Directed R&D Subsidies*

Gregory Casey† Ryo Horii‡

March 26, 2026

Abstract

Should R&D subsidies be targeted towards certain types of technologies? We build and calibrate a directed technical change growth model with both capital- and labor-augmenting technological change along the balanced growth path, circumventing the restrictions usually posed by the Uzawa steady-state theorem. We use the model to study how directed R&D policy affects the rate of economic growth and the labor share of income. For a given level of expenditure, capital-augmenting R&D subsidies provide the largest boost to economic growth and increase the labor share. Labor-augmenting R&D subsidies provide the smallest boost to growth and decrease the labor share. Uniform subsidies are an intermediate case for growth and are not distributionally neutral. In practice, however, perfectly directing the subsidy can be difficult. Our transitional analysis shows that perfect direction is neither necessary nor desirable. The lifetime equivalent variation of the representative household is maximized when the effective subsidy rate for capital-augmenting R&D is three times higher than the labor-augmenting one. The full benefits take more than 100 years to materialize, requiring sustained policy commitment.

*This paper was split out from an earlier working paper circulated under the title "A Multi-factor Uzawa Growth Theorem and Endogenous Capital-Augmenting Technological Change." We are grateful to Daron Acemoglu, Been-Lon Chen, Oded Galor, Andreas Irmen, Cecilia Garcia-Peñalosa, Miguel Leon-Ledesma, Esteban Rossi-Hansberg, Dean Scrimgeour, Chris Tonetti, Alain Venditti, Ping Wang, David Weil, and seminar participants at Aix-Marseille School of Economics, Brown University, Chulalongkorn University, Kobe University, Liberal Arts Macro Workshop, National Graduate Institute for Policy Studies (GRIPS), Shiga University, Society for Economic Dynamics, Osaka School of International Public Policy (OSIPP), Tongji University, for their helpful comments and suggestions. All remaining errors are our own. This study was financially supported by the JSPS Grant-in-Aid for Scientific Research (15H03329, 15H05729, 15H05728, 16K13353, 17K03788, 20H01477, 20H05631, 20H05633).

†Department of Economics, Williams College, Schapiro Hall, 24 Hopkins Hall Dr., Williamstown, MA, 01267. Email: gpc2@williams.edu.

‡Correspondence: Institute of Social and Economic Research (ISER), Osaka University, 6-1 Mihogaoka, Ibaraki city, Osaka 567-0047, Japan. Email: horii@iser.osaka-u.ac.jp.

Keywords: Balanced Growth, Directed Technical Change, Uzawa Theorem

JEL Classification Codes E13, E22, O33, O41

1 Introduction

Growth theory suggests that technological progress, generally driven by purposeful research and development (R&D), is the source of continued economic growth in developed economies (e.g., Jones and Romer, 2010). Firms under-supply R&D in the absence of policy intervention, because they do not internalize the benefits of intertemporal knowledge spillovers. Thus, R&D subsidies can boost growth and affect income distribution. How should these R&D subsidies be directed towards different types of technology?

In neoclassical models of economic growth, there are two broad types of technical change: labor-augmenting technical change (LATC) and capital-augmenting technical change (KATC). Unless the production function is Cobb-Douglas, these two types of technology have different impacts on economic growth and the labor-share of income. Evidence of both types of technical change are readily apparent in macroeconomic data (Greenwood et al., 1997; Grossman et al., 2017). In particular, a falling relative price of capital—i.e., investment-specific technical change (ISTC)—is a type of KATC.

We study the trade-off between subsidizing KATC and LATC, focusing on how directed R&D policy simultaneously shapes the rate of economic growth and the labor share of income. To study these outcomes, we construct an endogenous growth model with both types of technical change. We allow the elasticity of substitution between capital and labor to be different from one, implying that technology growth rates affect the labor share of income. We microfound the model using the task-based framework of Irmen and Tabaković (2017), which eliminates scale effects (Jones, 2002; Bloom et al., 2017). For the quantitative analysis, the model is calibrated to data from the United States.

Due to the restrictions posed by the Uzawa (1961) steady-state growth theorem, standard neoclassical growth models cannot simultaneously incorporate KATC and endogenous labor shares in the long-run. The Uzawa theorem implies that, in a two-factor neoclassical growth model, either (i) all technology is labor-augmenting or (ii) the production function is Cobb-Douglas, implying that technological progress does not affect the labor share of income (Jones and Scrimgeour, 2008; Grossman et al., 2017). To circumvent these restrictions, we add land as a third factor of production in the aggregate production function (Casey and Horii, 2024).

For a given level of expenditure, R&D subsidies for KATC boost economic growth more than labor-augmenting or uniform R&D subsidies. This result follows from two observations. First, there is a multiplier effect when KATC is subsidized. The KATC subsidy accelerates the growth of the effective amount of capital. It leads to faster output growth, which raises the supply of capital and, again, output. This process continues until the economy converges to the new BGP with considerably faster accumulation of effective capital. Second, since labor has a larger factor share than capital, firms tend to allocate more resources LATC. This means that, given the size of the budget for subsidies, a higher portion of R&D investment costs can be subsidized when the subsidies are targeted at capital-related technology. When the capital-augmenting technological change accelerates, the effective capital-labor ratio increases. Then, the R&D expenditure for LATC change also increases even without subsidies targeted to it. Since the elasticity of substitution is less than one, these targeted subsidies also increase the labor share of income. In our calibration, a small KATC subsidy is roughly three times more cost-effective than an LATC subsidy per unit of fiscal expenditure; even for a large program equal to 1% of GDP, the KATC subsidy still delivers roughly twice the growth gain of the LATC alternative, and it raises the labor share by nearly 12 percentage points. Together, these results suggest that well-targeted subsidies can simultaneously boost growth and raise the labor share. Our transitional analysis, however, shows that perfect direction is neither necessary nor desirable: the lifetime equivalent variation of the representative household is maximized at an effective subsidy-rate ratio of approximately $3:1$ between KATC and LATC.

Uniform R&D subsidies to all technologies are not necessarily neutral to factor shares. In our baseline calibration, they slightly raise the labor share of income while boosting growth, even when subsidies are not targeted to specific technology.[1]

As explained above, there is a multiplier effect from subsidizing KATC, but not from subsidizing LATC. Therefore, the uniform subsidy affects KATC more than LATC, resulting in a higher effective capital-labor ratio in the long run. This finding is absent from the existing endogenous-growth literature, which typically assumes

[1]Our analysis concerns the distribution of income across production factors. The implications for personal or wealth inequality, which depend on household heterogeneity not modeled here, are left for future research.

Cobb–Douglas production and therefore abstracts from studying how R&D subsidies affect income distribution.

While R&D subsidies are our main focus, we also consider the impacts of capital income taxation. Although the model features fully endogenous growth, capital income taxation does not affect the rate or direction of technical change in our framework, nor does it affect the pre-tax labor share of income. This is consistent with long-run U.S. data, in which the growth rate and the labor share have been relatively stable over long periods even as labor and capital tax rates have varied substantially (Jones, 2016).

After reviewing the related literature in Section 2, we build the model in Section 3. The key feature is a task-based production structure that allows firms to direct R&D toward capital- or labor-augmenting technologies, with equilibrium dynamics captured by three state variables. Section 4 derives the balanced growth path and shows analytically that KATC subsidies are more cost-effective than LATC subsidies in raising growth, and that they raise the labor share when the elasticity of substitution is below one. Section 5 calibrates the model to U.S. data, and Section 6 quantifies these effects through five policy experiments. Across scenarios and across the degree of subsidy direction, we compute the transitional dynamics under both a fixed-rate and a constant-budget policy design. Section 7 concludes.

2 Related literature

Our work also contributes to the literature on R&D policy and economic growth. In this literature, it is standard to use Cobb–Douglas production functions (e.g., Romer, 1990; Jones and Williams, 2000; Akcigit and Kerr, 2018), which assume that the labor share is constant and removes any fundamental difference between labor- and capital-augmenting technology. We build on this literature by showing that the effects of R&D policy depend on whether they target KATC or LATC and quantify the impact of these directed subsidies. We also show that R&D policies, even those that do not favor particular technologies, affect the labor share of income.

Our work also connects to the long literature showing that KATC (and ISTC in particular) is an important determinant of economic growth (e.g., Greenwood et al.,

1997) and the labor share of income (Grossman and Oberfield, 2022; Hubmer, 2023). We contribute to this literature by studying the consequences for R&D subsidies that are directed towards KATC and showing that they can boost economic growth. Güner (2023) is the most closely related work. He considers a two-sector model with separate consumption and investment goods with symmetric Cobb–Douglas production in both sectors, which makes the model consistent with the two-factor Uzawa theorem. He also shows that optimal subsidies favor the investment sector.[2] We build on this literature by studying a model that is consistent with both KATC and non-Cobb-Douglas production. In addition to improving quantitative realism, this allows us to simultaneously study growth and the labor share of income. We also provide more general conditions on when subsidizing KATC increases the rate of economic growth.

Our work is part of a growing literature that examines the relationship between endogenous growth and the Uzawa steady-state theorem. We build on this work in two ways. First, we present a directed technical-change growth model that converges to a BGP that has both KATC and non-Cobb-Douglas production. Previous work focuses on the case where the long-run production function is Cobb–Douglas (Jones, 2005; Leon-Ledesma and Satchi, 2019) or only labor-augmenting technology occurs along the BGP (Acemoglu, 2003; Irmen and Tabaković, 2017). Grossman et al. (2017, 2021) and Jones and Liu (2024) consider the case where one form of capital-augmenting technical change is offset by an opposing dis-augmenting force. Second, leveraging the unique features of our model, we study how policy affects the direction of technical change and the labor share of income—an exercise that is not possible without capital-augmenting technical change on the BGP and an EoS different from one.

Our work also contributes to the broader literature on the Uzawa theorem (e.g., Schlicht, 2006; Jones and Scrimgeour, 2008; Grossman et al., 2017). Casey and Horii (2024) show that, with exogenous technical change, a three-factor aggregate production function can deliver a BGP with positive capital-augmenting technical, but only under a knife-edge restriction on technology growth rates. We embed their results into a model with endogenous and directed technical change. We microfound the model using an extension of the task-based framework of Irmen and Tabaković (2017), which allows for fully endogenous growth without scale effects. This microfounded model

[2]See also Krusell (1998) who studies a model where all innovation happens via ISTC.

allows us to study the consequences of R&D policy, our primary goal. We also show that the knife-edge technology constraint identified by Casey and Horii (2024) is endogenously satisfied in the presence of directed technical change, implying that it does not impose extra restrictions on a fully specified model.

Our focus is on the direction of R&D subsidies when the only market failure is the externality from intertemporal knowledge spillovers. Several studies consider the optimal direction of R&D subsidies to address other externalities, such as pollution or labor market frictions (e.g., Acemoglu et al., 2012; Acemoglu and Restrepo, 2018; Acemoglu, 2023).[3] In addition, a long literature studies the relationship of natural resources scarcity and innovation (e.g., André and Smulders, 2014; Hassler et al., 2021; Casey, 2024). While these models are usually used to study resource conservation or climate change, we show that incorporating directed technical change and natural resources into growth models improves their ability to overcome the Uzawa steady state theorem, match macroeconomic data, and analyze policy question not directly related to natural resource economics.

3 Model

In this section, we build an endogenous growth model in which firms undertake R&D investments to improve factor-augmenting technologies. The government can subsidize R&D for capital- and labor-augmenting technologies at different rates, and we study how these directed subsidies affect growth and factor shares.

3.1 Structure

Tasks and production. There are non-overlapping generations of representative firms, each of which exists for only one period. A representative firm performs two types of tasks, M-tasks and N-tasks.[4] The number of M-tasks, as well as that of N-tasks, determines the amount of final output.

[3] See Hémous and Olsen (2021) for a review.

[4] This setting extends Irmen and Tabaković (2017) to incorporate a third production factor and directed R&D subsidies. A benefit of using the model of tasks is that it can incorporate R&D activity within a perfectly competitive economy. Another benefit is that the model is scale independent.

The M-tasks require effective capital and natural resources as inputs. The number of M-tasks the firm can complete is given by

$$M_t = (A_{K,t}K_t)^{\alpha} X_t^{1-\alpha}, \quad \alpha \in (0,1), \tag{1}$$

where $A_{K,t}$ is the capital-augmenting technology of the representative firm, K_t is the capital stock, and X_t represents natural resources and land. We refer to the RHS as the aggregate amount of the *capital composite*, which combines effective capital and effective natural resources with unit elasticity.[5] Natural resources grow exogenously at gross rate g_X, so that $X_t = X_0 \cdot g_X^t$.

An N-task uses only effective labor, $A_{L,t}L_t$, where $A_{L,t}$ is the labor-augmenting technology of the representative firm. The number of N-tasks is simply

$$N_t = A_{L,t}L_t. \tag{2}$$

Population (and labor supply, which is inelastic) grows at gross rate g_L, so that $L_t = L_0 \cdot g_L^t$.

By performing M_t and N_t tasks, the representative firm produces

$$Y_t = F(M_t, N_t) = F\big((A_{K,t}K_t)^{\alpha} X_t^{1-\alpha},\ A_{L,t}L_t\big) \tag{3}$$

units of output, where $F(\cdot)$ is a neoclassical production function with CRS, $f'(\mu) > 0$, $f''(\mu) < 0$, $f'(0^+) = \infty$, and $f'(\infty) = 0$.[6] A key feature of the production structure (3) is that capital and natural resources enter with unit elasticity of substitution within the capital composite M, while the elasticity between M and labor can differ from one. Casey and Horii (2024) show that this is precisely the structure under which capital-augmenting technological change is consistent with balanced growth when the production function is not Cobb–Douglas. The consistency requirement is a *technology condition*: the growth rates of M-tasks and N-tasks must be equal in

[5] We do not model R&D for the natural resource factor. We normalize $A_{X,t} = 1$ for all t, so that improvements in resource productivity are captured by the exogenous growth of X_t.

[6] We write $f(\mu) \equiv F(\mu, 1)$ for the intensive form, so that $F_M = f'(\mu)$ and $F_N = f(\mu) - \mu f'(\mu)$, where $\mu \equiv M/N$. These conditions are weaker than the standard Inada conditions: we do not require $f(\infty) = \infty$, which fails for CES production functions with elasticity of substitution below one.

the long run, which imposes a specific relationship between the rates of KATC and LATC. With exogenous technical change, this relationship is a knife-edge restriction on two independent parameters (Casey and Horii, 2024). In our model, however, the direction of R&D is chosen endogenously by firms, and we show in Section 4 that the technology condition is satisfied as an equilibrium outcome.

R&D technology. Technical knowledge can be kept within the firm for only one period, after which it becomes public. Thus, the representative firm at time t can freely use the technology of the period $t-1$ firm, $\{A_{K,t-1}, A_{L,t-1}\}$. In addition, the period t firm can improve each of the factor-augmenting technologies through R&D.

Tasks are differentiated and require separate R&D investments.[7] Specifically, to raise the capital-augmenting technology from $A_{K,t-1}$ to $A_{K,t}$ in one M-task, the firm must invest $\psi_K(\gamma_{K,t})$ units of final goods, where $\gamma_{K,t} \equiv A_{K,t}/A_{K,t-1} \geq 1$ is the gross rate of KATC. Similarly, $\psi_L(\gamma_{L,t})$ is the cost of improving the labor-augmenting technology in one N-task, where $\gamma_{L,t} \equiv A_{L,t}/A_{L,t-1}$.

We assume that the R&D cost functions $\psi_j(\cdot)$, $j = K, L$, satisfy

$$\psi_j(1) = 0, \quad \psi_j'(1) = 0, \quad \psi_j'(\gamma_j) > 0, \quad \psi_j''(\gamma_j) > 0 \quad \text{for all } \gamma_j > 1. \tag{4}$$

[8] Since the firm performs M_t M-tasks and N_t N-tasks, the total R&D investment costs are

$$R_{K,t} = M_t \cdot \psi_K(\gamma_{K,t}), \qquad R_{L,t} = N_t \cdot \psi_L(\gamma_{L,t}). \tag{5}$$

The one-period firm structure is a modeling device for finite appropriability of innovations. What matters for the equilibrium is that a firm can capture the returns to its R&D investment for a limited duration—one model period of χ years—after which the knowledge diffuses to all producers. An equivalent formulation would have longer-lived firms whose innovations are imitated after χ years, or whose patents expire after χ years; the equilibrium conditions would be identical. In Section 5, we

[7] From the symmetry of tasks within each group (M or N) and from the convexity of the R&D cost functions, it is always optimal to choose the same levels of $A_{K,t}$ and $A_{L,t}$ across individual tasks.

[8] The convexity of the R&D cost function, $\psi_j'' > 0$, can be explained by congestion in R&D activities. When many researchers are devoted to improvements in the same task at the same time, some of them will end up inventing the same innovation. The risk of duplication becomes more prominent as R&D inputs increase. See Horii and Iwaisako (2007) for a simple micro foundation.

calibrate $\chi \approx 4$ years.

R&D subsidies and profit. The government subsidizes R&D investments at rates $s_K \in [0,1)$ and $s_L \in [0,1)$ for capital- and labor-augmenting technologies, respectively. The post-subsidy costs borne by the firm are $(1-s_K)R_{K,t}$ and $(1-s_L)R_{L,t}$. The representative firm maximizes profit

$$\pi_t = F(M_t, N_t) - (1-s_K)R_{K,t} - (1-s_L)R_{L,t} - r_t K_t - \xi_t X_t - w_t L_t \tag{6}$$

with respect to $\{K_t, X_t, L_t, A_{K,t}, A_{L,t}\}$, taking as given factor prices $\{r_t, \xi_t, w_t\}$, lagged technology levels $\{A_{K,t-1}, A_{L,t-1}\}$, and the subsidy rates $\{s_K, s_L\}$.

Household. There is a representative household whose members supply labor inelastically. The household maximizes

$$U = \sum_{t=0}^{\infty} L_t \beta^t \frac{(C_t/L_t)^{1-\theta} - 1}{1-\theta}, \tag{7}$$

where C_t is aggregate consumption and $\beta \in (0,1)$ is the discount factor. The household owns capital K_t and receives factor income. The government levies taxes at rate τ_K on capital income and firm profits, and at rate τ_L on labor income, and provides a lump-sum transfer B_t. The budget constraint is

$$K_{t+1} = \big((1-\tau_K)(r_t - \delta) + 1\big)K_t + \xi_t X_t + (1-\tau_L)w_t L_t + (1-\tau_K)\pi_t + B_t - C_t. \tag{8}$$

[9] The household also satisfies the No-Ponzi-Game condition:

$$\lim_{T\to\infty} \frac{K_T}{\prod_{t=1}^{T}\big[1 + (1-\tau_K)(r_t - \delta)\big]} \geq 0. \tag{9}$$

Government. The government runs a balanced budget:

$$\tau_K(r_t - \delta)K_t + \tau_K \pi_t + \tau_L w_t L_t = s_K R_{K,t} + s_L R_{L,t} + B_t. \tag{10}$$

[9] The labor income tax τ_L is not distortionary because labor supply is inelastic.

Given the tax rates $\{\tau_K, \tau_L\}$ and subsidy rates $\{s_K, s_L\}$, the transfer B_t adjusts to balance the budget in every period.

Policy instruments. To summarize, the government has four policy instruments: R&D subsidies $\{s_K, s_L\}$ and factor income taxes $\{\tau_K, \tau_L\}$. We will show that R&D subsidies affect the direction of technological change and the growth rate, whereas factor income taxes affect the levels of capital and consumption on the BGP but not the direction of technology or the growth rate.

3.2 R&D by Firms and the Direction of Technological Change

Now we examine the behavior of the representative firm, focusing on the role of R&D. For convenience, we define $\mu_t \equiv M_t/N_t$, which is the relative task intensity in final good production. It also represents the ratio of the capital composite to effective labor. Because $F(\cdot)$ is CRS, we can write $F(M_t, N_t)/N_t = f(\mu_t)$, $F_M = f'(\mu_t)$, and $F_N = f(\mu_t) - \mu_t f'(\mu_t)$.

Factor demand. The first-order conditions for factor demand are (see Appendix A.1 for derivation)[10,11]

$$r_t = \frac{\alpha\mu_t}{k_t}\big(f'(\mu_t) - (1-s_K)\psi_K(\gamma_{K,t})\big), \tag{11}$$

$$\xi_t = \frac{(1-\alpha)M_t}{X_t}\big(f'(\mu_t) - (1-s_K)\psi_K(\gamma_{K,t})\big), \tag{12}$$

$$w_t = A_{L,t}\big(f(\mu_t) - \mu_t f'(\mu_t) - (1-s_L)\psi_L(\gamma_{L,t})\big), \tag{13}$$

where $k_t \equiv K_t/N_t$ is capital per effective labor.

R&D condition for KATC. Now, we turn to R&D, starting with the condition for improving the capital-augmenting technology $A_{K,t}$. The representative firm

[10]The RHS of (11) represents the net marginal product of K_t. It is the product of two parts: $\alpha M/K$, the marginal product of K_t in increasing M-tasks, and $f'(\mu) - (1-s_K)\psi_K(\gamma_K)$, the net marginal product of an M-task after the firm's R&D cost.

[11]By substituting (11), (12), (13), (14), and (15) into (6), it can be confirmed that the firm achieves zero profit, $\pi_t = 0$. This is due to the CRS property of the firm's problem.

chooses $\gamma_{K,t} \equiv A_{K,t}/A_{K,t-1} \geq 1$ to maximize profit. The first-order condition is (see Appendix A.1 for derivation)

$$(1 - s_K)\Big[\underbrace{\gamma_{K,t}\psi'_K(\gamma_{K,t})}_{\text{intensive margin}} + \underbrace{\alpha\,\psi_K(\gamma_{K,t})}_{\text{extensive margin}}\Big] = \alpha f'(\mu_t). \tag{14}$$

The left-hand side represents the marginal cost that the firm bears when raising $A_{K,t}$. It decomposes into the subsidy-adjusted cost $(1 - s_K)$ times the sum of two R&D cost components. The first, $\gamma_{K,t}\psi'_K(\gamma_{K,t})$, is the *intensive* marginal R&D cost: by intensifying R&D efforts in existing M-tasks to raise capital efficiency $A_{K,t}$, the firm can decrease capital inputs by just enough to perform one additional M-task. The second, $\alpha\,\psi_K(\gamma_{K,t})$, is the *extensive* marginal R&D cost: the saved capital is then used to perform a new M-task, which requires its own R&D investment.[12] The right-hand side, $\alpha f'(\mu_t)$, represents the marginal benefit from KATC: an increase in $A_{K,t}$ raises the number of M-tasks the firm can perform, and thereby increases output. The coefficient α appears because capital's share in the capital composite $M = (A_K K)^\alpha X^{1-\alpha}$ is α, so a 1% increase in A_K raises M by only α%.

As we formally prove in Proposition 1 below, condition (14) has a unique solution for $\gamma_{K,t}$ as a function of μ_t and s_K. Since $f'(\mu)$ is decreasing in μ (a higher μ means the capital composite is more abundant and hence less valuable), $\gamma_{K,t}$ is a decreasing function of μ_t. A higher subsidy rate s_K raises $\gamma_{K,t}$ by reducing the effective cost of R&D.

R&D condition for LATC. The first-order condition for $A_{L,t}$ has a parallel structure (see Appendix A.1):

$$(1 - s_L)\Big[\underbrace{\gamma_{L,t}\psi'_L(\gamma_{L,t})}_{\text{intensive margin}} + \underbrace{\psi_L(\gamma_{L,t})}_{\text{extensive margin}}\Big] = f(\mu_t) - \mu_t f'(\mu_t). \tag{15}$$

[12] The extensive margin carries the weight α because $dM/dA_K = \alpha M/A_K$: effective capital enters the capital composite with exponent α, so raising A_K frees only α units of the composite per unit of technology improvement.

The intensive and extensive margins have the same interpretation as in the KATC case, but applied to N-tasks.[13] The right-hand side, $f(\mu_t) - \mu_t f'(\mu_t)$, is the marginal product of an N-task: the increase in output from performing one additional N-task.

Condition (15) determines $\gamma_{L,t}$ as an increasing function of μ_t: Intuitively, a high value of $\mu_t \equiv M_t/N_t$ means that effective labor is relatively scarce, so the marginal product of an N-task is high, and the firm chooses a larger $\gamma_{L,t}$.

Direction of technological change. We summarize the R&D conditions in the following proposition.

Proposition 1 (Direction of Technological Change)**.** *Under assumption (4), conditions (14) and (15) uniquely determine the rates of technological progress as functions of the relative task intensity and the subsidy rates:*

$$\gamma_{K,t} = \widehat{\gamma}_K(\mu_t;\, s_K), \qquad \gamma_{L,t} = \widehat{\gamma}_L(\mu_t;\, s_L), \tag{16}$$

where $\widehat{\gamma}_K(\mu_t;\, s_K)$ is decreasing in μ_t and increasing in s_K, and $\widehat{\gamma}_L(\mu_t;\, s_L)$ is increasing in μ_t and increasing in s_L.

Proof. Define the composite marginal R&D cost functions $\Psi_K(\gamma_K) \equiv \gamma_K \psi_K'(\gamma_K) + \alpha\, \psi_K(\gamma_K)$ and $\Psi_L(\gamma_L) \equiv \gamma_L \psi_L'(\gamma_L) + \psi_L(\gamma_L)$. Under assumption (4), Ψ_K and Ψ_L are strictly increasing on $[1, \infty)$ with $\Psi_j(1) = 0$ and $\Psi_j(\infty) = \infty$. Therefore their inverses are well-defined, and the R&D conditions can be solved for $\gamma_{K,t}$ and $\gamma_{L,t}$ as functions of μ_t and the subsidy rates. The monotonicity properties follow from the monotonicity of $f'(\mu)$ and $f(\mu) - \mu f'(\mu)$. See Appendix A.2 for the full proof. □

Equilibrium Innovation Possibility Frontier. Given the subsidy rates $\mathbf{s} \equiv \{s_K, s_L\}$, the locus of equilibrium innovation rates (γ_K, γ_L) that firms choose as μ varies traces out what we call the *Equilibrium Innovation Possibility Frontier* (EIPF). Since $\widehat{\gamma}_K(\mu; s_K)$ is decreasing in μ and $\widehat{\gamma}_L(\mu; s_L)$ is increasing in μ, the EIPF is downward-sloping in (γ_K, γ_L) space.

The EIPF embodies the fundamental trade-off in the direction of technological change. When μ_t is high—meaning the capital composite is abundant relative to

[13] Unlike (14), the extensive margin in (15) does not carry a weight α. This is because $N_t = A_{L,t} L_t$ is linear in $A_{L,t}$, so $dN/dA_L = N/A_L$ without any exponent.

