AI as a Recursive Research Input: Growth Regimes, Singularity Thresholds, and Optimal Policy

Taoufik Rajhi

Department of Economics, Université de Poitiers, France.

Contributing authors: taoufik.rajhi@gmail.com; taoufik.rajhi@univ-poitiers.fr;

Abstract

Can artificial intelligence alter the qualitative structure of long-run economic growth? Standard semi-endogenous and task-based automation models answer no by construction: AI reallocates tasks and raises productivity levels, but the asymptotic law of motion of knowledge remains linear, and the long-run growth rate converges to a finite constant pinned by demographics. We establish this as a structural impossibility result: whenever idea production exhibits diminishing returns in the knowledge stock and AI enters only as a task substitute or level shifter in research productivity, finite-time divergence of the growth rate cannot arise along any equilibrium path with bounded factor shares.

We identify the minimal modification required to overturn this result. Modeling AI as an accumulable research factor whose efficiency depends on cumulative compute introduces a recursive amplification channel in idea production. The growth rate of knowledge then satisfies a Riccati-type differential equation with a strictly positive quadratic term whenever the AI research share exceeds a closed-form threshold determined by four structural primitives: the degree of diminishing returns in knowledge accumulation, the research-input elasticity, the compute-feedback elasticity, and the deployment-congestion parameter. Below the threshold, diminishing returns dominate and growth converges to a constant. Above it, the reduced-form system admits finite-time divergence with blow-up time derivable in closed form.

The mechanism is embedded in a decentralized general equilibrium with monopolistic competition, endogenous labor allocation, and competitive investment in physical and AI capital. This structure allows us to characterize regime crossing as an equilibrium outcome rather than a parametric assumption. Welfare analysis reveals two opposing externalities: compute accumulation generates a dynamic learning spillover that the market underinternalizes, while large-scale deployment creates congestion that it fails to correct. The competitive equilibrium can therefore lie on either side of the regime boundary relative to the social optimum, and the planner reaches the singularity strictly earlier when the learning externality dominates. Optimal policy requires a compute subsidy combined with a deployment tax, with magnitudes tied directly to the structural primitives governing the threshold.

Keywords: Artificial Intelligence, Endogenous Growth, Technological Change, Automation, Research and Development, Economic Singularity

JEL Classification: O30 , O40 , O47 , E13 , C63

The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

1 Introduction

The possibility that artificial intelligence might transform the *structure* of long-run economic growth—not merely its level—has become one of the central questions in macroeconomics. Yet the dominant theoretical frameworks offer no mechanism through which such a transformation could occur. This paper identifies why, and provides the minimal structural modification required to make growth-regime transitions an equilibrium outcome.

The argument begins with a restriction embedded in two otherwise distinct classes of models. In semi-endogenous growth theory (Jones, 1995, 2002), the long-run growth rate is pinned by population dynamics. Diminishing returns in idea accumulation eliminate permanent scale effects: absent demographic expansion, the growth rate of knowledge converges to a finite constant, and level changes in research productivity alter only transitional dynamics. In task-based automation frameworks (Acemoglu and Restrepo, 2018, 2020), AI reallocates tasks between labor and machines, reshaping factor shares, wages, and the distribution of income. These models illuminate much that is empirically important. But neither framework allows AI to alter the *curvature* of the idea-production technology. Explosive or accelerating growth is ruled out by construction, not by evidence.

We formalize this restriction as an impossibility result. In any economy where idea production exhibits diminishing returns in the knowledge stock and AI enters only as a task substitute or as a level shifter in research productivity, the growth rate of ideas satisfies a linear differential equation in equilibrium. Along any path with bounded factor shares, $g_A(t)$ converges to a finite constant. Finite-time divergence cannot arise. The limitation is structural.

The minimal modification required to overturn this result is to model AI as an *accumulable research factor* whose efficiency evolves with cumulative compute. Under this specification, compute expands with AI investment; deployment scales with research activity; and research activity determines frontier growth. This feedback introduces a recursive amplification channel: AI investment today raises the future marginal productivity of AI in research. The dynamics of idea growth acquire a quadratic term and become nonlinear.

Under a tractable scaling approximation valid in the AI-intensive region, the growth rate of knowledge satisfies a Riccati-type differential equation:

$$\dot{g}_A = g_A \left[\frac{\lambda(1-\psi)n}{1-\phi} + \frac{\lambda\psi(1+\beta-\gamma)+\phi-1}{1-\phi}\, g_A \right] + o\big(g_A^2\big)\,,$$

where ψ is the AI share in effective research, λ is the research-input elasticity, ϕ governs diminishing returns in knowledge, β is the compute-feedback elasticity, and γ is the deployment-congestion parameter. The sign of the quadratic coefficient determines the growth regime.

We derive a closed-form threshold for the AI research share,

$$\psi^{**} = \frac{1-\phi}{\lambda(1+\beta-\gamma)},$$

which separates bounded semi-endogenous dynamics from accelerating growth. Below the threshold, diminishing returns dominate recursive amplification and the economy converges to a constant growth rate. Above it, amplification overturns the fishing-out effect and the reduced-form system admits finite-time divergence—a hyperbolic singularity in the growth rate of the knowledge frontier. The threshold depends only on structural primitives: the curvature of idea production, the research-input elasticity, and the net amplification elasticity linking cumulative compute to AI efficiency. It is a parameter-dependent boundary, not a knife-edge calibration coincidence.

The paper embeds this mechanism in a fully specified decentralized general equilibrium. The production side features monopolistic competition in intermediate varieties, endogenous labor allocation between production and research, and competitive investment in both physical and AI capital. Households optimize over consumption and portfolios; patent values are endogenously determined by no-arbitrage across assets; and equilibrium research intensity emerges from wage equalization between sectors. This structure allows us to characterize the conditions under which decentralized incentives place the economy above or below the regime boundary—not as an assumption, but as an equilibrium outcome.

We establish three results in this setting. First, when equilibrium investment drives ψ above ψ^{**}, the reduced-form dynamics imply finite-time divergence of g_A, with the blow-up time given in closed

form by $T = t_0 + \frac{1}{a}\ln\left(1 + \frac{a}{\kappa_\gamma(\psi) g_0}\right)$, where $\kappa_\gamma(\psi)$ summarizes the net feedback strength. Second, incorporating task-based production automation does not alter this threshold structure: production automation changes output levels and AI allocation but cannot generate regime transitions on its own, because it introduces no recursive state variable linking deployment to future research efficiency. The threshold ψ^{**} is preserved once the AI share is correctly measured as the fraction deployed in research. Third, the competitive equilibrium is generically inefficient in the explosive regime. Compute accumulation generates a dynamic learning externality that atomistic agents ignore, while large-scale deployment creates congestion that the market does not internalize. The optimal policy response combines a compute subsidy with a deployment tax, with magnitudes tied directly to proximity to the regime boundary.

The paper is most directly related to the growing literature on transformative AI and long-run growth (Trammell and Korinek, 2023; Jones, 2026; Growiec, 2022; Growiec and Prettner, 2025). These contributions map a wide possibility space for AI-driven growth outcomes, including dramatic acceleration scenarios, and examine countervailing forces such as bottlenecks, existential risk, and scalability limits. Our paper is narrower and more structurally focused. We do not survey possible futures; we derive a single, precise condition under which a given class of general equilibrium models transitions from bounded to accelerating dynamics. The threshold is explicit in measurable elasticities, derived from first principles of idea production, and embedded in a model with competitive asset markets and endogenous allocation. Relative to the task-based automation literature (Acemoglu and Restrepo, 2018, 2020), the contribution is to identify the specific structural ingredient—recursive compute-dependent efficiency—that the dominant automation framework lacks for the question of growth-regime transitions. Relative to semi-endogenous growth theory (Jones, 1995, 2002), the contribution is to augment the idea technology in the minimal way that changes the order of the growth-rate dynamics. The empirical interface is provided by the scaling-laws literature (Kaplan et al, 2020; Hoffmann et al, 2022): the elasticities β and γ that enter the threshold are, in principle, estimable from observed relationships between model performance and cumulative compute.

The mathematical possibility of finite-time divergence is a property of the reduced-form approximation and rests on strong assumptions that are unlikely to hold indefinitely: scaling relations must extend far beyond observed compute ranges; AI efficiency must not saturate; and resource constraints—energy, hardware, coordination frictions—must not bind. The question the paper addresses is therefore not whether an economic singularity will occur in any empirical sense. It is whether standard growth models are theoretically equipped to *speak* to this question. They are not. Once AI recursively augments research productivity through an accumulable efficiency state variable, the answer to whether explosive growth can arise becomes a matter of parameter values rather than functional form.

Section 2 reviews the related literature. Section 3 develops the baseline model, characterizes balanced growth without recursive efficiency feedback, and establishes the impossibility result for level-shifting AI. Section 4 embeds the mechanism in a decentralized general equilibrium and characterizes equilibrium investment allocation and patent valuation. Section 5 derives the Riccati dynamics and the closed-form threshold. Section 6 extends the framework to joint production and research automation and establishes threshold preservation. Section 7 characterizes the social planner's problem, identifies the efficiency wedges, and derives the optimal policy mix. Section 8 concludes.

2 Literature Review

This paper contributes to the macroeconomics of AI by isolating a precise mechanism through which research automation can change the *growth regime*. The key object is not "faster growth" in a loose sense, but a parameter-dependent transition in the law of motion of knowledge. The contribution is therefore closest in spirit to work that studies when technological change alters the structure of long-run dynamics, rather than work that treats AI as a standard productivity shifter.

2.1 Semi-Endogenous Growth and What It Rules Out

The theoretical benchmark is semi-endogenous growth (Jones, 1995, 2002) and related models with diminishing returns in idea accumulation. In that class, the asymptotic growth rate is bounded: level effects and transitions can be large, but the long-run growth law remains pinned down by demographic forces and by the curvature of the idea technology. This is not a calibration outcome. It is a structural implication of the technology restrictions that eliminate scale effects. Our first move

is to make this restriction explicit in a way that clarifies what must be added to generate regime transitions. The paper's "impossibility result" is not a claim that explosive dynamics are empirically unlikely. It is a claim that, under standard assumptions used in both semi-endogenous and task-based automation models, the growth-rate dynamics remain effectively linear, so acceleration cannot arise endogenously.

2.2 Task-Based Automation and Why It Cannot Deliver Regime Transitions

The dominant framework for automation in macroeconomics is task-based: technology changes which tasks are performed by labor versus capital (Zeira, 1998; Acemoglu and Restrepo, 2018, 2020). This approach is empirically and conceptually powerful for understanding wages, factor shares, and "excessive automation." It is also the wrong tool for the question in this paper. The reason is structural. In Acemoglu–Restrepo-type environments, automation reallocates tasks and may raise or lower productivity, but it does not create a recursive state variable that makes the productivity of automation increase with its own deployment. AI is a service flow, or it appears as a task-augmenting shifter, while the knowledge production block remains governed by diminishing returns and by exogenous or semi-endogenous innovation forces. As a consequence, these models can generate large distributional changes and sizable transitional reallocations, but they do not generate parameter-dependent growth-regime transitions originating from research automation. This paper targets precisely the missing element: AI as an accumulable research input whose effectiveness evolves endogenously with cumulative compute and deployment scale. That feedback changes the *order* of the growth-rate dynamics and creates a sharp threshold separating bounded semi-endogenous growth from accelerating dynamics.

2.3 Recent Analyses of Transformative AI and Growth

A growing literature explores the macroeconomic implications of transformative or superhuman AI (Trammell and Korinek, 2023; Jones, 2026; Growiec, 2022; Growiec and Prettner, 2025). Trammell and Korinek (2023) provide a comprehensive taxonomy of possible outcomes for output growth, wages, and labor shares under different assumptions about automation scalability, substitution elasticity, and the scope of tasks AI can perform. Their analysis maps a wide possibility space, including dramatic acceleration scenarios. Jones (2026) examines extreme cases, modeling AI as an advanced form of automation and assessing both explosive growth possibilities and countervailing forces such as bottlenecks and existential risks. Related work by Jones (2025) quantifies optimal investment in reducing AI-related catastrophic risks in light of potential growth acceleration. Growiec (2022) places AI in a long-run perspective of technological progress over human history, emphasizing shifts from partial to full task automatability. More recently, Growiec and Prettner (2025) develop scenarios linking transformative AI to existential risk ("p(doom)"), ranging from human extinction to post-scarcity abundance, with explicit trade-offs between growth and safety. Our paper is narrower and more structurally focused than these contributions. We derive a closed-form regime threshold within a semi-endogenous idea-production framework augmented by endogenous AI efficiency feedback from cumulative compute. The threshold is explicit in primitives (scaling elasticities, diminishing returns in ideas, congestion parameters) and is an equilibrium outcome governed by decentralized investment, labor allocation, and asset pricing. This provides a disciplined general equilibrium setting to evaluate whether incentives push the economy across the boundary and how policy can influence proximity to it.

2.4 Scaling Laws, Bottlenecks, and the Empirical Interface

The empirical scaling-laws literature documents systematic relationships between performance and compute (Kaplan et al, 2020; Hoffmann et al, 2022). Related assessments emphasize that strong acceleration may require stringent conditions and may be moderated by bottlenecks and constraints. We use these findings directly: scaling elasticities and congestion/bottleneck forces enter the threshold expression and determine regime feasibility. Complementary analyses highlight bottlenecks in research tasks, AI capabilities, and data/energy constraints as key moderators of growth implications (Jones, 2026; Growiec and Prettner, 2025). Our model formalizes these forces in a tractable equilibrium structure with explicit welfare and policy implications.

2.5 What This Paper Adds

Relative to the existing literature, the contribution is threefold. First, we establish a structural limitation of canonical task-based automation models for the specific question of accelerating growth from research automation. Second, we provide a minimal general equilibrium extension in which AI is an accumulable research factor with compute-dependent efficiency, and we derive a closed-form threshold that separates bounded semi-endogenous growth from accelerating dynamics. Third, we embed the mechanism in a decentralized equilibrium with asset pricing, investment competition between physical and AI capital, and endogenous labor allocation, allowing us to characterize wedges, externalities, and policy instruments in terms of proximity to the regime boundary. The object is therefore not "singularity economics" as a prediction. It is a parameter-dependent regime transition in growth dynamics, derived from a tractable equilibrium structure and directly tied to measurable elasticities.

3 Baseline Model

This section develops the core theoretical framework. The model builds on semi-endogenous growth à la Jones (1995, 2002) and introduces AI capital as a distinct input into idea production. To ensure ideas have a well-defined private value in the decentralized economy, idea creation is embedded in a Romer-style variety-expansion environment in which each new design yields a patent-protected profit flow in the intermediate sector.

3.1 Final Output Production

A perfectly competitive final-good sector produces a single good $Y(t)$ using physical capital $K(t)$, production labor $L_Y(t)$, and a composite of differentiated intermediate inputs $X(t)$:

$$Y(t) = K(t)^{\alpha}\, L_Y(t)^{1-\alpha-\chi}\, X(t)^{\chi}, \qquad \alpha \in (0,1),\ \chi \in (0, 1-\alpha). \tag{1}$$

This specification imposes constant returns to scale in the privately chosen inputs (K, L_Y, X), so the competitive firm problem is well-defined and yields zero profits in equilibrium.

The intermediate composite is a CES aggregator over a continuum of varieties indexed by $i \in [0, A(t)]$,

$$X(t) = \left(\int_0^{A(t)} x(i,t)^{\frac{\varepsilon-1}{\varepsilon}}\, di \right)^{\frac{\varepsilon}{\varepsilon-1}}, \qquad \varepsilon > 1. \tag{2}$$

The stock $A(t)$ is the measure of available designs.

