures the slope of the Phillips curve where sector A is capacity-constrained and the central bank faces a genuine conflict between its objectives. Under augmentation it is 1.98, and under automation 1.95, nearly identical. Automation does not make it more costly to trade inflation for unemployment at the margin. Instead, it shifts the entire curve outward, confronting the central bank with higher inflation and higher unemployment at every point along the trade-off.

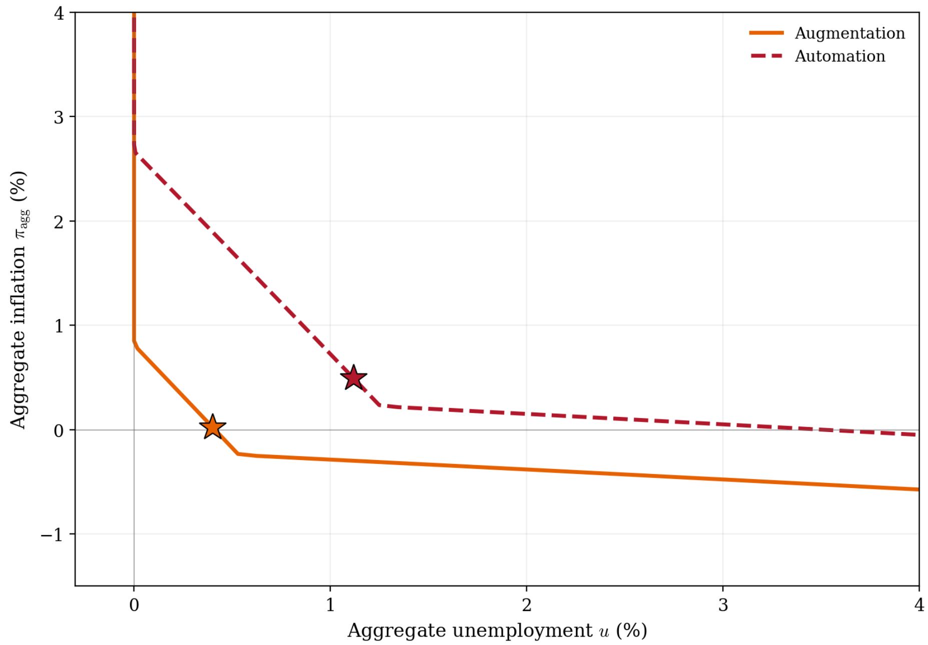

Figure 2: Aggregate Phillips curves under augmentation and automation at the baseline $\omega = 0.3$. The markers denotes the allocation selected under a standard Taylor rule with $\phi_\pi = 1.5$, $\phi_u = -1$.

Augmentation and the role of price stickiness. Augmentation is typically viewed as the benign case for AI, raising worker productivity in existing tasks rather than displacing them. In our framework, however, augmentation still generates unemployment—not because of the technology itself, but rather because of price stickiness. When prices are sticky, the fall in marginal cost passes only partly into prices, real demand falls short of full-employment output, labor demand contracts, and the DNWR binds.

Under flexible prices the augmentation trade-off disappears. The sector's price moves one-for-one with marginal cost, and since Cobb–Douglas preferences fix nominal spending on each good, real demand rises in exact proportion to the price fall. The expansion in demand raises labor demand by exactly enough to offset the decline from higher productivity, leaving employment unchanged. This is the sectoral, task-based counterpart of a familiar New Keynesian result: the contractionary employment response to a positive productivity shock is an artifact of nominal rigidity and vanishes once prices adjust (Galí 1999).

The automation shock, by contrast, is not mitigated by price flexibility. The displacement effect operates independently of prices, and flexible prices actually worsen the trade-off. Because automation raises rather than lowers marginal cost, rising costs pass directly into prices, demand contracts further, and employment falls by more. Figure 6 in the Appendix confirms this—the augmentation wedge falls to zero under flexible prices, while the automation wedge rises from 3.34pp to 5.99pp.

4.3 Sector size and inflation ambiguity

One of the most pressing open questions about AI is how broadly it will ultimately reshape production as adoption deepens and capabilities advance. Existing estimates vary widely, but some suggest that AI could eventually affect as much as 50% of economy-wide tasks (Eloundou et al. 2023). In our framework, we map the breadth of adoption directly to the expenditure share ω of the AI-affected sector. We therefore vary

ω and trace the aggregate macroeconomic consequences as the AI-adopting sector grows.[8] This exercise reveals that the inflationary consequences of AI are ambiguous and depend on both how AI manifests and the breadth of adoption. Automation is inflationary at any sector size, whereas augmentation ranges from mildly inflationary when the adopting sector is small to mildly deflationary when it is large.