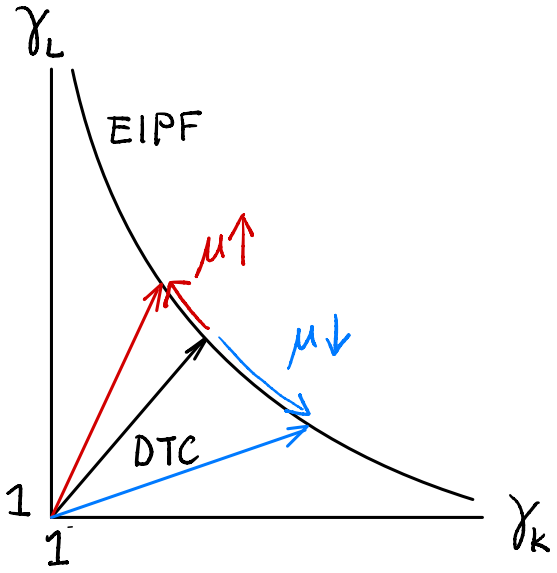


Figure 1: Equilibrium Innovation Possibility Frontier. As μ_t changes, the direction of technological change (γ_K, γ_L) moves along the EIPF curve.

effective labor—the marginal product of M-tasks ($f'(\mu)$) is low while that of N-tasks ($f(\mu) - \mu f'(\mu)$) is high. Firms respond by investing more in LATC and less in KATC. As Figure 1 illustrates, the DTC moves along the EIPF curve so that it enhances the factor that is relatively scarce.

Policy shifts of the EIPF. Directed R&D subsidies shift the EIPF. An increase in s_K stretches the EIPF horizontally: for a given μ, the equilibrium γ_K increases because the effective cost of KATC R&D falls. Similarly, an increase in s_L stretches the EIPF vertically (Figure 2). Note, however, that μ itself responds to the policy change, so the new equilibrium point on the shifted EIPF differs from the old one. A dynamic analysis is required to determine the long-run effects; this is the subject of Section 4.

3.3 Equilibrium Dynamics

The equilibrium path of this economy is given by the sequence of output, consumption, production factors, technologies, and R&D investments, $\{Y_t, C_t, K_t, X_t, L_t, A_{K,t}, A_{L,t}, R_{K,t}, R_{L,t}\}_{t=0}^{\infty}$, which satisfy the representative firm's optimization problem, the representative con-

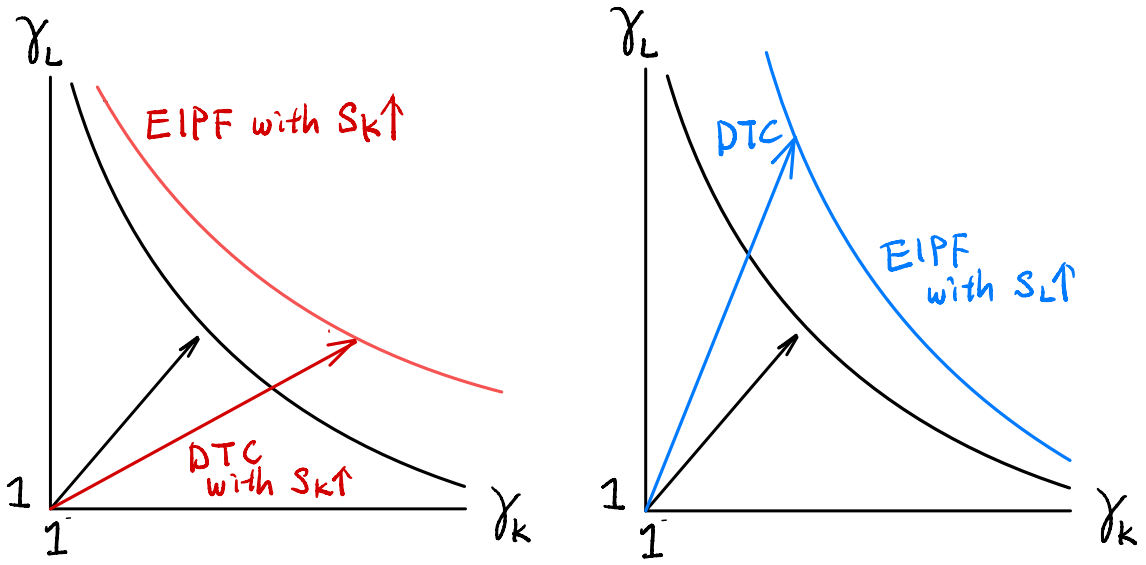


Figure 2: Policy shifts of the EIPF. Left: KATC subsidy s_K stretches the frontier horizontally. Right: LATC subsidy s_L stretches the frontier vertically.

sumer's utility maximization problem, and the market clearing conditions for output and production factors. The economy is endowed with K_0, X_0, and L_0 at time 0, as well as the initial levels of publicly available technologies, $A_{K,-1}$ and $A_{L,-1}$.

While the equilibrium involves many variables, we can analytically characterize its dynamic path in terms of only three: relative task intensity $\mu_t = M_t/N_t$, the amount of capital per effective labor $k_t \equiv K_t/(A_{L,t}L_t)$, and consumption per effective labor $c_t \equiv C_t/(A_{L,t}L_t)$. Below, we construct the equilibrium mapping from $\{\mu_t, k_t, c_t\}$ to $\{\mu_{t+1}, k_{t+1}, c_{t+1}\}$ for $t \geq 0$. The mapping and the initial conditions μ_0 and k_0, together with the transversality condition for c_t, will pin down the equilibrium path, from which all variables in the model can be recovered.

Net output. It is convenient to define the net aggregate output as $V_t = F(M_t, N_t) - R_{K,t} - R_{L,t}$, which means the aggregate output minus the total (pre-subsidy) R&D costs.[14] The net output per effective labor can be written as a function of μ_t

$$V_t/N_t = f(\mu_t) - \mu_t\, \psi_K\big(\widehat{\gamma}_K(\mu_t; s_K)\big) - \psi_L\big(\widehat{\gamma}_L(\mu_t; s_L)\big) \equiv v(\mu_t; \mathbf{s}). \tag{17}$$

[14]The resource constraint involves pre-subsidy R&D costs, not post-subsidy costs. The subsidy is a transfer from the government to the firm, which cancels out in the aggregate resource constraint.

Capital accumulation. Substituting profits (6) and the government budget (10) into the household budget constraint (8), all tax terms, subsidies, and factor payments cancel, yielding the aggregate resource constraint (see Appendix A.3)

$$K_{t+1} = V_t + (1-\delta)K_t - C_t. \tag{18}$$

The growth of aggregate capital can then be expressed as

$$\frac{K_{t+1}}{K_t} = \frac{v(\mu_t;\,\mathbf{s}) - c_t}{k_t} + 1 - \delta. \tag{19}$$

Dynamics for μ_{t+1}. The state variable $\mu_t \equiv M_t/N_t$ measures the relative abundance of the capital composite in production. Its evolution is shaped by two forces: factor accumulation (which shifts the task ratio through capital deepening and resource growth) and the equilibrium R&D response (which redirects innovation toward whichever factor is relatively scarce, Proposition 1). Because the R&D rates $\widehat{\gamma}_K(\mu_{t+1}; s_K)$ and $\widehat{\gamma}_L(\mu_{t+1}; s_L)$ are themselves functions of μ_{t+1}, the law of motion involves an implicit equation in μ_{t+1}. We show in Appendix A.3 that this equation can be solved explicitly as

$$\mu_{t+1} = \Gamma^{(-1)}\left(\frac{g_X^{1-\alpha}}{g_L}\left(\frac{v(\mu_t;\,\mathbf{s}) - c_t}{k_t} + 1 - \delta\right)^{\alpha}\mu_t\;;\;\mathbf{s}\right) \equiv \widehat{\mu}(\mu_t, k_t, c_t;\,\mathbf{s}), \tag{20}$$

Equation (20) has a natural two-step interpretation (see Appendix A.3 for the derivation). The argument of $\Gamma^{(-1)}$, which we denote μ_{t+1}^{pre}, is the relative task intensity that would obtain if technologies remained at their period-t levels; it is determined by the supply of production factors and by capital accumulation. The function $\Gamma^{(-1)}(\cdot\,;\mathbf{s})$ then describes how equilibrium R&D in period $t+1$ transforms μ_{t+1}^{pre} into the realized μ_{t+1}. Because firms direct R&D toward whichever factor is relatively scarce (Proposition 1), this R&D step partially offsets the supply-side movement of μ, providing a self-correcting force in the dynamics.

The capital accumulation multiplier. An important feature of the μ_{t+1} dynamics is the *capital accumulation multiplier*. The pre-R&D intensity μ_{t+1}^{pre} depends on K_{t+1}/K_t to the power α, reflecting the Cobb-Douglas structure of M-tasks. But on

the BGP, capital grows at the same rate as output. When the KATC rate γ_K increases permanently, faster growth of A_K raises output growth, which accelerates capital accumulation, which further raises output growth. This geometric series means that a permanent 1% increase in γ_K raises the growth rate of M-tasks by $\alpha/(1-\alpha)$%, not merely the α% that would follow from the direct effect of A_K on M alone. Since capital's share in the capital composite exceeds that of natural resources (i.e., $\alpha > 1/2$), the multiplier $\alpha/(1-\alpha)$ is greater than one. This amplification is central to why KATC subsidies are more cost-effective than LATC subsidies.

Capital and consumption dynamics. From (19), capital per effective labor evolves as

$$k_{t+1} = \frac{v(\mu_t;\, \mathbf{s}) - c_t + (1-\delta)k_t}{g_L\, \widehat{\gamma}_L(\widehat{\mu}(\mu_t, k_t, c_t;\, \mathbf{s});\, s_L)} \equiv \widehat{k}(\mu_t, k_t, c_t;\, \mathbf{s}), \tag{21}$$

where the numerator is net saving plus undepreciated capital; division by $g_L\, \widehat{\gamma}_L$ accounts for the growth of effective labor between periods t and $t+1$. The household's Euler equation is (see Appendix A.3)

$$\left(\frac{A_{L,t+1}\, c_{t+1}}{A_{L,t}\, c_t}\right)^{\theta} = \beta\big((1-\tau_K)(r_{t+1}-\delta)+1\big). \tag{22}$$

Using $A_{L,t+1}/A_{L,t} = \widehat{\gamma}_L(\mu_{t+1}; s_L)$ and substituting the rental rate (11) gives

$$c_{t+1} = \frac{\beta^{1/\theta}\, c_t}{\widehat{\gamma}_L(\mu_{t+1}; s_L)} \left\{(1-\tau_K)\left[\frac{\alpha\mu_{t+1}}{k_{t+1}}\Big(f'(\mu_{t+1}) - (1-s_K)\psi_K(\widehat{\gamma}_K(\mu_{t+1}; s_K))\Big) - \delta\right] + 1\right\}^{1/\theta} \tag{23}$$

where $\mu_{t+1} = \widehat{\mu}(\mu_t, k_t, c_t;\, \mathbf{s})$ and $k_{t+1} = \widehat{k}(\mu_t, k_t, c_t;\, \mathbf{s})$. We write $c_{t+1} = \widehat{c}(\mu_t, k_t, c_t;\, \mathbf{s}, \tau_K)$.

Equilibrium determination. The three laws of motion $\widehat{\mu}$, $\widehat{k}$, and $\widehat{c}$, together with the initial conditions for μ_0 and k_0 (determined by endowments and period-0 R&D; see Appendix A.3) and a standard transversality condition (Appendix A.3), fully characterize the equilibrium dynamics. A key feature is that μ_{t+1} depends only on (μ_t, k_t, c_t) and the R&D subsidies, not on τ_K; the capital tax enters only through the consumption dynamics $\widehat{c}$. This separation—technology direction first, then saving decisions—carries over to the BGP and underlies the policy results of Sections 4

and 6. In the next section, we analyze the long-run behavior of this economy by studying the balanced growth path.

4 Long-term Analysis

4.1 The Balanced Growth Path

Having characterized the equilibrium dynamics in Section 3.3, we now study the long-run behavior of this economy. We first derive the balanced growth path (BGP) and show that it features a positive rate of capital-augmenting technological change as an endogenous equilibrium outcome. We then conduct comparative statics to examine how R&D subsidies affect the direction of technological change, the long-run growth rate, and the cost-effectiveness of alternative policies.

Definition and intensive-form representation. A balanced growth path is an equilibrium path on which all variables grow at constant rates.

Lemma 1 (BGP characterization)**.** *On a BGP, the normalized state variables μ_t, k_t, and c_t are all constant.*

Proof. See Appendix B.1. □

We denote the BGP values by μ^*, k^*, and c^*. Since $\gamma_{K,t} = \widehat{\gamma}_K(\mu_t; s_K)$ and $\gamma_{L,t} = \widehat{\gamma}_L(\mu_t; s_L)$ by Proposition 1, these are also constant on the BGP: $\gamma_K^* \equiv \widehat{\gamma}_K(\mu^*; s_K)$ and $\gamma_L^* \equiv \widehat{\gamma}_L(\mu^*; s_L)$.

Determination of the direction of technological change. The BGP value of μ^* is pinned down by two steady-state conditions. From the μ-dynamics (20), the condition $\mu_{t+1}/\mu_t = 1$ requires that the growth of M-tasks and N-tasks are equalized. From the k-dynamics (21), the condition $k_{t+1}/k_t = 1$ pins down the BGP growth factor of the aggregate capital stock:

$$\left(\frac{K_{t+1}}{K_t}\right)^* = \widehat{\gamma}_L(\mu^*;\, s_L)\, g_L \equiv g^*. \tag{24}$$

Substituting this into the μ-stationarity condition and simplifying (see Appendix B.1):

$$g_L\,\widehat{\gamma}_L(\mu^*;\,s_L) = g_X\,\widehat{\gamma}_K(\mu^*;\,s_K)^{\alpha/(1-\alpha)} = g^*. \tag{25}$$

Equation (25) equates two expressions for the aggregate growth factor g^*. The left-hand side is the growth of N-tasks: $g_N \equiv g_L\,\gamma_L^*$. The right-hand side is the growth of M-tasks on the BGP: $g_M \equiv g_X\,(\gamma_K^*)^{\alpha/(1-\alpha)}$.[15]

The g_M expression and the capital accumulation multiplier. The exponent $\alpha/(1-\alpha)$ on γ_K^* in the g_M expression is the long-run manifestation of the capital accumulation multiplier introduced in Section 3.3. From the definition of M-tasks, $M_t = (A_{K,t}K_t)^{\alpha}X_t^{1-\alpha}$, a one-time increase in A_K raises M by only $\alpha\%$ on impact (holding K fixed). However, on the BGP, a permanent increase in the KATC *rate* γ_K raises not only A_K growth but also output growth, which accelerates capital accumulation. The resulting geometric series amplifies the growth effect: the elasticity of g_M with respect to γ_K is $\alpha/(1-\alpha)$, not α.[16]

The BGP condition (25) can be visualized by defining two curves as functions of μ:

$$g_M(\mu) \equiv g_X\,\widehat{\gamma}_K(\mu;\,s_K)^{\alpha/(1-\alpha)}, \qquad g_N(\mu) \equiv g_L\,\widehat{\gamma}_L(\mu;\,s_L). \tag{26}$$

Since γ_K is decreasing in μ (Proposition 1), g_M is downward-sloping. Since γ_L is increasing in μ, g_N is upward-sloping. When these two curves intersect, the intersection determines (μ^*, g^*), as shown in Figure 3.

Existence and uniqueness. The monotonicity of g_M and g_N guarantees that if the two curves cross, the intersection is unique. The following proposition establishes existence and uniqueness of the BGP.

Proposition 2 (BGP existence and uniqueness)**.** *Under assumption (4), suppose that*

$$\alpha\,f'(0^+) > \Psi_K(\bar{\gamma}_K) \qquad \textit{and} \qquad \lim_{\mu\to\infty}\left[f(\mu) - \mu f'(\mu)\right] > \Psi_L(\bar{\gamma}_L)\,, \tag{27}$$

[15] The per capita (or per worker) growth factor is $g^*/g_L = \gamma_L^*$.

[16] Even when $\gamma_K = 1$ (no KATC R&D is performed), capital still accumulates at the same rate as output on the BGP, so $g_M = g_X$, not $g_X^{1-\alpha}$. The multiplier converts the exogenous resource growth $g_X^{1-\alpha}$ into g_X through the capital accumulation channel.

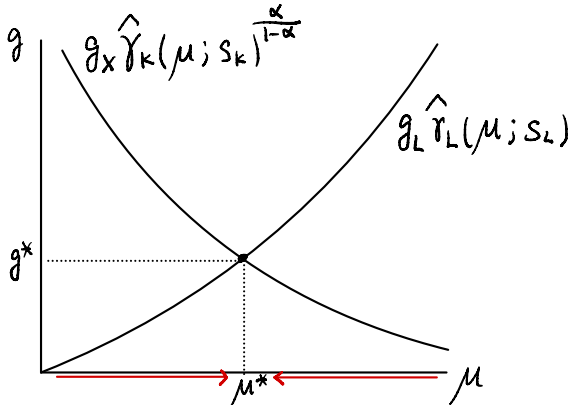


Figure 3: Determination of μ^* on the BGP. The downward-sloping curve $g_M(\mu)$ represents the growth of M-tasks and the upward-sloping curve $g_N(\mu)$ represents the growth of N-tasks. The intersection gives the BGP values (μ^*, g^*).

where $\bar{\gamma}_K \equiv (g_L/g_X)^{(1-\alpha)/\alpha}$ *and* $\bar{\gamma}_L \equiv g_X/g_L$. *Then there exists* $\bar{\beta} > 0$ *such that for all* $\beta \in (0, \bar{\beta})$, *there is a unique BGP with* $\mu^* > 0$, $k^* > 0$, $c^* > 0$, *and the transversality condition (Appendix A.3). On this BGP, both types of technological change are strictly positive:* $\widehat{\gamma}_K(\mu^*; s_K) > 1$ *and* $\widehat{\gamma}_L(\mu^*; s_L) > 1$.

Proof. See Appendix B.1. □

Condition (27) requires that the g_M and g_N curves cross at an interior point: the first inequality prevents the economy from staying at $\mu = 0$, and the second prevents μ from diverging.[17] The upper bound $\bar{\beta}$ rules out a household that is too patient: because per-capita income grows at rate γ_L^* on the BGP, too weak a discounting of the future causes the present value of lifetime income to diverge, violating the transversality condition and ruling out an equilibrium.

The technology condition and directed technological change. An important implication of Proposition 2 is that the technology condition discussed in Section 3.1 is satisfied endogenously. As noted there, Casey and Horii (2024) show that KATC is consistent with balanced growth under the production structure (3), but only if the rates of KATC and LATC satisfy a specific relationship—the technology condition.

[17] Under the standard Inada conditions (including $\lim_{\mu\to\infty}[f(\mu) - \mu f'(\mu)] = \infty$), condition (27) is automatically satisfied. For CES production functions with $\varepsilon < 1$, this limit is finite, so the second inequality in (27) must be verified directly at the calibrated values; we do so in Section 5.

With exogenous technical change, this would be a knife-edge restriction: an equality constraint on two independent parameters. The BGP condition (25) is precisely this technology condition. In our model, firms choose the direction of R&D to maximize profit, so the technology direction μ^* adjusts endogenously until (25) is satisfied. As Figure 3 illustrates, the monotonicity of $g_M(\mu)$ and $g_N(\mu)$ ensures that a unique μ^* is selected at which M-task growth and N-task growth are equalized. Directed technological change thus converts what would be an exogenous knife-edge restriction into an endogenous equilibrium outcome.

Factor income shares. The factor income shares in this economy can be expressed as functions of μ_t and the equilibrium R&D rates. Let $\varepsilon_f(\mu) \equiv \mu f'(\mu)/f(\mu)$ denote the output elasticity of M-tasks. From the R&D cost structure (5) and the equilibrium R&D rates in Proposition 1, the firm-borne R&D expenditure shares are

$$\begin{aligned} \sigma_{RK}(\mu;\, s_K) &\equiv \frac{(1-s_K)\,\mu\,\psi_K\big(\widehat{\gamma}_K(\mu; s_K)\big)}{f(\mu)}, \\ \sigma_{RL}(\mu;\, s_L) &\equiv \frac{(1-s_L)\,\psi_L\big(\widehat{\gamma}_L(\mu; s_L)\big)}{f(\mu)}, \end{aligned} \tag{28}$$

where the numerators are the R&D costs per effective labor borne by the firm, and the denominator is output per effective labor. Both depend on μ through the equilibrium R&D rates $\widehat{\gamma}_K(\mu; s_K)$ and $\widehat{\gamma}_L(\mu; s_L)$. Using the factor demand conditions (11)–(13) and the zero-profit condition, the capital, natural resource, and labor income shares are

$$\begin{aligned} &\sigma_K = \alpha\big(\varepsilon_f(\mu) - \sigma_{RK}(\mu; s_K)\big), \qquad \sigma_X = (1-\alpha)\big(\varepsilon_f(\mu) - \sigma_{RK}(\mu; s_K)\big), \\ &\sigma_L = 1 - \varepsilon_f(\mu) - \sigma_{RL}(\mu; s_L). \end{aligned} \tag{29}$$

These shares satisfy $\sigma_K + \sigma_X + \sigma_L + \sigma_{RK} + \sigma_{RL} = 1$ at every t, where $\sigma_R \equiv \sigma_{RK} + \sigma_{RL}$ denotes the private (post-subsidy) R&D expenditure share. The adding-up condition reflects the zero-profit condition: the government reimburses the subsidy portion of R&D costs, so only the firm-borne share σ_R appears in the income decomposition. The capital and natural resource shares maintain a constant ratio $\sigma_K/\sigma_X = \alpha/(1-\alpha)$,

reflecting the Cobb-Douglas structure of the capital composite.[18]

On the BGP, $\mu_t = \mu^*$ is constant, so the factor income shares are also constant and depend only on the first-stage variables. In particular, since μ^* determines γ_K^* and γ_L^*, the BGP shares σ_K^*, σ_X^*, and σ_L^* are fully determined by μ^* and the subsidy rates, independently of the household parameters β and θ or the capital income tax τ_K.

r^*, k^*, and c^* on the BGP. We now turn to the determination of the interest rate, capital, and consumption on the BGP. Unlike the first-stage variables (μ^*, g^*, and the factor shares), these depend on household preferences and capital taxation.

On the BGP, c_t is constant, so the Euler equation (22) reduces to a condition on the BGP interest rate:

$$(1 - \tau_K)(r^* - \delta) + 1 = \frac{(\gamma_L^*)^\theta}{\beta}. \tag{30}$$

The right-hand side is the intertemporal marginal rate of substitution: even though consumption per effective labor is constant on the BGP, per capita consumption grows at rate γ_L^*, so the household requires a gross return of $(\gamma_L^*)^\theta/\beta$.[19]

The capital intensity k^* is determined by the rental rate condition (11). Substituting (30) for r^* and solving for k^*:

$$k^* = \frac{(1 - \tau_K)\,\alpha\,\mu^*\left\{f'(\mu^*) - (1 - s_K)\,\psi_K\big(\widehat{\gamma}_K(\mu^*;\, s_K)\big)\right\}}{\dfrac{(\gamma_L^*)^\theta}{\beta} - 1 + (1 - \tau_K)\,\delta}. \tag{31}$$

The numerator is the after-tax net marginal product of capital, and the denominator is the effective discount rate. Since μ^* and γ_K^* are determined in the first stage by (25), the numerator depends on s_K and τ_K, while the denominator also depends on household parameters β and θ.

Consumption per effective labor is determined by the resource constraint on the BGP:

$$c^* = v(\mu^*;\, \mathbf{s}) - (g^* - 1 + \delta)\,k^*, \tag{32}$$

[18] For the derivation, see Appendix B.1.

[19] Condition (30) plays a key role in the robustness analysis of Section 5. When data moments pin down $r^* - \delta$, γ_L^*, τ_K, and δ, varying θ requires only β to adjust, and the first-stage variables (μ^*, g^*, γ_K^*, γ_L^*) are unaffected.

where $v(\mu^*; \mathbf{s})$ is the net output per effective labor defined in (17).

To summarize, the BGP variables are determined in a hierarchical fashion. In the first stage, μ^*, g^*, and the factor income shares are determined by the production and R&D primitives (f, ψ_j, α, g_X, g_L) and the R&D subsidies (s_K, s_L), independently of household preferences and capital taxation. In the second stage, the interest rate r^* and capital intensity k^* are determined by (30)–(31), bringing in β, θ, and τ_K. In the third stage, c^* follows as a residual from (32). This separation implies that R&D subsidies affect growth and factor shares (first stage), whereas capital income taxation affects only the levels of k^* and c^* (second and third stages).

4.2 Comparative Statics

We now examine how the two directed R&D subsidies, s_K and s_L, affect the balanced growth path. We present three analytical results: how each subsidy shifts the direction of technological change and the growth rate (Proposition 3), how a uniform subsidy tilts the technology direction (Proposition 5), and which type of subsidy is more cost-effective in raising growth (Proposition 6). Because the direction of technological change also affects factor income shares through Proposition 4, R&D subsidies have distributional consequences even when they are not explicitly directed. All proofs are in Appendix B.

Effect of directed subsidies on the direction of technological change. The first result characterizes how directed subsidies shift μ^* and g^*.