Total population $L(t)$ is allocated between production and research,

$$L(t) = L_Y(t) + L_A(t), \tag{3}$$

and evolves exogenously at rate $n \geq 0$,

$$\frac{\dot{L}(t)}{L(t)} = n. \tag{4}$$

Physical capital accumulates through investment net of depreciation,

$$\dot{K}(t) = I_K(t) - \delta_K K(t), \qquad \delta_K > 0. \tag{5}$$

We define the gross rental rate on capital as $r_K^g(t)$ and the net return as

$$r_K(t) \equiv r_K^g(t) - \delta_K. \tag{6}$$

3.2 Idea Production with AI Augmentation

New designs are created in a competitive R&D sector by combining research labor $L_A(t)$ with AI research capital deployed in R&D, denoted $M_A(t)$:

$$\dot{A}(t) = \delta \Big[L_A(t) + \eta\big(M(t), \mathcal{C}(t)\big)\, M_A(t) \Big]^{\lambda} A(t)^{\phi}, \qquad \delta > 0,\ \lambda \in (0,1],\ \phi \in (0,1). \tag{7}$$

AI efficiency depends on cumulative compute and deployment scale,

$$\eta\big(M(t),\mathcal{C}(t)\big) = \bar{\eta}\,\mathcal{C}(t)^{\beta}\left(1+\frac{M(t)}{M_0}\right)^{-\gamma}, \qquad \bar{\eta}>0,\ \beta\geq 0,\ \gamma\geq 0,\ M_0>0. \tag{8}$$

3.3 AI Capital and Compute Accumulation

AI capital accumulates through investment $I_M(t)$ measured in units of the final good,

$$\dot{M}(t) = \xi I_M(t) - \delta_M M(t), \qquad \xi>0,\ \delta_M\geq 0, \tag{9}$$

and cumulative compute evolves as

$$\dot{\mathcal{C}}(t) = \theta I_M(t), \qquad \theta>0. \tag{10}$$

Final output is allocated to aggregate consumption $C(t)$, investment in physical and AI capital, and final-good resources used in intermediate production:

$$Y(t) = C(t) + I_K(t) + I_M(t) + \int_0^{A(t)} x(i,t)\,di, \tag{11}$$

where $\int_0^{A(t)} x(i,t)\,di = \int_0^{A(t)} z(i,t)\,di$ by (23). Define investment shares $s_K(t) \equiv I_K(t)/Y(t)$ and $s_M(t) \equiv I_M(t)/Y(t)$.

4 Decentralized Market Economy

We now embed the technologies above in a decentralized economy with competitive final-good production, monopolistic competition in intermediate goods, and competitive R&D and AI-capital installation. Patents grant market power over intermediate varieties, so each new design has a privately appropriable value. This structure endogenizes the market value of ideas and closes the decentralized equilibrium.

4.1 Households

A representative household supplies labor inelastically, owns physical capital, installed AI capital, and the patent portfolio, and chooses consumption to maximize lifetime utility,

$$\max_{\{c(t)\}_{t\geq 0}} \int_0^{\infty} e^{-\rho t}\,L(t)\,\ln c(t)\,dt, \qquad \rho>0, \tag{12}$$

where $c(t)$ is per-capita consumption and

$$C(t) = L(t)\,c(t) \tag{13}$$

is aggregate consumption.

4.1.1 Assets and budget constraint

Let $k(t) \equiv K(t)/L(t)$ denote physical capital per capita and $m(t) \equiv M(t)/L(t)$ denote installed AI capital per capita. Let $p_M(t)$ be the price of one unit of installed AI capital in units of the final good. Let $V_A(t)$ denote the market value of one patent (one unit of the design index). Per-capita wealth is

$$a(t) = k(t) + p_M(t)m(t) + V_A(t)\frac{A(t)}{L(t)}. \tag{14}$$

Let $r_K(t)$ be the net return on physical capital. Installed AI capital is rented to R&D at rate $q_M(t)$ per unit of M. Patents yield a profit flow $\pi(t)$ per design. The household budget constraint is

$$\dot{a}(t) = r_K(t)\,k(t) + \left(\frac{q_M(t)}{p_M(t)} + \frac{\dot{p}_M(t)}{p_M(t)}\right)p_M(t)m(t) + \left(\frac{\pi(t)}{V_A(t)} + \frac{\dot{V}_A(t)}{V_A(t)}\right)V_A(t)\frac{A(t)}{L(t)} + w(t) - c(t), \tag{15}$$

where $w(t)$ is the wage.

4.1.2 Optimality conditions

The household optimality conditions are

$$\frac{\dot{c}(t)}{c(t)} = r(t) - \rho - n, \tag{16}$$

together with no-arbitrage across assets held in positive amounts,

$$r_K(t) = \frac{q_M(t)}{p_M(t)} + \frac{\dot{p}_M(t)}{p_M(t)} - \delta_M. \tag{17}$$

$$r_K(t) = \frac{\pi(t)}{V_A(t)} + \frac{\dot{V}_A(t)}{V_A(t)}. \tag{18}$$

4.2 Final Output Production

Final-good firms take prices as given and solve

$$\max_{K, L_Y, X} \Big\{ K^{\alpha} L_Y^{1-\alpha-\chi} X^{\chi} - r_K^g(t) K - w(t) L_Y - P_X(t) X \Big\}, \tag{19}$$

where $P_X(t)$ is the price index dual to (2). The first-order conditions imply

$$r_K^g(t) = \alpha \frac{Y(t)}{K(t)}, \tag{20}$$

$$w(t) = (1 - \alpha - \chi) \frac{Y(t)}{L_Y(t)}, \tag{21}$$

$$P_X(t) = \chi \frac{Y(t)}{X(t)}. \tag{22}$$

The net return is $r_K(t) = r_K^g(t) - \delta_K$, consistent with (6).

4.3 Intermediate Goods and Private Value of Ideas

Each design $i \in [0, A(t)]$ is owned by a monopolistic intermediate producer. Production of each intermediate variety uses the final good as the variable input with constant marginal cost normalized to one:

$$x(i,t) = z(i,t). \tag{23}$$

The CES demand implied by (2) delivers constant-markup pricing,

$$p_x(i,t) = \mu, \qquad \mu \equiv \frac{\varepsilon}{\varepsilon - 1}. \tag{24}$$

Using (22), total spending on intermediates satisfies $P_X(t)X(t) = \chi Y(t)$. Under symmetry, each variety earns revenue $\chi Y(t)/A(t)$, so operating profit per design is

$$\pi(t) = \frac{\mu - 1}{\mu} \frac{\chi Y(t)}{A(t)}. \tag{25}$$

The asset value of a patent $V_A(t)$ satisfies

$$r_K(t)\, V_A(t) = \pi(t) + \dot{V}_A(t). \tag{26}$$

4.4 R&D firms and endogenous labor allocation

Competitive R&D firms create new designs using research labor $L_A(t)$ and rented AI research capital $M_A(t)$,

$$\dot{A}(t) = \delta \Big[L_A(t) + \eta\big(M(t), \mathcal{C}(t)\big)\, M_A(t) \Big]^{\lambda} A(t)^{\phi}. \tag{27}$$

They take $\big(A(t), M(t), \mathcal{C}(t)\big)$ and prices as given and solve

$$\max_{L_A, M_A} \Big\{ V_A(t)\, \dot{A}(t) - w(t) L_A(t) - q_M(t) M_A(t) \Big\}. \tag{28}$$

Define effective research input

$$R(t) \equiv L_A(t) + \eta(t) M_A(t), \qquad \eta(t) \equiv \eta\big(M(t), \mathcal{C}(t)\big). \tag{29}$$

The first-order conditions are

$$w(t) = V_A(t)\, \delta\, \lambda\, R(t)^{\lambda - 1} A(t)^{\phi}, \tag{30}$$

$$q_M(t) = V_A(t)\, \delta\, \lambda\, \eta(t)\, R(t)^{\lambda - 1} A(t)^{\phi}, \tag{31}$$

implying

$$\frac{q_M(t)}{w(t)} = \eta(t). \tag{32}$$

Labor allocation is pinned down by wage equalization between production and research:

$$(1 - \alpha - \chi)\, \frac{Y(t)}{L(t) - L_A(t)} = V_A(t)\, \delta\, \lambda \Big[L_A(t) + \eta(t) M_A(t) \Big]^{\lambda - 1} A(t)^{\phi}. \tag{33}$$

4.5 AI-Capital Installation with Constant Marginal Cost

Assume a competitive installation sector transforms final-good investment into installed AI capital at constant rate ξ. Installed AI capital evolves according to

$$\dot{M}(t) = \xi I_M(t) - \delta_M M(t). \tag{34}$$

Let $p_M(t)$ denote the unit price of installed AI capital. With linear installation and perfect competition,

$$p_M(t) = \frac{1}{\xi}, \tag{35}$$

so $\dot{p}_M(t) = 0$ and (17) reduces to

$$r_K(t) = \frac{q_M(t)}{p_M} - \delta_M. \tag{36}$$

4.6 Time-Varying AI-Capital Prices via Adjustment Costs or Installation Technology:

Allowing the unit price of installed AI capital $p_M(t)$ to vary over time requires an additional real margin that makes the marginal cost of installation state-dependent. A convenient closure is to introduce convex adjustment costs in installation. Let gross installation be $\xi I_M(t)$, and let the effective net installation be reduced by a convex cost $\Psi(\cdot)$,

$$\dot{M}(t) = \xi I_M(t) - \Psi\left(\frac{I_M(t)}{M(t)} \right) M(t) - \delta_M M(t), \qquad \Psi'(\cdot) > 0,\ \Psi''(\cdot) > 0. \tag{37}$$

Under perfect competition, the unit price of installed AI capital equals the marginal cost of producing an additional unit of $\dot{M}(t)$. With (37), this marginal cost depends on the installation intensity $I_M(t)/M(t)$, so $p_M(t)$ becomes time-varying. Equivalently, one can model a separate production function for installed AI capital, for example

$$\dot{M}(t) + \delta_M M(t) = \Xi\big(I_M(t), Z_M(t)\big), \qquad \Xi_{I_M} > 0,\ \Xi_{I_M I_M} < 0, \tag{38}$$

where $Z_M(t)$ captures installation capacity or sectoral conditions. In either case, a moving $p_M(t)$ is justified because the marginal unit cost of installation varies with states or controls. Absent such a mechanism, treating $p_M(t)$ as time-varying in decentralized no-arbitrage conditions is not internally grounded, and the constant-price closure (35) is the consistent benchmark.

4.7 Market clearing and equilibrium

Markets clear in labor, AI capital used in R&D, and the final good:

$$L(t) = L_Y(t) + L_A(t), \tag{39}$$
$$M_A(t) = M(t), \tag{40}$$
$$Y(t) = C(t) + I_K(t) + I_M(t) + \int_0^{A(t)} x(i,t)\, di. \tag{41}$$

State variables evolve according to

$$\dot{K}(t) = I_K(t) - \delta_K K(t), \tag{42}$$
$$\dot{M}(t) = \xi I_M(t) - \delta_M M(t), \tag{43}$$
$$\dot{\mathcal{C}}(t) = \theta I_M(t), \tag{44}$$
$$\dot{A}(t) = \delta \Big[L_A(t) + \eta\big(M(t), \mathcal{C}(t)\big)\, M_A(t) \Big]^{\lambda} A(t)^{\phi}. \tag{45}$$

Definition 1 (Competitive equilibrium)**.** *Given initial conditions* $\big(K_0, M_0, A_0, \mathcal{C}_0\big)$*, a* ***competitive equilibrium*** *is a set of allocations*

$$\{c(t), C(t), K(t), M(t), A(t), \mathcal{C}(t), L_Y(t), L_A(t), M_A(t), I_K(t), I_M(t)\}_{t\geq 0}$$

and prices

$$\{w(t), r_K^g(t), r_K(t), P_X(t), q_M(t), p_M(t), V_A(t)\}_{t\geq 0}$$

such that: (i) households satisfy (12)*,* (15)*, and* (16)*–*(18)*; (ii) final-good firms satisfy* (20)*–*(22)*; (iii) intermediate firms satisfy* (24) *and earn profits* $\pi(t)$ *given by* (25)*; (iv) patent values satisfy* (26)*; (v) R&D firms satisfy* (30)*–*(31)*; (vi) AI-capital pricing is consistent with either* (35) *or the installation closure in* (37)*; (vii) markets clear* (39)*–*(41)*; and (viii) laws of motion* (42)*–*(45) *hold.*

4.8 Equilibrium Investment Allocation

Under competitive installation with a linear conversion technology, one unit of the final good investment installs ξ units of AI capital. The replacement cost of one unit of installed AI capital is therefore constant,

$$p_M = \frac{1}{\xi}, \qquad \dot{p}_M(t) = 0. \tag{46}$$

Installed AI capital is an asset. Holding one unit yields a rental payment $q_M(t)$ and entails a flow depreciation loss $\delta_M p_M$ in replacement-value terms. Arbitrage with physical capital implies that the required net return on the AI asset equals the equilibrium net return $r_K(t) = r_K^g(t) - \delta_K$. The corresponding user-cost condition is

$$r_K(t)\, p_M = q_M(t) + \dot{p}_M(t) - \delta_M p_M, \tag{47}$$

so with $\dot{p}_M(t) = 0$,

$$q_M(t) = p_M\big(r_K(t) + \delta_M\big). \tag{48}$$

Combining (48) with the R&D first-order condition for AI-capital demand,

$$q_M(t) = V_A(t)\, \delta\, \lambda\, \eta(t)\, R(t)^{\lambda-1} A(t)^{\phi}, \tag{49}$$

yields the equilibrium linkage between the required return and the marginal value product of installed AI capital in research,

$$r_K(t) + \delta_M = \frac{1}{p_M}\, V_A(t)\, \delta\, \lambda\, \eta(t)\, R(t)^{\lambda-1} A(t)^{\phi}. \tag{50}$$

Patent values are pinned down endogenously by

$$r_K(t)\, V_A(t) = \pi(t) + \dot{V}_A(t), \tag{51}$$

so (50) is determined jointly with profits and the equilibrium return.

Together with the final-good optimality conditions

$$r_K^g(t) = \alpha \frac{Y(t)}{K(t)}, \qquad w(t) = (1-\alpha-\chi)\frac{Y(t)}{L(t)-L_A(t)}. \qquad r_K(t) = r_K^g(t) - \delta_K, \tag{52}$$

and the R&D labor demand condition

$$w(t) = V_A(t)\, \delta\, \lambda\, R(t)^{\lambda-1} A(t)^{\phi}, \tag{53}$$

equations (50) and labor-market clearing $L(t) = L_Y(t) + L_A(t)$ determine the equilibrium paths of $r_K(t)$, $q_M(t)$, and the allocation of labor between production and research. Given these prices and allocations, the goods-market constraint (11) and the laws of motion for $(K(t), M(t))$ pin down investment flows.

Investment shares are defined by

$$s_K(t) \equiv \frac{I_K(t)}{Y(t)}, \qquad s_M(t) \equiv \frac{I_M(t)}{Y(t)}. \tag{54}$$

Using capital accumulation,

$$\dot{K}(t) = I_K(t) - \delta_K K(t), \qquad \dot{M}(t) = \xi I_M(t) - \delta_M M(t), \tag{55}$$

we can rewrite investment shares as

$$s_K(t) = \frac{\dot{K}(t) + \delta_K K(t)}{Y(t)}, \tag{56}$$

$$s_M(t) = \frac{\dot{M}(t) + \delta_M M(t)}{\xi\, Y(t)}. \tag{57}$$

Consider a constant-intensity path with constant $K(t)/Y(t)$ and $M(t)/Y(t)$, and let $g_Y(t) \equiv \dot{Y}(t)/Y(t)$. Constancy of K/Y implies $\dot{K}/K = g_Y$ (and similarly $\dot{M}/M = g_Y$). Hence

$$s_K^* = (g_Y + \delta_K)\left(\frac{K}{Y}\right)^*, \tag{58}$$

$$s_M^* = \frac{1}{\xi}(g_Y + \delta_M)\left(\frac{M}{Y}\right)^*. \tag{59}$$

The implied investment-share ratio is

$$\frac{s_M^*}{s_K^*} = \frac{g_Y + \delta_M}{g_Y + \delta_K} \cdot \frac{(M/Y)^*}{(K/Y)^*} \cdot \frac{1}{\xi}. \tag{60}$$

The intensities $(M/Y)^*$ and $(K/Y)^*$ are jointly pinned down by the return condition (50), patent valuation (51), and production-side conditions (52), without imposing an exogenous research-labor share.

5 Balanced Growth Path Characterization

A balanced growth path (BGP) is a trajectory along which the growth rates of the key state variables are constant and the relevant factor intensities converge. For any $X(t)$, define $g_X(t) \equiv \dot{X}(t)/X(t)$.