The asymmetry between the two shocks carries over to aggregate inflation, and the reason is visible from the aggregate price index. That index is a Cobb–Douglas aggregator of the two sectoral prices, with weights equal to the expenditure shares,

$$P = P_A^{1-\omega} P_B^{\omega}. \tag{25}$$

Since both sectoral prices equal one in the steady state, what matters for policy is the index's deviation from steady state rather than its level. Writing sector j's inflation as $\pi_j \equiv P_j - 1$, aggregate inflation $\pi \equiv P - 1$ is, to first order, the expenditure-weighted average of the two sectoral inflations,

$$\pi \approx (1-\omega)\pi_A + \omega\pi_B, \tag{26}$$

and aggregate unemployment is

$$u = \omega\, u_B, \tag{27}$$

since only sector B experiences unemployment (for policy rates at or below i^{ss}). The weights are exactly the expenditure shares, so ω governs both how much sector B's inflation contributes to the headline number and how much its unemployment registers in the aggregate. Now consider a policy rate inside the wedge, $i \in (i_B^{\text{full}}, i_A^{\text{full}})$, where the two sectors sit in different regimes. Sector A has reached capacity, so additional accommodation cannot raise its output and instead passes one-for-one into its price. Sector B, by contrast, remains in the unemployment regime, where its price is driven primarily by the production-side cost push the shock sets in motion rather than by demand.

This cost push has opposite signs under the two shocks. Augmentation lowers marginal cost, so $\pi_B < 0$ and sector B is disinflationary. Automation raises it, so $\pi_B > 0$ and sector B adds to inflation. Aggregate inflation thus inherits the sectoral asymmetry. At the same policy stance, sector B's inflation partly offsets sector A's under augmentation but reinforces it under automation. The same ω that scales aggregate unemployment therefore has opposing effects on aggregate inflation across the two shocks—under augmentation a larger AI sector buries the wedge, while under automation it amplifies both unemployment and inflation together.

This sign difference, propagating through the expenditure-share weights in $\pi \approx (1-\omega)\pi_A + \omega\pi_B$, generates qualitatively different relationships between i and π across shock types. Under augmentation, $\pi_A > 0$ and $\pi_B < 0$ inside the wedge, so the sign of aggregate inflation depends on ω. Panel (a) of Figure 3 traces this across $\omega \in \{0.3, 0.5, 0.7\}$. At $\omega = 0.3$, the Phillips curve crosses zero inflation at strictly positive unemployment. As ω rises, the negative π_B receives progressively more weight, shifting the curve downward—at $\omega = 0.5$ it crosses zero at a lower unemployment rate, and at $\omega = 0.7$ it lies at or below zero throughout, so full employment is consistent with essentially zero aggregate inflation. The wedge Δ_i is unchanged—the trade-off in interest-rate space is unaffected—but it becomes progressively invisible, and eventually inverted, in aggregate inflation.

Under automation, both π_A and π_B are positive inside the wedge, so there is no offsetting term. Panel (b) shows the curves remaining uniformly positive across all three values of ω. Achieving zero aggregate inflation

[8] Increasing ω scales the sector's endowments proportionally—$N_B = \omega$, $N_A = 1 - \omega$, and $\bar{K}_B$ proportional to N_B, with aggregate nominal expenditure $\bar{X}$ held fixed—so the economy remains at its symmetric full-employment steady state and ω jointly indexes sector B's shares of spending, labor, and capital.