Proposition 3 (Directed subsidies and technology direction)**.** *(i) An increase in s_K raises μ^*, raises g^*, raises γ_K^*, and lowers γ_L^*. (ii) An increase in s_L lowers μ^*, raises g^*, raises γ_L^*, and lowers γ_K^*.*

Proof. See Appendix B.2. □

Figure 4 illustrates the mechanism using the g_M–g_N diagram from Section 4.1. A KATC subsidy s_K reduces the effective cost of capital-augmenting R&D, which raises γ_K for each μ. Since $g_M(\mu) = g_X \widehat{\gamma}_K(\mu; s_K)^{\alpha/(1-\alpha)}$, the g_M curve shifts upward, while the g_N curve is unaffected. The new intersection lies to the right and above the original one: μ^* and g^* both increase. Symmetrically, an LATC subsidy s_L shifts

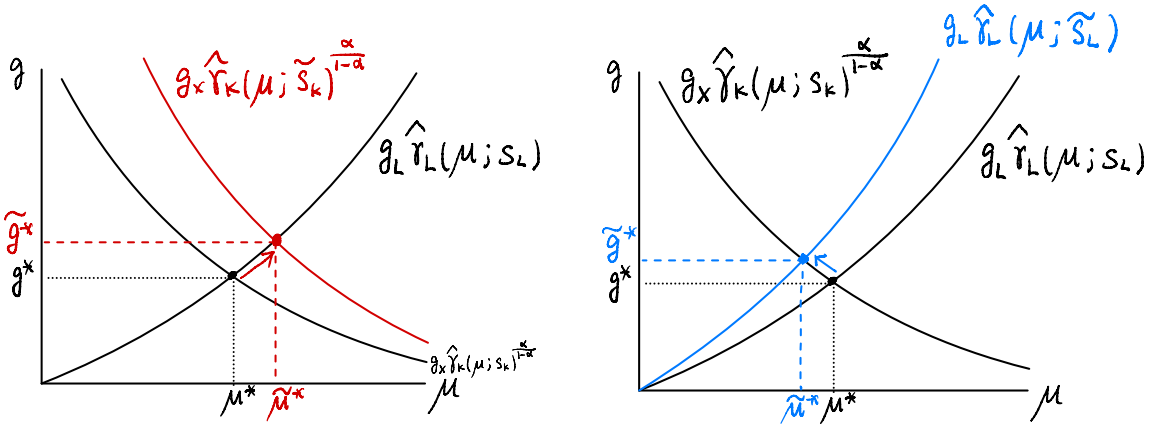


Figure 4: Policy effects on the BGP. Left: an increase in s_K shifts the g_M curve upward, raising μ^* and g^*. Right: an increase in s_L shifts the g_N curve upward, lowering μ^* while raising g^*.

the g_N curve upward, moving the intersection to the left and above: μ^* falls while g^* rises.

We can confirm that both subsidies raise g^*: directed R&D subsidies are always growth-enhancing regardless of which type of technology they target. However, the two subsidies have opposite effects on μ^*—the relative intensity of capital-composite tasks—and hence on the direction of technological change. A KATC subsidy tilts innovation toward the capital side, whereas an LATC subsidy tilts it toward the labor side. This asymmetry has consequences for factor income shares, as we explore next.

Labor share and the direction of technological change. To understand the distributional implications of Proposition 3, we need the relationship between μ^* and the labor income share σ_L^*. From (29), the labor share depends on μ^* through both the output elasticity $\varepsilon_f(\mu^*)$ and the R&D expenditure share σ_{RL}^*, and s_L also affects the labor share directly through γ_L^*. Under general functional forms, the relationship between the direction of technological change and factor shares can be complex. Under the CES and power-function specifications used in our calibration, however, the intuitive result obtains unconditionally: a KATC subsidy raises the labor share, and an LATC subsidy lowers it, with no additional parameter conditions

beyond $\lambda > 1$ and $\varepsilon < 1$.

Proposition 4 (Labor share and directed R&D subsidies)**.** *Suppose $f(\mu)$ is CES with elasticity of substitution $\varepsilon \in (0,1)$ and each $\psi_j(\gamma_j)$ is a power function of $(\gamma_j - 1)$ with exponent $\lambda > 1$, for $j = K, L$.*[20] *For any $s_K, s_L \in [0,1)$:*

(a) $d\sigma_L^*/ds_K > 0$: a KATC subsidy raises the labor share.

(b) $d\sigma_L^*/ds_L < 0$: an LATC subsidy lowers the labor share.

Proof. See Appendix B.3. □

Part (a) captures a paradox: subsidizing *capital*-augmenting technology raises the *labor* share. The mechanism is gross complementarity. A KATC subsidy raises μ^*, making M-tasks (the capital composite) relatively abundant and N-tasks (effective labor) relatively scarce. When capital and labor are gross complements ($\varepsilon < 1$), the resulting scarcity of effective labor raises its marginal product proportionally more, increasing the labor share.

Part (b) is the mirror image, reinforced by a second channel. An LATC subsidy lowers μ^* (Proposition 3), making labor relatively abundant and reducing the gross labor share $\sigma_L^G = \sigma_L^* + \sigma_{RL}^*$. In addition, the subsidy directly stimulates LATC R&D, raising γ_L^* and with it the private R&D expenditure σ_{RL}^* that is deducted from labor's gross share. Because both channels depress σ_L^*, the result requires no additional conditions. The bottom line is striking: capital-augmenting subsidies benefit labor, and labor-augmenting subsidies benefit capital.

Direction of uniform subsidies. Even when the subsidy is not deliberately directed, it need not be neutral with respect to the technology direction.

Proposition 5 (Uniform subsidy and technology direction)**.** *Under a uniform subsidy $s_K = s_L = s$, an increase in s raises μ^* whenever*

$$\frac{\varepsilon_{\Psi K}(\gamma_K^*)}{\varepsilon_{\Psi L}(\gamma_L^*)} < \frac{\alpha}{1-\alpha}, \tag{33}$$

where $\varepsilon_{\Psi j}(\gamma_j) \equiv \gamma_j \Psi_j'(\gamma_j)/\Psi_j(\gamma_j)$ is the elasticity of the composite marginal R&D cost.

[20] In Section 5, we parametrize the model with these functional forms and calibrate the parameters to U.S. data.

Proof. See Appendix B.4. □

The intuition can be seen by decomposing the growth effect of a uniform subsidy. Holding μ fixed, a uniform subsidy raises both γ_K and γ_L, shifting both the g_M and g_N curves upward. The question is which curve shifts more. Each R&D channel responds to the cost reduction in proportion to $1/\varepsilon_{\Psi j}$, but the K-side response feeds into growth through capital accumulation, amplifying it by the multiplier $\alpha/(1-\alpha) > 1$.[21] Whether g_M or g_N shifts more therefore depends on whether $[\alpha/(1-\alpha)] \times (1/\varepsilon_{\Psi K})$ exceeds $1/\varepsilon_{\Psi L}$, which is precisely condition (33).[22]

When condition (33) holds, the g_M curve shifts more than the g_N curve, so the intersection in the g_M–g_N diagram moves to the right, raising μ^*. Since the capital accumulation multiplier $\alpha/(1-\alpha)$ exceeds one whenever $\alpha > 1/2$, the condition is naturally satisfied when the two R&D cost functions are not too dissimilar. A uniform subsidy has two competing effects on the labor share. The first is a *pro-labor* effect: the directional tilt toward capital augmentation (higher μ^*) raises the marginal product of effective labor relative to that of the capital composite under gross complementarity. The second is an *anti-labor* effect: the subsidy directly stimulates LATC R&D, expanding the R&D cost share σ^*_{RL} and depressing the labor share. Whether the net effect is positive depends on whether the directional tilt is strong enough to overcome the direct LATC stimulus. In Appendix B.6, we provide a sufficient condition under which the pro-labor effect dominates at the margin, and verify that it holds with ample margin at the calibrated values.

Cost-effectiveness of directed subsidies. We now compare the fiscal cost of raising the BGP growth rate through each type of subsidy. Let σ_S denote the subsidy expenditure as a share of GDP. The cost-effectiveness of each subsidy is measured by $d\sigma^*_S/dg^*$: the increase in the subsidy-to-GDP ratio required to raise the BGP growth factor by one unit.

Proposition 6 (Cost-effectiveness of directed subsidies)**.** *Starting from $s_K = s_L = 0$, the KATC subsidy achieves a given increase in g^* at a lower subsidy-to-GDP cost*

[21] More precisely, holding μ constant, $d\log g_M = [\alpha/(1-\alpha)] \times (1/\varepsilon_{\Psi K}) \times ds/(1-s)$ and $d\log g_N = (1/\varepsilon_{\Psi L}) \times ds/(1-s)$. The factor $1/\varepsilon_{\Psi j}$ is the percentage increase in the R&D rate per percentage-point reduction in the effective cost, obtained by differentiating the FOC at fixed μ.

[22] We verify this condition at the calibrated values in Section 5.3.

(left-hand side) than the LATC subsidy (right-hand side):

$$\left.\frac{d\sigma_S^*}{dg^*}\right|_{\substack{s_K:variable\\ s_L:fixed}} < \left.\frac{d\sigma_S^*}{dg^*}\right|_{\substack{s_K:fixed\\ s_L:variable}}, \tag{34}$$

whenever

$$\frac{\varepsilon_{\psi K}(\gamma_K^*)}{\varepsilon_{\psi L}(\gamma_L^*)} > \alpha, \tag{35}$$

where $\varepsilon_{\psi j}(\gamma_j) \equiv \gamma_j \psi_j'(\gamma_j)/\psi_j(\gamma_j)$ *is the elasticity of the per-task R&D cost.*

More generally, for $s_K, s_L \geq 0$*, the KATC subsidy remains more cost-effective as long as*

$$\alpha \frac{\psi_K(\gamma_K^*)}{\Psi_K(\gamma_K^*)}\left(1 + \frac{s_K}{1-s_K}\frac{1-\varepsilon_f(\mu^*)}{\varepsilon_{f'}(\mu^*)}\right) < \frac{\psi_L(\gamma_L^*)}{\Psi_L(\gamma_L^*)}\left(1 + \frac{s_L}{1-s_L}\frac{1-\varepsilon_f(\mu^*)}{\varepsilon_{f'}(\mu^*)}\right), \tag{36}$$

where $\varepsilon_{f'}(\mu) \equiv -\mu f''(\mu)/f'(\mu) > 0$.

Proof. See Appendix B.5. □

The economic content of condition (35) can be seen by substituting the first-order conditions $\alpha f'(\mu^*) = \Psi_K(\gamma_K^*)$ and $f(\mu^*) - \mu^* f'(\mu^*) = \Psi_L(\gamma_L^*)$ (evaluated at $s_K = s_L = 0$), which yields the equivalent inequality

$$\frac{f'(\mu^*)}{\psi_K(\gamma_K^*)} > \frac{f(\mu^*) - \mu^* f'(\mu^*)}{\psi_L(\gamma_L^*)}. \tag{37}$$

The left-hand side is the marginal product of an M-task per unit of KATC R&D cost; the right-hand side is the marginal product of an N-task per unit of LATC R&D cost. The proposition says that the KATC subsidy is more cost-effective whenever the task-level output gain per unit of R&D cost is higher on the K-side.[23]

When each ψ_j is a power function of $(\gamma_j - 1)$ with exponent λ, we have $\varepsilon_{\psi j} = \lambda\gamma_j/(\gamma_j - 1)$, which is decreasing in γ_j. If $\gamma_K^* < \gamma_L^*$, as suggested by the data (see Section 5), then $\varepsilon_{\psi K} > \varepsilon_{\psi L} > 1 > \alpha$, so condition (35) is satisfied automatically.

The general condition (36) extends the result to economies with existing subsidies. When $s_K = s_L = 0$, it reduces to condition (35). For larger programs, the

[23] The threshold in (35) is α rather than one because the firm's FOC equates Ψ_K to $\alpha f'(\mu)$, not $f'(\mu)$: the marginal product of an M-task enters the KATC R&D decision scaled by α, so even when R&D cost elasticities are symmetric ($\varepsilon_{\psi K} = \varepsilon_{\psi L}$), the K-side ratio in (37) exceeds the L-side ratio by a factor of $1/\alpha$. This makes condition (35) easy to satisfy.

cost-effectiveness advantage of KATC subsidies diminishes as a growing share of subsidy expenditure goes toward R&D that firms would have undertaken at lower rates. If subsidy rates become sufficiently large, condition (36) can be violated and the KATC advantage disappears; this is especially likely when the subsidy mix is heavily tilted toward KATC, since a high s_K relative to s_L raises the left-hand side of (36) disproportionately.

Long-run neutrality of capital income taxation. We also note that, in our framework, the capital income tax τ_K is neutral with respect to the BGP growth rate and the direction of technological change. From the hierarchical structure of the BGP (end of Section 4.1), the first-stage variables μ^*, g^*, γ_K^*, γ_L^*, and the factor income shares are all determined by the R&D conditions (14)–(15) and the BGP condition (25), which do not involve τ_K. The capital income tax enters only through the Euler equation (30), so its long-run effects are confined to the levels of k^* and c^*.

The reason τ_K is absent from the R&D conditions is that, under U.S. tax law, R&D expenditures are deductible as a current operating expense in the year incurred (IRC §174A), and R&D is predominantly financed from internal funds (Hall and Lerner, 2010). Physical capital expenditures, by contrast, must generally be capitalized and recovered over multi-year depreciation schedules (5–7 years for equipment, up to 39 years for structures under MACRS). Our model captures this asymmetry: R&D expenditure is paid out of pre-tax earnings in the profit function (6), so the R&D decision is unaffected by τ_K. The return on physical capital, by contrast, is subject to τ_K in the household budget constraint (8): although depreciation is excluded from the tax base, the net return $(r_t - \delta)$ is taxed, creating a wedge in the saving decision. Along the transitional path, τ_K does affect the speed of capital accumulation and hence the dynamics of all variables; we discuss this in Section 6.2.

How much directed R&D subsidies raise the growth rate and shift the labor share depends on parameter values, which we now estimate by matching U.S. macroeconomic data.

5 Calibration

5.1 Functional Forms and Data

Functional forms. We adopt the CES specification for the aggregate production function and power-function R&D costs:

$$F(M_t, N_t) = B\left[\eta\, M_t^{\frac{\varepsilon-1}{\varepsilon}} + (1-\eta)\, N_t^{\frac{\varepsilon-1}{\varepsilon}}\right]^{\frac{\varepsilon}{\varepsilon-1}}, \quad \varepsilon > 0, \quad \eta \in (0,1), \tag{38}$$

$$\psi_K(\gamma_K) = z_K(\gamma_K - 1)^\lambda, \qquad \psi_L(\gamma_L) = z_L(\gamma_L - 1)^\lambda, \qquad \lambda > 1. \tag{39}$$

The CES form nests Cobb–Douglas ($\varepsilon = 1$) as a special case. The power-function R&D cost satisfies condition (4) and is a standard assumption in the endogenous growth literature (Acemoglu et al., 2018). We allow the cost level parameters z_K and z_L to differ across technologies but assume the same degree of convexity λ for both.

Parameters and the period length. With these functional forms, the model has 15 parameters: 12 structural parameters—B, ε, η, α, λ, z_K, z_L, β, θ, δ, g_L, g_X—and 3 policy parameters—s_K, s_L, τ_K. To calibrate the model, we also need to determine the *period length* χ, measured in years.

The period length carries an important economic meaning. As described in Section 3.1, technical knowledge can be kept within the firm for only one period, after which it becomes public. Thus χ represents the duration for which a firm can monopolize the benefit from its R&D investments.

Moreover, χ plays a central role in the calibration because the model's one-period quantities must be mapped to annual data. Throughout this section, we use a superscript a to denote *annualized rates*: for example, $\gamma_L^{*a} = 1.86\%$ is the annualized growth rate of per-capita output on the BGP, while the corresponding per-period variable $\gamma_L^* = (1+\gamma_L^{*a})^\chi$ is the growth factor over one model period of χ years. Similarly, δ^a and β^a denote the annualized depreciation rate and discount rate, respectively. All per-period variables used in the model—growth factors, depreciation rates, discount factors, and R&D cost parameters—depend on χ through such compounding relationships (see Appendix C.1 for the full list of conversion formulas). As we show in Section 5.2, χ is not set exogenously but is itself pinned down by the data through

an equilibrium condition.

Exogenous parameters. The parameter λ governs the convexity of the R&D cost function. Two lines of empirical evidence support quadratic costs ($\lambda = 2$) as the baseline. First, the empirical literature on the elasticity of patents with respect to R&D expenditures finds values around 0.5, which implies quadratic curvature of the underlying cost function (Griliches, 1990; Hall and Ziedonis, 2001; Blundell et al., 2002). Second, studies of the impact of R&D tax credits on firm-level R&D expenditure find a tax price elasticity of approximately unity, which again corresponds to a quadratic cost function (Hall, 1993; Bloom et al., 2002; Wilson, 2009). Both Acemoglu et al. (2018) and Akcigit and Kerr (2018) adopt quadratic R&D costs based on these estimates. We also examine $\lambda = 3/2$ and $\lambda = 5/2$ as alternate values to confirm the robustness of the results.

We set $\theta = 1$ (log utility) as the baseline and examine $\theta \in \{1/2, 3/2\}$ as alternate values. Indeed, the predictions for growth and the direction of technological change are completely invariant to θ when the model is recalibrated to the same data moments.

We also assume that the pre-policy capital income tax rate is $\tau_K = 15\%$, a conservative benchmark based on the U.S. long-term capital gains tax rate.[24] We verify robustness at $\tau_K = 23.8\%$ (20% federal capital gains plus the 3.8% Net Investment Income Tax) in Section 5.3. We also assume that there are no R&D subsidies ($s_K = s_L = 0$) at the baseline. Note that this does not ignore the tax deductibility of R&D: as discussed in Section 4.2, R&D expenditure is deducted from the firm's pre-tax revenue in the profit function (6), consistent with its treatment as a current operating expense under U.S. tax law. Table 1 summarizes the exogenous parameter values.

Target moments. Including χ, 5 of the 16 unknowns are set exogenously (Table 1). The remaining 11 are pinned down by 9 empirical moments from U.S. data, the BGP technology condition $g_M = g_N$, and the factor share adding-up condition $\sigma_K^* + \sigma_X^* +$

[24] In the model, the return r encompasses capital gains, dividends, and interest payments, so the effective rate on total capital income may be higher.

Table 1: Exogenous parameters

Parameter	Symbol	Baseline	Alternatives
R&D cost convexity	λ	2	3/2, 5/2
CRRA parameter	θ	1	1/2, 3/2
KATC R&D subsidy	s_K	0%	Varied in Section 6
LATC R&D subsidy	s_L	0%	Varied in Section 6
Capital income tax	τ_K	15%	23.8%; varied in Section 6

Table 2: Target moments

Moment	Variable	Value	Source / method
Capital–output ratio	$(K/Y)^{*a}$	2.9	BEA/FRED, 1960–2019 mean
Labor income share	σ_L^*	62.5%	BEA/FRED, 1960–2020 mean
Private R&D expenditure share	σ_R^*	2.7%	BEA/FRED, 1960–2020 mean
Depreciation / GDP	$(\delta K/Y)^{*\dagger}$	14%	BEA/FRED, 1960–2020 mean
Per-capita growth rate	γ_L^{*a}	1.86%	Log-linear est., 1960–2024
Relative investment price decline	γ_K^{*a}	0.80%	Log-linear est., 1960–2020
Net return on capital	$(r^* - \delta)^a$	4%	McGrattan et al. (2003)
Population growth rate	g_L^a	1.02%	Log-linear est., 1960–2020
K–L elasticity of substitution	EoS_{KL}	0.6 (0.8)	Oberfield & Raval (2014)

†Since depreciation δK and output Y are both flow variables measured over the same period, the period length scales both equally and their ratio does not require annualization. The same applies to σ_L^*, σ_R^*, and EoS_{KL}.

$\sigma_L^* + \sigma_R^* = 1$, yielding an exactly identified system.[25] Table 2 lists the 9 moments; all except the elasticity of substitution are constructed from Bureau of Economic Analysis (BEA) data accessed through FRED (Federal Reserve Bank of St. Louis).

Per-capita growth and the consumption deflator. On the BGP, the per-capita growth rate of consumption equals γ_L^{*a}, the annualized rate of LATC. In standard macroeconomic practice, real GDP growth is measured using the GDP deflator, and most calibration exercises equate this with the growth rate of final output in the model. However, the GDP deflator incorporates the prices of both consumption

[25] Among the 11 unknowns, g_X is not directly targeted but recovered from the BGP technology condition (equation 25). In addition, not all income shares are independent observables: σ_K^* is inferred from the return, capital–output, and depreciation moments, while σ_X^* and α are jointly recovered from the adding-up condition and the production structure $\sigma_K^*/\sigma_X^* = \alpha/(1-\alpha)$. See the "untargeted moments" discussion below.

goods and investment goods. Since the model's output is denominated in consumption units, it does not reflect the decline in the relative price of investment. Using the GDP deflator to measure real output growth would therefore confound two distinct phenomena: the growth of real output measured in consumption units, and the cheapening of investment goods due to KATC.

To obtain a measure consistent with the model's accounting, we deflate nominal GDP by the *Personal Consumption Expenditure (PCE) price index* rather than the GDP deflator. Log-linear estimation over 1960–2024 yields $\gamma_L^{*a} = 1.86\%$. The ability to make this distinction is an advantage of our framework: because we explicitly model the decline in the relative price of investment through endogenous KATC, we can correctly separate consumption-good output growth from the investment price effect.

Measuring KATC from relative investment prices. In our model, the resource constraint (18) measures output, consumption, and investment in the same unit (the consumption good). One unit of the consumption good can be converted into one unit of physical capital K, which represents A_K efficiency units of effective capital. As a result, the price of one efficiency unit of capital in terms of the consumption good is $1/A_K$, and this price declines at rate γ_K^{*a} per year as A_K grows. We measure γ_K^{*a} from the ratio of the BEA investment price deflator to the consumption price deflator, applying log-linear estimation over 1960–2020, which yields $\gamma_K^{*a} = 0.80\%$ per year.

All growth-rate moments are estimated by log-linear estimation (OLS of $\log y$ on time), which uses all data points and reduces sensitivity to endpoint outliers. Level moments are arithmetic averages over the indicated sample period.

The elasticity of substitution between capital and labor. The ninth moment, the elasticity of substitution between capital and labor (EoS_{KL}), determines the CES parameter ε. Oberfield and Raval (2021) estimate this elasticity using micro-level variation in capital shares across U.S. manufacturing plants, obtaining a value of approximately 0.6; Antras et al. (2004) report a similar estimate using aggregate time-series methods. A meta-regression analysis of 77 U.S. studies finds that the long-run aggregate elasticity lies in the range 0.45–0.87, with substantial heterogeneity driven

by modeling choices for technological dynamics (Knoblach et al., 2020). Our baseline of $\text{EoS}_{KL} = 0.6$ sits near the center of this range; we also examine $\text{EoS}_{KL} = 0.8$, near the upper end, as a robustness check in Section 5.3. Because of the nested production structure, ε differs from EoS_{KL}; we explain this mapping in Section 5.2.

5.2 Parameter Estimation

Overview of the calibration procedure. Given the 5 exogenous parameters (Table 1) and 9 target moments (Table 2), we determine the 11 remaining unknowns—B, ε, η, α, z_K, z_L, δ, β, g_L, g_X, and the period length χ. Most parameters can be expressed as closed-form functions of χ and the target moments. The key step is therefore pinning down χ.

The period length χ is determined first, by an equilibrium condition that links the R&D share to the other moments (described below). Note that several intermediate quantities—such as $\alpha(\chi)$ and $\sigma_K^*(\chi)$—depend on χ and are used in this condition. Once χ is pinned down, the remaining parameters are determined sequentially: (i) g_L and other per-period growth factors from the annualization formulas (e.g., $g_L = (1 + g_L^a)^\chi$); (ii) k^*, $f(\mu^*)$, and μ^* from the normalization; (iii) δ and r^* from the depreciation and return targets; (iv) β from the Euler equation; (v) g_X from the BGP technology condition (25); (vi) $\varepsilon_f(\mu^*)$ from the factor share and R&D first-order conditions; (vii) ε from the elasticity of substitution condition; (viii) η from the definition of ε_f; (ix) B from the production function; (x) z_K and z_L from the R&D first-order conditions. The complete derivation of each step is in Appendix C.1.

Determination of the period length. By substituting the R&D first-order conditions into the factor share equations, we can express the R&D expenditure share as a function of factor shares, technology growth rates, and the R&D cost curvature (see Appendix C.1 for the derivation):

$$\sigma_R^* = \underbrace{\frac{\sigma_K^*(\chi)}{\lambda}\frac{\gamma_K^*(\chi)-1}{\gamma_K^*(\chi)}}_{\sigma_{RK}^*(\chi)\text{: KATC R\&D share}} + \underbrace{\frac{\sigma_L^*}{\lambda}\frac{\gamma_L^*(\chi)-1}{\gamma_L^*(\chi)}}_{\sigma_{RL}^*(\chi)\text{: LATC R\&D share}}. \tag{40}$$

The right-hand side decomposes the aggregate R&D share into its capital- and labor-augmenting components, both of which depend on the period length χ through the per-period growth factors $\gamma_K^*(\chi) = (1+\gamma_K^{*a})^\chi$ and $\gamma_L^*(\chi) = (1+\gamma_L^{*a})^\chi$, and (for the KATC component) the capital share $\sigma_K^*(\chi)$.[26]

Since a longer period implies larger per-period innovation rates and hence higher R&D expenditure shares, the right-hand side of (40) is strictly increasing in χ, while the left-hand side is fixed at $\sigma_R^* = 2.7\%$. The period length is therefore uniquely pinned down; at the baseline, $\chi \approx 4.10$ years.

The elasticity of substitution between capital and labor. The CES parameter ε governs the elasticity of substitution between M-tasks and N-tasks. However, the empirically estimated elasticity of substitution between capital and labor, EoS_{KL}, differs from ε because of the nested production structure: capital enters M-tasks through the Cobb–Douglas capital composite $(A_K K)^\alpha X^{1-\alpha}$ with exponent $\alpha < 1$. As shown in Casey and Horii (2024), the relationship between the two is

$$\text{EoS}_{KL} = \frac{1 + \dfrac{KF_K}{LF_L}}{\dfrac{1}{\varepsilon}\left(\dfrac{KF_K}{LF_L} + \alpha\right) + (1-\alpha)}, \qquad \text{where} \quad \frac{KF_K}{LF_L} = \frac{\alpha\,\varepsilon_f(\mu^*)}{1-\varepsilon_f(\mu^*)}. \tag{41}$$

When $\alpha = 1$, M-tasks consist solely of effective capital and this formula reduces to $\text{EoS}_{KL} = \varepsilon$. When $\alpha < 1$, the presence of natural resources in the capital composite dampens the substitutability between K and L, so $\text{EoS}_{KL} < \varepsilon$. We target $\text{EoS}_{KL} = 0.6$ and solve for ε from (41). At the calibrated values ($\alpha = 0.757$, $\varepsilon_f(\mu^*) = 0.352$), we obtain $\varepsilon = 0.554$.

Baseline calibration. The system is exactly identified: each unknown is uniquely determined and all 9 target moments are matched exactly by construction. Table 3 summarizes the calibrated parameter values under the baseline specification ($\lambda = 2$, $\theta = 1$).