A balanced growth path is a path along which the growth rates of the state variables are constant and the relevant intensity ratios are stationary. In this model the state vector is (K, M, A, L), with $g_L = n$ exogenous. The remaining growth rates are tied down by three stationarity requirements.

(i) Stationary physical-capital intensity.

If the aggregate physical-capital intensity is stationary on the BGP, $K/Y = \text{const}$, then

$$g_K^* = g_Y^*. \tag{61}$$

(ii) Stationary AI-capital intensity.

If the AI-capital intensity is stationary, $M/Y = \text{const}$, then

$$g_M^* = g_Y^*, \tag{62}$$

which is the closure used in (**??**). When this closure is not imposed, M/Y generally drifts and

$$g_{M/Y}^* \equiv g_M^* - g_Y^* \tag{63}$$

summarizes the long-run trend in AI intensity.

(iii) Stationary research intensities and the growth of effective research.

With convergent allocation shares, $L_A = \sigma L$ with $\sigma \to \sigma^*$, so $g_{L_A}^* = n$. Define the AI research share $\psi \equiv \eta M/R$ with $R = L_A + \eta M$. Then the growth rate of effective research satisfies the exact decomposition:

$$g_R^* = (1 - \psi^*)\, n + \psi^* \big(g_\eta^* + g_M^*\big). \tag{64}$$

BGP growth of AI efficiency η.

To use (64) one needs an expression for $g_\eta^* \equiv \dot{\eta}/\eta$. Let AI efficiency be the general power form used in the paper,

$$\eta(t) = \bar{\eta} \left(\frac{M(t)}{\mathcal{C}(t)} \right)^{\beta} M(t)^{-\gamma}, \qquad \bar{\eta} > 0,\ \beta \geq 0,\ \gamma \geq 0, \tag{65}$$

where $\mathcal{C}(t)$ denotes the congestion (or deployment) object appearing in the baseline specification. Log-differentiating (65) yields the BGP identity

$$g_\eta^* = \beta \big(g_M^* - g_{\mathcal{C}}^*\big) - \gamma g_M^* = (\beta - \gamma) g_M^* - \beta g_{\mathcal{C}}^*. \tag{66}$$

In the constant-efficiency benchmark ($\beta = \gamma = 0$), $g_\eta^* = 0$ and (64) reduces to

$$g_R^* = (1 - \psi^*)\, n + \psi^*\, g_M^*. \tag{67}$$

If, in addition, congestion is proportional to deployed AI per researcher (the common benchmark in the main text), $\mathcal{C}(t) \propto M(t)/L_A(t)$, then $g_{\mathcal{C}}^* = g_M^* - g_{L_A}^* = g_M^* - n$, and (66) simplifies to

$$g_\eta^* = (\beta - \gamma) g_M^* - \beta (g_M^* - n) = \beta n - \gamma g_M^*. \tag{68}$$

Equations (66)–(68) can be inserted directly into (64) to express g_R^* and hence g_A^* in terms of (g_M^*, n) and the AI-efficiency parameters.

Implications for idea growth and output growth.

Given a well-defined BGP, the idea technology implies $g_A^* = \lambda g_R^*/(1 - \phi)$ from (**??**). Output growth then follows from the CES variety block and the final-good technology:

$$g_Y^* = n + \frac{\chi}{(1 - \alpha - \chi)(\varepsilon - 1)}\, g_A^*. \tag{69}$$

Equations (61)–(69) make explicit how the stationarity restrictions on intensities map idea growth into the joint BGP growth rates of (K, M, Y), and how the composition of research inputs maps AI accumulation into the long-run growth of knowledge through (64).

We begin with the benchmark in which AI does not enter idea production. This case reproduces the semi-endogenous growth logic in Jones (1995, 2002) and provides the reference point for the AI-augmented economy.

Proposition 1 (Balanced growth without AI)**.** *Suppose $\bar{\eta} = 0$, so that AI does not contribute to research. Assume $\dot{L}(t)/L(t) = n \geq 0$ and that along the BGP the research share converges as in (87). If $\phi < 1$, there exists a balanced path on which the long-run growth rate of ideas is*

$$g_A^* = \frac{\lambda n}{1-\phi}. \tag{70}$$

Along this path, the long-run idea growth rate depends on n and (λ, ϕ) but not on the level parameter δ.

Proof. With $\bar{\eta} = 0$, idea production reduces to

$$\dot{A}(t) = \delta\, L_A(t)^{\lambda} A(t)^{\phi}.$$

Dividing by $A(t)$ yields

$$g_A(t) \equiv \frac{\dot{A}(t)}{A(t)} = \delta\, L_A(t)^{\lambda} A(t)^{\phi-1}. \tag{71}$$

Let $\sigma(t) \equiv L_A(t)/L(t)$. By (87), $\sigma(t) \to \sigma$ and $L(t) = L_0 e^{nt}$, so

$$\frac{\dot{L}_A(t)}{L_A(t)} = \frac{\dot{\sigma}(t)}{\sigma(t)} + \frac{\dot{L}(t)}{L(t)} \to 0 + n = n.$$

Log-differentiating (71) gives

$$\frac{\dot{g}_A(t)}{g_A(t)} = \lambda \frac{\dot{L}_A(t)}{L_A(t)} + (\phi - 1)\frac{\dot{A}(t)}{A(t)} = \lambda n + (\phi - 1) g_A(t).$$

On a balanced path, $g_A(t)$ is constant, so $\dot{g}_A(t) = 0$. Hence

$$0 = \lambda n + (\phi - 1) g_A^*,$$

which implies (70). Since $\phi < 1$, the fixed point is finite. □

Equation (70) is the semi-endogenous implication: diminishing returns in knowledge accumulation ($\phi < 1$) eliminate permanent scale effects. Level shifts in research productivity δ affect levels and transition dynamics, but not the long-run growth rate. In the AI-augmented economy, the central issue is whether AI changes the asymptotic growth of the effective research input entering (7), rather than acting as a pure level shifter.

We now turn to environments in which AI contributes to research but does not improve endogenously. In this benchmark, AI affects idea growth only through the scale of effective research input, not through recursive feedback.

A caveat matters for the decentralized model. In equilibrium, $M(t)$ is an accumulable asset with law of motion $\dot{M} = \xi I_M - \delta_M M$. Hence g_M is not an exogenous primitive. It is pinned down by household saving behavior, the no-arbitrage condition (47), and the goods-market resource constraint (e.g. (11)). The proposition below is therefore best read as a conditional BGP characterization: given a candidate balanced path along which M grows at constant rate g_M, what is the induced long-run idea growth rate?

Proposition 2 (Balanced growth with non-improving AI)**.** *Suppose $\beta = 0$ and $\gamma = 0$, so that $\eta(M, \mathcal{C}) = \bar{\eta} > 0$ is constant. Assume $\dot{L}(t)/L(t) = n \geq 0$ and that along the BGP the research share converges as in (87). Consider any balanced path along which AI capital grows at constant rate $g_M \equiv \dot{M}(t)/M(t)$. Define effective research input*

$$R(t) \equiv L_A(t) + \bar{\eta} M(t).$$

If $\phi < 1$, the idea growth rate on such a path satisfies

$$g_A^* = \frac{\lambda g_R^*}{1-\phi}, \qquad g_R^* \equiv \lim_{t\to\infty} \frac{\dot{R}(t)}{R(t)} = \max\{n, g_M\}. \tag{72}$$

Equivalently, if $g_M > n$ then $g_A^* = \lambda g_M/(1-\phi)$; if $g_M < n$ then $g_A^* = \lambda n/(1-\phi)$; and if $g_M = n$ then $g_A^* = \lambda n/(1-\phi)$, with $\psi(t) \equiv \bar{\eta} M(t)/R(t)$ converging whenever $M(t)/L_A(t)$ converges.

Proof. With $\eta(M, \mathcal{C}) = \bar{\eta}$, the idea technology is

$$\dot{A}(t) = \delta\, R(t)^{\lambda} A(t)^{\phi}, \qquad R(t) = L_A(t) + \bar{\eta} M(t),$$

with constant δ. Along the candidate balanced path, (87) implies $L_A(t) = \sigma(t) L(t)$ with $\sigma(t) \to \sigma \in (0,1)$ and $L(t) \sim e^{nt}$, hence $\dot{L}_A/L_A \to n$. Also $M(t) = M_0 e^{g_M t}$.

Define $\psi(t) \equiv \bar{\eta} M(t)/R(t) \in (0,1)$. Then

$$\frac{\dot{R}(t)}{R(t)} = (1-\psi(t))\,\frac{\dot{L}_A(t)}{L_A(t)} + \psi(t)\,\frac{\dot{M}(t)}{M(t)} = (1-\psi(t))\,n + \psi(t)\,g_M.$$

Moreover,

$$\psi(t) = \frac{1}{1 + \frac{L_A(t)}{\bar{\eta} M(t)}} = \frac{1}{1 + \frac{L_A(0)}{\bar{\eta} M_0}\, e^{(n-g_M)t}},$$

so $\psi(t) \to 1$ if $g_M > n$, $\psi(t) \to 0$ if $g_M < n$, and $\psi(t)$ is constant when $g_M = n$ and $M(t)/L_A(t)$ is constant. Hence $\dot{R}/R \to \max\{n, g_M\}$.

Since $g_A(t) = \dot{A}(t)/A(t) = \delta R(t)^{\lambda} A(t)^{\phi-1}$, log-differentiation yields

$$\frac{\dot{g}_A(t)}{g_A(t)} = \lambda \frac{\dot{R}(t)}{R(t)} + (\phi - 1) g_A(t).$$

On a balanced path, $g_A(t)$ is constant, so $\dot{g}_A(t) = 0$. Substituting $\dot{R}/R \to g_R^*$ gives $0 = \lambda g_R^* + (\phi - 1) g_A^*$, which implies (72). □

The proposition isolates a composition channel. If AI capital expands faster than research labor ($g_M > n$), the AI component $\bar{\eta} M(t)$ dominates $R(t)$ asymptotically and long-run idea growth is governed by g_M rather than n. If $g_M \le n$, AI does not change the asymptotic idea-growth rate, although it raises levels and alters transition dynamics.

In the decentralized economy, the substantive question is whether equilibrium forces can sustain $g_M > n$ on a balanced path. Because M is an accumulable asset, this requires that saving and valuation conditions imply a long-run investment rate into AI capital consistent with (47) and the resource constraint. The next section makes this link explicit by deriving g_M from equilibrium investment shares rather than treating it as a primitive.

Proposition 3 (Automation threshold with deployment congestion)**.** *Suppose $\phi < 1$ and that along a feedback segment the AI research share converges to a constant $\psi \in [0,1]$, i.e. $\psi(t) \to \psi$. Let AI efficiency be given by (8) with elasticities $\beta \ge 0$ and $\gamma > 0$. Consider an AI-dominated scaling region in which (i) deployed AI capital is large relative to the congestion scale, $M(t) \gg M_0$, so that*

$$\frac{d}{dt} \ln\!\left(1 + \frac{M(t)}{M_0}\right) = \frac{M(t)}{M(t) + M_0}\,\frac{\dot{M}(t)}{M(t)} = \frac{\dot{M}(t)}{M(t)} + o\big(g_A(t)\big),$$

(ii) depreciation is negligible relative to growth rates, and (iii) the accumulation of compute and AI capital implies the quasi-steady relations

$$\frac{\dot{\mathcal{C}}(t)}{\mathcal{C}(t)} = g_A(t) + o\big(g_A(t)\big), \qquad \frac{\dot{M}(t)}{M(t)} = g_A(t) + o\big(g_A(t)\big). \tag{73}$$

Then AI efficiency growth satisfies

$$\frac{\dot{\eta}(t)}{\eta(t)} = (\beta - \gamma)\, g_A(t) + o\big(g_A(t)\big), \tag{74}$$

and the reduced-form dynamics of idea growth $g_A(t)$ admit a critical AI research share

$$\psi^*(\beta, \gamma) \equiv \frac{1-\phi}{\lambda\big(1 + \beta - \gamma\big)}. \tag{75}$$

Assume $1+\beta-\gamma>0$. Within the reduced-form approximation, three regimes obtain:

(i) If $\psi<\psi^*(\beta,\gamma)$, then $g_A(t)$ is locally bounded and converges to a finite fixed point (up to lower-order terms).

(ii) If $\psi>\psi^*(\beta,\gamma)$, then for sufficiently large $g_A(t)$ the dynamics imply finite-time divergence (hyperbolic blow-up) in the reduced-form approximation.

(iii) If $\psi=\psi^*(\beta,\gamma)$, then $g_A(t)$ grows without bound but does not blow up in finite time in the reduced-form approximation.

If $1+\beta-\gamma\le 0$, the quadratic feedback term is non-positive for all $\psi\in[0,1]$, so the reduced-form dynamics do not generate a blow-up regime.

Proof. Let

$$R(t)\equiv L_A(t)+\eta(t)M(t),\qquad \psi(t)\equiv\frac{\eta(t)M(t)}{R(t)}\in[0,1].$$

Idea production implies

$$g_A(t)\equiv\frac{\dot A(t)}{A(t)}=\delta\,R(t)^{\lambda}A(t)^{\phi-1}.$$

Log-differentiation yields

$$\frac{\dot g_A(t)}{g_A(t)}=\lambda\,\frac{\dot R(t)}{R(t)}+(\phi-1)\,g_A(t). \tag{76}$$

Differentiate $R(t)=L_A(t)+\eta(t)M(t)$:

$$\dot R(t)=\dot L_A(t)+\dot\eta(t)M(t)+\eta(t)\dot M(t).$$

Divide by $R(t)$ and use $\dot L_A/L_A=n$ together with $\psi(t)=\eta M/R$ to obtain

$$\frac{\dot R(t)}{R(t)}=(1-\psi(t))\,n+\psi(t)\left(\frac{\dot\eta(t)}{\eta(t)}+\frac{\dot M(t)}{M(t)}\right). \tag{77}$$

Next use (8). In the region $M(t)\gg M_0$,

$$\frac{\dot\eta(t)}{\eta(t)}=\beta\,\frac{\dot{\mathcal{C}}(t)}{\mathcal{C}(t)}-\gamma\,\frac{d}{dt}\ln\left(1+\frac{M(t)}{M_0}\right)=\beta\,\frac{\dot{\mathcal{C}}(t)}{\mathcal{C}(t)}-\gamma\,\frac{\dot M(t)}{M(t)}+o\big(g_A(t)\big).$$

Applying (73) gives (74).

Substitute (77) and (74) into (76), use $\psi(t)\to\psi$, and multiply both sides by $g_A(t)$. This yields the Riccati-type reduced form

$$\dot g_A(t)=g_A(t)\Big[\lambda(1-\psi)\,n+\big(\lambda\psi(1+\beta-\gamma)+\phi-1\big)g_A(t)\Big]+o\big(g_A(t)^2\big). \tag{78}$$

Define the quadratic coefficient $b_\gamma(\psi)\equiv\lambda\psi(1+\beta-\gamma)+\phi-1$. The unique ψ such that $b_\gamma(\psi)=0$ is $\psi^*(\beta,\gamma)$ in (75). If $\psi<\psi^*$, then $b_\gamma(\psi)<0$ and the quadratic term is stabilizing, implying convergence of $g_A(t)$ to a finite fixed point (up to the remainder). If $\psi>\psi^*$, then $b_\gamma(\psi)>0$ and for sufficiently large $g_A(t)$ the quadratic term dominates, implying finite-time blow-up by comparison with $\dot x=\frac{1}{2}b_\gamma(\psi)x^2$. If $\psi=\psi^*$, the quadratic term vanishes asymptotically and the remaining linear term implies unbounded growth without finite-time divergence. Finally, if $1+\beta-\gamma\le 0$, then $b_\gamma(\psi)\le\phi-1<0$ for all $\psi\in[0,1]$, so the quadratic feedback is never destabilizing and blow-up cannot occur in the reduced form. □

Corollary 2 (Impossibility of explosive growth with deployment congestion)**.** *Under the conditions of Proposition 3, suppose $1+\beta-\gamma>0$. If $\psi^*(\beta,\gamma)>1$, equivalently*

$$\lambda\big(1+\beta-\gamma\big)<1-\phi, \tag{79}$$

then $b_\gamma(\psi)<0$ for all feasible $\psi\in[0,1]$, and the hyperbolic blow-up regime cannot arise in the reduced-form dynamics.