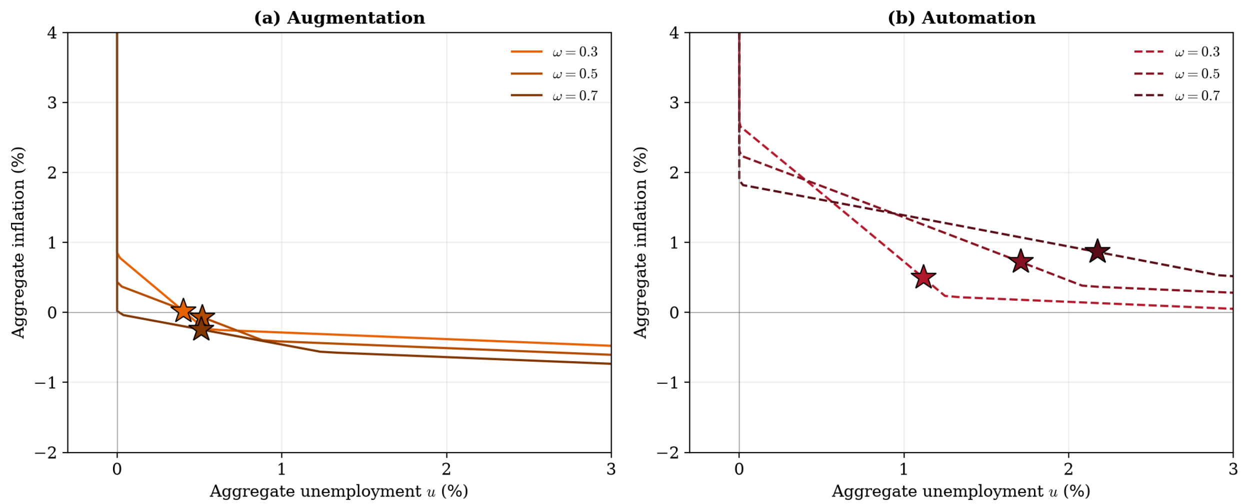

Figure 3: Aggregate Phillips curves under augmentation (panel a) and automation (panel b) at three values of the AI-adopting sector size: $\omega = 0.3, 0.5, 0.7$, plotted in graduated shades from light to dark within each panel. Markers denote the allocation selected under a standard Taylor rule ($\phi_\pi = 1.5$, $\phi_u = -1$).

would require raising i above i^{ss}, pushing sector A into the demand-constrained regime and generating unemployment in both sectors. As ω rises, a larger adopting sector scales up both its unemployment and inflation imprint, so the curves rotate—flatter, with lower inflation at full employment and higher inflation at the do-nothing rate—but price stability requires unemployment regardless of sector size.

The two panels share a nearly identical slope at every ω—the cross-ω counterpart of their near-equal sacrifice ratios. What separates them is the level. Augmentation's offsetting $\pi_B < 0$ shifts the menu down through zero as ω grows, making full employment consistent with zero or negative inflation, whereas automation's reinforcing $\pi_B > 0$ holds it uniformly positive. The sign of the menu, not its slope, is the shock-dependent signal.

The Taylor rule contrast. The markers in Figure 3 make the asymmetry concrete. Under augmentation, the Taylor rule clusters at low unemployment across all three values of ω, with realized aggregate inflation slipping from near zero at $\omega = 0.3$ to mildly negative at $\omega = 0.7$—giving no signal that a meaningful sectoral trade-off exists. Under automation, the same rule produces allocations that deteriorate on both dimensions as ω rises, with aggregate inflation and unemployment increasing together. A central bank watching only headline inflation would conclude that augmentation becomes increasingly benign as the AI-adopting sector grows, while automation becomes progressively more severe—even though the underlying sectoral wedge is identical across sector sizes under both shocks.

The implication is sharp. Low aggregate inflation under augmentation is consistent with two structurally distinct situations: a small AI-adopting sector with a genuinely modest trade-off, or a large adopting sector whose severe trade-off is masked by offsetting sectoral prices. Aggregate inflation alone cannot distinguish them. Under automation, by contrast, aggregate inflation is a reliable indicator because both sectoral prices move in the same direction. The signal value of aggregate inflation for monetary policy therefore depends on the nature of the underlying shock—it is high under automation and low under augmentation when the adopting sector is large.

The two-sector structure is essential for this result; a one-sector model cannot generate the offsetting force. A model-implied empirical implication follows directly: economies with larger AI-adopting sectors should exhibit weaker pass-through from sectoral cost pressures to aggregate inflation under augmentation, but little to no weakening under automation.

5 Extensions

5.1 Labor and capital mobility

Our baseline holds sectoral factor allocations fixed. We now allow households to reallocate labor or capital across sectors subject to quadratic adjustment costs $\psi_L, \psi_K \geq 0$, with the baseline corresponding to $\psi_L = \psi_K = \infty$ and costless mobility to $\psi_L = 0$ or $\psi_K = 0$.[9]

The top row of Figure 4 shows Phillips curves under labor mobility. Under both shocks, the wage floor binds in sector B and, at any given policy rate inside the wedge, the expected real wage differential $\rho_A W_A/P > \rho_B W_B/P$ induces workers to flow out of sector B. As labor mobility costs fall, the Phillips curve flattens and shifts inward. The flattening operates through sector A: as workers move in, the sector can absorb additional accommodation through higher employment rather than prices, reducing the inflation cost per unit of accommodation. The inward shift operates through sector B: as workers leave, the labor pool that must be absorbed at full employment shrinks, requiring less accommodation to close the unemployment gap. Moving from fixed to costless labor mobility, Δ_i narrows from 1.48pp to 0.44pp under augmentation and from 3.34pp to 1.05pp under automation.