The calibrated model generates a per-capita output growth rate of $g_{\text{pc}}^a = \gamma_L^{*a} = 1.86\%$ and an aggregate output growth rate of $g^a = 2.90\%$. The BGP values of the

[26] The capital share $\sigma_K^*(\chi)$ is itself a function of χ because it is inferred from the return and depreciation targets: $\sigma_K^*(\chi) = [(1.04)^\chi - 1] \cdot (2.9/\chi) + 0.14$.

Table 3: Calibrated parameters (baseline: $\lambda = 2$, $\theta = 1$)

Parameter	Symbol	Value	Annualized
Period length	χ	4.10	—
Share of K in M	α	0.757	—
Depreciation rate	δ	0.198	5.24%
Discount factor	β	0.939	1.54%
KATC R&D cost	z_K	7.03	28.8
LATC R&D cost	z_L	15.1	62.1
CES elasticity (M vs. N)	ε	0.554	—
CES distribution	η	0.510	—
TFP constant	B	2.90	0.707
Natural resource growth	g_X	1.016	0.38%

The Value column reports per-period quantities. The Annualized column shows the corresponding annualized rates or the per-year equivalents $z_j^a = \chi^{\lambda-1} z_j$ and $B^a = B/\chi$. Parameters that do not depend on χ (such as α, ε, and η) have no separate annualized form.

state variables are $\mu^* = 2.24$, $k^* = 2.9$, and $c^* = 3.06$ (in per-effective-labor units).

Untargeted moments. In addition to the 9 targeted moments, the calibrated model implies values for the capital income share and the natural resource income share. Since $\sigma_L^* = 62.5\%$ and $\sigma_R^* = 2.7\%$ are targeted, the non-labor, non-R&D income $1 - \sigma_L^* - \sigma_R^* = 34.8\%$ is split into $\sigma_K^* = 26.3\%$ and $\sigma_X^* = 8.5\%$ by the calibrated $\alpha = 0.757$. The implied σ_X^* is consistent with Caselli and Feyrer (2007), who estimate the share of all non-reproducible factors at about 8% for the U.S.

The implied $g_X^a = 0.38\%$ is derived from the BGP technology condition (25), not set directly. To assess its plausibility, we decompose X into its components. Valentinyi and Herrendorf (2008) estimate that land accounts for about 5% of GDP, which represents roughly five-eighths of the Caselli–Feyrer non-reproducible share. The remaining three-eighths is attributable to fossil fuels and other natural resources, whose physical consumption in the U.S. has grown at roughly 1% per year over 1960–2019.[27] Since the supply of land is essentially fixed, the weighted growth rate of the composite resource is approximately $(1.01)^{3/8} - 1 \approx 0.37\%$, which is close to our calibrated value. Table 4 summarizes these untargeted implications.

[27] Computed from the U.S. Energy Information Administration, *Monthly Energy Review*, Table 1.3 (Primary Energy Consumption by Source).

Table 4: Untargeted moments

Moment	Symbol	Value	Annualized
Capital income share	σ_K^*	26.3%	—
Natural resource share	σ_X^*	8.5%	—
Natural resource growth	$g_X - 1$	1.56%	0.38%

5.3 Long-Run Implications

Saddle-path stability and convergence speed. The BGP is saddle-path stable for all parameter configurations considered in this section (Table 5): the linearized dynamics around (μ^*, k^*, c^*) have exactly two stable eigenvalues and one unstable eigenvalue, as required for a well-determined equilibrium with two predetermined variables (μ, k) and one jump variable (c).

The two stable eigenvalues reveal two distinct time scales of convergence with important policy implications. The slow mode (half-life $\approx$ 93 years) governs the adjustment of the technology direction μ; since γ_K, γ_L, and the factor income shares are all determined by μ in the first stage (Section 4.1), this mode sets the pace at which growth and income distribution converge to their new long-run values. The fast mode (half-life $\approx$ 5 years) governs the Ramsey-type adjustment of k and c, which track μ's slow movement closely. The full growth and distributional benefits of R&D subsidies therefore materialize only over more than 100 years—a point we quantify through the transitional dynamics in Section 6.[28]

Quantitative comparative statics. We now evaluate the analytical conditions at the calibrated values to assess the quantitative policy implications.

The key question for policy design is how much growth each type of R&D subsidy delivers per unit of fiscal cost. Proposition 4 requires no parameter conditions beyond $\lambda > 1$ and $\varepsilon < 1$, both of which hold by assumption. The sufficient conditions for Propositions 5 and 6 are also satisfied with ample margin.[29] As a result, all the qualitative predictions of Section 4.2 are confirmed: KATC subsidies raise the labor share while LATC subsidies lower it (Proposition 4); a uniform subsidy tilts the

[28] See Appendix C.2 for the eigenvalue computation and eigenvector decomposition.

[29] Condition (33) requires $\varepsilon_{\Psi K}/\varepsilon_{\Psi L} < \alpha/(1-\alpha)$; at the calibrated values, $2.26 < 3.11$. Condition (35) requires $\varepsilon_{\psi K}/\varepsilon_{\psi L} > \alpha$; we obtain $2.26 > 0.757$.

Table 5: Saddle-path stability and convergence speed across parameter configurations

					Half-life (years)		
Configuration	θ	λ	τ_K	EoS_{KL}	Slow (μ)	Fast (k)	Stability
Baseline	1	2	15%	0.6	93.3	5.2	Saddle
$\theta = 1/2$	1/2				90.2	3.8	Saddle
$\theta = 3/2$	3/2				93.3	6.3	Saddle
$\lambda = 3/2$		3/2			50.7	4.8	Saddle
$\lambda = 5/2$		5/2			136.1	5.5	Saddle
$\tau_K = 23.8\%$			23.8%		93.3	5.5	Saddle
$\text{EoS}_{KL} = 0.8$				0.8	100.4	6.0	Saddle

Blank entries denote baseline values ($\theta = 1$, $\lambda = 2$, $\tau_K = 15\%$, $\text{EoS}_{KL} = 0.6$). Each row varies only the parameter indicated in its label; all others are held at baseline. Half-life $= \chi \ln 2/|\ln \lambda_i|$, where λ_i is the eigenvalue of the linearized system. Slow: technology direction μ. Fast: capital intensity k and consumption c. Eigenvalues are reported in Appendix C.2.

technology direction toward KATC (Proposition 5); and the KATC subsidy is more cost-effective (Proposition 6). The marginal cost-effectiveness ratio—the growth gain per unit of fiscal cost from a KATC subsidy relative to an LATC subsidy—is

$$\frac{dg^a/d\sigma_S|_{s_K}}{dg^a/d\sigma_S|_{s_L}} = \frac{0.376}{0.129} \approx 2.92. \tag{42}$$

That is, a marginal KATC subsidy raises the per-capita growth rate roughly *three times* as much as a marginal LATC subsidy per unit of fiscal expenditure.

Robustness to the R&D cost convexity parameter. Among the exogenous parameters, λ is the only one to which the quantitative results are sensitive—θ and τ_K do not affect any first-stage variable (see below). We examine $\lambda \in \{3/2, 2, 5/2\}$, re-calibrating all other parameters to match the same 9 data moments for each λ.

All qualitative conclusions survive: both directed subsidies raise the BGP growth rate, with KATC having a larger effect; KATC subsidies raise the labor share while LATC subsidies lower it; and a uniform subsidy tilts the technology direction toward KATC. Saddle-path stability is maintained for all λ (Table 5), and the sufficient conditions hold with ample margin throughout. The only dimension in which λ matters is the absolute magnitude of policy effects: a lower λ makes the R&D cost function less convex, amplifying firms' innovation response to subsidies. Importantly,

however, the marginal cost-effectiveness ratio of KATC over LATC is highly stable at approximately 3 (ranging from 2.88 to 2.97 across λ). The full set of recalibrated parameters is reported in Table C.3 (Appendix C.3).

Robustness to θ and τ_K. By the hierarchical structure of the BGP (Section 4.1), θ and τ_K enter only through the Euler equation (30) and do not affect any first-stage variable. When either parameter changes, only β adjusts to match the interest rate target, so the growth rate, factor shares, cost-effectiveness ratios, and all comparative statics results are *exactly the same* for $\theta \in \{1/2, 1, 3/2\}$ and for $\tau_K = 23.8\%$.[30] The only policy-relevant difference at $\tau_K = 23.8\%$ is the after-tax return, which falls from 3.40% to 3.05% at the baseline, and the tax rate in S5, which rises from 30.4% to 40.8%.

Robustness to the elasticity of substitution. We now recalibrate the model to $\text{EoS}_{KL} = 0.8$, near the upper end of the empirical range discussed in Section 5.1, keeping all other target moments unchanged. Because EoS_{KL} enters only at the final stage of the calibration—when the CES parameters (ε, η, B) are determined from the production function (Appendix C.1)—all other calibrated quantities are completely unaffected, and all *growth* effects of R&D subsidies are essentially invariant to EoS_{KL}. Quantitatively, a 1%-of-GDP KATC subsidy raises per-capita growth by +0.79 percentage points at $\text{EoS}_{KL} = 0.8$, compared with +0.81 pp at the baseline, and the marginal cost-effectiveness ratio of KATC over LATC subsidies remains at 2.92. The *distributional* effects, by contrast, depend on EoS_{KL} because factor income shares are governed by the curvature of the CES production function; we discuss the quantitative implications in Section 6.

6 Policy Experiments

The calibrated model allows us to quantify the effects of alternative R&D policies. We consider five policy scenarios—KATC subsidy, LATC subsidy, uniform subsidy,

[30] Under policy changes, θ and τ_K do affect the levels of k^* and c^* (and hence K/Y) through the savings response, but these level differences do not affect the growth rate, the direction of technological change, or the ranking of policies.

Table 6: Five policy scenarios (BGP fiscal shares, % of GDP)

#	Scenario	σ^*_{SK}	σ^*_{SL}	σ^*_S	σ^*_τ
—	Baseline (no intervention)	0	0	0	1.85
S1	KATC subsidy	1.00	0	1.00	(1.28)
S2	LATC subsidy	0	1.00	1.00	(2.11)
S3	Uniform subsidy	(0.15)	(0.85)	1.00	(1.95)
S4	Capital tax reduction	0	0	0	0.85
S5	KATC subsidy + tax increase	1.00	0	1.00	2.85

Values in parentheses are equilibrium outcomes, not scenario-defining targets. σ^*_{SK}: KATC subsidy expenditure / GDP. σ^*_{SL}: LATC subsidy expenditure / GDP. $\sigma^*_S = \sigma^*_{SK} + \sigma^*_{SL}$: total subsidy expenditure / GDP. σ^*_τ: capital tax revenue / GDP. S4 foregoes 1% of GDP in tax revenue relative to baseline (σ^*_τ: 1.85% $\to$ 0.85%). S5 finances S1's subsidy by raising τ_K from 15% to 30.4%.

capital tax reduction, and budget-neutral KATC subsidy—each calibrated to the same fiscal cost of 1% of GDP. We first compare their long-run balanced growth paths and then examine the transitional dynamics from one BGP to another.

6.1 Five Scenarios

Scenario definitions. We consider five policy scenarios, summarized in Table 6. All scenarios impose the same long-run fiscal burden: the total subsidy expenditure equals 1% of GDP on the new BGP. The capital tax reduction (S4) is calibrated so that the revenue foregone also equals 1% of GDP, and the budget-neutral scenario (S5) finances the KATC subsidy by raising the capital income tax rate.

The budget composition in Table 6 reflects the large baseline size difference between the two R&D sectors: K-augmenting R&D accounts for only about 16% of total R&D spending, while L-augmenting accounts for 84%.

Long-run outcomes. Table 7 summarizes the key BGP outcomes. The KATC subsidy (S1) is the most effective growth policy, raising per-capita growth by 0.81 percentage points—roughly twice the effect of the LATC subsidy (S2, +0.41 pp) and 50% more than the uniform subsidy (S3, +0.53 pp). While Proposition 6 predicts a marginal cost-effectiveness ratio of approximately 3, the discrete advantage for a 1%-of-GDP budget is approximately 2. This gap arises because the KATC sector accounts for only 16% of baseline R&D spending, so the subsidy rate s_K must reach

Table 7: Balanced growth path comparison across policy scenarios

Variable	Baseline	S1 KATC	S2 LATC	S3 Uniform	S4 Tax cut	S5 KATC+tax
s_K (%)	0	74.4	0	22.8	0	74.4
s_L (%)	0	0	27.6	22.8	0	0
τ_K (%)	15.0	15.0	15.0	15.0	7.2	30.4
g^a_{pc} (%)	1.86	2.67	2.27	2.39	1.86	2.67
Δg^a_{pc} (pp)	—	+0.81	+0.41	+0.53	0	+0.81
σ^*_L (%)	62.5	74.3	59.5	62.8	62.5	74.3
$\widetilde{\sigma}^*_K$ (%)	10.5	7.3	12.0	11.0	10.9	6.5
σ^*_R (%)	2.70	4.15	3.15	3.39	2.70	4.15
μ^*	2.24	5.10	1.95	2.35	2.24	5.10
K/Y	2.90	1.60	2.93	2.62	3.02	1.44
$(1-\tau_K)(r-\delta)^a$ (%)	3.40	4.20	3.81	3.93	3.42	4.14

s_K, s_L: R&D subsidy rates. τ_K: capital income tax rate. g^a_{pc}: annualized per-capita growth rate. $\widetilde{\sigma}^*_K \equiv (1-\tau_K)(\sigma^*_K - \delta K/Y)$: after-tax, after-depreciation capital income / GDP. $(1-\tau_K)(r-\delta)^a$: after-tax net return on capital. σ^*_R: private (post-subsidy) R&D expenditure / GDP. K/Y: annualized capital–output ratio.

74% to allocate 1% of GDP; at such a high rate, most of the subsidy expenditure goes toward R&D that firms would have undertaken even without the subsidy, and only the additional R&D induced at the margin generates extra growth. Nevertheless, a twofold growth advantage for the same fiscal cost represents a substantial policy gain. The capital tax reduction (S4) has zero long-run growth effect, while the budget-neutral KATC subsidy (S5) achieves the same growth as S1, confirming that capital taxation is growth-neutral (Section 4.1). On the distributional side, the KATC subsidy raises the labor share by nearly 12 pp while the LATC subsidy lowers it by 3 pp (Proposition 4), and even the uniform subsidy is not distributionally neutral (Proposition 5). Appendix C.4 confirms that these rankings are robust to the choice of the R&D cost convexity parameter λ. At $\text{EoS}_{KL} = 0.8$, the growth effects and policy rankings are virtually unchanged, but the distributional effects are substantially attenuated: the labor share response to KATC and LATC subsidies is roughly halved, and the small positive labor share effect of the uniform subsidy reverses sign (Appendix C.4). We explore the mechanisms behind these results through the transitional dynamics below.

6.2 Transitional Dynamics

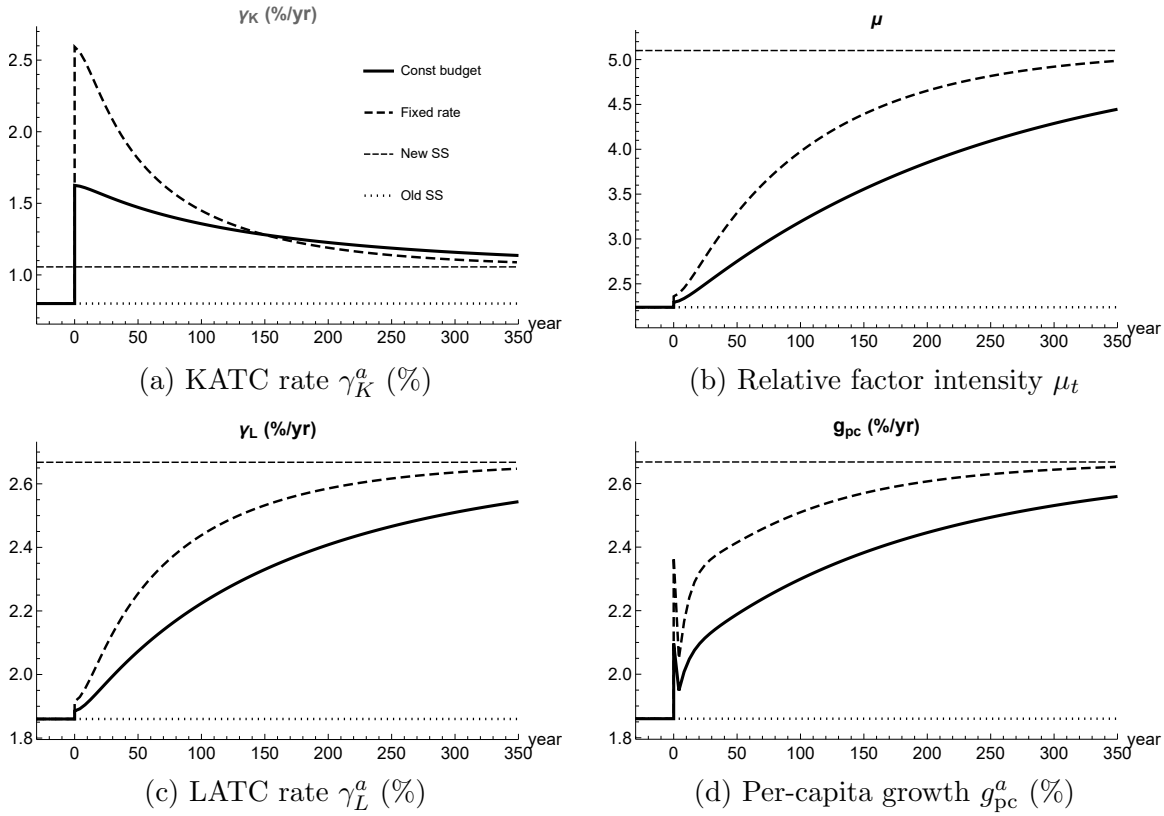
6.2.1 Scenario 1: KATC Subsidy

We begin with the *fixed-rate* design in which the KATC subsidy rate is set permanently to $s_K = 0.744$ from period 0 onward. Figure 5 reveals the layered multiplier mechanism through which the subsidy raises long-run growth. The fixed-rate transition path (dashed curves) displays two distinct effects, both governed by the slow adjustment of μ_t (half-life ≈ 93 years; Section 5.3).

Direct effect: the KATC rate spike and its gradual decay. The KATC subsidy triggers an immediate burst of capital-augmenting innovation: γ_K^a spikes from 0.80% to approximately 2.6% at impact (panel a). As μ_t rises and $f'(\mu)$ falls, the profitability of KATC R&D declines, and γ_K^a gradually decays toward its new BGP value of 1.06%. The speed of this decay is set by μ's adjustment (the slow eigenvalue), not by the capital accumulation process. In the μ-dynamics (A.1), the term $(K_{t+1}/K_t)^\alpha$ exceeds its BGP value whenever k_t is tracking a moving target $k^*(\mu_t)$, amplifying the rise of μ_t (panel b) beyond what the R&D incentive alone would produce. This capital-accumulation reinforcement persists throughout the transition: as long as μ_t is rising and $k^*(\mu_t)$ is shifting, the feedback from capital deepening continues to push μ upward.

Induced effect: the rise of labor-augmenting innovation. As μ_t rises, the marginal product of N-tasks increases, making labor-augmenting R&D more profitable. Firms respond by increasing LATC investment even without any change in s_L: γ_L^a gradually rises from 1.86% toward 2.67% (panel c). This induced response is slower than the direct effect because γ_L depends on $f(\mu) - \mu f'(\mu)$, which requires μ to accumulate a substantial shift before the LATC profitability change becomes appreciable. Since per-capita growth equals γ_L^{*a} on the BGP, the induced LATC acceleration is the engine of the permanent growth gain.

The growth path. Panel (d) makes the direct–induced sequence visible. At impact, per-capita growth jumps to about 2.35%, driven by the γ_K spike. It then dips to about 2.05% as the initial γ_K surge subsides while γ_L has not yet risen appreciably. From this trough, growth recovers monotonically over roughly 200 years toward the new BGP of 2.67%, propelled by the slow rise of γ_L. The dip-then-recovery is the signature



(a) KATC rate γ_K^a (%)

(b) Relative factor intensity μ_t

(c) LATC rate γ_L^a (%)

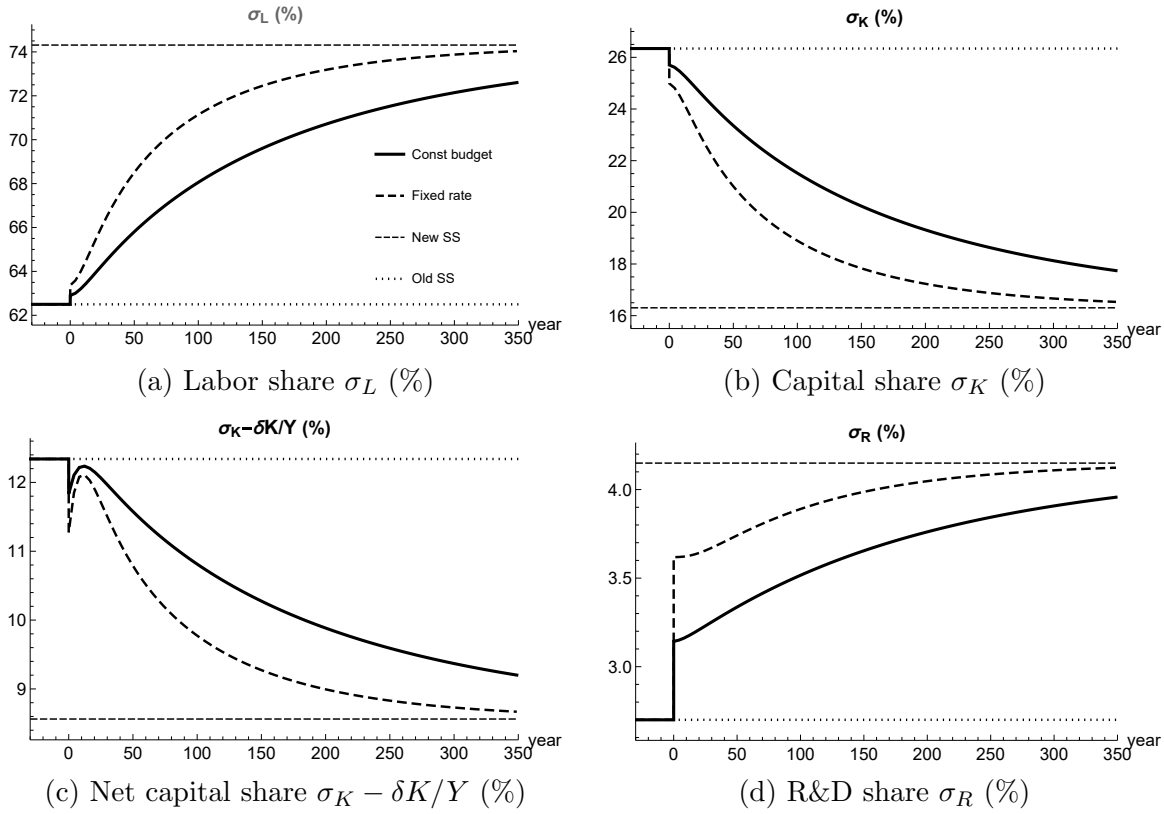
(d) Per-capita growth g_{pc}^a (%)

Figure 5: Scenario 1: KATC subsidy—mechanism. (a) γ_K^a spikes from 0.80% to ∼2.6% at impact, then declines to the new BGP of 1.06%. (b) μ_t rises from 2.24 toward 5.10. (c) γ_L^a rises from 1.86% toward 2.67%, induced by higher μ_t. (d) Per-capita growth jumps, dips, then recovers toward 2.67%.

of the induced effect taking over from the direct effect.[31]

The KATC subsidy produces a dramatic redistribution: the labor income share rises from 62.5% to 74.3%, a nearly 12 percentage point increase (Figure 6, panel a). This is the distributional paradox established in Proposition 4: under gross complementarity, subsidizing capital-augmenting technology raises the labor share by making effective labor relatively scarcer and hence more valuable. The capital share falls from 26.3% to 16.3% (panel b). The large magnitude reflects the mutually reinforcing nature of the two transition channels: the direct KATC stimulus and the capital-accumulation feedback both push μ upward, compounding the distributional

[31]The initial spike in γ_K^a (panel a) reflects the profit incentive to exploit the new subsidy immediately. As μ_t rises and $f'(\mu)$ falls, the KATC rate gradually declines toward its new BGP value $\gamma_K^{*a} = 1.06\%$, which is still above the baseline of 0.80%.



(a) Labor share σ_L (%)

(b) Capital share σ_K (%)

(c) Net capital share $\sigma_K - \delta K/Y$ (%)

(d) R&D share σ_R (%)

Figure 6: Scenario 1: KATC subsidy—distributional effects. (a) σ_L rises from 62.5% to 74.3%. (b) σ_K falls from 26.3% to 16.3%. (c) Net capital share declines more moderately because lower K/Y reduces depreciation. (d) σ_R rises from 2.7% to 4.15%, primarily from induced LATC R&D.

shift over the slow convergence of μ_t.[32]

The net capital income share $\sigma_K - \delta K/Y$ (panel c) declines less steeply than the gross capital share because the fall in K/Y from 2.90 to 1.60 reduces the depreciation burden ($\delta K/Y$ falls from 14.0% to 7.7%). The private R&D expenditure share (panel d) rises from 2.7% to 4.15%, with the increase primarily driven by the induced expansion of LATC R&D (ψ_L^* nearly doubles from 0.093 to 0.197).

Fiscal cost and the constant-budget alternative. Even under the fixed-rate design, the transition is slow: the full growth and distributional effects take roughly

[32]The KATC subsidy achieves both faster growth and a more equal factor income distribution, avoiding the usual growth–equity trade-off. This property is specific to the case $\varepsilon < 1$ and would be reversed if M-tasks and N-tasks were gross substitutes (some macro-level estimates using long-run trends yield $\sigma > 1$; see, e.g., Karabarbounis and Neiman, 2014).