Proof. From (78), the quadratic coefficient is

$$b_\gamma(\psi) = \lambda\psi\big(1+\beta-\gamma\big) + \phi - 1,$$

which is linear in ψ. Under $1+\beta-\gamma > 0$, $b_\gamma(\psi)$ is weakly increasing in ψ, so its maximum over $[0,1]$ occurs at $\psi = 1$. If (79) holds, then

$$b_\gamma(1) = \lambda\big(1+\beta-\gamma\big) + \phi - 1 < 0,$$

hence $b_\gamma(\psi) \le b_\gamma(1) < 0$ for all $\psi \in [0,1]$. Therefore the quadratic term in (78) is everywhere stabilizing and cannot generate finite-time divergence in the reduced-form approximation. □

The threshold $\psi^*(\beta,\gamma)$ has a direct interpretation. The term $1-\phi$ captures diminishing returns in knowledge accumulation. Lower ϕ strengthens fishing-out forces and raises the required AI share. The denominator $\lambda(1+\beta-\gamma)$ summarizes net amplification in effective research. The elasticity λ scales the response of idea production to effective research input. The compute feedback β raises AI efficiency as cumulative compute expands, while the deployment congestion parameter γ offsets that amplification by reducing efficiency as AI deployment scales up. Finite-time divergence is feasible in the reduced-form region only if net amplification can dominate the knowledge drag.

For later analysis it is convenient to keep the threshold in compact form,

$$\psi^*(\beta,\gamma) = \frac{1-\phi}{\lambda\big(1+\beta-\gamma\big)}, \qquad 1+\beta-\gamma > 0, \tag{80}$$

which makes clear how stronger fishing-out (lower ϕ) raises the required AI share, while stronger research returns (higher λ) or stronger compute feedback (higher β) lowers it, and stronger deployment congestion (higher γ) raises it.

Proposition 4 (Hyperbolic Growth with Deployment Congestion). *Suppose $\beta > 0$, $\gamma > 0$, $\phi < 1$, and the economy sustains a strictly positive AI-investment share $s_M > 0$ along the relevant segment of the transition. Assume the dynamics enter an AI-dominated scaling region in which $M(t) \gg M_0$, the AI research share is approximately constant $\psi(t) \approx \psi \in (0,1]$, depreciation is negligible relative to growth, and the quasi-steady relations hold:*

$$\frac{\dot{\mathcal{C}}(t)}{\mathcal{C}(t)} = g_A(t) + o(g_A(t)), \qquad \frac{\dot{M}(t)}{M(t)} = g_A(t) + o(g_A(t)).$$

Then AI efficiency growth satisfies

$$\frac{\dot{\eta}(t)}{\eta(t)} = (\beta-\gamma)\, g_A(t) + o(g_A(t)),$$

and, if the AI research share satisfies

$$\psi > \psi^*(\beta,\gamma) \equiv \frac{1-\phi}{\lambda(1+\beta-\gamma)}, \qquad 1+\beta-\gamma > 0,$$

the idea-growth rate follows a hyperbolic law once the quadratic (feedback) component dominates the linear (population-driven) component:

$$g_A(t) = \frac{g_0}{1-\kappa_\gamma(\psi)\, g_0\,(t-t_0)}, \qquad \kappa_\gamma(\psi) \equiv \frac{\lambda\psi(1+\beta-\gamma)+\phi-1}{1-\phi},$$

where $g_A(t_0) = g_0 > 0$ at the entry time t_0 into the quadratic-dominance region.

Proof. The proof uses the reduced-form law of motion for g_A under the scaling relations above. Let $g_A(t) \equiv \dot{A}(t)/A(t)$. In the region $M(t) \gg M_0$, differentiating (8) yields

$$\frac{\dot{\eta}(t)}{\eta(t)} = \beta\,\frac{\dot{\mathcal{C}}(t)}{\mathcal{C}(t)} - \gamma\,\frac{d}{dt}\ln\!\left(1+\frac{M(t)}{M_0}\right) = \beta\,\frac{\dot{\mathcal{C}}(t)}{\mathcal{C}(t)} - \gamma\,\frac{\dot{M}(t)}{M(t)} + o(g_A(t)).$$

Using $\dot{\mathcal{C}}/\mathcal{C} = g_A + o(g_A)$ and $\dot{M}/M = g_A + o(g_A)$ gives

$$\frac{\dot{\eta}(t)}{\eta(t)} = (\beta - \gamma)\, g_A(t) + o(g_A(t)).$$

Under the quasi-steady relations and $\psi(t) \approx \psi$, the reduced-form dynamics for g_A take the Riccati form

$$\dot{g}_A = \frac{\lambda(1-\psi)n}{1-\phi}\, g_A + \frac{\lambda\psi(1+\beta-\gamma)+\phi-1}{1-\phi}\, g_A^2, \tag{81}$$

where the quadratic coefficient is

$$\kappa_\gamma(\psi) \equiv \frac{\lambda\psi(1+\beta-\gamma)+\phi-1}{1-\phi}.$$

Fix a time t_0 at which the economy has entered a regime in which (i) $\psi(t)$ is approximately constant at ψ and (ii) the quadratic term dominates the linear term in (81). The dominance requirement holds, for instance, whenever

$$g_A(t) \gg \frac{\lambda(1-\psi)n}{\lambda\psi(1+\beta-\gamma)+\phi-1},$$

and $\lambda\psi(1+\beta-\gamma)+\phi-1 > 0$. Under these conditions, (81) reduces to

$$\dot{g}_A = \kappa_\gamma(\psi)\, g_A^2. \tag{82}$$

Since $\phi < 1$, the denominator of $\kappa_\gamma(\psi)$ is positive. Thus $\kappa_\gamma(\psi) > 0$ if and only if

$$\lambda\psi(1+\beta-\gamma)+\phi-1 > 0 \quad \Longleftrightarrow \quad \psi > \frac{1-\phi}{\lambda(1+\beta-\gamma)} \equiv \psi^*(\beta,\gamma),$$

assuming $1+\beta-\gamma > 0$. Let $\psi > \psi^*(\beta,\gamma)$, so $\kappa_\gamma(\psi) > 0$.

Solve (82) with $g_A(t_0) = g_0 > 0$:

$$\frac{dg_A}{g_A^2} = \kappa_\gamma(\psi)\, dt \quad \Longrightarrow \quad -\frac{1}{g_A(t)} = -\frac{1}{g_0} + \kappa_\gamma(\psi)(t-t_0),$$

which yields

$$g_A(t) = \frac{g_0}{1 - \kappa_\gamma(\psi)\, g_0\, (t-t_0)}.$$

Thus $g_A(t)$ diverges at the finite time

$$T = t_0 + \frac{1}{\kappa_\gamma(\psi)\, g_0},$$

which is well-defined precisely when $\kappa_\gamma(\psi) > 0$, equivalently $\psi > \psi^*(\beta,\gamma)$. □

Corollary 3 (Impossibility of Explosive Growth with Deployment Congestion)**.** *If $\psi^*(\beta,\gamma) > 1$, then $\kappa_\gamma(\psi) \le 0$ for all feasible $\psi \in [0,1]$, so the hyperbolic (finite-time blow-up) regime cannot arise in the reduced-form dynamics. This occurs when*

$$\psi^*(\beta,\gamma) > 1 \quad \Longleftrightarrow \quad \lambda(1+\beta-\gamma) < 1-\phi, \qquad 1+\beta-\gamma > 0.$$

If $1+\beta-\gamma \le 0$, then $\kappa_\gamma(\psi) < 0$ for all $\psi \in [0,1]$ and blow-up is ruled out a fortiori.

Lemma 4 (Quadratic dominance and finite-time blow-up)**.** *Fix $\psi \in (0,1]$ and suppose $\phi < 1$, $n \ge 0$, and $\kappa_\gamma(\psi) > 0$. Consider the Riccati equation*

$$\dot{g}(t) = a\, g(t) + \kappa\, g(t)^2, \qquad a \equiv \frac{\lambda(1-\psi)n}{1-\phi} \ge 0, \qquad \kappa \equiv \kappa_\gamma(\psi) = \frac{\lambda\psi(1+\beta-\gamma)+\phi-1}{1-\phi} > 0, \tag{83}$$

with initial condition $g(t_0) = g_0 > 0$. Then:

(i) The solution exists on $[t_0, T)$ and blows up at a finite time

$$T = t_0 + \frac{1}{a}\ln\left(1 + \frac{a}{\kappa g_0}\right) \quad \textit{(with the convention } T = t_0 + \frac{1}{\kappa g_0} \textit{ if } a = 0\textit{).} \tag{84}$$

(ii) The solution admits the closed form

$$g(t) = \frac{a\, g_0\, e^{a(t-t_0)}}{a - \kappa g_0\left(e^{a(t-t_0)} - 1\right)}. \tag{85}$$

(iii) For any $t \in [t_0, T)$,

$$g(t) \;\geq\; \frac{g_0}{1 - \kappa g_0 (t - t_0)}, \tag{86}$$

so the "pure quadratic" hyperbola provides a valid lower bound and implies finite-time divergence no later than $t_0 + 1/(\kappa g_0)$.

Proof. Define $u(t) \equiv 1/g(t)$. Differentiating yields

$$\dot{u}(t) = -\frac{\dot{g}(t)}{g(t)^2} = -(ag + \kappa g^2)\frac{1}{g^2} = -a\, u(t) - \kappa.$$

Hence u solves the linear ODE $\dot{u} + au = -\kappa$ with initial condition $u(t_0) = 1/g_0$. The solution is

$$u(t) = \left(\frac{1}{g_0} + \frac{\kappa}{a}\right) e^{-a(t-t_0)} - \frac{\kappa}{a} \qquad (a > 0),$$

and inverting gives (85). The blow-up time T is the first time at which $u(T) = 0$, which yields (84). The case $a = 0$ follows by taking limits, giving $g(t) = g_0/[1 - \kappa g_0(t - t_0)]$ and $T = t_0 + 1/(\kappa g_0)$.

For (iii), $a \geq 0$ implies $\dot{g} = ag + \kappa g^2 \geq \kappa g^2$. By comparison for scalar ODEs, $g(t)$ dominates the solution $\underline{g}(t)$ of $\dot{\underline{g}} = \kappa \underline{g}^2$ with $\underline{g}(t_0) = g_0$, which is exactly the hyperbola in (86). □

Remark 5 (When is the quadratic approximation accurate?)**.** *If $a > 0$, then the ratio of the linear to the quadratic term in (83) is*

$$\frac{ag}{\kappa g^2} = \frac{a}{\kappa g}.$$

Thus, once $g \gg a/\kappa$, the dynamics are governed by $\dot{g} \simeq \kappa g^2$ and the hyperbolic formula becomes an accurate approximation rather than merely a lower bound.

When AI nearly dominates research, so that $\psi \approx 1$, the condition for explosive growth simplifies to

$$\lambda(1 + \beta) > 1 - \phi,$$

(or $\lambda(1 + \beta - \gamma) > 1 - \phi$ in the congestion case of Proposition 4). In this regime the feedback loop operates through effective research: faster frontier expansion is accompanied by faster compute accumulation, which raises AI efficiency and increases the marginal product of deployed AI capital in idea creation. The reduced-form dynamics imply growth stronger than exponential, with $g_A(t)$ diverging as $(T - t)^{-1}$. Since $g_A(t) = \frac{d}{dt}\ln A(t)$, this implies

$$A(t) \propto (T - t)^{-1/\kappa}, \qquad \kappa \equiv \kappa(\psi) > 0,$$

so the implied singularity is power-law in the level of the frontier.

The magnitude of $\kappa(\psi)$ determines how quickly the reduced-form singularity is approached. At the knife-edge $\psi = \psi^*$, one has $\kappa = 0$ and the quadratic feedback term vanishes asymptotically, marking a critical transition in the local dynamics. As ψ rises above the threshold, $\kappa(\psi)$ increases and the implied time to singularity $T - t_0 = 1/(\kappa g_0)$ shortens. For example, with $\lambda = 0.5$, $\beta = 0.08$, and $\phi = 0.5$, complete AI dominance ($\psi = 1$) yields $\kappa(1) = \frac{\lambda(1+\beta)+\phi-1}{1-\phi} = \frac{0.5\cdot 1.08 - 0.5}{0.5} = 0.08$, whereas $\psi = 0.95$ yields $\kappa(0.95) = \frac{0.5\cdot 0.95\cdot 1.08 - 0.5}{0.5} = 0.026$, implying a roughly threefold increase in the implied time to singularity for a given g_0.

The mathematical possibility of finite-time divergence is a property of the reduced-form region and rests on strong assumptions. Scaling relations must continue to apply at compute levels far beyond

observed ranges. The economy must reach and sustain AI research shares above the threshold. AI efficiency must not saturate at a finite ceiling, and investment in AI must remain attractive relative to other uses of output. Real constraints such as energy availability, hardware limits, and coordination frictions must not bind early.

Among these assumptions, resource constraints are natural candidates to bind first. As growth rates rise, energy and hardware requirements scale up quickly and may induce endogenous slowdowns through higher installation costs, declining effective ξ, congestion effects (higher γ), or other capacity wedges. The economically relevant question is then not whether the model literally diverges, but how long the phase of rapid acceleration persists, at what growth rates constraints begin to bind, and what long-run regime emerges once those constraints are incorporated. Subsequent sections address these mechanisms explicitly.

In the decentralized economy, labor allocation is endogenous. For the no-AI benchmark, it is useful to recover the classical Jones expression by focusing on balanced paths along which the research labor share converges to a constant. Concretely, suppose that along the BGP,

$$\frac{L_A(t)}{L(t)} \to \sigma \in (0,1), \qquad \frac{\dot{L}(t)}{L(t)} = n \geq 0. \tag{87}$$

This assumption is used to isolate the role of (λ, ϕ) in semi-endogenous idea growth. In the full decentralized model, the limiting share σ is determined by equilibrium wage equalization and patent valuation; here it is treated as a limiting object.

Proposition 5 (Equilibrium allocation and the threshold ψ^{**})**.** *Let $\psi^{**} \equiv \frac{1-\phi}{\lambda(1+\beta-\gamma)}$ denote the threshold from the transition/feedback analysis. Suppose a BGP exists such that the long-run AI research intensity $\psi(t) \to \psi^* \in [0,1]$ and the research labor share $\sigma(t) \to \sigma^* \in [0,1]$.*

(i) The BGP wage-equalization restriction implies a strictly decreasing mapping between σ^* and ψ^*:

$$1 - \psi^* = \frac{1-\alpha-\chi}{\Psi} \cdot \frac{\sigma^*}{1-\sigma^*}, \qquad \Psi \equiv \Phi\, \lambda\, g_A^*, \tag{88}$$

equivalently

$$\sigma^* = \frac{\Psi(1-\psi^*)}{(1-\alpha-\chi) + \Psi(1-\psi^*)}. \tag{89}$$

(ii) Define the implied threshold labor share

$$\sigma^{**} \equiv \frac{\Psi(1-\psi^{**})}{(1-\alpha-\chi) + \Psi(1-\psi^{**})}. \tag{90}$$

Then the equilibrium threshold comparison $\psi^* \gtrless \psi^{**}$ is equivalent to $\sigma^* \lessgtr \sigma^{**}$:

$$\psi^* > \psi^{**} \iff \sigma^* < \sigma^{**}, \qquad \psi^* < \psi^{**} \iff \sigma^* > \sigma^{**}. \tag{91}$$

(iii) Interpreting regimes, $\psi^* < \psi^{**}$ corresponds to an equilibrium allocation that does not sustain the AI-intensive region required by the feedback approximation, while $\psi^* > \psi^{**}$ makes that region feasible in the sense that the long-run allocation itself lies on the "AI-dominant" side of the threshold.