The bottom row shows Phillips curves under capital mobility with labor fixed. Under augmentation, the productivity gain lowers R_B and, at impact, capital flows out of sector B. As capital enters sector A, rental rates fall, reducing marginal costs and P_A. Because sector A dominates aggregate output, this decline shifts the curve inward, narrowing Δ_i from 1.48pp to 1.06pp. Under automation, the opposite occurs: the task-share shift raises R_B and, at impact, pulls capital into sector B. Higher K_B partially restores labor demand in sector B but makes capital scarce in sector A, raising marginal costs and P_A. Because sector A is larger, this cost-push dominates and the curve shifts outward, widening Δ_i from 3.34pp to 4.09pp. Capital mobility thus narrows the wedge under augmentation but widens it under automation, amplifying the qualitative ranking $\Delta_i^{\text{auto}} > \Delta_i^{\text{aug}}$ relative to the immobile baseline.

Automation, the labor share, and the tension between short-run stabilization and long-run factor allocations. In the benchmark setup with immobile factors, long-run factor allocations—the period after the shock when both prices and wages are flexible—coincide with the pre-shock steady state. With factor mobility, however, the long-run allocation differs from the steady state, and this highlights a tension between short-run stabilization and long-run structural adjustment. Automation permanently shifts the task boundary, reducing the efficient allocation of labor to sector B and requiring workers to reallocate toward sector A. Monetary accommodation that holds employment in sector B at its pre-shock level may delay this reallocation, propping up employment in a sector where the efficient labor input has structurally declined. A central bank focused on short-run unemployment may therefore impede the very adjustment that raises long-run productivity. This tension does not arise under augmentation, where the efficient labor allocation is unchanged and restoring full employment in sector B is consistent with the long-run equilibrium.

[9] We omit the joint-mobility case because the Phillips curves are essentially identical to the labor-only case.