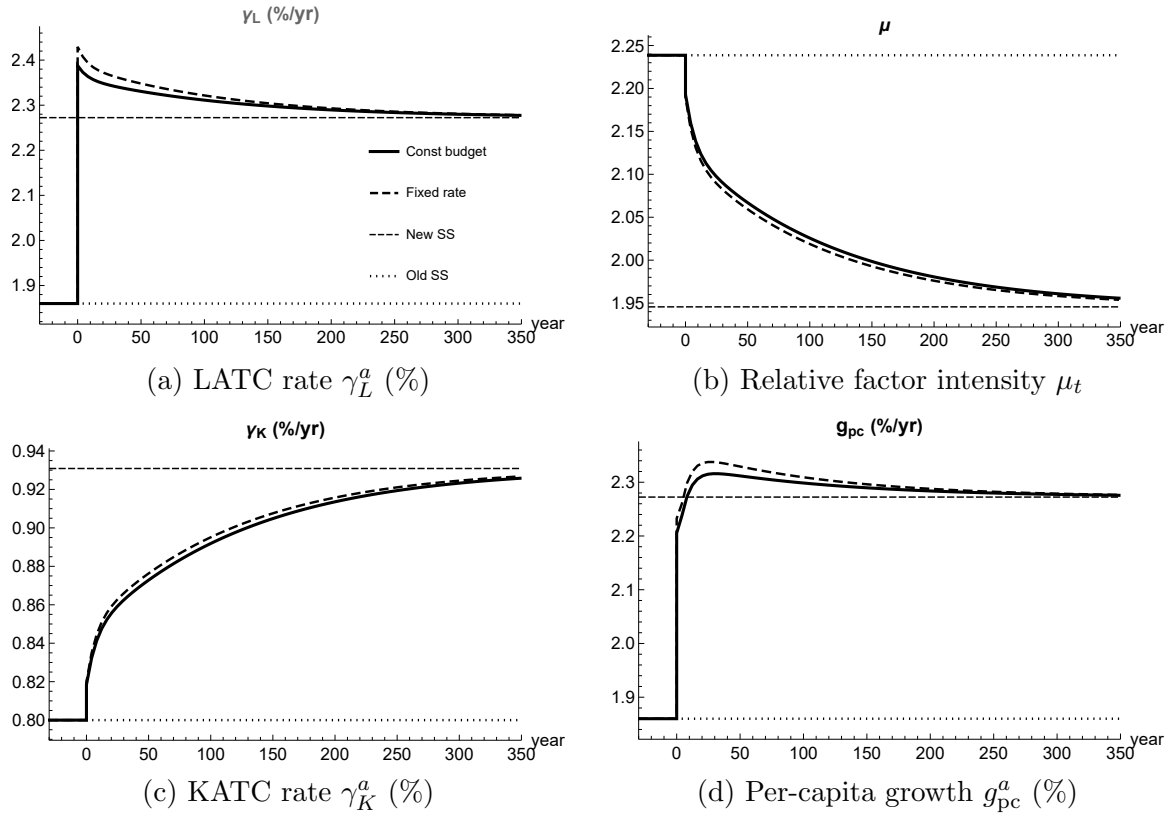
200 years to materialize, governed by the $\approx$ 93-year half-life of μ_t (Section 5.3). A further complication is fiscal. On the new BGP, μ^* is high and KATC R&D is correspondingly modest ($\gamma_K^{*a} = 1.06\%$), so a subsidy rate of 74% costs only 1% of GDP. At time 0, however, μ_0 is still low and the high subsidy rate elicits a much larger KATC response ($\gamma_K^a \approx 2.6\%$), pushing fiscal expenditure to approximately three times its long-run level before it gradually declines as μ_t rises. Such a sharp spike in government expenditure may not be tolerated fiscally or politically.

We therefore also consider a *constant-budget* design in which the subsidy rate adjusts each period to hold fiscal expenditure at exactly 1% of GDP (solid curves in all figures; see Appendix D.2 for the algorithm). Under this design, s_K starts well below its long-run value and rises only as μ_t increases. The initial KATC stimulus is therefore much weaker than under the fixed-rate path. On top of the inherently slow adjustment of μ_t, the budget constraint delays the onset of the R&D stimulus itself, stretching the transition over several centuries (Figure 5, panel d). Measured by the lifetime equivalent variation (LEV; Appendix D.3), the constant-budget path delivers an LEV gain of 70.6%—a substantial return on a policy that costs only 1% of GDP at all times. Even though both paths converge to the same BGP, the fixed-rate LEV is much higher (120.6%). This large difference reflects the faster realization of R&D benefits when the budget constraint is more flexible.

6.2.2 Scenario 2: LATC Subsidy

The LATC subsidy ($s_L = 0.276$) directly stimulates labor-augmenting innovation: γ_L^a rises immediately, overshooting the new BGP temporarily (Figure 7, panel a). Meanwhile, μ_t falls from 2.24 to 1.95 (panel b)—the economy shifts toward more labor-intensive techniques, the opposite direction from S1.

The transition is the mirror image of S1. The *direct effect* is the γ_L spike (panel a): as μ_t falls and $f(\mu) - \mu f'(\mu)$ declines, the profitability of LATC R&D gradually falls, and γ_L^a decays toward its new BGP value. The *induced effect* raises γ_K^a from 0.80% to 0.93% (panel c) as the lower μ_t increases $f'(\mu)$ and makes KATC R&D more profitable. A notable asymmetry with S1 arises through the capital-accumulation feedback. Higher effective labor raises output, saving, and capital accumulation, so $(K_{t+1}/K_t)^\alpha$ exceeds its BGP value throughout the transition, exerting *upward*



(a) LATC rate γ_L^a (%)

(b) Relative factor intensity μ_t

(c) KATC rate γ_K^a (%)

(d) Per-capita growth g_{pc}^a (%)

Figure 7: Scenario 2: LATC subsidy. (a) γ_L^a rises immediately, overshooting the new BGP temporarily. (b) μ_t falls from 2.24 to 1.95. (c) γ_K^a is induced upward from 0.80% to 0.93%. (d) Growth rises to 2.27%.

pressure on μ—the opposite direction from the direct effect. This persistent buffering explains why μ falls by only 13% (from 2.24 to 1.95) and why σ_L declines by only 3 percentage points—far less than the 12 percentage point shift under S1, where the same feedback *reinforces* rather than buffers the direct effect.

The LATC subsidy reduces the labor share from 62.5% to 59.5%. This is the mirror image of S1's "distribution paradox": subsidizing *labor*-augmenting technology hurts the labor share under gross complementarity. The LATC subsidy delivers less growth (+0.41 pp vs +0.81 pp) while worsening income distribution—an unambiguously inferior outcome on both dimensions compared to S1.

The roughly 2× growth advantage of S1 over S2 (0.81 pp vs 0.41 pp per 1% GDP) arises from two reinforcing mechanisms. The first is a "narrow and thick" allocation effect. K-augmenting R&D accounts for only 16% of baseline R&D spending, so a 1%

GDP budget concentrated on the small KATC sector produces a much larger shift in γ_K and μ^* than the same budget spread across the LATC sector.

The second is the capital accumulation multiplier (Section 3.3), which amplifies the K-side response by a factor of $\alpha/(1-\alpha) \approx 3.1$. Moreover, the capital-accumulation feedback *reinforces* the direct effect under S1 (both push μ upward) but *buffers* it under S2 (the two forces partially cancel), contributing to the large distributional asymmetry ($|\Delta\sigma_L|$ of 12 pp under S1 vs 3 pp under S2).

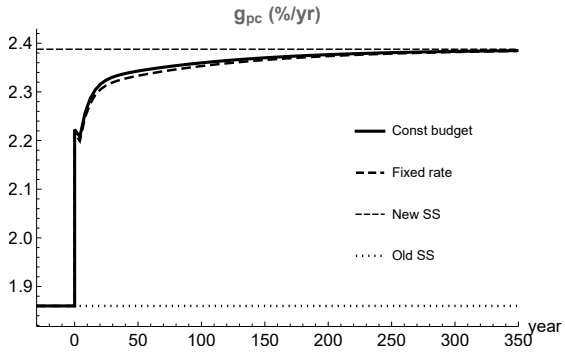
Unlike S1, the LATC subsidy does not produce fiscal front-loading: the moderate long-run rate ($s_L^* \approx 28\%$) does not elicit an outsized R&D response at the initial μ_0, and the fiscal cost stays near 1% of GDP throughout the transition. The convergence speed is accordingly similar under both designs, and the constant-budget and fixed-rate paths are nearly indistinguishable. The LEV is 66.0% (constant-budget) versus 68.7% (fixed-rate)—a substantial LEV gain, but more modest than S1's 70.6% under the same constant-budget constraint.

6.2.3 Scenario 3: Uniform Subsidy

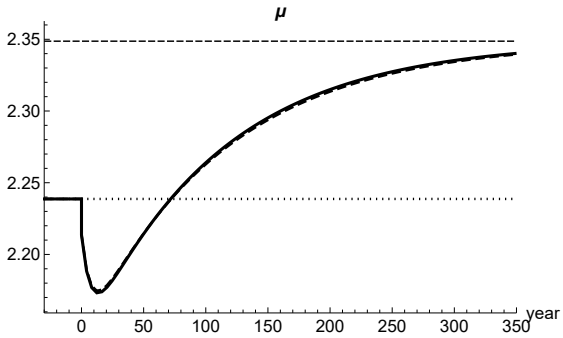
The uniform subsidy ($s_K = s_L = 0.228$) appears to be a direction-neutral policy, but it is not. The new $\mu^* = 2.35$ exceeds the baseline $\mu^* = 2.24$ (Figure 8, panel b), as predicted by Proposition 5: the capital accumulation multiplier $\alpha/(1-\alpha) \approx 3.1$ amplifies the K-side response disproportionately.

The most striking feature of S3 is the non-monotonic path of the labor share (panel c). While the long-run change is small (+0.3 pp, from 62.5% to 62.8%), the transition path reveals a meaningful temporary decline of approximately 1 percentage point, lasting roughly 30 years, before reversing and converging to the new BGP.

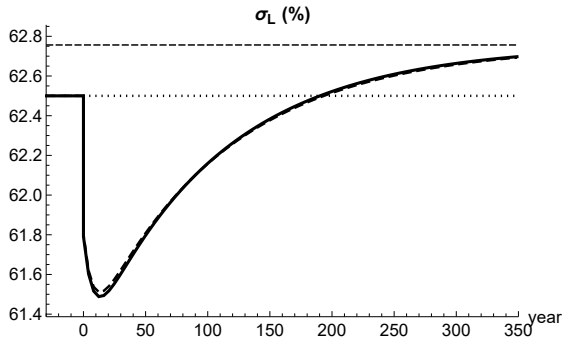
This dip-then-recovery pattern arises because the two components of the subsidy affect μ at different speeds. In the short-to-medium term (0–40 years), the LATC component dominates: it accounts for ∼85% of the subsidy budget and its direct effect on γ_L is immediate and large, pushing μ_t downward and—under $\varepsilon < 1$—lowering the labor share. In the long term (40+ years), the KATC component's multiplier effect catches up: the $\alpha/(1-\alpha) \approx 3.1$ amplification in (25) means that even a modest initial boost to γ_K eventually generates a large upward shift in μ^*. As μ_t reverses and rises



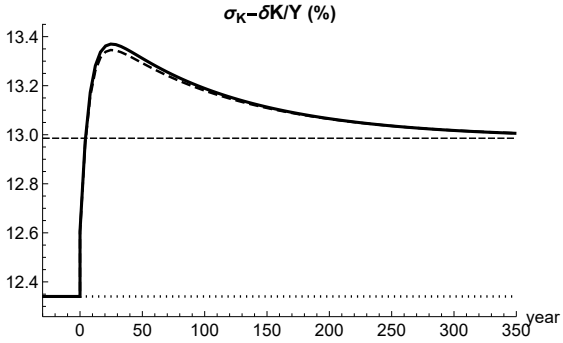
(a) Per-capita growth g^a_{pc} (%)



(b) Relative factor intensity μ_t



(c) Labor share σ_L (%)



(d) Net capital share $\sigma_K - \delta K/Y$ (%)

Figure 8: Scenario 3: Uniform subsidy. (a) Growth rises gradually to 2.39%. (b) μ_t initially dips, then reverses and converges from below to $\mu^* = 2.35$. (c) σ_L exhibits non-monotonic dynamics: a ∼1 pp dip for ∼30 years before rising to 62.8%. (d) Net capital share also shows non-monotonic behavior.

above baseline toward $\mu^* = 2.35$, the labor share recovers.[33]

The constant-budget and fixed-rate paths are again nearly identical, as the moderate subsidy rate ($s \approx 23\%$) keeps the fiscal cost close to 1% of GDP throughout the transition. The LEV is 79.4% (constant-budget) versus 77.8% (fixed-rate).

6.2.4 Scenario 4: Capital Tax Reduction

Scenario 4 provides the most visually striking confirmation of the model's theoretical structure. The capital tax rate is nearly halved (from 15.0% to 7.2%), at a fiscal cost

[33]This non-monotonic labor share path under uniform subsidies has not, to our knowledge, been discussed in the prior literature. It illustrates the slow adjustment of μ_t (Section 5.3) and suggests that policymakers who adopt uniform R&D subsidies expecting distributional neutrality may face unexpected political opposition during the first several decades of the transition.

of 1% of GDP. The result: *zero long-run growth effect* (Figure 9). This is consistent with the empirical observation by Jones (2016) that long-run growth rates have been remarkably stable across countries despite large variation in capital tax policies. The structural explanation is that τ_K does not enter the first-stage R&D conditions that determine μ^*, g^*, and the factor income shares (Section 4.1); only k^* and c^* shift, leaving the long-run growth rate unchanged regardless of the tax level.

In panel (a), the old and new steady-state lines overlap—the per-capita growth rate returns to 1.86% after a tiny transitory spike. The only permanent effect is on the capital–output ratio, which rises modestly from 2.90 to 3.02 (panel b). The labor share and factor income distribution are virtually unchanged in the long run (panels c, d).

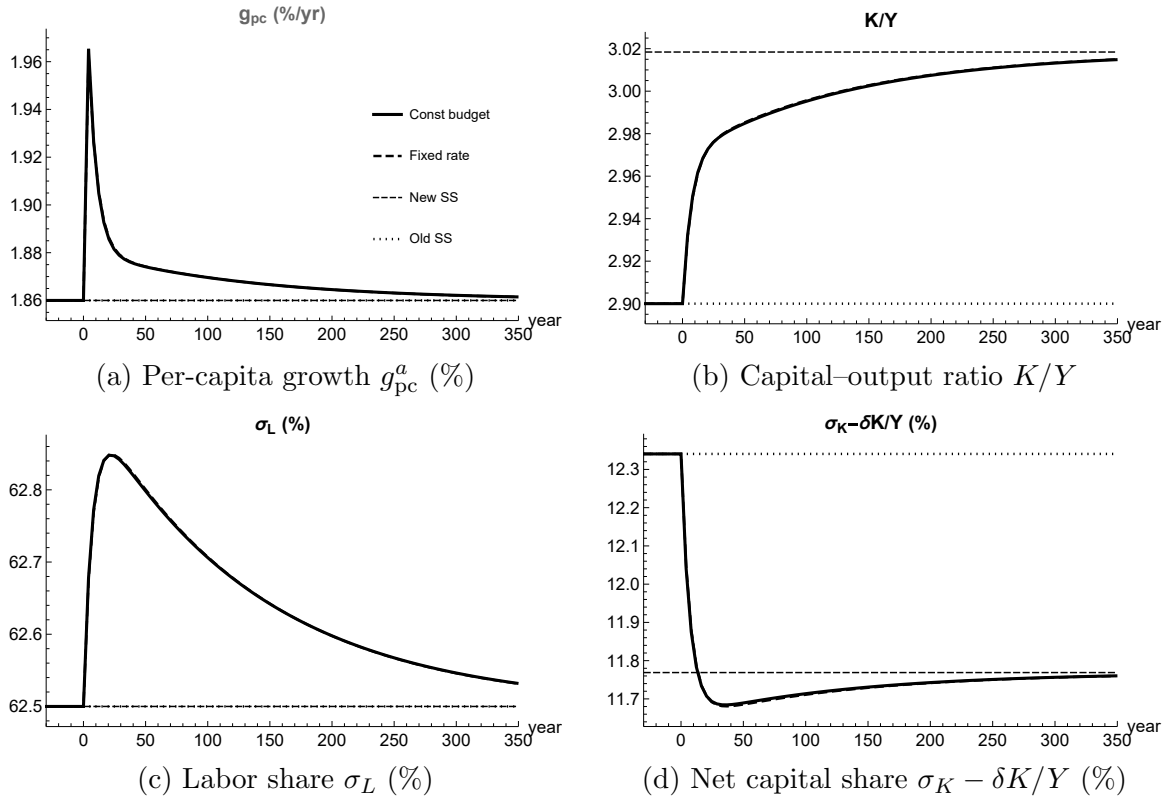
S4 thus isolates the fast eigenvalue mode identified in Section 5.3: the only adjustment is the Ramsey-type savings response that moves k_t toward the new k^*, with the bulk of the adjustment occurring within ∼5 years (panel b). The technology direction μ_t is temporarily perturbed but returns to μ^* once capital accumulation stabilizes.

A perhaps surprising finding is that the capital tax cut delivers remarkably little benefit even to capital owners: capital deepening raises the depreciation burden and reduces the gross return, almost entirely offsetting the lower tax rate. The after-tax net return rises by only 0.03 percentage points despite the near-halving of τ_K, and even a uniform R&D subsidy (S3) delivers both higher growth and a higher net capital income share than the tax cut (Table 7).

The constant-budget and fixed-rate paths coincide exactly, because the capital tax rate required to forgo 1% of GDP in revenue is nearly independent of the technology state μ_t. The LEV is almost zero (0.9%) under both designs, suggesting that the capital income tax cut does not benefit the representative household even in the transition.

6.2.5 Scenario 5: Budget-Neutral KATC Subsidy

Scenario 5 combines the KATC subsidy of S1 with a capital tax increase (τ_K: 15% → 30.4%) that finances it. The result (Figure 10) confirms the fiscal equivalence property: all growth-related BGP values—μ^*, g^a_{pc}, γ^{*a}_K, γ^{*a}_L, σ^*_L—are exactly identical to S1. The transition paths for these variables are also nearly identical (panels a,



(a) Per-capita growth g^a_{pc} (%)

(b) Capital–output ratio K/Y

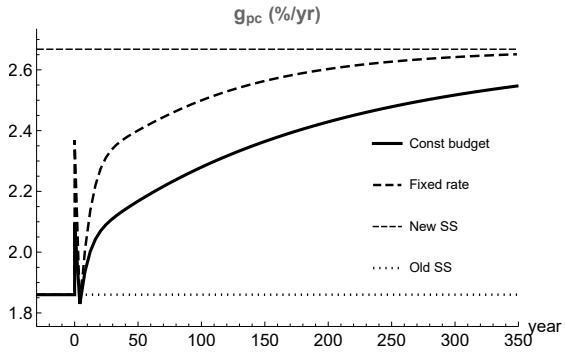
(c) Labor share σ_L (%)

(d) Net capital share $\sigma_K - \delta K/Y$ (%)

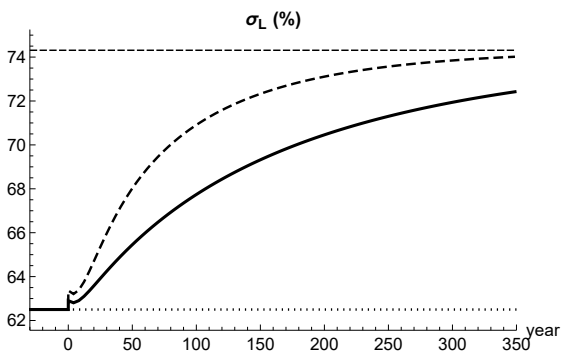
Figure 9: Scenario 4: Capital tax reduction. (a) Growth shows a tiny transitory spike, then returns to 1.86%—old and new BGP lines overlap. (b) The only permanent effect is K/Y rising from 2.90 to 3.02. (c) Labor share is virtually unchanged. (d) Net capital share increases only marginally despite the near-halving of the tax rate.

b), because the S4 isolation experiment shows that the capital tax channel shifts μ by only $\sim$0.05 (about 2% of the baseline $\mu^* = 2.24$), negligible compared to the R&D-driven shift of $\Delta\mu^* = 2.86$ under S1.

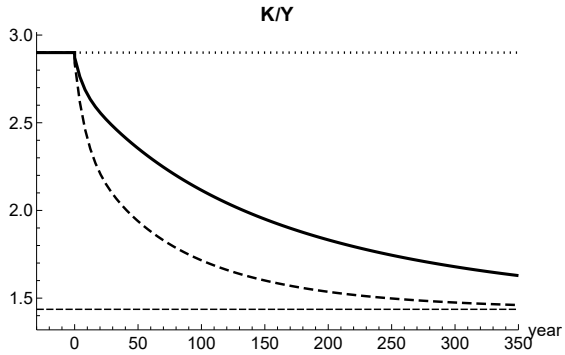
The principal difference from S1 lies on the capital side. The after-tax return $(1-\tau_K)(r-\delta)^a$ drops sharply at impact—from the baseline 3.40% to approximately 2.7% (panel d). The after-tax return then recovers through two channels. First, the lower after-tax return reduces saving, and K/Y declines from 2.90 toward 1.44 (panel c); the resulting capital scarcity raises the gross return r. Second, as μ_t rises and A_K improves, each unit of capital becomes more productive, further raising r. The first channel, driven by the fast mode, accounts for the bulk of the early recovery; the second, governed by the slow movement of μ_t, continues long after K/Y has largely stabilized. The after-tax return recovers past the baseline after roughly 20–25 years.



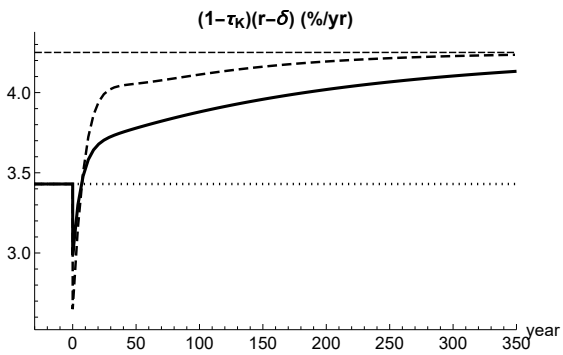
(a) Per-capita growth g^a_{pc} (%)



(b) Labor share σ_L (%)



(c) Capital–output ratio K/Y



(d) After-tax return $(1-\tau_K)(r-\delta)^a$ (%)

Figure 10: Scenario 5: Budget-neutral KATC subsidy. (a) Growth path is nearly identical to S1, rising to 2.67%. (b) Labor share is nearly identical to S1, rising to 74.3%. (c) K/Y falls from 2.90 to 1.44, driven by the fast mode. (d) After-tax return initially drops sharply, then recovers to 4.14%.

On the new BGP, it converges to 4.14%, only 6 basis points below S1's 4.20%.[34]

A notable feature of S5 is the "depreciation dividend." With K/Y falling to 1.44 (half the baseline), the depreciation share of GDP drops from 14.0% to 6.9%. This 7.1 percentage point reduction in depreciation costs exceeds the 1% GDP cost of the R&D subsidy by a factor of seven. The saved resources flow into consumption: the consumption–output ratio C/Y rises from 0.745 to 0.823, the highest of all scenarios.

Remarkably, financing the KATC subsidy through a capital tax increase rather than new spending delivers nearly the same LEV gain as S1: the constant-budget LEV is 67.6%, compared with 70.6% for S1, despite the fact that S5 requires no net

[34]The near-convergence of after-tax returns across S1 and S5 follows from the Euler equation (30): in steady state, the after-tax return is pinned by the growth rate and preference parameters, both of which are identical across S1 and S5. The tax increase is almost entirely offset by the rise in the gross return.

fiscal outlay while S1 costs 1% of GDP. The fixed-rate LEV is 117.4%, and the gap between the two designs is almost the same magnitude as in S1, reflecting the same budget-constrained delay of the initial KATC stimulus.

A common theme across the R&D subsidy scenarios is the slow convergence of the technology direction μ_t (half-life ≈ 90 years), which governs the pace at which growth rates and factor shares adjust. R&D subsidies must therefore be viewed as long-term commitments: the full growth benefits materialize only after several decades, or even centuries when the budget is less flexible, and premature withdrawal would strand the economy on a suboptimal transition path.

6.3 Incompletely Directed Subsidies

The preceding experiments assume that the government can perfectly direct subsidies toward either KATC or LATC R&D. In practice, classifying R&D activities as purely capital- or labor-augmenting is difficult. To assess how much targeting precision matters, we parameterize the degree of direction by $\rho \equiv s_K/(s_K + s_L) \in [0, 1]$, the fraction of total R&D subsidy directed toward KATC, and compute the BGP for each ρ while holding the total fiscal cost at 1% of GDP. At $\rho = 0$ only LATC R&D is subsidized (S2), at $\rho = 1/2$ both types receive equal subsidy rates (S3, the uniform benchmark), and at $\rho = 1$ only KATC R&D is subsidized (S1).

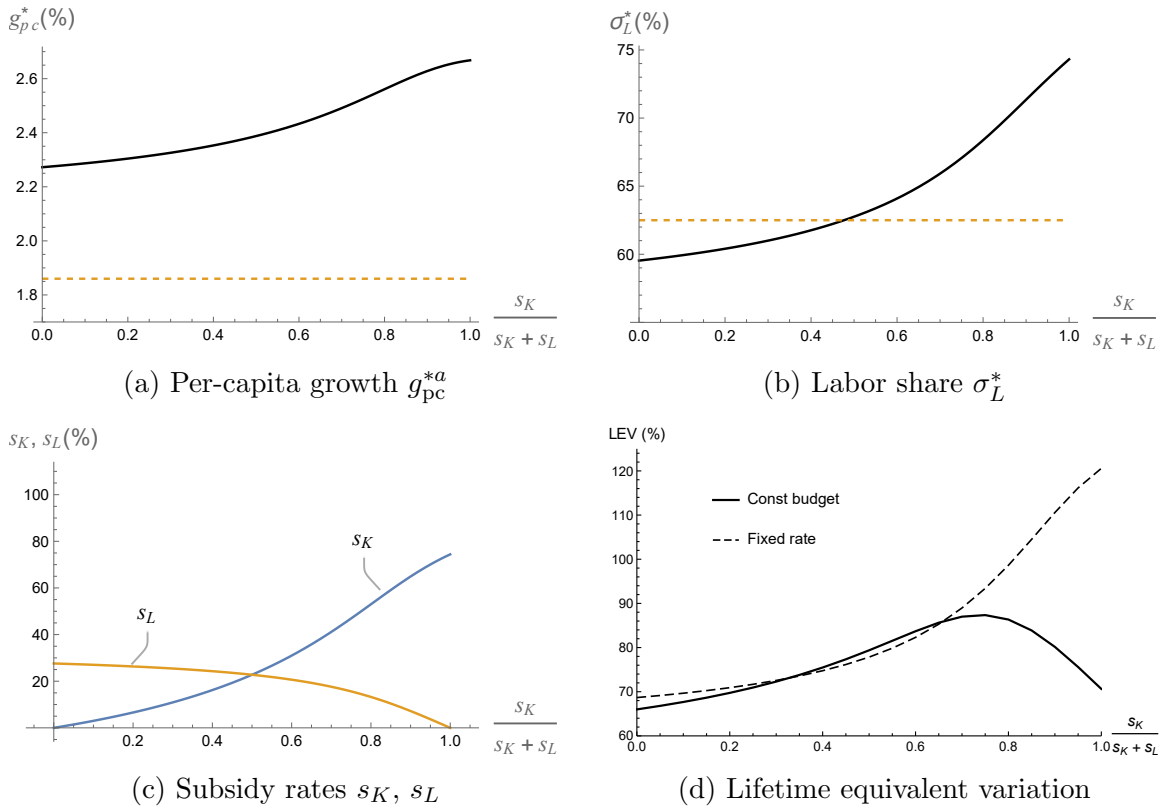
Figure 11 reports the results. Panels (a) and (b) show that both per-capita growth and the labor share increase monotonically in ρ, confirming numerically the comparative statics of Proposition 3: every incremental shift of the subsidy mix toward KATC raises both the BGP growth rate and the labor share. Panel (c) reveals the cost of full direction: because the KATC R&D base is small relative to the LATC base, concentrating the 1% GDP budget on KATC alone ($\rho = 1$) requires a subsidy rate of $s_K \approx 74\%$. Such a high rate creates two problems. First, under the fixed-rate design, the high subsidy rate elicits an outsized R&D response during the transition, pushing fiscal expenditure to roughly three times its long-run level (Section 6.2.1). Imposing a constant-budget constraint resolves the fiscal spike but forces the initial subsidy rate well below s_K^*, severely weakening the early R&D stimulus and stretching the transition over several centuries. Second, at such a high rate, most of the subsidy expenditure goes toward R&D that firms would have undertaken without the subsidy,

and the marginal cost-effectiveness of the KATC subsidy deteriorates (Proposition 6, condition (36)).