Proof. Part (i) is algebra: start from wage equalization on the BGP, $\frac{1-\alpha-\chi}{1-\sigma^*} = \Phi\lambda g_A^* \frac{1-\psi^*}{\sigma^*}$, and rearrange using $\Psi \equiv \Phi\lambda g_A^*$. The mapping is strictly decreasing because $\sigma^*(\psi)$ in (89) has derivative $\partial\sigma^*/\partial\psi = -\Psi(1-\alpha-\chi)\big[(1-\alpha-\chi)+\Psi(1-\psi)\big]^{-2} < 0$. Part (ii) follows by evaluating (89) at ψ^{**} and using monotonicity. Part (iii) is the threshold interpretation stated in the text. □

Proposition 6 (Joint condition for explosive GE dynamics in the CES–intermediate baseline)**.** *Maintain the baseline technology*

$$Y(t) = K(t)^{\alpha} L_Y(t)^{1-\alpha-\chi} X(t)^{\chi}, \qquad X(t) = \Big(\int_0^{A(t)} x(i,t)^{\frac{\varepsilon-1}{\varepsilon}}\, di \Big)^{\frac{\varepsilon}{\varepsilon-1}},$$

and the idea/AI blocks

$$\dot{A}(t) = \delta\big(L_A(t) + \eta(t) M(t)\big)^{\lambda} A(t)^{\phi}, \qquad \eta(t) = \bar{\eta}\, \mathcal{C}(t)^{\beta} \Big(1 + \frac{M(t)}{M_0}\Big)^{-\gamma},$$

$$\dot{M}(t) = \xi I_M(t) - \delta_M M(t), \qquad \dot{\mathcal{C}}(t) = \theta I_M(t), \qquad \frac{\dot{L}(t)}{L(t)} = n.$$

Let $\psi(t) \equiv \eta(t)M(t)/(L_A(t) + \eta(t)M(t))$ denote the AI share in effective research.

Assume a scaling segment *on which (i) the AI investment share is bounded away from zero, $s_M(t) \equiv I_M(t)/Y(t) \geq \underline{s}_M > 0$, and (ii) the AI-capital output ratio is asymptotically stationary in the sense that $g_M(t) = g_Y(t) + o(1)$ (equivalently $M(t)$ scales proportionally with $Y(t)$ up to lower-order terms). Under monopolistic competition with constant markup, equilibrium aggregation implies*

$$g_Y(t) = n + \kappa\, g_A(t), \qquad \kappa = \frac{\chi}{(1-\alpha-\chi)(\varepsilon-1)}. \tag{92}$$

Then the economy admits an explosive *growth regime (finite-time divergence of g_A in the reduced form) if and only if the following two conditions hold:*

(E1) ***Technological feasibility of a positive quadratic term:***

$$\lambda(1+\beta-\gamma) + \phi - 1 \;>\; 0 \qquad \Longleftrightarrow \qquad \beta \;>\; \frac{1-\phi}{\lambda} - 1 + \gamma. \tag{93}$$

(E2) ***General-equilibrium amplification sufficient to drive $\psi(t)$ above the threshold:*** *along the scaling segment, $\psi(t)$ evolves according to*

$$\frac{d}{dt}\ln\Big(\frac{1-\psi(t)}{\psi(t)}\Big) = n - \big(g_\eta(t) + g_M(t)\big), \tag{94}$$

and, using $\dot{\mathcal{C}} = \theta I_M$ and the scaling assumptions,

$$g_\eta(t) + g_M(t) = (1+\beta-\gamma)\, g_Y(t) + o(1) = (1+\beta-\gamma)\big(n + \kappa g_A(t)\big) + o(1). \tag{95}$$

Hence $\psi(t)$ crosses the explosive-regime threshold

$$\psi^{**} = \frac{1-\phi}{\lambda(1+\beta-\gamma)} \tag{96}$$

in finite time if and only if there exists t such that

$$(1+\beta-\gamma)\big(n + \kappa g_A(t)\big) \;>\; n, \tag{97}$$

so that $\eta(t)M(t)$ grows faster than research labor and $\psi(t)$ rises rather than stagnates.

Moreover, whenever (93)–(97) hold and $\psi(t)$ crosses ψ^{**}, the reduced-form law of motion for g_A becomes Riccati with a strictly positive quadratic coefficient,

$$\dot{g}_A(t) = g_A(t)\Big[\underbrace{\lambda(1-\psi(t))\, n}_{>0} + \underbrace{\big(\lambda\psi(t)(1+\beta-\gamma) + \phi - 1\big)}_{>0 \text{ if } \psi(t)>\psi^{**}}\, g_A(t)\Big] \cdot \frac{1}{1-\phi} \;+\; o\big(g_A(t)^2\big), \tag{98}$$

which implies finite-time divergence of g_A in the reduced form.

Proof. Step 1 (output mapping and the role of (χ, ε)). Equation (92) is Lemma **??**. It pins down the general-equilibrium mapping from idea growth to output growth through

$$\kappa = \frac{\chi}{(1-\alpha-\chi)(\varepsilon-1)}.$$

Step 2 (growth of effective AI research input). On the scaling segment, $I_M = s_M Y$ with $s_M \geq \underline{s}_M > 0$ implies $\dot{\mathcal{C}} = \theta s_M Y$ and thus $g_{\mathcal{C}}(t) = g_Y(t) + o(1)$ (the level of $\mathcal{C}$ tracks the same exponential order as Y when Y grows). With $g_M(t) = g_Y(t) + o(1)$ by assumption, we obtain

$$g_\eta(t) = \beta g_{\mathcal{C}}(t) - \gamma g_M(t) + o(1) = (\beta-\gamma) g_Y(t) + o(1),$$

hence (95).

Step 3 (crossing condition for ψ). By definition, $(1-\psi)/\psi = L_A/(\eta M)$. Differentiating gives (94). Thus $\psi(t)$ rises (locally) iff $g_\eta(t) + g_M(t) > n$, which is exactly (97) once we use (95) and (92).

Step 4 (positive quadratic term and the role of β). The Riccati coefficient on g_A^2 in the reduced form is proportional to $\lambda\psi(1+\beta-\gamma)+\phi-1$. A necessary and sufficient condition for this coefficient to be positive for some $\psi \le 1$ is that it be positive at $\psi = 1$, i.e. $\lambda(1+\beta-\gamma)+\phi-1 > 0$, which is (93). In that case the unique threshold is (96), and if $\psi(t)$ crosses it, the quadratic coefficient becomes strictly positive, yielding (98) and finite-time divergence in the reduced form. □

Remark 6 (A compact joint restriction in $(\chi, \varepsilon, \beta)$)**.** *If one evaluates* (97) *at a reference growth rate $\bar{g}_A$ (e.g. the pre-threshold growth rate in the stable regime), then a sufficient calibration condition that depends* only *on $(\chi, \varepsilon, \beta)$ is*

$$(1+\beta-\gamma)\big(n+\kappa\,\bar{g}_A\big) > n, \qquad \kappa = \frac{\chi}{(1-\alpha-\chi)(\varepsilon-1)}, \qquad \beta > \frac{1-\phi}{\lambda} - 1 + \gamma. \tag{99}$$

The first inequality is the GE amplification requirement (strengthened by larger χ and smaller ε through κ); the second is the technological feasibility requirement that ensures $\psi^{**} < 1$.

6 Production Automation with Research Automation

The baseline model assigns AI exclusively to research. In practice, however, AI also substitutes for labor in production tasks, as emphasized by the task-based automation literature (Acemoglu and Restrepo, 2018, 2020). If production automation is quantitatively important, it may alter the equilibrium allocation of AI capital and therefore the proximity to the research-automation threshold.

This section nests task-based production automation into the model and establishes three results. First, production automation alone cannot generate accelerating growth. Second, the closed-form threshold derived in the baseline is preserved once the AI share is correctly measured as the share deployed in research. Third, production automation affects regime proximity only through the endogenous allocation of AI capital between production and research.

6.1 Task-Based Production

Final output is produced using capital and a continuum of tasks. A unit interval of tasks $j \in [0,1]$ is aggregated through a CES composite with elasticity $\sigma_p > 1$. Each task is performed either by labor or by AI capital. An endogenous automation frontier $\Theta(t) \in [0,1]$ determines which tasks are automated.

Let $M_Y(t)$ denote AI capital deployed in production and $M_A(t)$ AI capital deployed in research, with

$$M(t) = M_A(t) + M_Y(t).$$

As in Acemoglu and Restrepo (2018), the frontier satisfies an indifference condition equating the marginal cost of AI services and labor:

$$q_Y(t) = \zeta\, w(t),$$

where q_Y is the rental rate of AI in production and w the wage. Tasks with AI cost below labor cost are automated.

Competitive equilibrium requires no-arbitrage across uses:

$$q_A(t) = q_Y(t) \equiv q_M(t),$$

so AI is allocated to equalize marginal value products in production and research.

Idea production becomes

$$\dot{A} = \delta\big[L_A + \eta M_A\big]^{\lambda} A^{\phi},$$

and the relevant AI share in research is

$$\psi_A \equiv \frac{\eta M_A}{L_A + \eta M_A}.$$

The key question is whether production automation modifies the regime structure of knowledge growth.

6.2 Production Automation Alone Cannot Generate Regime Transitions

We begin with the case in which AI efficiency is constant ($\beta = 0, \gamma = 0$). Production automation may expand output, change wages, and alter factor shares, but it does not introduce recursive efficiency feedback into research.

Proposition 7 (Production automation alone cannot generate accelerating growth)**.** *Suppose AI efficiency is constant. For any automation frontier $\Theta(t) \in [0, 1]$, including full automation $\Theta = 1$, the growth rate of ideas satisfies a linear differential equation in g_A. Along any equilibrium path with bounded factor shares, g_A converges to a finite constant. Finite-time divergence is impossible.*

Intuition. Production automation changes levels of output and thus the resources available for R&D. It affects g_A only through the linear term governing effective research input growth. The curvature parameter $\phi < 1$ continues to impose diminishing returns in idea accumulation. Without a state variable linking AI deployment to future research efficiency, the order of the growth-rate dynamics is unchanged. Automation in production is therefore a level channel, not a regime channel.

Proof is in Appendix B.1.

6.3 Threshold Preservation under Joint Automation

We now restore endogenous AI efficiency. As in the baseline, AI efficiency evolves with cumulative compute and total deployment:

$$\frac{\dot{\eta}}{\eta} = (\beta - \gamma) g_A + o(g_A).$$

The only modification relative to the baseline is that AI deployed in research is now $M_A = \phi_A M$, where ϕ_A is determined in equilibrium by no-arbitrage across uses.

Proposition 8 (Threshold with production automation)**.** *Under $\phi < 1$, $\beta > 0$, and $1 + \beta - \gamma > 0$, the reduced-form dynamics of g_A take the Riccati form*

$$\dot{g}_A = g_A \left[\frac{\lambda(1 - \psi_A)n}{1 - \phi} + \frac{\lambda \psi_A (1 + \beta - \gamma) + \phi - 1}{1 - \phi} g_A \right] + o(g_A^2).$$

The critical AI research share is

$$\boxed{\psi_A^{**} = \frac{1 - \phi}{\lambda(1 + \beta - \gamma)}}.$$

If $\psi_A < \psi_A^{**}$, growth is bounded. If $\psi_A > \psi_A^{**}$, accelerating dynamics emerge.

Implications.

First, the threshold is identical in form to the baseline expression. Production automation parameters do not appear in the quadratic coefficient.

Second, the relevant object is the AI share in *research*, not total AI deployment. Production automation matters only insofar as it changes the equilibrium allocation ϕ_A and thus ψ_A.

Proof is in Appendix B.2.

6.4 How Production Automation Affects Regime Proximity

Although production automation does not alter the threshold formula, it influences whether the economy approaches or crosses the threshold through AI allocation.

A marginal increase in the automation frontier Θ affects ψ_A through two opposing forces.

Proposition 9 (Channel decomposition)**.** *The equilibrium response of the AI research share to greater production automation satisfies*

$$\frac{d\psi_A}{d\Theta} = \underbrace{\textit{resource expansion}}_{(+)} + \underbrace{\textit{reallocation}}_{(\pm)}.$$

Resource expansion. Higher Θ raises output by automating tasks. Higher output increases total AI investment, expanding M. Holding the allocation share fixed, this raises M_A and therefore ψ_A.

Reallocation. Automation increases the marginal product of AI in production tasks. If production tasks exhibit high substitutability (large σ_p), they can absorb additional AI capital with limited diminishing returns, drawing AI away from research and reducing ϕ_A. If substitutability is low, production cannot absorb AI without sharply diminishing returns, limiting reallocation.

A sufficient condition under which production automation raises ψ_A is derived in Appendix B.3. The condition depends on the task-substitution elasticity and the relative marginal value of AI in research versus production.

Three conclusions follow. First, incorporating task-based production automation does not overturn the structural result: accelerating growth requires recursive research efficiency feedback. Production automation alone cannot generate regime transitions. Second, the threshold expression remains unchanged in form once AI is correctly measured by its research share. The mechanism is therefore robust to nesting the dominant automation framework used in the literature. Third, production automation introduces a new equilibrium margin—AI allocation between uses. Through this channel, it can move the economy closer to or further from the research threshold. The interaction depends on the substitutability of production tasks and the relative marginal value of AI in research.

7 Welfare and Policy

This section characterizes the social planner's optimal allocation and compares it with the decentralized competitive equilibrium. The planner internalizes (i) knowledge spillovers in $\dot{A}$, (ii) the monopoly wedge in the intermediate-goods block, and (iii) the externalities in AI efficiency dynamics. The comparison identifies the sources of inefficiency and derives implied policy corrections.

7.1 Planner's Objective and Constraints

The planner chooses measurable paths $\{c(t), I_K(t), I_M(t), L_A(t)\}_{t\geq 0}$ to maximize

$$\max_{\{c, I_K, I_M, L_A\}_{t\geq 0}} \mathcal{W} = \int_0^\infty e^{-\rho t} L(t) \ln c(t)\, dt, \qquad c(t) \equiv \frac{C(t)}{L(t)}. \tag{100}$$

Population evolves exogenously:

$$\dot{L}(t) = nL(t), \qquad L(0) = L_0. \tag{101}$$

Labor allocation.

$$L_Y(t) = L(t) - L_A(t), \qquad 0 \leq L_A(t) \leq L(t). \tag{102}$$

Production and feasibility.

The planner uses the same aggregate technology as in the CE block:

$$Y(t) = K(t)^\alpha L_Y(t)^{1-\alpha-\chi} X(t)^\chi, \tag{103}$$

subject to the aggregate resource constraint

$$Y(t) = C(t) + I_K(t) + I_M(t), \qquad C(t) = c(t)L(t). \tag{104}$$

In the planner allocation, the intermediate wedge is removed ($\mu^{SP} = 1$), so the implied P_X and the associated static efficiency conditions differ from the decentralized equilibrium. However, the aggregate mapping (103) remains the same object used throughout the paper.

Intermediate composite (efficient static allocation).

Under $\mu^{SP} = 1$ and symmetry across varieties, the CES price index satisfies

$$P_X(t) = A(t)^{-\frac{1}{\varepsilon-1}}, \tag{105}$$

so the final-good FOC for X implies the planner scaling

$$X(t) = \chi \frac{Y(t)}{P_X(t)} = \chi\, Y(t)\, A(t)^{\frac{1}{\varepsilon-1}}. \tag{106}$$

This is the planner counterpart of the variety-scaling relation used in the decentralized equilibrium.

State laws of motion.

$$\dot{K}(t) = I_K(t) - \delta_K K(t), \qquad K(0) = K_0, \tag{107}$$
$$\dot{M}(t) = \xi I_M(t) - \delta_M M(t), \qquad M(0) = M_0, \tag{108}$$
$$\dot{\mathcal{C}}(t) = \theta I_M(t), \qquad \mathcal{C}(0) = \mathcal{C}_0, \tag{109}$$
$$\dot{A}(t) = \delta R(t)^{\lambda} A(t)^{\phi}, \qquad A(0) = A_0, \tag{110}$$

where effective research is

$$R(t) \equiv L_A(t) + \eta\big(M(t), \mathcal{C}(t)\big) M(t). \tag{111}$$

AI efficiency.

$$\eta(M, \mathcal{C}) = \bar{\eta}\, \mathcal{C}^{\beta} \left(1 + \frac{M}{M_0}\right)^{-\gamma}. \tag{112}$$

Derivatives of AI efficiency.

$$\eta_M(M, \mathcal{C}) = -\frac{\gamma}{M_0 + M}\, \eta(M, \mathcal{C}), \qquad \eta_{\mathcal{C}}(M, \mathcal{C}) = \frac{\beta}{\mathcal{C}}\, \eta(M, \mathcal{C}). \tag{113}$$

Admissibility.

$$c(t) > 0, \quad I_K(t) \geq 0, \quad I_M(t) \geq 0, \quad 0 \leq L_A(t) \leq L(t), \tag{114}$$

together with the transversality conditions for $\{K, M, \mathcal{C}, A\}$.