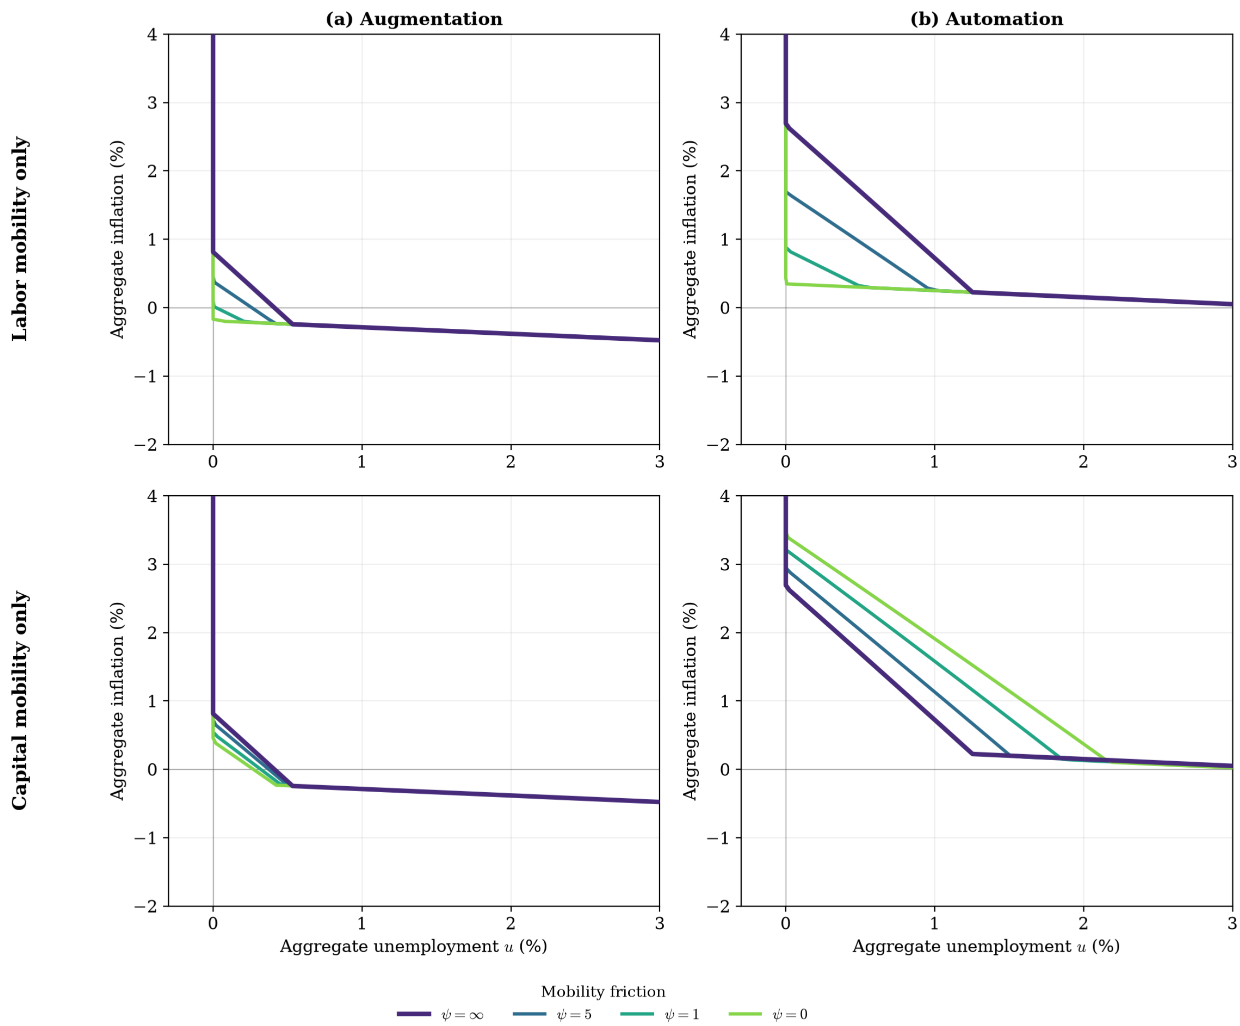

Figure 4: Aggregate Phillips curves under augmentation (left column) and automation (right column) as the factor-mobility cost varies. The top row varies the labor-mobility cost and the bottom row the capital-mobility cost; within each panel the four curves correspond to $\psi = \{0, 1, 5, \infty\}$, plotted in graduated shades from light (full mobility) to dark (immobile baseline).

5.2 CES preferences

Our baseline Cobb–Douglas preferences fix nominal expenditure shares at $(1 - \omega, \omega)$ regardless of equilibrium prices, so the only general-equilibrium link between sectors runs through monetary policy. We generalize to CES preferences, which nest Cobb–Douglas at $\sigma = 1$ and make expenditure shares $s_j = \omega_j (P_j/P)^{1-\sigma}$ endogenous to relative prices.

The key observation is that monetary accommodation raises P_A/P_B within the wedge region: sector A is already at full employment, so additional nominal expenditure raises prices there rather than output, while sector B absorbs its share mainly through output. Under augmentation, where P_A/P_B rises monotonically with accommodation, substitutes ($\sigma > 1$) rotate demand toward the cheaper good B, aiding the central bank, while complements ($\sigma < 1$) divert demand into the capacity-constrained sector A, intensifying the leak and widening the wedge.[10] Under automation the sign of the relative-price gap flips along the accommodation path. The shock initially pushes P_B above P_A, but as accommodation increases and P_A is bid up, P_A/P_B rises above one and the same demand-diversion channel reasserts itself.

Figure 5 plots the Phillips curves at $\sigma \in \{0.5, 1, 2\}$ for each shock. Under augmentation (left panel), the curves fan out monotonically: the complements economy lies furthest from the origin and the substitutes economy closest, reflecting the widening and narrowing of the wedge described above. Under automation (right panel), the ordering at the Taylor-rule equilibria is reversed—complements deliver the lowest unemployment and inflation and substitutes the highest—because the shock itself raises P_B/P_A and the CES interaction initially aids rather than resists stabilization. As accommodation increases and the relative price flips, however, the curves cross and the complements economy becomes the most costly to stabilize, consistent with the demand-diversion channel that dominates in the upper portion of the trade-off.

CES preferences also shift the position of the Phillips curve at a given level of accommodation, since the initial relative price movement redistributes demand across sectors absent any policy response.

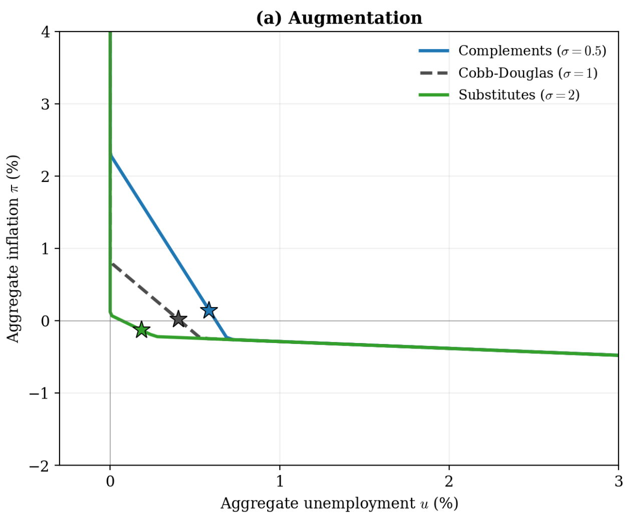

Figure 5: Aggregate Phillips curves under augmentation (panel a) and automation (panel b) at three values of the elasticity of substitution between sectoral goods: $\sigma = 0.5$ (complements, blue), $\sigma = 1$ (Cobb–Douglas, dashed grey), and $\sigma = 2$ (substitutes, green). Markers denote the allocation selected under a standard Taylor rule ($\phi_\pi = 1.5$, $\phi_u = -1$).