Welfare comparison across ρ. Panel (d) reports the lifetime equivalent variation (LEV) over the 350-year transition path for each ρ, comparing the fixed-rate (dashed) and constant-budget (solid) policy designs (Appendix D.3). Under the fixed-rate design, the LEV increases monotonically in ρ, reaching 120.6% at $\rho = 1$: without a budget constraint, concentrating subsidies on the small KATC sector always strengthens the initial R&D stimulus, and the fiscal "overshoot" early in the transition raises the representative household's lifetime consumption rather than penalizing it. Under the constant-budget design, however, the LEV is non-monotonic: it peaks at $\rho \approx 0.75$ (LEV $\approx 87\%$) and declines toward $\rho = 1$, where it falls to 70.6%.

The decline of the constant-budget LEV from $\rho = 0.75$ to $\rho = 1$ in panel (d) is a direct consequence of the two problems identified above. Reducing ρ below 1 lowers the required s_K^*, which simultaneously eases the budget constraint—allowing a stronger initial R&D stimulus and faster convergence—and restores fiscal efficiency by shifting expenditure toward the steeper part of the R&D cost curve. At $\rho \approx 0.75$ (a subsidy-rate ratio of $s_K : s_L \approx 3 : 1$), the LEV gain from faster convergence and improved cost-effectiveness just offsets the lower long-run growth rate, producing the LEV peak.

A separate consideration is that, in practice, many R&D projects simultaneously advance both capital- and labor-augmenting capabilities, making a clean classification into KATC or LATC inherently difficult. Perfectly directed subsidies ($\rho = 1$) are therefore not only suboptimal under a constant budget, but also difficult to implement. Fortunately, perfect direction is not required: under the constant-budget design, the LEV-maximizing allocation corresponds to an effective subsidy-rate ratio of $s_K : s_L \approx 3 : 1$ (LEV $\approx 87\%$), well above the uniform benchmark ($\rho = 1/2$, LEV $= 79\%$). Even a policy that only partially tilts R&D subsidies toward capital-augmenting technologies captures the bulk of the achievable gains.



(a) Per-capita growth g^{*a}_{pc}

(b) Labor share σ^*_L

(c) Subsidy rates s_K, s_L

(d) Lifetime equivalent variation

Figure 11: Incompletely directed subsidies. Horizontal axis: $\rho \equiv s_K/(s_K + s_L)$. (a)–(b) Long-run BGP growth rate and labor share; dashed lines mark the no-subsidy baseline. (c) Subsidy rates s_K and s_L. (d) LEV over the 350-year transition; solid = constant-budget (peaks at $\rho \approx 0.75$), dashed = fixed-rate (monotonically increasing).

7 Conclusion

This paper builds a directed technical change model with both capital- and labor-augmenting technological change along the balanced growth path, circumventing the Uzawa steady-state theorem through a three-factor production structure. The main message is that R&D subsidies tilted toward capital-augmenting technologies can simultaneously accelerate growth and raise the labor share of income. In our calibration, a small KATC subsidy is roughly three times more cost-effective than an LATC subsidy; even for a large program equal to 1% of GDP, it delivers roughly twice the growth gain and raises the labor share by nearly 12 percentage points. Capital income taxation, by contrast, has no long-run effect on growth or the direction of technological change in our framework, because R&D is paid out of pre-tax firm earnings and

its optimality conditions do not involve τ_K.

An important caveat is that the transition to the new balanced growth path is slow: in our simulations, it takes roughly 90 years for the technology ratio μ_t to close half the gap to its new long-run value, so realizing the full benefits of R&D policy requires sustained commitment spanning more than a century.

In practice, perfect targeting is neither required nor, under a realistic budget constraint, even desirable. Our analysis of the transitional dynamics shows that the lifetime equivalent variation (LEV) is maximized when the effective subsidy-rate ratio between KATC and LATC R&D is approximately 3 : 1, rather than at full KATC direction. The practical lesson is clear: a policy that tilts R&D subsidies toward capital-augmenting technologies—even when the two types cannot be perfectly separated—captures the bulk of the potential gains in both growth and the labor share.

References

Acemoglu, D. (2003): "Labor-and capital-augmenting technical change," *Journal of the European Economic Association*, 1, 1–37.

——— (2023): "Distorted innovation: does the market get the direction of technology right?" in *AEA Papers and Proceedings* Volume 113, 1–28, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Acemoglu, D., P. Aghion, L. Bursztyn, and D. Hemous (2012): "The environment and directed technical change," *American Economic Review*, 102, 131–166.

Acemoglu, D., U. Akcigit, H. Alp, N. Bloom, and W. Kerr (2018): "Innovation, reallocation, and growth," *American Economic Review*, 108, 3450–91.

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.

Akcigit, U., and W. R. Kerr (2018): "Growth through heterogeneous innovations," *Journal of Political Economy*, 126, 1374–1443.

André, F. J., and S. Smulders (2014): "Fueling growth when oil peaks: Directed technological change and the limits to efficiency," *European Economic Review*, 69, 18–39.

Antras, P. et al. (2004): "Is the US aggregate production function Cobb-Douglas? New estimates of the elasticity of substitution," *Contributions to Macroeconomics*, 4, 1–34.

Bloom, N., C. I. Jones, J. Van Reenen, and M. Webb (2017): "Are ideas getting harder to find?"Technical Report 23782.

Bloom, N., R. Griffith, and J. Van Reenen (2002): "Do R&D Tax Credits Work? Evidence from a Panel of Countries 1979–1997," *Journal of Public Economics*, 85, 1–31.

Blundell, R., R. Griffith, and F. Windmeijer (2002): "Individual Effects and Dynamics in Count Data Models," *Journal of Econometrics*, 108, 113–131.

Caselli, F., and J. Feyrer (2007): "The marginal product of capital," *Quarterly Journal of Economics*, 122, 535–568.

Casey, G. (2024): "Energy efficiency and directed technical change: implications for climate change mitigation," *Review of Economic Studies*, 91, 192–228.

Casey, G., and R. Horii (2024): "A Generalized Uzawa Growth Theorem," *Journal of Political Economy Macroeconomics*, 2, 336–373.

Greenwood, J., Z. Hercowitz, and P. Krusell (1997): "Long-run implications of investment-specific technological change," *American Economic Review*, 87, 342–362.

Griliches, Z. (1990): "Patent Statistics as Economic Indicators: A Survey," *Journal of Economic Literature*, 28, 1661–1707.

Grossman, G. M., E. Helpman, E. Oberfield, and T. Sampson (2017): "Balanced Growth Despite Uzawa," *American Economic Review*, 107, 1293–1312.

——— (2021): "Endogenous Education and Long-Run Factor Shares," *American Economic Review: Insights*, 3, 215–32.

Grossman, G. M., and E. Oberfield (2022): "The elusive explanation for the declining labor share," *Annual Review of Economics*, 14, 93–124.

Güner, I. (2023): "Growth and welfare implications of sector-specific innovations," *Review of Economic Dynamics*, 47, 204–245.

Hall, B. H. (1993): "R&D Tax Policy During the 1980s: Success or Failure?" in *Tax Policy and the Economy* Volume 7: MIT Press, 1–35.

Hall, B. H., and J. Lerner (2010): "The Financing of R&D and Innovation," in *Handbook of the Economics of Innovation* Volume 1: North-Holland, 609–639.

Hall, B. H., and R. H. Ziedonis (2001): "The Patent Paradox Revisited: An Empirical Study of Patenting in the U.S. Semiconductor Industry, 1979–1995," *RAND Journal of Economics*, 32, 101–128.

Hassler, J., P. Krusell, and C. Olovsson (2021): "Directed technical change as a response to natural resource scarcity," *Journal of Political Economy*, 129, 3039–3072.

Hémous, D., and M. Olsen (2021): "Directed technical change in labor and environmental economics," *Annual Review of Economics*, 13, 571–597.

Horii, R., and T. Iwaisako (2007): "Economic growth with imperfect protection of intellectual property rights," *Journal of Economics*, 90, 45–85.

Hubmer, J. (2023): "The race between preferences and technology," *Econometrica*, 91, 227–261.

Irmen, A., and A. Tabaković (2017): "Endogenous capital-and labor-augmenting technical change in the neoclassical growth model," *Journal of Economic Theory*, 170, 346–384.

Jones, B. F., and X. Liu (2024): "A framework for economic growth with capital-embodied technical change," *American Economic Review*, 114, 1448–1487.

Jones, C. I. (2002): "Sources of US economic growth in a world of ideas," *American Economic Review*, 92, 220–239.

——— (2005): "The shape of production functions and the direction of technical change," *Quarterly Journal of Economics*, 120, 517–549.

——— (2016): "The facts of economic growth," *Handbook of Macroeconomics*, 2, 3–69.

Jones, C. I., and P. M. Romer (2010): "The new Kaldor facts: Ideas, institutions, population, and human capital," *American Economic Journal: Macroeconomics*, 2, 224–245.

Jones, C. I., and D. Scrimgeour (2008): "A new proof of Uzawa's steady-state growth theorem," *Review of Economics and Statistics*, 90, 180–182.

Jones, C. I., and J. C. Williams (2000): "Too much of a good thing? The economics of investment in R&D," *Journal of economic growth*, 5, 65–85.

Karabarbounis, L., and B. Neiman (2014): "The global decline of the labor share," *Quarterly Journal of Economics*, 129, 61–103.

Knoblach, M., M. Roessler, and P. Zwerschke (2020): "The elasticity of substitution between capital and labour in the US economy: A meta-regression analysis," *Oxford Bulletin of Economics and Statistics*, 82, 62–82.

Krusell, P. (1998): "Investment-specific R &D and the decline in the relative price of capital," *Journal of Economic Growth*, 3, 131–141.

Leon-Ledesma, M. A., and M. Satchi (2019): "Appropriate Technology and Balanced Growth," *Review of Economic Studies*, 86, 807–835.

Oberfield, E., and D. Raval (2021): "Micro data and macro technology," *Econometrica*, 89, 703–732.

Romer, P. M. (1990): "Endogenous technological change," *Journal of political Economy*, 98, S71–S102.

Schlicht, E. (2006): "A Variant of Uzawa's Theorem," *Economics Bulletin*, 5, 1–5.

Uzawa, H. (1961): "Neutral inventions and the stability of growth equilibrium," *Review of Economic Studies*, 28, 117–124.

Valentinyi, A., and B. Herrendorf (2008): "Measuring factor income shares at the sectoral level," *Review of Economic Dynamics*, 11, 820–835.

Wilson, D. J. (2009): "Beggar Thy Neighbor? The In-State, Out-of-State, and Aggregate Effects of R&D Tax Credits," *Review of Economics and Statistics*, 91, 431–436.

A Derivations for the Model

A.1 First-Order Conditions

FOC for K_t. The firm maximizes (6) with respect to K_t. The first-order condition $\partial\pi_t/\partial K_t = 0$ gives $F_M \cdot \frac{dM_t}{dK_t} - (1 - s_K)\frac{dR_{K,t}}{dK_t} - r_t = 0$. Since $M_t = (A_{K,t}K_t)^\alpha X_t^{1-\alpha}$, we have $dM_t/dK_t = \alpha M_t/K_t = \alpha\mu_t/k_t$. The R&D cost contribution is $dR_{K,t}/dK_t = (dM_t/dK_t) \cdot \psi_K(\gamma_{K,t}) = \alpha(\mu_t/k_t)\psi_K(\gamma_{K,t})$. Substituting and using $F_M = f'(\mu_t)$: $r_t = \frac{\alpha\mu_t}{k_t}\big(f'(\mu_t) - (1 - s_K)\psi_K(\gamma_{K,t})\big)$. □

FOC for X_t. Similarly, $dM_t/dX_t = (1 - \alpha)M_t/X_t$, and $\xi_t = \frac{(1-\alpha)M_t}{X_t}\big(f'(\mu_t) - (1 - s_K)\psi_K(\gamma_{K,t})\big)$. □

FOC for L_t. Since $N_t = A_{L,t}L_t$, we have $dN_t/dL_t = A_{L,t}$ and $dR_{L,t}/dL_t = A_{L,t}\psi_L(\gamma_{L,t})$: $w_t = A_{L,t}\big(f(\mu_t) - \mu_t f'(\mu_t) - (1 - s_L)\psi_L(\gamma_{L,t})\big)$. □

FOC for $A_{K,t}$. From $\partial\pi_t/\partial A_{K,t} = 0$: $\frac{\partial F}{\partial A_{K,t}} = (1-s_K)\frac{dR_{K,t}}{dA_{K,t}}$. The LHS is $\partial F/\partial A_{K,t} = f'(\mu_t)\cdot\alpha M_t/A_{K,t}$. The RHS is $\frac{dR_{K,t}}{dA_{K,t}} = \frac{dM_t}{dA_{K,t}}\psi_K(\gamma_{K,t}) + M_t\cdot\frac{1}{A_{K,t-1}}\psi'_K(\gamma_{K,t}) = \frac{M_t}{A_{K,t}}\big[\alpha\,\psi_K(\gamma_{K,t}) + \gamma_{K,t}\,\psi'_K(\gamma_{K,t})\big]$. Canceling $M_t/A_{K,t}$: $\alpha f'(\mu_t) = (1 - s_K)\big[\alpha\,\psi_K(\gamma_{K,t}) + \gamma_{K,t}\,\psi'_K(\gamma_{K,t})\big] = (1 - s_K)\,\Psi_K(\gamma_{K,t})$. □

FOC for $A_{L,t}$. From $\partial\pi_t/\partial A_{L,t} = 0$: The LHS is $[f(\mu_t) - \mu_t f'(\mu_t)] \cdot L_t$. The RHS is $\frac{dR_{L,t}}{dA_{L,t}} = L_t\big[\psi_L(\gamma_{L,t}) + \gamma_{L,t}\,\psi'_L(\gamma_{L,t})\big]$. Canceling L_t: $f(\mu_t) - \mu_t f'(\mu_t) = (1 - s_L)\big[\psi_L(\gamma_{L,t}) + \gamma_{L,t}\,\psi'_L(\gamma_{L,t})\big] = (1 - s_L)\,\Psi_L(\gamma_{L,t})$. □

Zero profit. Substituting (11), (12), (13), (14), and (15) into (6), one can verify that $\pi_t = 0$ due to the CRS property.

A.2 Proof of Proposition 1

Properties of Ψ_K and Ψ_L. For $\Psi_K(\gamma_K) = \gamma_K\psi'_K(\gamma_K) + \alpha\psi_K(\gamma_K)$: $\Psi_K(1) = 1 \cdot \psi'_K(1) + \alpha\psi_K(1) = 0$ (from (4)); $\Psi_K(\infty) = \infty$ (since $\psi'_K > 0$ for $\gamma_K > 1$); and $\Psi'_K(\gamma_K) = (1 + \alpha)\psi'_K(\gamma_K) + \gamma_K\psi''_K(\gamma_K) > 0$ for $\gamma_K > 1$. Therefore Ψ_K is a strictly increasing bijection from $[1, \infty)$ to $[0, \infty)$, and Ψ_K^{-1} exists with $\Psi_K^{-1\prime} > 0$. The same argument applies to Ψ_L with α replaced by 1.

Existence and uniqueness. From (14): $\gamma_{K,t} = \Psi_K^{-1}\left(\frac{\alpha f'(\mu_t)}{1-s_K}\right) \equiv \widehat{\gamma}_K(\mu_t;\, s_K)$. Since $f'(\mu)$ is decreasing in μ (concavity and Inada conditions) and Ψ_K^{-1} is increasing, $\widehat{\gamma}_K(\mu; s_K)$ is decreasing in μ. Since $\alpha f'(\mu)/(1-s_K)$ is increasing in s_K, $\widehat{\gamma}_K(\mu; s_K)$ is increasing in s_K.

Similarly: $\gamma_{L,t} = \Psi_L^{-1}\left(\frac{f(\mu_t)-\mu_t f'(\mu_t)}{1-s_L}\right) \equiv \widehat{\gamma}_L(\mu_t;\, s_L)$. Since $f(\mu) - \mu f'(\mu)$ is increasing in μ, $\widehat{\gamma}_L(\mu; s_L)$ is increasing in μ and in s_L. □

EIPF: limiting behavior. When $\mu \to 0$: $f'(\mu) \to \infty$ (Inada) $\Rightarrow \widehat{\gamma}_K(\mu; s_K) \to \infty$; $f(\mu) - \mu f'(\mu) \to 0 \Rightarrow \widehat{\gamma}_L(\mu; s_L) \to 1$. When $\mu \to \infty$: $f'(\mu) \to 0 \Rightarrow \widehat{\gamma}_K(\mu; s_K) \to 1$; $f(\mu) - \mu f'(\mu) \to \infty \Rightarrow \widehat{\gamma}_L(\mu; s_L) \to \infty$.

A.3 Derivation of Equilibrium Dynamics

Resource constraint. Substituting the profit π_t and the government transfer B_t into the household budget (8): the profit tax terms cancel, $(1-\tau_K)\pi_t + \tau_K \pi_t = \pi_t$; capital income terms yield $[(1-\tau_K)(r_t - \delta) + 1]K_t - r_t K_t + \tau_K(r_t - \delta)K_t = (1-\delta)K_t$; resource income: $\xi_t X_t - \xi_t X_t = 0$; wage income: $(1-\tau_L)w_t L_t - w_t L_t + \tau_L w_t L_t = 0$; R&D costs: $-(1-s_K)R_K - s_K R_K = -R_K$ and similarly $-R_L$. Therefore: $K_{t+1} = F(M_t, N_t) - R_{K,t} - R_{L,t} + (1-\delta)K_t - C_t = V_t + (1-\delta)K_t - C_t$. □

Dynamics of μ_{t+1}. The growth factor of μ_{t+1} is $\mu_{t+1}/\mu_t = (M_{t+1}/M_t)/(N_{t+1}/N_t)$. By using (1), (2), and (19), its value in equilibrium can be written as

$$\frac{\mu_{t+1}}{\mu_t} = \frac{g_X^{1-\alpha}}{g_L\, \widehat{\gamma}_L(\mu_{t+1}; s_L)} \left(\widehat{\gamma}_K(\mu_{t+1}; s_K)\left(\frac{v(\mu_t;\, \mathbf{s}) - c_t}{k_t} + 1 - \delta\right)\right)^{\alpha}. \tag{A.1}$$

Both sides depend on μ_{t+1}, so we decompose the dynamics into two steps. Define the pre-R&D relative factor intensity:

$$\mu_{t+1}^{\text{pre}} \equiv \frac{\left(A_{K,t}K_{t+1}\right)^{\alpha} X_{t+1}^{1-\alpha}}{A_{L,t}L_{t+1}} = \frac{g_X^{1-\alpha}}{g_L}\left(\frac{v(\mu_t;\, \mathbf{s}) - c_t}{k_t} + 1 - \delta\right)^{\alpha} \mu_t. \tag{A.2}$$

The post-R&D value of μ_{t+1} is related to μ_{t+1}^{pre} by

$$\mu_{t+1}^{\text{pre}} = \frac{\widehat{\gamma}_L(\mu_{t+1}; s_L)}{\widehat{\gamma}_K(\mu_{t+1}; s_K)^{\alpha}} \mu_{t+1} \equiv \Gamma(\mu_{t+1};\, \mathbf{s}). \tag{A.3}$$

Proposition 1 implies that $\Gamma(\mu; \mathbf{s})$ is strictly increasing with $\lim_{\mu\to 0}\Gamma(\mu) = 0$ and $\lim_{\mu\to\infty}\Gamma(\mu) = \infty$. Therefore $\Gamma^{(-1)}$ is well-defined, and inverting gives equation (20). □

Euler equation. The Lagrangian FOCs yield: $\lambda_t = \beta^t (C_t/L_t)^{-\theta}$, $\lambda_t = \lambda_{t+1}\big[(1-\tau_K)(r_{t+1}-\delta)+1\big]$. Combining and using $C_t/L_t = A_{L,t}c_t$: $(A_{L,t}c_t)^{-\theta} = \beta\big[(1-\tau_K)(r_{t+1}-\delta)+1\big](A_{L,t+1}c_{t+1})^{-\theta}$. This gives equation (22). Rearranging with $A_{L,t+1}/A_{L,t} = \widehat{\gamma}_L(\mu_{t+1}; s_L)$ and substituting (11) for r_{t+1} gives equation (23). □

Initial conditions. Given endowments $\{K_0, X_0, L_0, A_{K,-1}, A_{L,-1}\}$, the pre-R&D relative factor intensity at $t = 0$ is: $\mu_0^{\text{pre}} = \frac{(A_{K,-1}K_0)^\alpha X_0^{1-\alpha}}{A_{L,-1}L_0}$. Then $\mu_0 = \Gamma^{(-1)}(\mu_0^{\text{pre}}; \mathbf{s})$, and $k_0 = K_0/(\widehat{\gamma}_L(\mu_0; s_L)\cdot A_{L,-1}\cdot L_0)$.

Transversality condition. From the Euler equation and the NPG condition, the transversality condition in terms of the normalized state variables is:

$$\lim_{T\to\infty}(\beta g_L)^T \left(\prod_{t=1}^{T}\widehat{\gamma}_L(\mu_t; s_L)\right)^{1-\theta} c_T^{-\theta}k_T = 0. \tag{A.4}$$

B BGP and Comparative Statics Proofs

We introduce auxiliary notation used throughout this appendix. For given BGP values $(\mu^*, \gamma_K^*, \gamma_L^*)$, define

$$\phi_K \equiv \frac{\alpha}{1-\alpha}\cdot\frac{g_X\,(\gamma_K^*)^{(2\alpha-1)/(1-\alpha)}}{(1-s_K)\,\Psi_K'(\gamma_K^*)}, \quad \phi_L \equiv \frac{g_L}{(1-s_L)\,\Psi_L'(\gamma_L^*)}, \quad \Phi \equiv \alpha\,\phi_K + \mu^*\,\phi_L. \tag{B.1}$$

These coefficients measure the sensitivity of the BGP growth factor to changes in the KATC and LATC rates, respectively. Since $\Psi_j' > 0$ for $\gamma_j > 1$ (Appendix A.2), ϕ_K, ϕ_L, and Φ are strictly positive.

B.1 Balanced Growth Path

Proof of Lemma 1. On a BGP, all growth factors are constant. From (21), $k_{t+1}/k_t = 1$ requires $K_{t+1}/K_t = \widehat{\gamma}_L(\mu_{t+1}; s_L)\,g_L$, which must be constant. Since

g_L is a parameter and $\widehat{\gamma}_L(\cdot\,; s_L)$ is strictly monotone (Proposition 1), constancy of $\widehat{\gamma}_L(\mu_{t+1}; s_L)$ requires μ_t to be constant. With $\mu_t = \mu^*$, equations (21) and (23) then imply that k_t and c_t are also constant. □

Derivation of the BGP condition. With $\mu_t = \mu^*$ constant, the stationarity of k_t gives $(K_{t+1}/K_t)^* = \widehat{\gamma}_L(\mu^*; s_L)\, g_L \equiv g^*$. Substituting this into the μ-dynamics stationarity condition $\mu_{t+1}/\mu_t = 1$ from (A.1) and simplifying yields (25).

Proof of Proposition 2: existence and uniqueness. Condition (27) is a sufficient condition for the existence of a BGP. It is stated in terms of the primitives $(f, \psi_j, \alpha, g_X, g_L)$ only and does not involve the subsidy rates (s_K, s_L). This is because $s_j > 0$ reduces the effective R&D cost $(1 - s_j)\Psi_j$ in each FOC, thereby raising $\widehat{\gamma}_j$ for every μ; hence any $s_K, s_L \in [0, 1)$ only relaxes the endpoint inequalities, and (27) at $s_K = s_L = 0$ is sufficient for all subsidy levels.

Define the two curves in (26). By Proposition 1, $g_M(\mu)$ is strictly decreasing in μ and $g_N(\mu)$ is strictly increasing.

Endpoints. As $\mu \to 0^+$: the KATC FOC gives $\widehat{\gamma}_K(0^+; s_K) = \Psi_K^{-1}(\alpha f'(0^+)/(1 - s_K))$. Since Ψ_K is strictly increasing with $\Psi_K(1) = 0$, the first inequality of (27) ensures that $\widehat{\gamma}_K(0^+; s_K) > \bar{\gamma}_K$ and hence $g_M(0^+) > g_L = g_N(0^+)$. As $\mu \to \infty$: since $f'(\infty) = 0$, $\widehat{\gamma}_K(\infty; s_K) = 1$ and $g_M(\infty) = g_X$. The second inequality of (27) ensures $g_N(\infty) > g_X = g_M(\infty)$.