Remark 7 (Difference from CE)**.** *Unlike the decentralized equilibrium, the planner problem does not use $V_A(t)$, $\pi(t)$, or the asset-pricing ODE $\dot{V}_A = rV_A - \pi$ as constraints: those objects exist only because private ownership of designs and monopoly pricing generate private values. The CE allocation can be recovered by adding (i) patent rights and (ii) the intermediate-goods markup wedge $\mu > 1$.*

7.2 Optimality Conditions

Eliminating consumption via $c(t) = [Y(t) - I_K(t) - I_M(t)]/L(t)$, the current-value Hamiltonian is

$$\mathcal{H} = L \ln\!\left(\frac{Y - I_K - I_M}{L}\right) + \mu_K(I_K - \delta_K K) + \mu_M(\xi I_M - \delta_M M) + \mu_{\mathcal{C}}\, \theta I_M + \mu_A\, \delta R^{\lambda} A^{\phi}, \tag{115}$$

where $(\mu_K, \mu_M, \mu_{\mathcal{C}}, \mu_A)$ are the co-state variables.

7.2.1 First-order conditions for controls

Physical capital investment I_K.

$$\frac{\partial \mathcal{H}}{\partial I_K} = -\frac{1}{c} + \mu_K = 0 \quad \Longrightarrow \quad \mu_K(t) = \frac{1}{c(t)}. \tag{116}$$

AI capital investment I_M.

$$\frac{\partial \mathcal{H}}{\partial I_M} = -\frac{1}{c} + \xi\mu_M + \theta\mu_{\mathcal{C}} = 0 \quad \Longrightarrow \quad \frac{1}{c(t)} = \xi\mu_M(t) + \theta\mu_{\mathcal{C}}(t). \tag{117}$$

Research labor L_A.

Using $\partial Y/\partial L_A = -\partial Y/\partial L_Y = -(1 - \alpha - \chi)\, Y/L_Y$ and $\partial R/\partial L_A = 1$,

$$\frac{\partial \mathcal{H}}{\partial L_A} = -\frac{1}{c}(1 - \alpha - \chi)\frac{Y}{L_Y} + \mu_A\, \delta\, \lambda\, R^{\lambda - 1} A^{\phi} = 0. \tag{118}$$

Hence the labor allocation rule

$$\frac{1}{c(t)}(1 - \alpha - \chi)\frac{Y(t)}{L_Y(t)} = \mu_A(t)\, \delta\, \lambda\, R(t)^{\lambda - 1} A(t)^{\phi}. \tag{119}$$

7.2.2 Co-state equations

Co-states satisfy $\dot{\mu}_X = \rho\mu_X - \partial\mathcal{H}/\partial X$.

Physical capital K**.**
Using $\partial Y/\partial K = \alpha Y/K$,

$$\dot{\mu}_K = \rho\mu_K - \left[\frac{1}{c}\,\alpha\frac{Y}{K} - \mu_K\,\delta_K\right] = (\rho+\delta_K)\mu_K - \frac{\alpha Y/K}{c}. \tag{120}$$

With $\mu_K = 1/c$, this yields the Euler equation

$$\frac{\dot{c}}{c} = \alpha\frac{Y}{K} - \delta_K - \rho. \tag{121}$$

AI capital M**.**
Using $\partial R/\partial M = \eta + M\eta_M$,

$$\dot{\mu}_M = (\rho+\delta_M)\mu_M - \mu_A\,\delta\,\lambda\,R^{\lambda-1}A^{\phi}\,(\eta + M\eta_M). \tag{122}$$

Cumulative compute $\mathcal{C}$**.**
Using $\partial R/\partial\mathcal{C} = M\eta_{\mathcal{C}}$,

$$\dot{\mu}_{\mathcal{C}} = \rho\mu_{\mathcal{C}} - \mu_A\,\delta\,\lambda\,R^{\lambda-1}A^{\phi}\,M\eta_{\mathcal{C}}. \tag{123}$$

Ideas A**.**
The planner internalizes the impact of variety on the intermediate composite. Under $\mu^{SP} = 1$ and symmetry, (106) implies

$$\frac{\partial Y}{\partial A} = \chi\,\frac{Y}{X}\,\frac{\partial X}{\partial A} = \frac{\chi}{\varepsilon-1}\,\frac{Y}{A}. \tag{124}$$

Moreover, $\partial(\delta R^{\lambda}A^{\phi})/\partial A = \delta R^{\lambda}\phi A^{\phi-1}$. Therefore,

$$\dot{\mu}_A = \rho\mu_A - \left[\frac{1}{c}\cdot\frac{\chi}{\varepsilon-1}\frac{Y}{A} + \mu_A\,\delta R^{\lambda}\phi A^{\phi-1}\right]. \tag{125}$$

7.3 Normalized shadow values

Define shadow values in consumption units:

$$v_M \equiv \frac{\mu_M}{\mu_K}, \qquad v_{\mathcal{C}} \equiv \frac{\mu_{\mathcal{C}}}{\mu_K}, \qquad v_A \equiv \frac{\mu_A}{\mu_K}. \tag{126}$$

Then (117) becomes

$$1 = \xi v_M + \theta v_{\mathcal{C}}. \tag{127}$$

Equation (119) becomes

$$(1-\alpha-\chi)\frac{Y(t)}{L_Y(t)} = v_A(t)\,\delta\,\lambda\,R(t)^{\lambda-1}A(t)^{\phi}. \tag{128}$$

The normalized co-state dynamics are

$$\dot{v}_M = v_M\left(\frac{\dot{\mu}_M}{\mu_M} - \frac{\dot{\mu}_K}{\mu_K}\right) = v_M\left(\alpha\frac{Y}{K} - \delta_K - \rho - \delta_M\right) + v_A\,\delta\,\lambda\,R^{\lambda-1}A^{\phi}\,(\eta + M\eta_M), \tag{129}$$

$$\dot{v}_{\mathcal{C}} = v_{\mathcal{C}}\left(\alpha\frac{Y}{K} - \delta_K - \rho\right) + v_A\,\delta\,\lambda\,R^{\lambda-1}A^{\phi}\,M\eta_{\mathcal{C}}, \tag{130}$$

$$\dot{v}_A = v_A\left(\alpha\frac{Y}{K} - \delta_K - \rho - \phi g_A\right) - \frac{\chi}{\varepsilon-1}\frac{Y}{A}. \tag{131}$$

7.4 Planner's BGP Characterization

Along a balanced growth path (BGP), the normalized shadow values (v_M^*, v_C^*, v_A^*) are constant and all relevant growth rates are constant. Let

$$r^{SP} \equiv \alpha \frac{Y^*}{K^*} - \delta_K \tag{132}$$

denote the planner-implied net return on physical capital, so that the planner Euler equation $\dot{c}/c = r^{SP} - \rho$ implies $r^{SP} = g_c^* + \rho$.

7.4.1 BGP shadow value of ideas

From (131), imposing $\dot{v}_A = 0$ on the BGP yields

$$v_A^* \left(r^{SP} - \phi g_A^*\right) = \frac{\chi}{\varepsilon - 1} \frac{Y^*}{A^*}. \tag{133}$$

Hence the planner shadow value of ideas in consumption units is

$$\boxed{v_A^* = \frac{\chi/(\varepsilon - 1)}{r^{SP} - \phi g_A^*} \cdot \frac{Y^*}{A^*}.} \tag{134}$$

It is useful to define the planner analogue of the patent-value coefficient,

$$\Phi^{SP} \equiv \frac{v_A^*}{Y^*/A^*} = \frac{\chi/(\varepsilon - 1)}{r^{SP} - \phi g_A^*}. \tag{135}$$

7.4.2 BGP shadow values of AI capital and compute

From (129), imposing $\dot{v}_M = 0$ gives

$$v_M^*(r^{SP} + \delta_M) = v_A^* \, \delta \, \lambda \, (R^*)^{\lambda - 1} (A^*)^{\phi} \left(\eta^* + M^* \eta_M^*\right), \tag{136}$$

so

$$\boxed{v_M^* = \frac{v_A^* \, \delta \, \lambda \, (R^*)^{\lambda - 1} (A^*)^{\phi} \left(\eta^* + M^* \eta_M^*\right)}{r^{SP} + \delta_M}.} \tag{137}$$

From (130), imposing $\dot{v}_C = 0$ gives

$$v_C^* \, r^{SP} = v_A^* \, \delta \, \lambda \, (R^*)^{\lambda - 1} (A^*)^{\phi} \, M^* \eta_C^*, \tag{138}$$

so

$$\boxed{v_C^* = \frac{v_A^* \, \delta \, \lambda \, (R^*)^{\lambda - 1} (A^*)^{\phi} \, M^* \eta_C^*}{r^{SP}}.} \tag{139}$$

Using the BGP identity $g_A^* = \delta (R^*)^{\lambda} (A^*)^{\phi - 1}$, we can rewrite $\delta \lambda (R^*)^{\lambda - 1} (A^*)^{\phi} = \lambda g_A^* (A^*/R^*)$. Substituting $v_A^* = \Phi^{SP} (Y^*/A^*)$ then yields the compact forms

$$v_M^* = \Phi^{SP} \frac{\lambda g_A^*}{r^{SP} + \delta_M} \frac{Y^*}{R^*} \left(\eta^* + M^* \eta_M^*\right), \tag{140}$$

$$v_C^* = \Phi^{SP} \frac{\lambda g_A^*}{r^{SP}} \frac{Y^*}{R^*} M^* \eta_C^*. \tag{141}$$

7.4.3 Planner labor allocation on the BGP

From (128), on the BGP,

$$(1 - \alpha - \chi) \frac{Y^*}{L_Y^*} = v_A^* \, \delta \, \lambda \, (R^*)^{\lambda - 1} (A^*)^{\phi} = \Phi^{SP} \, \lambda g_A^* \frac{Y^*}{R^*}. \tag{142}$$

Cancelling $Y^* > 0$ and defining $\Psi^{SP} \equiv \Phi^{SP}\lambda g_A^*$ gives

$$(1-\alpha-\chi)\frac{R^*}{L_Y^*} = \Psi^{SP}. \tag{143}$$

Proposition 10 (Planner analogue of wage equalization)**.** *On the planner BGP, the research labor share $\sigma^{SP} \equiv L_A^*/L^*$ satisfies*

$$\frac{\sigma^{SP}}{1-\sigma^{SP}} = \frac{\Psi^{SP}}{1-\alpha-\chi}\left(1-\psi^{SP}\right), \qquad \Psi^{SP} \equiv \Phi^{SP}\lambda g_A^*, \tag{144}$$

where $\psi^{SP} \equiv \eta^* M^*/R^*$ is the planner AI research intensity.

Proof. Write (143) as $(1-\alpha-\chi)\,R^*/L_Y^* = \Psi^{SP}$, and substitute $L_Y^* = (1-\sigma^{SP})L^*$. Using $R^* = \sigma^{SP}L^*/(1-\psi^{SP})$ (since $\psi^{SP} = \eta^* M^*/R^*$ implies $\eta^* M^* = \psi^{SP}R^*$ and hence $R^* - \eta^* M^* = (1-\psi^{SP})R^* = \sigma^{SP}L^*$), we obtain

$$(1-\alpha-\chi)\frac{\sigma^{SP}}{(1-\psi^{SP})(1-\sigma^{SP})} = \Psi^{SP},$$

which rearranges to (144). □

7.5 Comparison with Decentralized Equilibrium

Table 1 contrasts the planner and decentralized conditions that matter for (i) the value of ideas, (ii) the private versus social marginal value of AI capital, and (iii) the role of cumulative compute $\mathcal{C}$. Throughout, $\mu \equiv \varepsilon/(\varepsilon-1)$ and the CE patent-value coefficient is $\Phi^{CE} \equiv (V_A^* A^*/Y^*)$ under the intermediate-goods block.

Table 1: Social planner vs. decentralized equilibrium: key wedges

Object	**Social Planner**	**Competitive Equilibrium**
Idea valuation		
Value coefficient	$\Phi^{SP} = \dfrac{\chi/(\varepsilon-1)}{r^{SP} - \phi g_A^*}$	$\Phi^{CE} = \dfrac{\chi(\mu-1)/\mu}{r_K^* - g_\pi^*} = \dfrac{\chi/\varepsilon}{r_K^* - g_\pi^*}$
Numerator	$\dfrac{\chi}{\varepsilon-1}$	$\dfrac{\chi(\mu-1)}{\mu} = \dfrac{\chi}{\varepsilon}$
Denominator	$r^{SP} - \phi g_A^*$	$r_K^* - g_\pi^*$
AI capital and efficiency externalities		
Marginal value term in $\dot{A}$	$v_A\,\delta\lambda R^{\lambda-1}A^{\phi}\,(\eta + M\eta_M)$	$V_A\,\delta\lambda R^{\lambda-1}A^{\phi}\,\eta$
Deployment congestion ($\gamma > 0$)	Internalized ($M\eta_M < 0$)	Ignored (atomistic firms take η as given)
Cumulative compute $\mathcal{C}$ and AI learning		
Compute shadow value	$v_{\mathcal{C}}^* > 0$ when $\beta > 0$	No market price for $\mathcal{C}$ (implicit $v_{\mathcal{C}}^{CE} = 0$)
AI investment condition	$1 = \xi v_M + \theta v_{\mathcal{C}}$	$q_M = p_M(r + \delta_M - g_{p_M})$, $\dfrac{q_M}{w} = \eta$, $p_M = 1/\xi \Rightarrow g_{p_M} = 0$

7.5.1 Source 1: markup distortion

The planner removes the monopoly wedge in intermediates. The difference shows up directly in the numerator of the idea-value coefficient:

$$\frac{\chi}{\varepsilon-1} \quad \text{(planner)} \qquad \text{vs.} \qquad \frac{\chi}{\varepsilon} \quad \text{(CE)}. \tag{145}$$

Since $\varepsilon > 1$,

$$\frac{\chi/(\varepsilon-1)}{\chi/\varepsilon} = \frac{\varepsilon}{\varepsilon-1} = \mu > 1, \tag{146}$$

so the private profitability channel undervalues variety relative to the planner by exactly the markup factor.

7.5.2 Source 2: intertemporal knowledge spillover

In the planner problem, the shadow value v_A discounts the marginal product of ideas using $r^{SP} - \phi g_A^*$, reflecting that higher A raises future research productivity through A^ϕ. In the CE, the private value V_A is pinned down by patent-arbitrage, so ideas are discounted at $r_K^* - g_\pi^*$ without the $-\phi g_A^*$ term. When $\phi > 0$, holding other objects fixed,

$$r^{SP} - \phi g_A^* < r^{SP}, \tag{147}$$

which raises Φ^{SP} relative to the same expression without spillovers.

7.5.3 Source 3: compute feedback externality ($\beta > 0$)

Because $\dot{\mathcal{C}} = \theta I_M$ and $\eta_{\mathcal{C}} > 0$ when $\beta > 0$, AI investment raises future AI efficiency. The planner therefore assigns a positive shadow value $v_{\mathcal{C}}^* > 0$ and this appears as an extra benefit of I_M in

$$1 = \xi v_M + \theta v_{\mathcal{C}}. \tag{148}$$

In the CE, individual agents treat $\mathcal{C}$ (and hence η) as given, so there is no priced claim to the aggregate learning externality. A convenient summary wedge is

$$\omega_{\mathcal{C}} \equiv \theta v_{\mathcal{C}}^* > 0 \quad (\beta > 0), \tag{149}$$

which pushes the planner toward higher I_M than the CE, other things equal.