6 Discussion

The central message of our paper is that when AI adoption is sector-specific, monetary policy faces a stabilization challenge even when technological change is productivity-enhancing. Because monetary policy is a blunt instrument, accommodating unemployment in the affected sector generates inflationary pressure in unaffected sectors that cannot immediately expand capacity. The severity of this trade-off depends on whether AI primarily augments or automates labor, but also on the breadth of adoption, the mobility of factors across sectors, and how household demand responds to relative price changes.

[10] Guerrieri et al. (2022) find a similar role for cross-sector complementarity in propagating sector-specific supply shocks, though their mechanism operates through household income effects under incomplete insurance rather than the relative-price reallocation studied here.

Policy Implications

Several implications follow for central banks navigating AI-driven technological change.

Diagnosis matters. The appropriate monetary response depends on whether AI primarily augments labor or automates it. Observing higher productivity growth is not sufficient; policymakers must assess how technological change affects labor demand across sectors, and how large those sectors are. Misinterpreting automation as neutral productivity growth risks underestimating both the employment consequences and the inflation cost of the required accommodation. Equally, observing low aggregate inflation when the AI-adopting sector is large does not imply the absence of a stabilization problem—the sectoral wedge and associated unemployment may persist even when aggregate price pressures are muted. Furthermore, to the extent that AI-driven technological change requires structural adjustment in the labor market, excessive monetary accommodation risks delaying that adjustment by sustaining employment in contracting sectors beyond what is efficient.

Monetary policy has limits. Central banks operate through aggregate demand and cannot directly reallocate spending toward sectors experiencing labor-market disruption. Eliminating unemployment in affected sectors may therefore require accepting inflation in unaffected ones. A central bank balancing both objectives may have to tolerate unemployment that reflects the limits of an aggregate instrument rather than a deficiency of demand.

Complementary policies reduce the wedge. Labor reallocation subsidies and retraining programs narrow the inflation–employment trade-off available to monetary authorities and are complements to, not substitutes for, monetary accommodation in the short run. Policies that facilitate capital mobility are more ambiguous: while greater capital mobility narrows the wedge under augmentation, it widens it under automation, as capital drawn toward the AI-adopting sector raises rental rates in the unaffected sector, pushing up marginal costs and adding an additional inflationary dimension to the stabilization problem.

Structural parameters are empirically central. The relative size of AI-exposed sectors and the degree of nominal rigidities jointly determine the magnitude of the trade-off. Sector size shapes both the distribution of inflationary pressure and whether aggregate inflation is a reliable signal of sectoral stress; greater price stickiness slows relative price adjustment. Improved measurement of these parameters should be a priority, particularly as AI capabilities expand and diffuse across sectors.

7 Conclusion

When AI adoption is sector-specific and nominal rigidities inhibit the immediate adjustment of prices and wages, monetary policy faces a stabilization challenge because of the aggregate nature of its instrument. Whether that challenge is mild or severe depends critically on the nature of the shock. Augmentation opens a wedge between the policy rates required to clear the two labor markets, but falling prices in the AI-adopting sector can offset inflation elsewhere, allowing the central bank to target low aggregate inflation while addressing sectoral unemployment. For output-equivalent shocks, automation opens a wider wedge through a direct displacement effect compounded by a capital-substitution externality, and both sectoral prices rise, so no such offset exists. The central bank can choose a point along the resulting inflation–unemployment trade-off, but cannot eliminate it.

The breadth of AI adoption governs not whether this trade-off exists but how it appears in aggregate data. The sectoral wedge is invariant to sector size, but its imprint on both aggregate inflation and aggregate unemployment is not. A larger adopting sector scales up the unemployment imprint under both shocks, since aggregate unemployment is the labor force weighted average of sectoral unemployment rates. The inflation imprint, however, diverges across shocks. Under augmentation, the falling sectoral price carries progressively more weight in the aggregate index as the sector grows, so a large adopting sector can exhibit zero or negative aggregate inflation even as the wedge and sectoral unemployment persist. The apparent stability is cancellation, not balance, and a central bank reading only headline inflation would judge augmentation increasingly benign precisely as its underlying disruption widens. Under automation both sectoral prices rise, no offset exists, and aggregate inflation reliably tracks the trade-off at any sector size. Distinguishing the two therefore requires sectoral rather than aggregate data.