By the intermediate value theorem and the strict monotonicity of both curves, a unique intersection μ^* exists, satisfying (25) with $g^* > \max\{g_X, g_L\}$. Since both $\widehat{\gamma}_K(\mu^*; s_K)$ and $\widehat{\gamma}_L(\mu^*; s_L)$ exceed 1 at the intersection, both types of technological change are strictly positive.

Capital and consumption. Given μ^* and g^*, k^* is determined by (31). The numerator is positive because $f'(\mu^*) > (1 - s_K)\psi_K(\gamma_K^*)$ (which follows from the fact that the net marginal product of an M-task, $f'(\mu) - (1 - s_K)\psi_K(\gamma_K)$, must be positive for factor prices to be positive). The denominator is $(\gamma_L^*)^\theta/\beta - 1 + (1 - \tau_K)\delta$, which is positive for all β below $\bar{\beta} \equiv (\gamma_L^*)^\theta/(1 - (1 - \tau_K)\delta)$. Since μ^* does not depend on β (it is determined by (25)), $\bar{\beta}$ is an explicit function of the first-stage variables. Given $k^* > 0$, c^* in (32) is the residual of net output over gross investment and is positive when β is sufficiently small.

Transversality condition. On the BGP, the transversality condition (A.4) reduces to $\lim_{T\to\infty}(\beta g_L(\gamma_L^*)^{1-\theta})^T c^{*-\theta}k^* = 0$, which holds if and only if $\beta g_L(\gamma_L^*)^{1-\theta} < 1$. This is satisfied for all $\beta \in (0, \bar{\beta})$ after tightening $\bar{\beta}$ if necessary. □

Factor income shares. From the factor demand conditions (11)–(13), the zero-profit condition, and the definitions of $\varepsilon_f(\mu) \equiv \mu f'(\mu)/f(\mu)$ and σ_{Rj} in (28), one can verify (29) by direct computation. Since μ^* is constant on the BGP, all factor shares are constant and determined solely by μ^* and the subsidy rates, independently of β, θ, and τ_K. □

B.2 Proof of Proposition 3

We use the total differential method. Differentiating the first-order conditions (14), (15), and the BGP condition (25) simultaneously:

$$\alpha f''(\mu)\, d\mu = -\Psi_K(\gamma_K)\, ds_K + (1 - s_K)\, \Psi_K'(\gamma_K)\, d\gamma_K, \tag{B.2}$$

$$-\mu f''(\mu)\, d\mu = -\Psi_L(\gamma_L)\, ds_L + (1 - s_L)\, \Psi_L'(\gamma_L)\, d\gamma_L, \tag{B.3}$$

$$g_L\, d\gamma_L = \tfrac{\alpha}{1-\alpha}\, g_X\, \gamma_K^{(2\alpha-1)/(1-\alpha)}\, d\gamma_K. \tag{B.4}$$

Solving (B.2) and (B.3) for $d\gamma_K$ and $d\gamma_L$, substituting into (B.4), and using the notation in (B.1):

$$d\mu = \frac{1}{f''(\mu)\,\Phi}\left(\phi_L\, \Psi_L(\gamma_L)\, ds_L - \phi_K\, \Psi_K(\gamma_K)\, ds_K\right). \tag{B.5}$$

Since $f''(\mu) < 0$ and $\Phi > 0$: (i) $ds_K > 0$, $ds_L = 0 \Rightarrow d\mu > 0$; (ii) $ds_K = 0$, $ds_L > 0 \Rightarrow d\mu < 0$.

For the growth effect, substituting $d\gamma_L$ from (B.3) into $dg = g_L\, d\gamma_L$ and simplifying:

$$dg = \frac{\phi_K \phi_L}{\Phi}\left(\alpha\, \Psi_L(\gamma_L)\, ds_L + \mu\, \Psi_K(\gamma_K)\, ds_K\right) > 0 \tag{B.6}$$

for either $ds_K > 0$ or $ds_L > 0$. The effects on γ_K^* and γ_L^* follow from (B.2)–(B.3) and the sign of $d\mu$. □

B.3 Proof of Proposition 4

From (29) and (28), the labor share can be written as

$$\sigma_L^* = \underbrace{\big(1-\varepsilon_f(\mu^*)\big)}_{\equiv\,\sigma_L^G(\mu^*)} \cdot \underbrace{\frac{\varepsilon_{\psi L}(\gamma_L^*)}{1+\varepsilon_{\psi L}(\gamma_L^*)}}_{\equiv\,\omega(\gamma_L^*)}, \tag{B.7}$$

where $\varepsilon_{\psi L} \equiv \gamma_L \psi_L'/\psi_L = \lambda\gamma_L/(\gamma_L-1)$ for the power-function specification. The factor $\sigma_L^G \equiv 1-\varepsilon_f(\mu^*)$ is the gross labor share, inclusive of LATC R&D expenditure ($\sigma_L^G = \sigma_L^* + \sigma_{RL}^*$); it is increasing in μ^* since ε_f is decreasing in μ when $\varepsilon < 1$. The factor ω is the fraction of gross labor income retained after LATC R&D costs; it is decreasing in γ_L^* since $\varepsilon_{\psi L}$ is decreasing in γ_L.

We prove parts (a) and (b) by computing the total derivatives $d\sigma_L^*/ds_K$ and $d\sigma_L^*/ds_L$ directly. The signs $d\mu^*/ds_K > 0$ and $d\mu^*/ds_L < 0$ and the formula (B.5) are established in Proposition 3. What remains is to derive $d\sigma_L^*/d\mu^*$ and the direct effect of s_L on γ_L^* at fixed μ.

Derivatives of σ_L^G, ω, and the direct channel. Recall from Proposition 6 the notation $\varepsilon_{f'}(\mu) \equiv -\mu f''(\mu)/f'(\mu) > 0$; for CES, $\varepsilon_{f'} = (1-\varepsilon_f)(1-\varepsilon)/\varepsilon$. Differentiating $\sigma_L^G = 1-\varepsilon_f(\mu)$ using the CES property $d\varepsilon_f/d\mu = -\varepsilon_f(1-\varepsilon_f)(1-\varepsilon)/(\varepsilon\mu)$:

$$\frac{d\sigma_L^G}{d\mu} = \frac{\varepsilon_f\,\varepsilon_{f'}}{\mu} > 0. \tag{B.8}$$

Differentiating $\omega(\gamma_L) = \varepsilon_{\psi L}/(1+\varepsilon_{\psi L})$ with $\varepsilon_{\psi L} = \lambda\gamma_L/(\gamma_L-1)$:

$$\omega'(\gamma_L) \equiv \frac{d\omega}{d\gamma_L} = -\frac{1}{(1+\varepsilon_{\psi L})\,\gamma_L(\gamma_L-1)} < 0. \tag{B.9}$$

Finally, differentiating the LATC first-order condition $(1-s_L)\Psi_L(\gamma_L) = f(\mu) - \mu f'(\mu)$ with respect to s_L, holding μ fixed:

$$\left.\frac{\partial\gamma_L}{\partial s_L}\right|_{\mu:\text{fixed}} = \frac{\Psi_L(\gamma_L)}{(1-s_L)\,\Psi_L'(\gamma_L)} = \frac{\gamma_L}{\varepsilon_{\Psi L}} > 0. \tag{B.10}$$

This direct channel—s_L raises γ_L^* for given μ—does not arise in Proposition 3, because there γ_L adjusts only through μ.

Proof of (a): $d\sigma_L^*/ds_K > 0$. Since s_K enters only the KATC first-order condition, it affects σ_L^* only through μ^*:

$$\frac{d\sigma_L^*}{ds_K} = \left[\frac{d\sigma_L^G}{d\mu}\cdot\omega + \sigma_L^G\cdot\omega'\cdot\left.\frac{\partial\gamma_L}{\partial\mu}\right|_{s_L:\text{fixed}}\right]\cdot\frac{d\mu^*}{ds_K}.$$

From (B.5), $d\mu^*/ds_K > 0$, and $d\mu^*/ds_K$ contains the factor $1/f''(\mu) \propto 1/\varepsilon_{f'}$ in its denominator. The bracket involves (B.8) and $\partial\gamma_L/\partial\mu|_{s_L:\text{fixed}}$, both proportional to $\varepsilon_{f'}$. In the product, $\varepsilon_{f'}$ cancels between the numerator (from the bracket) and the denominator (from $d\mu^*/ds_K$), leaving only ε-free factors. All remaining quantities—ε_f, ω, $f'(\mu)$, ϕ_K, Ψ_K, Φ, $(1-\varepsilon_f)$—are strictly positive for $\varepsilon \in (0,1)$ and $\lambda > 1$. Hence $d\sigma_L^*/ds_K > 0$. □

Proof of (b): $d\sigma_L^*/ds_L < 0$. Unlike s_K, the subsidy s_L affects σ_L^* both through μ^* and directly through γ_L^*:

$$\frac{d\sigma_L^*}{ds_L} = \underbrace{\left[\frac{d\sigma_L^G}{d\mu}\cdot\omega + \sigma_L^G\cdot\omega'\cdot\left.\frac{\partial\gamma_L}{\partial\mu}\right|_{s_L:\text{fixed}}\right]\cdot\frac{d\mu^*}{ds_L}}_{\mu\text{-path}} + \underbrace{\sigma_L^G\cdot\omega'\cdot\left.\frac{\partial\gamma_L}{\partial s_L}\right|_{\mu:\text{fixed}}}_{\text{direct channel}}.$$

The bracket in the μ-path is the same positive expression as in part (a). Multiplied by $d\mu^*/ds_L < 0$ (Proposition 3), the μ-path is negative (again with $\varepsilon_{f'}$ cancelling). The direct channel is negative because $\sigma_L^G > 0$, $\omega' < 0$ from (B.9), and $\partial\gamma_L/\partial s_L|_{\mu:\text{fixed}} > 0$ from (B.10). Since both terms are strictly negative, $d\sigma_L^*/ds_L < 0$. □

B.4 Proof of Proposition 5

Set $s_K = s_L = s$ and $ds_K = ds_L = ds$ in (B.5):

$$d\mu = \frac{ds}{f''(\mu)\,\Phi}\big(\phi_L\,\Psi_L(\gamma_L) - \phi_K\,\Psi_K(\gamma_K)\big).$$

Since $f''(\mu) < 0$ and $\Phi > 0$, $d\mu > 0$ if and only if $\phi_K \Psi_K(\gamma_K) > \phi_L \Psi_L(\gamma_L)$. Using the BGP condition $g^* = g_X \gamma_K^{\alpha/(1-\alpha)} = g_L \gamma_L$ to simplify the ratio ϕ_K/ϕ_L, and the FOCs to substitute $\Psi_K = \alpha f'/(1-s)$ and $\Psi_L = (f - \mu f')/(1-s)$, the condition reduces to $\frac{\alpha}{1-\alpha} \cdot \frac{1}{\varepsilon_{\Psi K}(\gamma_K^*)} > \frac{1}{\varepsilon_{\Psi L}(\gamma_L^*)}$, which rearranges to (33). □

B.5 Proof of Proposition 6

We reparametrize the BGP equilibrium by the common growth factor g. On the BGP, $\gamma_L = g/g_L$ and $\gamma_K = (g/g_X)^{(1-\alpha)/\alpha}$ are determined by g alone.

K-subsidy path (s_L fixed). Holding s_L fixed, the LATC first-order condition $(1 - s_L)\Psi_L(g/g_L) = f(\mu) - \mu f'(\mu)$ implicitly defines $\mu = \mu_L(g)$. Differentiating: $\mu_L'(g) = -1/(\mu\, \phi_L\, f''(\mu)) > 0$. The KATC first-order condition $\alpha f'(\mu_L(g)) = (1 - s_K)\Psi_K(\gamma_K(g))$ then pins down $s_K = s_K(g)$ with $s_K'(g) = \Phi/(\mu\, \phi_K \phi_L\, \Psi_K(\gamma_K)) > 0$. The subsidy expenditure per effective labor is $e^{(K)}(g) = s_K(g)\, \mu_L(g)\, \psi_K(\gamma_K(g)) + s_L\, \psi_L(\gamma_L(g))$. Differentiating and simplifying:

$$\left.\frac{de}{dg}\right|_{\substack{s_K:\text{var.}\\ s_L:\text{fix.}}} = \frac{\Phi}{\phi_K \phi_L} \frac{\psi_K}{\Psi_K} + \frac{s_K}{1 - s_K} \frac{\mu}{\phi_K} \frac{\psi_K'}{\Psi_K'} - \frac{s_K}{\phi_L} \frac{\psi_K}{\mu f''(\mu)} + \frac{s_L}{1 - s_L} \frac{1}{\phi_L} \frac{\psi_L'}{\Psi_L'}. \tag{B.11}$$

L-subsidy path (s_K fixed). By the symmetric argument, $\mu = \mu_K(g)$ is defined by the KATC FOC with $\mu_K'(g) = 1/(\alpha\, \phi_K\, f''(\mu)) < 0$, and $s_L = s_L(g)$ with $s_L'(g) = \Phi/(\alpha\, \phi_K \phi_L\, \Psi_L(\gamma_L)) > 0$. Differentiating $e^{(L)}(g)$:

$$\left.\frac{de}{dg}\right|_{\substack{s_K:\text{fix.}\\ s_L:\text{var.}}} = \frac{1}{\alpha} \frac{\Phi}{\phi_K \phi_L} \frac{\psi_L}{\Psi_L} + \frac{s_L}{1 - s_L} \frac{1}{\phi_L} \frac{\psi_L'}{\Psi_L'} + \frac{s_K}{1 - s_K} \frac{\mu}{\phi_K} \frac{\psi_K'}{\Psi_K'} + \frac{s_K}{\alpha\, \phi_K} \frac{\psi_K}{f''(\mu)}. \tag{B.12}$$

Comparison. Subtracting (B.12) from (B.11):

$$\left.\frac{de}{dg}\right|_{\substack{s_K:\text{var.}\\ s_L:\text{fix.}}} - \left.\frac{de}{dg}\right|_{\substack{s_K:\text{fix.}\\ s_L:\text{var.}}} = \frac{\Phi}{\phi_K \phi_L} \left[\frac{\psi_K}{\Psi_K} - \frac{1}{\alpha} \frac{\psi_L}{\Psi_L} - \frac{s_K}{\mu \alpha} \frac{\psi_K}{f''(\mu)} \right]. \tag{B.13}$$

Since $\Phi/(\phi_K \phi_L) > 0$, the KATC subsidy achieves lower de/dg when the bracketed expression is negative.

Condition at $s_K = s_L = 0$. When $s_K = 0$, the last term in the bracket vanishes and the condition reduces to $\psi_K/\Psi_K < (1/\alpha)(\psi_L/\Psi_L)$. Using $\Psi_K/\psi_K = \alpha + \varepsilon_{\psi K}$ and $\Psi_L/\psi_L = 1 + \varepsilon_{\psi L}$, this is equivalent to $\varepsilon_{\psi K}/\varepsilon_{\psi L} > \alpha$, giving condition (35).

The rewritten form (37) follows by substituting the FOCs $\Psi_K = \alpha f'/(1 - s_K)$ and $\Psi_L = (f - \mu f')/(1 - s_L)$, evaluated at $s_K = s_L = 0$.

Conversion to the subsidy-to-GDP metric. Since $\sigma_S = e/f(\mu)$, $d\sigma_S/dg = (1/f(\mu))(de/dg - e\, f'(\mu)\,(d\mu/dg)/f(\mu))$. At $s_K = s_L = 0$, $e = 0$, so $d\sigma_S/dg = (1/f(\mu))\, de/dg$ and condition (35) applies directly to the σ_S metric as stated in (34).

For the general case $s_K, s_L \geq 0$, the additional term $-ef'/(f^2) \cdot d\mu/dg$ modifies the comparison. The third term in the bracket of (B.13) is replaced by $(s_K\psi_K - e\, f'/f) \cdot (-1/(\mu\alpha f''))$. Substituting $e = s_K\mu\psi_K + s_L\psi_L$, using the FOCs to express Ψ_K/f' and $\Psi_L/(f - \mu f')$, and defining $\varepsilon_{f'}(\mu) \equiv -\mu f''(\mu)/f'(\mu) > 0$, the condition $d\sigma_S^*/dg^*|_{s_K:\text{var.}} < d\sigma_S^*/dg^*|_{s_L:\text{var.}}$ becomes (36). □

B.6 Uniform Subsidy and the Labor Share

Lemma 2 (Uniform subsidy and the labor share)**.** *Under the maintained assumptions of Proposition 4 and condition (33), if*

$$\lambda(\lambda - 1)\left(\frac{\alpha}{1-\alpha}\,\frac{\varepsilon_{\Psi L}}{\varepsilon_{\Psi K}} - 1\right) > (\gamma_L^* - 1)\left(1 + \frac{\alpha}{1-\alpha}\,\frac{\varepsilon_{\Psi L}(1 - \varepsilon_f)}{\varepsilon_f\,\varepsilon_{\Psi K}}\right), \tag{B.14}$$

then a marginal uniform subsidy ($ds_K = ds_L = ds$) raises the labor share: $d\sigma_L^*/ds|_{\text{uniform}} > 0$.

Proof. Setting $ds_K = ds_L = ds$ in the results of Proposition 4:

$$\left.\frac{d\sigma_L^*}{ds}\right|_{\text{uniform}} = \underbrace{\left[\frac{d\sigma_L^G}{d\mu}\cdot\omega + \sigma_L^G\cdot\omega'\cdot\left.\frac{\partial\gamma_L}{\partial\mu}\right|_{s_L:\text{fixed}}\right]\cdot\left(\frac{d\mu^*}{ds_K} + \frac{d\mu^*}{ds_L}\right)}_{\mu\text{-path}} + \underbrace{\sigma_L^G\cdot\omega'\cdot\left.\frac{\partial\gamma_L}{\partial s_L}\right|_{\mu:\text{fixed}}}_{\text{direct LATC channel}<0}. \tag{B.15}$$

The bracket in the μ-path is strictly positive (as shown in the proof of part a of Proposition 4). From (B.5), $d\mu^*/ds_K + d\mu^*/ds_L = \mu(\phi_K\Psi_K - \phi_L\Psi_L)/(\varepsilon_{f'}f'\Phi)$, which is positive under condition (33). Hence the μ-path is positive, while the direct channel is negative. We need to show that the μ-path dominates.

Substituting (B.8), (B.9), and (B.5) into (B.15), the $\varepsilon_{f'}$-cancellation applies to both the μ-path and the bracket, giving

$$\left.\frac{d\sigma_L^*}{ds}\right|_{\text{uniform}} = (1-\varepsilon_f)\Bigg[\underbrace{\frac{\varepsilon_f\,\omega + (1-\varepsilon_f)\,\frac{\alpha}{1-\alpha}\,|\omega'|\,\frac{\gamma_L}{\varepsilon_{\Psi K}}}{1+\frac{\alpha}{1-\alpha}\,\frac{\varepsilon_{\Psi L}(1-\varepsilon_f)}{\varepsilon_f\,\varepsilon_{\Psi K}}}\cdot\left(\frac{\alpha}{1-\alpha}\,\frac{\varepsilon_{\Psi L}}{\varepsilon_{\Psi K}}-1\right)}_{\mu\text{-path}} + \underbrace{\omega'\cdot\frac{\gamma_L}{\varepsilon_{\Psi L}}}_{\text{direct}}\Bigg].$$

Since $(1-\varepsilon_f) > 0$, the sign depends on the bracketed expression. Dividing the μ-path by $|\omega'|\cdot\gamma_L/\varepsilon_{\Psi L}$, the condition $d\sigma_L^*/ds|_{\text{uniform}} > 0$ reduces to

$$\frac{\varepsilon_f\,\omega/|\omega'| + (1-\varepsilon_f)\,\frac{\alpha}{1-\alpha}\,\frac{\gamma_L}{\varepsilon_{\Psi K}}}{1+\frac{\alpha}{1-\alpha}\,\frac{\varepsilon_{\Psi L}(1-\varepsilon_f)}{\varepsilon_f\,\varepsilon_{\Psi K}}}\cdot\frac{\varepsilon_{\Psi L}}{\gamma_L}\cdot\left(\frac{\alpha}{1-\alpha}\,\frac{\varepsilon_{\Psi L}}{\varepsilon_{\Psi K}}-1\right) > 1.$$

For the power-function specification, $\omega/|\omega'| = \varepsilon_{\psi L}\gamma_L(\gamma_L-1) = \lambda\gamma_L^2$. Substituting and simplifying, the left-hand side factors as

$$\frac{(\lambda\gamma_L-1)\big((1+\lambda)\gamma_L-1\big)}{\gamma_L-1}\cdot\frac{\frac{\alpha}{1-\alpha}\,\frac{\varepsilon_{\Psi L}}{\varepsilon_{\Psi K}}-1}{1+\frac{\alpha}{1-\alpha}\,\frac{\varepsilon_{\Psi L}(1-\varepsilon_f)}{\varepsilon_f\,\varepsilon_{\Psi K}}} > 1.$$

The first factor depends only on γ_L^* and λ. Since $(\lambda\gamma_L-1)((1+\lambda)\gamma_L-1)/(\gamma_L-1) > \lambda(\lambda-1)/(\gamma_L-1)$ for all $\gamma_L > 1$ and $\lambda > 1$, the inequality holds whenever (B.14) is satisfied. □

All quantities in (B.14) are determined before EoS_{KL} enters the calibration (Section 5.2), so the condition is independent of ε. The left-hand side of (B.14) is a product of the R&D cost convexity margin $\lambda(\lambda-1)$ and the excess of the K-side multiplier over the L-side (condition 33). The right-hand side is small because it contains the per-period LATC innovation rate $\gamma_L^* - 1 \approx 0.079$.

Numerical verification. At the baseline ($\lambda = 2$), the left-hand side of (B.14) equals 0.903 and the right-hand side equals 0.288, giving a margin of 3.13×. At $\lambda = 3/2$ and $\lambda = 5/2$, the margins are 1.54× and 5.03×, respectively. Since all quantities are ε-invariant, these margins apply identically at $\text{EoS}_{KL} = 0.8$.

Table C.1: Conversion between annualized and per-period variables

Quantity	Annualized → per-period	Per-period → annualized
Growth factor γ	$\gamma = (1+\gamma^a)^\chi$	$\gamma^a = \gamma^{1/\chi} - 1$
Discount factor β	$\beta = (1+\beta^a)^{-\chi}$	$\beta^a = \beta^{-1/\chi} - 1$
Depreciation δ	$\delta = 1-(1-\delta^a)^\chi$	$\delta^a = 1-(1-\delta)^{1/\chi}$
R&D cost z_j	$z_j = z_j^a/\chi^{\lambda-1}$	$z_j^a = \chi^{\lambda-1} z_j$
TFP B	$B = \chi \cdot B^a$	$B^a = B/\chi$
Capital–output ratio	$(K/Y)^* = (K/Y)^a/\chi$	$(K/Y)^a = \chi \cdot (K/Y)^*$

The following quantities do not depend on χ: σ_L^*, σ_R^*, σ_K^*, α, η, ε, θ, λ, $(\delta K/Y)^*$.

C Calibration and Robustness

C.1 Calibration Details

This subsection provides the complete derivation of the calibration procedure summarized in Section 5.2. We first present the formulas for converting between annualized and per-period variables, then derive each calibrated parameter step by step.

Normalization. We normalize the initial technology levels and factor endowments at $t = -1$ (the last period before the model begins) as follows:

$$A_{K,-1} = A_{L,-1} = X_{-1} = L_{-1} = 1, \qquad Y_{-1}/\chi = 1.$$

The last condition says that annualized output equals one in the initial period. Since $Y_{-1} = F(M_{-1}, N_{-1}) = N_{-1} f(\mu^*) = f(\mu^*)$, this implies $f(\mu^*) = \chi$. Similarly, $K_{-1}/(Y_{-1}/\chi) = (K/Y)^a = 2.9$ with $Y_{-1}/\chi = 1$ gives $K_{-1} = 2.9$, so $k^* = K_{-1}/(A_{L,-1}L_{-1}) = 2.9$.

Conversion between annualized and per-period variables. One model period spans χ years. Throughout, a superscript a denotes annualized net rates and unadorned symbols denote per-period gross factors (or levels). Table C.1 lists the conversion formulas.

The growth factor formula follows from compounding: if the annualized growth rate is γ^a, then the growth factor over χ years is $(1+\gamma^a)^\chi$. The depreciation formula reflects survival compounding: the fraction of capital surviving one year is $(1-\delta^a)$,

so the fraction surviving χ years is $(1-\delta^a)^\chi = 1-\delta$. The R&D cost scaling follows from $\psi_K(\gamma_K) = z_K(\gamma_K - 1)^\lambda$: when the period is χ times longer, $\gamma_K - 1 \approx \chi \cdot \gamma_K^a$ is approximately χ times larger, so the cost scales as χ^λ; to keep the annualized cost comparable, z_K must scale as $1/\chi^{\lambda-1}$. The TFP scaling follows from $f(\mu^*) = \chi$ and the fact that the CES aggregate within f does not depend on χ, so B scales linearly with χ.

Step 1: Exogenous parameters. The following parameters are set exogenously: $\theta = 1$ (log utility), $\lambda = 2$ (quadratic R&D cost), $\tau_K = 0.15$ (capital income tax), $s_K = s_L = 0$ (no R&D subsidies at baseline), and $g_L^a = 1.02\%$ (population growth rate).

Step 2: Period length χ. The period length is determined by the R&D share condition (40). We derive this condition by eliminating z_K and z_L from the factor share equations.

From (29), the capital share satisfies $\sigma_K^*/\alpha = \varepsilon_f(\mu^*) - \sigma_{RK}^*$, where $\sigma_{RK}^* = (1 - s_K)\mu^*\psi_K(\gamma_K^*)/f(\mu^*)$. Substituting the R&D first-order condition (14), which implies $(1-s_K)z_K = \alpha f'(\mu^*)/\Psi_{K,\text{body}}(\gamma_K^*, \lambda, \alpha)$, we can express σ_{RK}^* in terms of σ_K^*, α, γ_K^*, and λ alone.[35] Here $\Psi_{K,\text{body}}(\gamma, \lambda, \alpha) \equiv \Psi_K(\gamma)/z_K = \alpha(\gamma-1)^\lambda + \gamma\lambda(\gamma-1)^{\lambda-1}$ is the z_K-free part of the composite marginal R&D cost, obtained by substituting the power-function specification (39) into the definition of Ψ_K. Summing the KATC and LATC components yields the R&D share condition (40).