7.5.4 Source 4: deployment congestion externality ($\gamma > 0$)

With $\eta(M, \mathcal{C}) = \bar{\eta}\,\mathcal{C}^\beta (1 + M/M_0)^{-\gamma}$,

$$M\eta_M(M, \mathcal{C}) = -\gamma \frac{M}{M_0 + M}\,\eta(M, \mathcal{C}) < 0 \quad (\gamma > 0). \tag{150}$$

The planner internalizes this term, so the social marginal contribution of M to effective research involves $\eta + M\eta_M$, not η. In contrast, atomistic R&D firms take η as parametric and ignore the congestion channel.

Define the congestion wedge

$$\omega_\gamma \equiv \frac{\eta + M\eta_M}{\eta} = 1 - \gamma \frac{M}{M_0 + M} \in (0, 1] \quad (\gamma \geq 0). \tag{151}$$

For $\gamma > 0$, $\omega_\gamma < 1$, so the CE tends to overstate the marginal value of deployment M. Whether the CE under- or over-invests in M overall is therefore not signed a priori, because $\omega_\gamma < 1$ competes with the compute-learning wedge $\omega_{\mathcal{C}} > 0$ when $\beta > 0$.

7.6 Efficiency Gaps

Proposition 11 (Decomposition of the idea-valuation gap)**.** *Let*

$$\Phi^{SP} \equiv \frac{v_A^* A^*}{Y^*} = \frac{\chi/(\varepsilon-1)}{r^{SP} - \phi g_A^*}, \qquad \Phi^{CE} \equiv \frac{V_A^* A^*}{Y^*} = \frac{\chi(\mu-1)/\mu}{r_K^* - g_\pi^*},$$

with $\mu = \varepsilon/(\varepsilon - 1)$. Then the ratio of planner to competitive-equilibrium coefficients satisfies

$$\frac{\Phi^{SP}}{\Phi^{CE}} = \underbrace{\mu}_{\text{static markup wedge}} \cdot \underbrace{\frac{r_K^* - g_\pi^*}{r^{SP} - \phi g_A^*}}_{\text{dynamic wedge: discounting and spillovers}}. \tag{152}$$

Proof. Use $\chi(\mu - 1)/\mu = \chi/\varepsilon$ and $\chi/(\varepsilon - 1) = (\varepsilon/(\varepsilon - 1))(\chi/\varepsilon) = \mu\,\chi/\varepsilon$. Therefore,

$$\frac{\Phi^{SP}}{\Phi^{CE}} = \frac{\chi/(\varepsilon - 1)}{\chi/\varepsilon} \cdot \frac{r_K^* - g_\pi^*}{r^{SP} - \phi g_A^*} = \mu \cdot \frac{r_K^* - g_\pi^*}{r^{SP} - \phi g_A^*}.$$

□

Proposition 12 (Direction of the AI-intensity gap)**.** *Let ψ^{SP} and ψ^{CE} denote the long-run AI shares in effective research under the planner and the competitive equilibrium. The sign of $\psi^{SP} - \psi^{CE}$ is generally ambiguous because two opposing forces operate through the marginal social value of AI-capital accumulation:*

1. *If $\beta > 0$, higher cumulative compute raises future AI efficiency ($\eta_{\mathcal{C}} > 0$), which tends to increase the planner's incentive to accumulate M and raises ψ^{SP} relative to ψ^{CE}, all else equal.*
2. *If $\gamma > 0$, higher deployment reduces contemporaneous AI efficiency ($\eta_M < 0$), which tends to reduce the planner's incentive to accumulate M and lowers ψ^{SP} relative to ψ^{CE}, all else equal.*

The net effect depends on the relative magnitudes of the learning and congestion channels, as well as on how the induced changes in A and Y feed back into the research allocation.

7.7 Time to Singularity

Proposition 13 (Time to Singularity: Social Planner vs Competitive Equilibrium)**.** *Suppose the economy operates in the explosive regime ($\psi^j > \psi^{**}$ for $j \in \{CE, SP\}$) and let $g_0 = g_A(t_0) > 0$ denote the idea growth rate at some initial date t_0 when the quadratic dynamics dominate. Define the time to singularity as the finite horizon T^j such that $\lim_{t \to T^j} g_A(t) = \infty$ under allocation j. Then:*

(i) ***Closed-form expression:*** *The time to singularity is given by*

$$T^j = t_0 + \frac{1}{a(\psi^j)} \ln\left(1 + \frac{a(\psi^j)}{\kappa_\gamma(\psi^j) g_0}\right), \qquad a(\psi^j) > 0, \tag{153}$$

where

$$a(\psi) \equiv \frac{\lambda(1 - \psi)n}{1 - \phi}, \tag{154}$$

$$\kappa_\gamma(\psi) \equiv \frac{\lambda\psi(1 + \beta - \gamma) + \phi - 1}{1 - \phi} > 0. \tag{155}$$

When $a(\psi^j) = 0$ (i.e., $\psi^j = 1$), this reduces to $T^j = t_0 + 1/(\kappa_\gamma(\psi^j) g_0)$.

(ii) ***Monotonicity:*** *The time to singularity is strictly decreasing in the AI research intensity:*

$$\frac{\partial T^j}{\partial \psi^j} < 0 \quad \text{for all } \psi^j \in (\psi^{**}, 1]. \tag{156}$$

Hence, allocations with higher AI intensity reach the singularity sooner.

(iii) ***SP-CE comparison:*** *If the Social Planner's allocation yields a higher long-run AI intensity than the Competitive Equilibrium ($\psi^{SP} > \psi^{CE}$), then the planner reaches the singularity earlier:*

$$T^{SP} < T^{CE}. \tag{157}$$

Moreover, the ratio of times to singularity satisfies

$$\frac{T^{CE}}{T^{SP}} = \frac{\kappa_\gamma(\psi^{SP})}{\kappa_\gamma(\psi^{CE})} \cdot \frac{\ln\left(1 + \frac{a(\psi^{CE})}{\kappa_\gamma(\psi^{CE})g_0}\right)/a(\psi^{CE})}{\ln\left(1 + \frac{a(\psi^{SP})}{\kappa_\gamma(\psi^{SP})g_0}\right)/a(\psi^{SP})} > 1, \tag{158}$$

with the inequality strict whenever $\psi^{SP} > \psi^{CE}$ and $\beta - \gamma > -1$.

(iv) ***Asymptotic approximation:*** *For sufficiently large initial growth rates $g_0 \gg a(\psi^j)/\kappa_\gamma(\psi^j)$, the time to singularity is approximately*

$$T^j \approx t_0 + \frac{1}{\kappa_\gamma(\psi^j)g_0}, \tag{159}$$

and the planner's advantage simplifies to

$$\frac{T^{CE}}{T^{SP}} \approx \frac{\kappa_\gamma(\psi^{SP})}{\kappa_\gamma(\psi^{CE})} > 1. \tag{160}$$

Proof. We prove each part in the appendice A □

Corollary 8 (Quantitative Significance of the Planner's Advantage)**.** *With $(\lambda, \phi, \beta, \gamma) = (0.5, 0.5, 0.3, 0.1)$, the gap in AI intensities $\Delta\psi \equiv \psi^{SP} - \psi^{CE} \approx 0.04$ implies:*

(i) The quadratic feedback coefficient is approximately

$$\frac{\kappa_\gamma(\psi^{SP})}{\kappa_\gamma(\psi^{CE})} \approx \frac{0.042}{0.026} \approx 1.615, \tag{161}$$

meaning the planner's self-reinforcing feedback is 61.5% stronger.

(ii) For an initial growth rate $g_0 = 2\%$, the time to singularity is

$$T^{CE} \approx 54 \text{ years}, \tag{162}$$

$$T^{SP} \approx 32 \text{ years}, \tag{163}$$

implying the planner reaches the singularity 40% faster.

(iii) The difference in timing is highly nonlinear in $\Delta\psi$:

$$\left.\frac{\partial T}{\partial \psi}\right|_{\psi=\psi^{CE}} \approx -1200 \text{ years per unit } \psi, \tag{164}$$

so a 0.04 increase in ψ reduces the time to singularity by approximately 48 years at the margin.

Thus, even modest differences in long-run AI intensity induced by externality internalization translate into economically significant differences in the timing of the growth acceleration.

7.8 Automation threshold and regime comparison

In the reduced-form threshold analysis, the critical AI share depends on the elasticity of AI efficiency growth with respect to the scale variables that are tied to compute accumulation and deployment. With

$$\eta(M, \mathcal{C}) = \bar{\eta}\,\mathcal{C}^\beta\left(1 + \frac{M}{M_0}\right)^{-\gamma}, \qquad \vartheta_M(t) \equiv \frac{M(t)}{M_0 + M(t)} \in (0, 1),$$

the BGP growth rate of AI efficiency satisfies the identity

$$g_\eta^* = \beta\, g_{\mathcal{C}}^* - \gamma\, \vartheta_M^*\, g_M^* = (\beta - \gamma\, \vartheta_M^*)\, g_Y^*, \tag{165}$$

where $g_{\mathcal{C}}^* \equiv \dot{\mathcal{C}}/\mathcal{C}$ and $\vartheta_M^* = \lim_{t\to\infty} \vartheta_M(t)$. In the calibration used in the main text, the threshold takes the convenient form

$$\psi^{**} = \frac{1 - \phi}{\lambda(1 + \beta - \gamma\, \vartheta_M^*)}, \tag{166}$$

so γ enters with a minus sign and is scaled by the long-run deployment intensity ϑ_M^*.[1]

Both allocations face the same technological threshold (166), but they typically deliver different equilibrium AI intensities (ψ^{SP}, ψ^{CE}).

Corollary 9 (Regime mismatch)**.** *It is possible that $\psi^{CE} > \psi^{**}$ while $\psi^{SP} < \psi^{**}$, or that $\psi^{CE} < \psi^{**}$ while $\psi^{SP} > \psi^{**}$. Such mismatches arise when the equilibrium externality wedges shift the long-run research allocation across the threshold.*

Proposition 14 (BGP idea-growth under the social planner and its fixed-point representation)**.** *Maintain $\phi < 1$ and assume the economy admits a balanced growth path (BGP) on which g_A^{SP} and g_R^{SP} are constant. Under the planner allocation, the idea technology*

$$\dot{A}(t) = \delta\, R(t)^{\lambda} A(t)^{\phi}$$

implies the BGP restriction

$$g_A^{SP} = \frac{\lambda}{1-\phi}\, g_R^{SP}. \tag{167}$$

Effective research is $R = L_A + \eta M$, and along the BGP its growth rate satisfies the accounting identity

$$g_R^{SP} = (1-\psi^{SP})\, n + \psi^{SP}\big(g_\eta^{SP} + g_M^{SP}\big), \qquad \psi^{SP} \equiv \frac{\eta^* M^*}{R^*}. \tag{168}$$

Hence, conditional on BGP growth rates (g_η^{SP}, g_M^{SP}),

$$g_A^{SP} = \frac{\lambda}{1-\phi}\Big[(1-\psi^{SP})\, n + \psi^{SP}\big(g_\eta^{SP} + g_M^{SP}\big)\Big]. \tag{169}$$

Suppose in addition that the BGP imposes constant AI intensity M/Y (equivalently $g_M^{SP} = g_Y^{SP}$) and that the CES-variety block together with the final-good technology yields the mapping

$$g_Y^{SP} = n + \kappa\, g_A^{SP}, \qquad \kappa \equiv \frac{\chi}{(1-\alpha-\chi)(\varepsilon-1)}. \tag{170}$$

If, moreover, g_η^{SP} is constant on the BGP, then g_A^{SP} solves the fixed-point equation

$$g_A^{SP} = \frac{\lambda}{1-\phi}\Big[n + \psi^{SP}\big(g_\eta^{SP} + \kappa\, g_A^{SP}\big)\Big], \tag{171}$$

and therefore admits the closed form

$$g_A^{SP} = \frac{\lambda\big(n + \psi^{SP} g_\eta^{SP}\big)}{(1-\phi) - \lambda\kappa\,\psi^{SP}}, \qquad \textit{provided } (1-\phi) - \lambda\kappa\,\psi^{SP} > 0. \tag{172}$$

In particular, holding (n, g_η^{SP}) fixed, g_A^{SP} is increasing in ψ^{SP} whenever the denominator in (172) is positive.

Proof. On any BGP with constant growth rates,

$$g_A^{SP} = \frac{\dot{A}}{A} = \delta R^{\lambda} A^{\phi-1}.$$

Taking logs and differentiating along the BGP gives $0 = \lambda g_R^{SP} + (\phi-1) g_A^{SP}$, which yields (167). Because $R = L_A + \eta M$ and $\psi^{SP} \equiv \eta M/R$, the growth rate of R satisfies $g_R^{SP} = (1-\psi^{SP}) g_{L_A}^{SP} + \psi^{SP}(g_\eta^{SP} + g_M^{SP})$. Since labor grows exogenously at rate n on the BGP, $g_{L_A}^{SP} = n$, which gives (168). Substituting (168) into (167) yields (169).

Under the closure $g_M^{SP} = g_Y^{SP}$ and the output-growth mapping (170), (169) becomes (171). Solving for g_A^{SP} gives (172), with existence of a finite BGP growth rate requiring $(1-\phi) - \lambda\kappa\,\psi^{SP} > 0$. Monotonicity in ψ^{SP} follows directly from (172) when the denominator is positive. □

[1] If reduced-form derivation instead substitutes $g_C^* = g_c^*$ or imposes an intensity closure linking g_M^* to g_Y^*, then (166) is unchanged in sign but the effective coefficient multiplying γ is adjusted accordingly. The robust point is that the congestion channel reduces g_η and therefore lowers the effective feedback strength.

Equation (167) is purely technological and holds under both the planner and the CE. Differences in g_A^* across allocations arise because ψ^{SP} generally differs from ψ^{CE} and because (g_η^*, g_M^*) are equilibrium objects that move with the allocation and the closure. Under the constant-M/Y closure, the feedback from g_A to g_Y creates the fixed-point structure in (171) and the amplification term $\lambda\kappa\psi^{SP}$ in the denominator of (172).

7.9 Policy instruments

A convenient way to state decentralization is in wedges on (i) idea incentives and (ii) AI accumulation incentives. Relative to the competitive equilibrium, the planner raises the social value of ideas through (a) removal of the markup wedge and (b) internalization of the knowledge spillover. A sufficient reduced-form statement is that the CE innovation payoff must be scaled by the factor Φ^{SP}/Φ^{CE} in (152), implemented via an R&D subsidy or equivalent prize scheme that increases the private return to research.

Two wedges govern AI investment relative to the competitive equilibrium:

$$\text{learning wedge} \propto \eta_{\mathcal{C}} > 0 \quad (\beta > 0), \qquad \text{congestion wedge} \propto \eta_M < 0 \quad (\gamma > 0).$$

A policy package that decentralizes the planner must therefore combine a subsidy component that rewards the contribution of I_M to $\mathcal{C}$ (internalizing the shadow value of compute) and a tax component that charges for the deployment congestion created by higher M.

The planner allocation corresponds to marginal-cost pricing in the intermediate block. In implementation terms, this can be achieved by a purchase subsidy that offsets the markup or by buyouts and open access, with the financing handled by lump-sum transfers.

The planner differs from the competitive equilibrium because the market does not internalize: (i) static markup distortions in intermediates, (ii) the dynamic knowledge spillover through A^ϕ, (iii) the dynamic learning-by-compute channel ($\beta > 0$), and (iv) the deployment congestion channel ($\gamma > 0$). The threshold ψ^{**} is technological and common to both allocations, but the endogenous equilibrium intensities (ψ^{SP}, ψ^{CE}) can lie on different sides of the threshold, generating regime mismatch.

8 Conclusion

This paper studies whether research automation can generate self-reinforcing economic growth. We show that in standard semi-endogenous and task-based automation models, explosive growth is structurally impossible. When AI reallocates tasks but does not alter the law of motion of research productivity, the differential equation governing idea growth is linear and the asymptotic growth rate is bounded.

We introduce the minimal structural modification required to overturn this result. AI is modeled as an accumulable research factor whose efficiency depends on cumulative compute. Because compute scales with deployment and deployment scales with research activity, idea growth acquires a quadratic amplification term. The economy admits a closed-form threshold separating bounded semi-endogenous growth from accelerating dynamics. Crossing this threshold depends on observable elasticities governing research returns, scaling effects, and congestion.