These conclusions sharpen rather than soften when factors are mobile and expenditure shares are endogenous. Labor mobility narrows the wedge under both shocks by moving rationed workers toward the unaffected sector. Capital mobility narrows it under augmentation but widens it under automation, as mobile capital drawn toward the AI-adopting sector raises the cost of capital in the unaffected one, amplifying the cost-push. When households view consumption goods A and B as complements, monetary accommodation increasingly rotates demand away from the affected sector, widening the sectoral wedge and amplifying the central bank's problem. Substitutability instead rotates demand toward the affected sector, narrowing the wedge and easing the problem. Across all of these extensions, the ranking holds—automation remains the more challenging shock to stabilize than augmentation.

The short-run macroeconomic consequences of AI cannot be inferred from aggregate data alone. They depend on whether AI augments or displaces labor, the breadth of adoption, the mobility of factors, how household preferences shape demand reallocation, and how quickly prices and wages adjust. As AI diffuses, recognizing these structural channels—and the limits of an aggregate instrument in addressing what are fundamentally sectoral imbalances—will be central to effective macroeconomic stabilization.

In practice, AI is neither pure automation nor pure augmentation but a combination of the two, and which force dominates will differ across sectors and occupations—and shift over time as the technology diffuses and capabilities expand. Whether a given wave of adoption proves easier or harder to stabilize therefore depends on the relative strength of these two forces. Our framework does not settle the empirical question, but it offers a structured way to reason about it. By isolating how automation and augmentation each shape the sectoral wedge, the slope and position of the Phillips curve, and the trade-offs facing monetary policy, it provides building blocks for understanding their combined effect.

References

Acemoglu, D.: 2024, The simple macroeconomics of AI, *Working Paper 32487*, National Bureau of Economic Research.

Acemoglu, D. and Restrepo, P.: 2018, The race between man and machine: Implications of technology for growth, factor shares, and employment, *American Economic Review* **108**(6), 1488–1542.

Acemoglu, D. and Restrepo, P.: 2019, Automation and new tasks: How technology displaces and reinstates labor, *Journal of Economic Perspectives* **33**(2), 3–30.

Aghion, P. and Bunel, S.: 2024, AI and growth: Where do we stand?, *Technical report*, Federal Reserve Bank of San Francisco.

Aghion, P., Jones, B. F. and Jones, C. I.: 2019, Artificial intelligence and economic growth, *in* A. Agrawal, J. Gans and A. Goldfarb (eds), *The Economics of Artificial Intelligence: An Agenda*, NBER and University of Chicago Press, pp. 237–282.

Basso, H. S. and Rachedi, O.: 2025, Robot adoption and inflation dynamics, *American Economic Review* .

Eloundou, T., Manning, S., Mishkin, P. and Rock, D.: 2023, GPTs are GPTs: An early look at the labor market impact potential of large language models, *Technical Report 2303.10130*, arXiv.

Felten, E. W., Raj, M. and Seamans, R.: 2021, Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses, *Strategic Management Journal* **42**(12), 2195–2217.

Fornaro, L. and Wolf, M.: 2022, The scars of supply shocks: Implications for monetary policy, *Journal of Monetary Economics* **131**, 18–36.

Fornaro, L. and Wolf, M.: 2026, Macroeconomic policies for AI, *Working paper*, CREI.

Fueki, T., Katsuki, S., Muto, I. and Sugisaki, Y.: 2023, Automation and nominal rigidities, *IMES Discussion Paper 2023-E-1*, Bank of Japan, Institute for Monetary and Economic Studies.

Galí, J.: 1999, Technology, employment, and the business cycle: Do technology shocks explain aggregate fluctuations?, *American Economic Review* **89**(1), 249–271.

Guerrieri, V., Lorenzoni, G., Straub, L. and Werning, I.: 2021, Monetary policy in times of structural reallocation, *Jackson Hole Economic Policy Symposium*.

Guerrieri, V., Lorenzoni, G., Straub, L. and Werning, I.: 2022, Macroeconomic implications of COVID-19: Can negative supply shocks cause demand shortages?, *American Economic Review* **112**(5), 1437–1474.

Lenzu, S.: 2026, Artificial intelligence and monetary policy: A framework and perspective on cyclical transmission, structural transition, and financial stability, *Working paper*, Federal Reserve Bank of New York.

Rotemberg, J. J.: 1982, Sticky prices in the United States, *Journal of Political Economy* **90**(6), 1187–1211.

Zeira, J.: 1998, Workers, machines, and economic growth, *Quarterly Journal of Economics* **113**(4), 1091–1117.