All per-period quantities in (40) depend on χ: $\gamma_K^*(\chi) = (1+\gamma_K^{*a})^\chi$, $\gamma_L^*(\chi) = (1+\gamma_L^{*a})^\chi$, and $\sigma_K^*(\chi) = [(1.04)^\chi - 1] \cdot 2.9/\chi + 0.14$ (the last expression is derived in Step 3 below). Since the left-hand side $\sigma_R^* = 2.7\%$ is fixed by data and the right-hand side is strictly increasing in χ, there is a unique solution, obtained numerically as $\chi \approx 4.10$ years.

[35] The key simplification is $\psi_{\text{body}}(\gamma, \lambda)/\Psi_{K,\text{body}}(\gamma, \lambda, \alpha) = (\gamma-1)/(\alpha(\gamma-1) + \gamma\lambda)$, which, after substitution, yields $\sigma_{RK}^* = (\sigma_K^*/\lambda)(\gamma_K^* - 1)/\gamma_K^*$. A parallel derivation gives $\sigma_{RL}^* = (\sigma_L^*/\lambda)(\gamma_L^* - 1)/\gamma_L^*$.

Step 3: Capital share σ_K^*, capital-composite share α, and resource share σ_X^*. The capital income share is the sum of the net return to capital and depreciation:

$$\sigma_K^* = (r^* - \delta) \cdot \left(\frac{K}{Y}\right)^* + \left(\frac{\delta K}{Y}\right)^* = \left[(1.04)^\chi - 1\right] \cdot \frac{2.9}{\chi} + 0.14, \tag{C.1}$$

where $r^* - \delta = (1 + 0.04)^\chi - 1$ is the per-period net return (from the annualization formula in Table C.1), $(K/Y)^* = (K/Y)^a/\chi = 2.9/\chi$ is the per-period capital–output ratio, and $(\delta K/Y)^* = 0.14$ does not require annualization. Note that σ_K^* is not targeted directly but inferred from the return, capital–output, and depreciation targets.

From (29), the income shares satisfy $\sigma_K^*/\sigma_X^* = \alpha/(1-\alpha)$ and $\sigma_K^* + \sigma_X^* = 1 - \sigma_L^* - \sigma_R^*$. Solving:

$$\alpha = \frac{\sigma_K^*}{1 - \sigma_L^* - \sigma_R^*}, \qquad \sigma_X^* = (1 - \alpha)(1 - \sigma_L^* - \sigma_R^*). \tag{C.2}$$

Step 4: k^*, $f(\mu^*)$, δ, and r^*. The normalization yields $k^* = (K/Y)^a = 2.9$ and $f(\mu^*) = \chi$ as described above. The depreciation rate follows from the target $\delta K/Y$:

$$\delta = \frac{\left(\frac{\delta K}{Y}\right)^*}{\left(\frac{K}{Y}\right)^*} = \frac{0.14}{2.9/\chi} = \frac{0.14\,\chi}{2.9}. \tag{C.3}$$

The gross interest rate follows: $r^* = [(1.04)^\chi - 1] + \delta$.

Step 5: μ^*. From the normalization, $M_{-1} = (A_{K,-1}K_{-1})^\alpha X_{-1}^{1-\alpha} = K_{-1}^\alpha$ and $N_{-1} = A_{L,-1}L_{-1} = 1$. Thus:

$$\mu^* = M_{-1}/N_{-1} = k^{*\alpha} = 2.9^\alpha. \tag{C.4}$$

Step 6: β from the Euler equation. On the BGP, the Euler equation (30) gives:

$$\beta = \frac{(\gamma_L^*)^\theta}{(1 - \tau_K)(r^* - \delta) + 1}, \tag{C.5}$$

where $(r^* - \delta) = (1.04)^\chi - 1$ is the per-period net return. Note that β depends on θ and τ_K, but this does not affect μ^*, g^*, or any first-stage variable.

Step 7: g_X from the BGP technology condition. The natural resource growth factor g_X is not directly targeted by data. Instead, it is derived from the BGP condition (25), which equates the growth of M-tasks and N-tasks:

$$g_X = \frac{g_L \cdot \gamma_L^*}{\gamma_K^{*\alpha/(1-\alpha)}}. \tag{C.6}$$

Step 8: $\varepsilon_f(\mu^*)$ from the factor shares and R&D first-order conditions. The output elasticity of M-tasks at the BGP is obtained by eliminating z_K from the capital share equation. From (29), $\sigma_K^*/\alpha = \varepsilon_f(\mu^*) - \sigma_{RK}^*$. Substituting the expression for σ_{RK}^* derived in Step 2 and solving for ε_f:

$$\varepsilon_f(\mu^*) = \frac{\sigma_K^*}{\alpha}\left(\frac{\alpha(\gamma_K^* - 1)}{\gamma_K^*\lambda} + 1\right). \tag{C.7}$$

The derivation uses the identity[36] $\sigma_{RK}^* = (\sigma_K^*/\alpha) \cdot \alpha\psi_{\text{body}}/\Psi_{K,\text{body}}$.

Step 9: $f'(\mu^*)$. From the definition $\varepsilon_f(\mu) \equiv \mu f'(\mu)/f(\mu)$:

$$f'(\mu^*) = \frac{\chi \cdot \varepsilon_f(\mu^*)}{\mu^*}. \tag{C.8}$$

Step 10: z_K and z_L from the R&D first-order conditions. Solving (14) and (15) for z_K and z_L at the baseline ($s_K = s_L = 0$):

$$z_K = \frac{\alpha\, f'(\mu^*)}{\Psi_{K,\text{body}}(\gamma_K^*, \lambda, \alpha)} = \frac{\alpha\, f'(\mu^*)}{\alpha(\gamma_K^* - 1)^\lambda + \gamma_K^*\lambda(\gamma_K^* - 1)^{\lambda-1}}, \tag{C.9}$$

$$z_L = \frac{f(\mu^*) - \mu^* f'(\mu^*)}{\Psi_{L,\text{body}}(\gamma_L^*, \lambda)} = \frac{f(\mu^*) - \mu^* f'(\mu^*)}{(\gamma_L^* - 1)^\lambda + \gamma_L^*\lambda(\gamma_L^* - 1)^{\lambda-1}}. \tag{C.10}$$

With subsidies $s_K, s_L > 0$, the right-hand sides are divided by $(1 - s_K)$ and $(1 - s_L)$, respectively.

Step 11: ε from the elasticity of substitution condition. As described in Section 5.2, the CES parameter ε differs from EoS_{KL} due to the nested production

[36] This is the same identity used in Step 2. For the power-function specification $\psi_K(\gamma) = z_K(\gamma - 1)^\lambda$, we have $\psi_K/\Psi_K = (\gamma-1)^\lambda/[\alpha(\gamma-1)^\lambda + \gamma\lambda(\gamma-1)^{\lambda-1}] = (\gamma-1)/[\alpha(\gamma-1) + \gamma\lambda]$.

structure. Solving (41) for ε:

$$\varepsilon = \frac{\alpha}{\big(1-(1-\alpha)\,\varepsilon_f(\mu^*)\big)\,\mathrm{EoS}_{KL}^{-1} - (1-\alpha)\big(1-\varepsilon_f(\mu^*)\big)}. \tag{C.11}$$

Step 12: η from the definition of ε_f. The CES intensive form (38) has the elasticity $\varepsilon_f(\mu) = \eta\,\mu^{(\varepsilon-1)/\varepsilon}\big/\big[\eta\,\mu^{(\varepsilon-1)/\varepsilon} + (1-\eta)\big]$. Solving for η:

$$\eta = \frac{1}{\mu^{*\frac{\varepsilon-1}{\varepsilon}}\left(\dfrac{1}{\varepsilon_f(\mu^*)} - 1\right) + 1}. \tag{C.12}$$

Step 13: B from the production function. The normalization $f(\mu^*) = \chi$ and the CES form (38) give:

$$B = \chi\left[\eta\,(\mu^*)^{\frac{\varepsilon-1}{\varepsilon}} + (1-\eta)\right]^{-\frac{\varepsilon}{\varepsilon-1}}. \tag{C.13}$$

C.2 Local Stability Analysis

The dynamical system (20)–(23) defines a mapping $(\mu_{t+1}, k_{t+1}, c_{t+1}) = (\widehat{\mu}, \widehat{k}, \widehat{c})(\mu_t, k_t, c_t)$. We evaluate the 3×3 Jacobian matrix numerically at the BGP (μ^*, k^*, c^*) using a finite-difference approximation.[37]

For the baseline calibration ($\lambda = 2$, $\theta = 1$), the Jacobian is

$$J = \begin{pmatrix} 1.2305 & 0.4733 & -0.2806 \\ -0.1568 & 0.7217 & -0.1917 \\ -0.4860 & -0.8630 & 1.3075 \end{pmatrix},$$

with eigenvalues 0.970, 0.581, and 1.710. Since μ_t and k_t are predetermined and c_t is a jump variable, saddle-path stability requires exactly two stable eigenvalues ($|\lambda_i| < 1$) and one unstable eigenvalue ($|\lambda_i| \geq 1$). This condition is satisfied: the two stable eigenvalues govern convergence of the state variables, while the unstable eigenvalue pins down the unique initial consumption c_0 on the convergent saddle path.

[37] We use the `ND` function from the Mathematica `NumericalCalculus` package with step scale 10^{-4}.

Since one model period spans χ years, a per-period eigenvalue λ_i corresponds to an annual convergence rate of $|\ln \lambda_i|/\chi$ and a half-life of $\chi \ln 2/|\ln \lambda_i|$. The two stable eigenvalues thus correspond to half-lives of approximately 93 years (slow, ≈ 0.97) and 5 years (fast, ≈ 0.58). The normalized eigenvectors in the (μ, k, c) space are $(1.00\ -0.59\ -0.06)$ for the slow mode and $(-0.30\ 1.00\ 0.98)$ for the fast mode. The slow mode thus loads almost entirely on μ, while the fast mode loads jointly on k and c. Because the separation in time scales is large (a factor of roughly 19 in half-lives), k_t tracks its instantaneous target $k^*(\mu_t)$ closely, and the observed convergence speed of K/Y, factor shares, and growth rates is governed by μ's slow dynamics, as discussed in Section 5.3.

Sensitivity to θ, λ, τ_K, and EoS_{KL}. The slow half-life is nearly invariant to θ, τ_K, and EoS_{KL}, confirming that these parameters do not substantively affect the first-stage dynamics governing the technology direction. The R&D cost convexity λ is the key determinant: lower convexity ($\lambda = 3/2$) roughly halves the slow half-life to 51 years, while higher convexity ($\lambda = 5/2$) extends it to 136 years. The fast mode remains in the range of 4–6 years throughout. Table C.2 reports the underlying eigenvalues.

Table C.2: Eigenvalues of the linearized system

Configuration	Stable (slow)	Stable (fast)	Unstable
Baseline	0.970	0.581	1.710
$\theta = 1/2$	0.969	0.474	2.097
$\theta = 3/2$	0.970	0.635	1.561
$\lambda = 3/2$	0.959	0.641	1.504
$\lambda = 5/2$	0.974	0.524	1.924
$\tau_K = 23.8\%$	0.970	0.594	1.670
$\text{EoS}_{KL} = 0.8$	0.972	0.624	1.624

C.3 Robustness: Calibrated Parameters

Table C.3 reports the calibrated parameters under each alternative specification considered in Section 5.3. All calibrations match the same 9 empirical moments listed in Table 2. The baseline column ($\lambda = 2$, $\theta = 1$, $\tau_K = 15\%$, $\text{EoS}_{KL} = 0.6$) is shown

Table C.3: Calibrated parameters under alternative specifications

Parameter	Baseline	$\lambda=3/2$	$\lambda=5/2$	$\theta=1/2$	$\theta=3/2$	$\tau_K=23.8\%$	$\text{EoS}_{KL}=0.8$
χ (years)	4.10	3.06	5.17				
α	0.757	0.750	0.765				
δ^a	5.24%	5.09%	5.40%				
β^a	1.54%	1.53%	1.55%	2.48%	0.61%	1.21%	
z_K^a	28.8	2.60	318.7				
z_L^a	62.1	8.74	436.8				
ε	0.554	0.552	0.556				0.768
η	0.510	0.509					0.410
B^a	0.707	0.710	0.705				0.736
g_X^a	0.38%	0.47%	0.27%				

Blank entries denote values identical to the baseline. When θ or τ_K varies, only β adjusts (through the Euler equation); all other parameters are unchanged. When λ varies, most parameters change because the period length χ is re-determined; the R&D cost parameters z_K^a and z_L^a absorb the bulk of the adjustment. When EoS_{KL} varies, only the CES parameters (ε, η, B) change, because EoS_{KL} enters the calibration only at the final stage (Section 5.2).

on the left; each subsequent column varies only the indicated parameter, and blank entries denote values identical to the baseline.

The most notable feature is that the R&D cost parameters z_K^a and z_L^a vary by two orders of magnitude across λ, while all other parameters change by less than 2%. This is because z_j and λ enter the R&D cost function $\psi_j(\gamma_j) = z_j(\gamma_j - 1)^\lambda$ multiplicatively with λ in the exponent: a change in the curvature λ is absorbed almost entirely by the level z_j, leaving the equilibrium R&D rates γ_K^* and γ_L^* (and hence all other calibrated quantities) nearly unchanged.

C.4 Robustness: Policy Scenarios

Table C.4 reports the policy scenario results across the three values of λ and the alternative $\text{EoS}_{KL} = 0.8$ calibration. The five scenarios are defined in Section 6.1; each allocates a budget equal to 1% of GDP to R&D subsidies or tax reductions.

The key qualitative results are robust across all specifications: (i) the ranking KATC > Uniform > LATC > Tax reduction holds for growth effects; (ii) the KATC subsidy delivers roughly twice the growth effect of the LATC subsidy for the same 1% GDP budget (growth ratio 1.90–2.04 across λ); (iii) KATC subsidies raise the labor share while LATC subsidies lower it; (iv) tax reductions have zero growth effect.

Table C.4: Policy scenario results across λ and EoS_{KL} (1% of GDP budget)

	Baseline		$\lambda = 3/2$		$\lambda = 5/2$		$\text{EoS}_{KL} = 0.8$	
Scenario	g^a_{pc}	σ^*_L	g^a_{pc}	σ^*_L	g^a_{pc}	σ^*_L	g^a_{pc}	σ^*_L
Baseline	1.86	62.5	—		—		—	
S1. KATC subsidy	2.67	74.3	3.37	72.0	2.41	75.3	2.65	67.8
S2. LATC subsidy	2.27	59.5	2.60	59.4	2.15	59.6	2.27	60.8
S3. Uniform subsidy	2.39	62.8	2.80	62.2	2.23	63.1	2.39	62.4
S4. Tax reduction	1.86	62.5	—		—		—	
S5. KATC + tax	2.67	74.3	3.37	72.0	2.41	75.3	2.65	67.8
Growth ratio (S1/S2)	1.98		2.04		1.90		1.93	

All values in percentages. g^a_{pc}: per-capita growth rate. σ^*_L: BGP labor income share. Growth ratio: $\Delta g^a_{\text{pc}}(\text{S1})/\Delta g^a_{\text{pc}}(\text{S2})$ for the same 1% GDP budget. The gap between this discrete ratio (≈ 2) and the marginal cost-effectiveness ratio of ≈ 3 (Proposition 6) is explained in Section 5.3. Baseline uses $\lambda = 2$, $\text{EoS}_{KL} = 0.6$; the $\text{EoS}_{KL} = 0.8$ column uses $\lambda = 2$. "—" indicates values identical to baseline. Since θ and τ_K do not affect any first-stage variable, all entries are also invariant to these parameters.

The parameter λ affects the *absolute magnitude* of policy effects. A lower λ makes the R&D cost function less convex, so firms' innovation choices are more responsive to subsidies: the KATC growth gain ranges from +0.55 percentage points ($\lambda = 5/2$) to +1.51 percentage points ($\lambda = 3/2$). This amplification applies symmetrically to all subsidy types, preserving the relative rankings and cost-effectiveness ratios.

The parameter EoS_{KL} affects the *distributional* effects but leaves growth effects essentially unchanged. Comparing the $\lambda = 2$ and $\text{EoS}_{KL} = 0.8$ columns, the growth rates are nearly identical (e.g., 2.67% vs 2.65% for KATC), while the labor share responses are roughly halved: the KATC subsidy raises σ^*_L by +5.3 pp (vs +11.8 pp at the baseline), and the LATC subsidy lowers it by −1.7 pp (vs −3.0 pp). Most strikingly, the uniform subsidy's labor share effect reverses sign, from +0.3 pp to −0.1 pp. This sensitivity reflects the nonlinear dependence of distributional effects on the CES curvature: the *marginal* effect $d\sigma^*_L/ds|_{\text{uniform}}$ is positive and ε-invariant (Appendix B.6), but higher-order terms become adverse when ε is closer to unity. In the Cobb–Douglas limit ($\varepsilon \to 1$), distributional effects vanish entirely while growth effects survive.

D Computational Algorithm for Transition Dynamics

This appendix describes the shooting algorithm used to compute the transition dynamics in Section 6.2. We first describe the algorithm for the fixed-rate policy (where s_K, s_L, τ_K are fixed at their new steady-state values for all $t \geq 0$), then explain how it is modified for the constant-budget policy (where rates adjust each period to maintain fiscal expenditure at 1% of GDP).

D.1 Fixed-Rate Transition

The algorithm relies on the mapping $\Gamma(\mu;\, \mathbf{s})$ defined in Appendix A.3 (equation (A.3)), which links the pre-R&D relative factor intensity μ^{pre} to the post-R&D value μ through the equilibrium R&D rates $\widehat{\gamma}_K$ and $\widehat{\gamma}_L$. Its inverse $\Gamma^{(-1)}$ appears in the law of motion for μ_t (20) and is used throughout the computation.

Step 1: Compute the new BGP. Given the new policy parameters (s_K, s_L, τ_K), compute the new balanced growth path (μ^*, k^*, c^*) following Section 4.1: μ^* and g^* from (25), k^* from (31), and c^* from (32).

Step 2: Compute initial conditions. At $t = 0$, the economy transitions from the old BGP $(\mu^*_{\text{old}}, k^*_{\text{old}}, c^*_{\text{old}})$ to the new policy regime. The initial state (μ_0, k_0) is determined as follows. Under the old BGP, capital grows at rate g^*_{old}, so $K_1/K_0 = g^*_{\text{old}}$. The pre-R&D intensity at $t = 0$ is computed from (A.2):

$$\mu_0^{\text{pre}} = \frac{(g^*_{\text{old}})^{\alpha}\, g_X^{1-\alpha}}{g_L}\, \mu^*_{\text{old}}, \tag{D.1}$$

$$\mu_0 = \Gamma^{(-1)}(\mu_0^{\text{pre}};\, \mathbf{s}), \tag{D.2}$$

$$k_0 = \frac{v(\mu^*_{\text{old}};\, \mathbf{s}_{\text{old}}) - c^*_{\text{old}} + (1-\delta)\, k^*_{\text{old}}}{g_L\, \widehat{\gamma}_L(\mu_0;\, s_L)}. \tag{D.3}$$

Step 3: Shooting (bisection over c_0). Search for $c_0 \in [0, c_{\max}]$ where $c_{\max} = v(\mu_0;\, \mathbf{s}) + (1-\delta)\, k_0$, such that the forward-iterated path $\{(\mu_t, k_t, c_t)\}_{t=0}^{T}$ converges to (μ^*, k^*, c^*).

For each candidate c_0, iterate $(\mu_{t+1}, k_{t+1}, c_{t+1}) = (\widehat{\mu}, \widehat{k}, \widehat{c})(\mu_t, k_t, c_t)$ forward using the three laws of motion (20), (21), and (23). The bisection classifies each path: if $k_t < 0$ (overshooting consumption), decrease c_0; if $c_t < 0$ (undershooting consumption), increase c_0. Each path is iterated for $T = 400$ years ($400/\chi$ periods) and is required to reach the neighborhood of the new steady state. Only the first 350 years are retained; the tail is discarded because the divergent component of the saddle path, though exponentially small, is amplified over time and may contaminate the later periods.

Precision requirement. Due to the saddle-path instability, the algorithm requires approximately 200 digits of working precision to resolve c_0 with sufficient accuracy over 400-year paths.

D.2 Constant-Budget Transition

Under the fixed-rate policy, holding the subsidy and tax rates constant causes the fiscal cost to deviate from 1% of GDP during the transition (see the footnote in Section 6.2). Under the constant-budget policy, the policy instruments adjust each period to maintain the fiscal target specified in Table 6: for Scenarios 1–3 the subsidy rates adjust to keep $\sigma_S = 1\%$ of GDP, for Scenario 4 the tax rate adjusts to keep the revenue reduction at 1% of GDP relative to the old BGP, and for Scenario 5 both the subsidy rate and the tax rate adjust. The shooting algorithm is identical to Section D.1 except that the policy rates and hence μ_{t+1} (and, for S4–S5, $\tau_{K,t+1}$) must be determined endogenously at each step.

Simultaneous determination of $\mathbf{s}_{t+1}$ and μ_{t+1}. Under the fixed-rate policy, the subsidy rates $\mathbf{s} = (s_K, s_L)$ are constant, so the period-$(t+1)$ R&D rates $\gamma_{K,t+1} = \widehat{\gamma}_K(\mu_{t+1}; s_K)$ and $\gamma_{L,t+1} = \widehat{\gamma}_L(\mu_{t+1}; s_L)$ (equations (14)–(15)) depend only on μ_{t+1}. Under the constant-budget policy, the subsidy rates $\mathbf{s}_{t+1}$ are themselves endogenous, which changes the equilibrium R&D rates $\gamma_{K,t+1} = \widehat{\gamma}_K(\mu_{t+1}; s_{K,t+1})$ and $\gamma_{L,t+1} = \widehat{\gamma}_L(\mu_{t+1}; s_{L,t+1})$, and hence μ_{t+1} through the Γ mapping (A.3). Therefore the subsidy rates and the technology state must be determined simultaneously. The key observation is that μ_{t+1}^{pre} from (A.2) depends only on (μ_t, k_t, c_t) and the current-period rates

$\mathbf{s}_t$, not on $\mathbf{s}_{t+1}$. Define the subsidy expenditure share of GDP as

$$\sigma_S(\mu;\, s_K, s_L) \equiv \frac{s_K\, \mu\, \psi_K\big(\widehat{\gamma}_K(\mu;\, s_K)\big) + s_L\, \psi_L\big(\widehat{\gamma}_L(\mu;\, s_L)\big)}{f(\mu)}, \tag{D.4}$$

and the capital tax revenue share of GDP as

$$\sigma_\tau(\mu, k;\, \tau_K) \equiv \tau_K \left(\sigma_K(\mu) - \frac{\delta\, k}{f(\mu)}\right), \tag{D.5}$$

where $\sigma_K(\mu)$ is the capital income share (29).

Subsidy rates and μ_{t+1} (simultaneous). In each subsidy scenario, only one degree of freedom remains (see below), so given μ_{t+1}^{pre} the budget constraint

$$\sigma_S\big(\Gamma^{(-1)}(\mu_{t+1}^{\text{pre}};\, \mathbf{s}_{t+1}),\ \mathbf{s}_{t+1}\big) = 1\% \tag{D.6}$$

is a single-variable root-finding problem. Its solution determines $\mathbf{s}_{t+1}$, from which $\mu_{t+1} = \Gamma^{(-1)}(\mu_{t+1}^{\text{pre}};\, \mathbf{s}_{t+1})$ follows immediately, and then k_{t+1} from (21).

Tax rate and c_{t+1} (sequential). Unlike the subsidy rates, the capital tax rate $\tau_{K,t+1}$ does not require simultaneous determination: it does not affect μ_{t+1} or k_{t+1}, and enters only through the Euler equation (23). For Scenarios 4 and 5, $\tau_{K,t+1}$ is therefore determined from the tax-revenue target $\sigma_\tau(\mu_{t+1}, k_{t+1};\, \tau_{K,t+1}) = \sigma_\tau^{\text{target}}$ before evaluating c_{t+1}.[38] For S1–S3, τ_K is constant and c_{t+1} is computed directly.

Scenario-specific budget constraints.

- **(Incompletely) directed subsidies** (Section 6.3): for a fixed $\rho \equiv s_K/(s_K + s_L) \in [0, 1]$, find $s_{K,t}$ and $s_{L,t}$ such that $s_{K,t} : s_{L,t} = \rho : (1-\rho)$ and $\sigma_S(\mu_t;\, s_{K,t},\, s_{L,t}) = 1\%$.

- S1 (KATC), S2 (LATC), and S3 (uniform) are special cases with $\rho = 1$, 0, and 1/2, respectively.

[38] Because σ_τ is not monotone in τ_K (Laffer curve), two tax rates may achieve the same target revenue. We select the lower root and impose an upper bound $\tau_{K,t+1} \le \bar{\tau}_K$ where $\bar{\tau}_K$ is the revenue-maximizing rate.

- **S4** (Tax cut): no subsidy; $\sigma_\tau^{\text{target}} = \sigma_\tau^{\text{old}} - 1\%$ (Table 6).[38]
- **S5** (Subsidy + tax): $s_{K,t}$ from the subsidy constraint (as $\rho = 1$); $\sigma_\tau^{\text{target}} = \sigma_\tau^{\text{old}} + 1\%$.[38]

D.3 Welfare Comparison: Lifetime Equivalent Variation

To compare welfare across policy scenarios, we compute the lifetime equivalent variation (LEV). The LEV λ is the constant proportional increase in per-capita consumption along the baseline path that yields the same lifetime utility as the policy path. From the household objective (7), the defining equation is

$$\sum_{t=0}^{T} L_t \, \beta^t \, u\left((1+\lambda) \, \frac{C_t^{\text{base}}}{L_t} \right) = \sum_{t=0}^{T} L_t \, \beta^t \, u\left(\frac{C_t^{\text{policy}}}{L_t} \right), \tag{D.7}$$

where $u(x) = (x^{1-\theta} - 1)/(1-\theta)$.[39]

For $\theta = 1$ (log utility), equation (D.7) admits the closed-form solution

$$\lambda = \exp\left(\frac{U^{\text{policy}} - U^{\text{base}}}{\sum_{t=0}^{T} L_t \, \beta^t} \right) - 1, \tag{D.8}$$

where $U = \sum_{t=0}^{T} L_t \, \beta^t \log(C_t/L_t)$. For general θ, the LEV is computed numerically from (D.7).

[39] The household objective (7) is defined over an infinite horizon. Because the transition path is computed over a finite horizon for computational reasons (Section D.1), we evaluate the LEV over the first $T = 350$ years of the simulated path. The discounting ($\beta < 1$) and the near-convergence to the new BGP ensure that the truncation error is negligible.