A Appendix

Table 1: Baseline Calibration

Parameter	Symbol	Value
Sector-B expenditure share	ω	0.30
Baseline automation threshold	α_{baseline}	0.30
Variety elasticity	ε	6
Rotemberg adjustment cost	φ	11
Discount factor	β	0.96
Sectoral labor endowments	N_A, N_B	$0.70, 0.30$
Sectoral capital endowments	$\bar{K}_A, \bar{K}_B$	$0.55, 0.24$
Pre-shock nominal expenditure	$\bar{X}$	1.71
Augmentation shock magnitude	ΔA_B	$1.00 \to 1.03$
Automation shock magnitude	$\Delta \alpha_B$	$0.30 \to 0.34$

Notes: $\bar{K}_A, \bar{K}_B$ and $\bar{X}$ are not free parameters. $\bar{K}_j$ is calibrated so that $R_j^{ss} = \mu^{-1/\alpha_j}$ clears the capital market at full employment with $W_j = 1$, and $\bar{X}$ so that demand absorbs full-employment output at these unit prices.

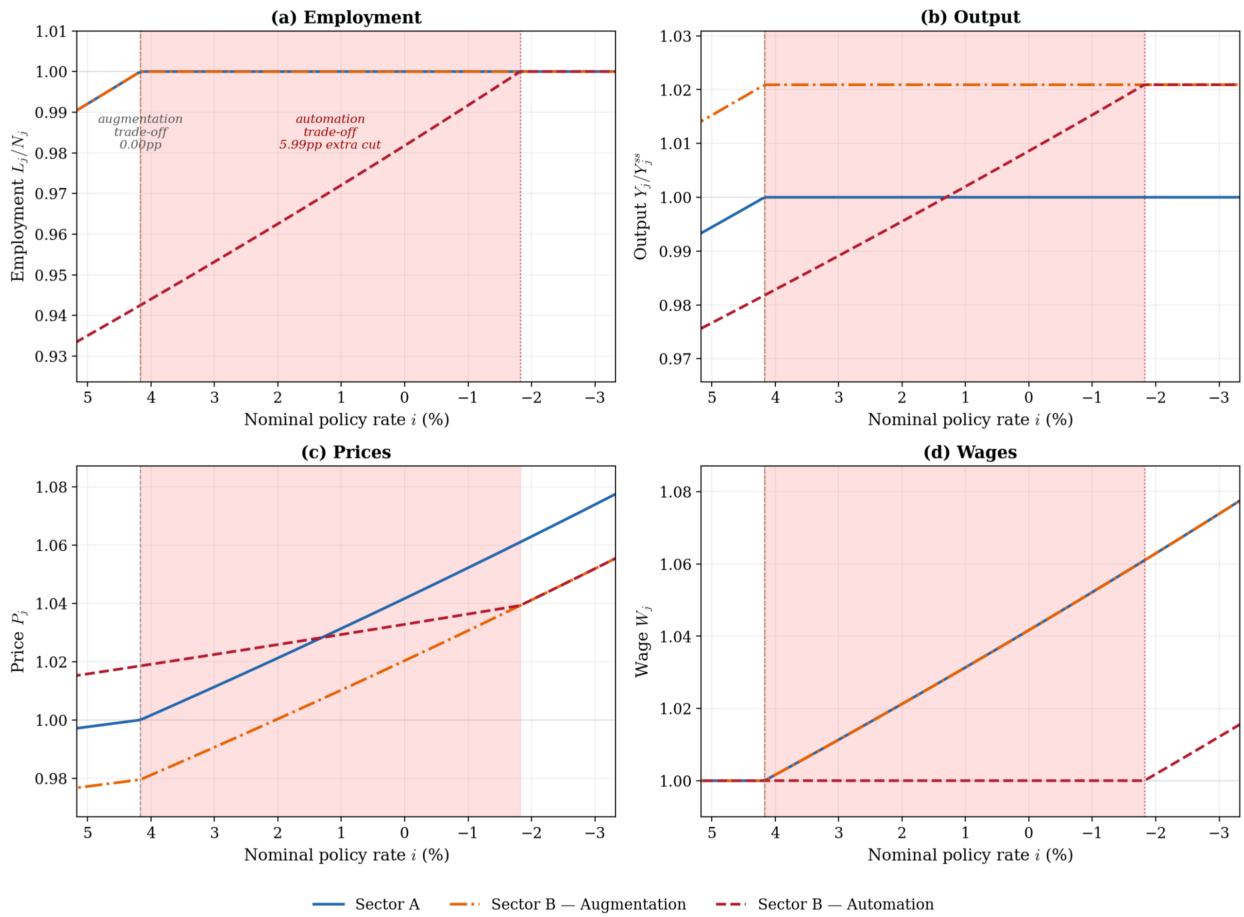

Figure 6: Sectoral employment, output, prices, and nominal wages as functions of the nominal interest rate i under the flexible-price limit ($\varphi = 0$). Employment and output are normalized by their respective steady-state values.