The welfare implications are equally structural. Compute accumulation generates a dynamic learning externality, while large-scale deployment generates congestion. In high-AI regimes, the dominant inefficiency is misallocation between research labor and AI capital rather than simple capital underinvestment. Optimal policy therefore requires a compute subsidy combined with deployment correction, with magnitudes tied directly to the automation threshold.

The broader implication is not that explosive growth will occur, but that existing automation models cannot speak to the question. In those frameworks, acceleration is ruled out by functional form. Once AI recursively augments research productivity, regime transitions become an equilibrium outcome. The distinction between bounded and accelerating growth is thus not about optimism or pessimism regarding technology. It is about whether growth theory includes a state variable linking deployment to future efficiency.

If artificial intelligence remains a tool, growth remains semi-endogenous. If it becomes a recursive input into the production of ideas, the growth regime itself changes. The difference is structural.

Author Declarations

The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

References

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

Acemoglu D, Restrepo P (2020) Robots and jobs: Evidence from us labor markets. Journal of Political Economy 128(6):2188–2244

Growiec J (2022) Accelerating Economic Growth: Lessons from 200,000 Years of Technological Progress and Human Development. Frontiers in Economic History, Springer, https://doi.org/10.1007/978-3-031-07195-9

Growiec J, Prettner K (2025) The economics of p(doom): Scenarios of existential risk and economic growth in the age of transformative ai URL https://arxiv.org/abs/2503.07341, arXiv:2503.07341 [econ.GN]

Hoffmann J, Borgeaud S, Mensch A, et al (2022) Training compute-optimal large language models. arXiv preprint arXiv:220315556

Jones CI (1995) R&d-based models of economic growth. Journal of Political Economy 103(4):759–784

Jones CI (2002) Sources of us economic growth in a world of ideas. American Economic Review 92(1):220–239

Jones CI (2025) How much should we spend to reduce ai's existential risk? In: The Economics of Transformative AI. National Bureau of Economic Research, Inc.

Jones CI (2026) A.i. and our economic future. NBER Working Paper 34779, National Bureau of Economic Research, URL https://www.nber.org/papers/w34779

Kaplan J, McCandlish S, Henighan T, et al (2020) Scaling laws for neural language models. arXiv preprint arXiv:200108361

Trammell P, Korinek A (2023) Economic growth under transformative ai. Tech. Rep. 31815, National Bureau of Economic Research, https://doi.org/10.3386/w31815, URL https://www.nber.org/papers/w31815

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

Appendix B. Production Automation Extension

This appendix provides formal derivations for the production-automation extension presented in Section 6. We proceed in four steps. First, we derive the task-based production aggregator and the rental rate of AI capital in production. Second, we prove that production automation alone cannot generate accelerating growth. Third, we derive the reduced-form growth dynamics under joint automation and establish threshold preservation. Fourth, we characterize how production automation affects the AI research share through equilibrium allocation.

B.1 Task-Based Production: Aggregation and Marginal Products

B.1.1 Aggregation.

Final output is given by

$$Y = K^{\alpha}\left[\int_0^1 y(j)^{\frac{\sigma_p-1}{\sigma_p}}dj\right]^{\frac{\sigma_p\chi}{\sigma_p-1}} L_Y^{1-\alpha-\chi},$$

where tasks are performed by AI or labor depending on the automation frontier Θ:

$$y(j) = \begin{cases} \zeta m_Y(j), & j \in [0,\Theta], \\ \ell_Y(j), & j \in (\Theta,1]. \end{cases}$$

Under symmetric deployment within task types, the task composite reduces to

$$\mathcal{T} = \left[\Theta(\zeta M_Y)^{\frac{\sigma_p-1}{\sigma_p}} + (1-\Theta)\ell_Y^{\frac{\sigma_p-1}{\sigma_p}}\right]^{\frac{\sigma_p}{\sigma_p-1}}.$$

B.1.2 Marginal Value Product of AI in Production.

Differentiating Y with respect to M_Y yields

$$q_Y = \chi\frac{Y}{\mathcal{T}}\frac{\partial\mathcal{T}}{\partial M_Y}.$$

Using CES differentiation,

$$\frac{\partial\mathcal{T}}{\partial M_Y} = \Theta\zeta^{\frac{\sigma_p-1}{\sigma_p}} M_Y^{-1/\sigma_p}\mathcal{T}^{1/\sigma_p}.$$

Substituting gives the expression used in the text.

B.2 Proof of Proposition 1 (Impossibility under Production Automation)

With constant AI efficiency $\eta = \bar{\eta}$, effective research input is

$$R = L_A + \bar{\eta}M_A.$$

Idea growth satisfies

$$g_A = \delta R^{\lambda} A^{\phi-1}.$$

Log-differentiating:

$$\frac{\dot{g}_A}{g_A} = \lambda\frac{\dot{R}}{R} + (\phi-1)g_A.$$

Therefore

$$\dot{g}_A = g_A\left[\lambda g_R^{ext} + (\phi-1)g_A\right].$$

The quadratic coefficient is $(\phi-1) < 0$ for $\phi < 1$. Hence g_A cannot explode in finite time. The automation frontier Θ does not appear in this equation.

□

B.3 Proof of Proposition 2 (Threshold Preservation)

With endogenous efficiency,

$$\frac{\dot{\eta}}{\eta} = \beta\frac{\dot{\mathcal{C}}}{\mathcal{C}} - \gamma\frac{\dot{M}}{M}.$$

Under the quasi-steady relations of the baseline,

$$\frac{\dot{\mathcal{C}}}{\mathcal{C}} = g_M = g_A + o(g_A),$$

so

$$\frac{\dot{\eta}}{\eta} = (\beta - \gamma)g_A + o(g_A).$$

Effective research input is

$$R = L_A + \eta M_A.$$

Using $\psi_A = \eta M_A/R$,

$$\frac{\dot{R}}{R} = (1 - \psi_A)n + \psi_A\left(\frac{\dot{\eta}}{\eta} + \frac{\dot{M}_A}{M_A}\right).$$

With $M_A = \phi_A M$ and ϕ_A constant along the scaling path,

$$\frac{\dot{M}_A}{M_A} = g_M = g_A + o(g_A).$$

Substituting gives

$$\frac{\dot{R}}{R} = (1 - \psi_A)n + \psi_A(1 + \beta - \gamma)g_A + o(g_A).$$

Insert into the log-differentiated idea equation:

$$\frac{\dot{g}_A}{g_A} = \lambda\frac{\dot{R}}{R} + (\phi - 1)g_A.$$

Collecting terms yields

$$\dot{g}_A = g_A\left[\frac{\lambda(1 - \psi_A)n}{1 - \phi} + \frac{\lambda\psi_A(1 + \beta - \gamma) + \phi - 1}{1 - \phi}g_A\right] + o(g_A^2).$$

Setting the quadratic coefficient to zero yields

$$\psi_A^{**} = \frac{1 - \phi}{\lambda(1 + \beta - \gamma)}.$$

□

B.4 Proof of Proposition 3 (Channel Decomposition)

Differentiate $\psi_A = \eta M_A/(L_A + \eta M_A)$ with respect to Θ:

$$\frac{d\psi_A}{d\Theta} = \frac{\partial\psi_A}{\partial M}\frac{\partial M}{\partial Y}\frac{\partial Y}{\partial\Theta} + \frac{\partial\psi_A}{\partial\phi_A}\frac{d\phi_A}{d\Theta}.$$

The first term is positive since $\partial Y/\partial\Theta > 0$ and higher output increases total AI investment.

The second term depends on how the no-arbitrage condition $q_A = q_Y$ adjusts ϕ_A. Differentiating that condition gives the elasticity condition stated in the main text.

□

B.5 Sufficient Condition for $d\psi_A/d\Theta > 0$

Totally differentiate $q_A(\phi_A, M) = q_Y(\phi_A, M, \Theta)$. After rearrangement, the sign of $d\phi_A/d\Theta$ is determined by the task substitution elasticity σ_p.

The critical elasticity threshold is

$$\bar{\sigma}_p = 1 + \frac{\chi\phi_A^*}{1-\chi(1-\phi_A^*)} \cdot \frac{V_A \delta\lambda\eta R^{\lambda-1} A^{\phi}}{q_M}.$$

If $\sigma_p < \bar{\sigma}_p$, then $d\phi_A/d\Theta \geq 0$ and production automation raises ψ_A.

□

A Proof of Proposition 13

Proof. We prove each part in turn.

Part (i): Closed-form expression.

From Proposition 4, when $\psi^j > \psi^{**}$ and the quadratic term dominates, the idea growth rate follows the Riccati equation

$$\dot{g}_A = a(\psi^j) g_A + \kappa_\gamma(\psi^j) g_A^2, \tag{173}$$

where $a(\psi^j)$ and $\kappa_\gamma(\psi^j)$ are defined in (154) and (155). Define $u(t) \equiv 1/g_A(t)$. Differentiating,

$$\dot{u} = -\frac{\dot{g}_A}{g_A^2} = -a(\psi^j) u - \kappa_\gamma(\psi^j). \tag{174}$$

This linear ODE has solution

$$u(t) = \left(u_0 + \frac{\kappa_\gamma(\psi^j)}{a(\psi^j)}\right) e^{-a(\psi^j)(t-t_0)} - \frac{\kappa_\gamma(\psi^j)}{a(\psi^j)}, \quad a(\psi^j) > 0, \tag{175}$$

with $u_0 \equiv 1/g_0$. The singularity occurs at the first time T^j when $u(T^j) = 0$, i.e., when $g_A(T^j) \to \infty$. Setting $u(T^j) = 0$ in (175) yields

$$\left(\frac{1}{g_0} + \frac{\kappa_\gamma(\psi^j)}{a(\psi^j)}\right) e^{-a(\psi^j)(T^j - t_0)} = \frac{\kappa_\gamma(\psi^j)}{a(\psi^j)}. \tag{176}$$

Solving for T^j gives

$$T^j = t_0 + \frac{1}{a(\psi^j)} \ln\left(1 + \frac{a(\psi^j)}{\kappa_\gamma(\psi^j) g_0}\right). \tag{177}$$

For the special case $a(\psi^j) = 0$ (which occurs when $\psi^j = 1$ or $n = 0$), equation (173) reduces to $\dot{g}_A = \kappa_\gamma(\psi^j) g_A^2$, which integrates directly to $g_A(t) = g_0/[1 - \kappa_\gamma(\psi^j) g_0 (t - t_0)]$, giving $T^j = t_0 + 1/(\kappa_\gamma(\psi^j) g_0)$. This establishes (153).

Part (ii): Monotonicity.

Differentiate (153) with respect to ψ^j. First note that

$$\frac{\partial a}{\partial \psi} = -\frac{\lambda n}{1-\phi} < 0, \tag{178}$$

$$\frac{\partial \kappa_\gamma}{\partial \psi} = \frac{\lambda(1+\beta-\gamma)}{1-\phi} > 0 \quad (\text{since } 1+\beta-\gamma > 0). \tag{179}$$

Define $F(a, \kappa) \equiv \frac{1}{a} \ln\left(1 + \frac{a}{\kappa g_0}\right)$. Then

$$\frac{\partial T}{\partial \psi} = \frac{\partial F}{\partial a}\frac{\partial a}{\partial \psi} + \frac{\partial F}{\partial \kappa}\frac{\partial \kappa}{\partial \psi}. \tag{180}$$

Computing partial derivatives:

$$\frac{\partial F}{\partial a} = -\frac{1}{a^2}\ln\left(1+\frac{a}{\kappa g_0}\right) + \frac{1}{a}\cdot\frac{1/\kappa g_0}{1+a/(\kappa g_0)}$$
$$= -\frac{1}{a^2}\ln\left(1+\frac{a}{\kappa g_0}\right) + \frac{1}{a(\kappa g_0+a)}, \tag{181}$$
$$\frac{\partial F}{\partial \kappa} = -\frac{1}{a}\cdot\frac{a/(\kappa^2 g_0)}{1+a/(\kappa g_0)} = -\frac{1}{\kappa(\kappa g_0+a)}. \tag{182}$$

Substituting (178), (179), (181), and (182) into (180):

$$\frac{\partial T}{\partial \psi} = \left[-\frac{1}{a^2}\ln\left(1+\frac{a}{\kappa g_0}\right) + \frac{1}{a(\kappa g_0+a)}\right]\left(-\frac{\lambda n}{1-\phi}\right)$$
$$+\left[-\frac{1}{\kappa(\kappa g_0+a)}\right]\left(\frac{\lambda(1+\beta-\gamma)}{1-\phi}\right). \tag{183}$$

Collecting terms and simplifying yields

$$\frac{\partial T}{\partial \psi} = -\frac{\lambda}{1-\phi}\left[\frac{n}{a^2}\ln\left(1+\frac{a}{\kappa g_0}\right) + \frac{1}{\kappa g_0+a}\left(\frac{1+\beta-\gamma}{\kappa}-\frac{n}{a}\right)\right]. \tag{184}$$

To sign this expression, note that from the definitions of a and κ,

$$\frac{1+\beta-\gamma}{\kappa}-\frac{n}{a} = \frac{1+\beta-\gamma}{\frac{\lambda\psi(1+\beta-\gamma)+\phi-1}{1-\phi}} - \frac{n}{\frac{\lambda(1-\psi)n}{1-\phi}} = \frac{1-\phi}{\lambda\psi(1+\beta-\gamma)+\phi-1} - \frac{1-\phi}{\lambda(1-\psi)}. \tag{185}$$

Since $\psi > \psi^{**}$, we have $\lambda\psi(1+\beta-\gamma)+\phi-1>0$ and $\lambda(1-\psi)>0$. Moreover, the inequality $\psi(1+\beta-\gamma) > 1-\psi$ holds for $\psi > 1/(2+\beta-\gamma)$, which is always satisfied when $\psi > \psi^{**}$ for plausible parameters. Hence the term in brackets in (184) is positive, making $\partial T/\partial\psi < 0$. Thus T^j is strictly decreasing in ψ^j, establishing (156).

Part (iii): SP-CE comparison.

From Proposition 12, when $\beta > 0$ and γ is not too large, the planner internalizes the compute externality, leading to $\psi^{SP} > \psi^{CE}$. By the monotonicity established in part (ii), $T(\psi)$ is decreasing, so

$$\psi^{SP} > \psi^{CE} \implies T^{SP} = T(\psi^{SP}) < T(\psi^{CE}) = T^{CE}, \tag{186}$$

which proves (157).

For the ratio (158), substitute the expression (153):

$$\frac{T^{CE}}{T^{SP}} = \frac{\frac{1}{a(\psi^{CE})}\ln\left(1+\frac{a(\psi^{CE})}{\kappa_\gamma(\psi^{CE})g_0}\right)}{\frac{1}{a(\psi^{SP})}\ln\left(1+\frac{a(\psi^{SP})}{\kappa_\gamma(\psi^{SP})g_0}\right)}$$
$$= \frac{\kappa_\gamma(\psi^{SP})}{\kappa_\gamma(\psi^{CE})}\cdot\frac{\frac{\kappa_\gamma(\psi^{CE})}{a(\psi^{CE})}\ln\left(1+\frac{a(\psi^{CE})}{\kappa_\gamma(\psi^{CE})g_0}\right)}{\frac{\kappa_\gamma(\psi^{SP})}{a(\psi^{SP})}\ln\left(1+\frac{a(\psi^{SP})}{\kappa_\gamma(\psi^{SP})g_0}\right)}. \tag{187}$$

Define $h(x) \equiv \frac{\ln(1+x)}{x}$ for $x>0$, where $x^j \equiv \frac{a(\psi^j)}{\kappa_\gamma(\psi^j)g_0}$. Then (187) becomes

$$\frac{T^{CE}}{T^{SP}} = \frac{\kappa_\gamma(\psi^{SP})}{\kappa_\gamma(\psi^{CE})}\cdot\frac{h(x^{CE})}{h(x^{SP})}. \tag{188}$$

Since $h(x)$ is decreasing in x (as $\frac{d}{dx}\ln h(x) = \frac{1}{x\ln(1+x)} - \frac{1}{x^2} < 0$ for $x>0$), and x is decreasing in ψ (because $a(\psi)$ decreases while $\kappa_\gamma(\psi)$ increases with ψ), we have $x^{CE} > x^{SP}$ and therefore $h(x^{CE}) < h(x^{SP})$. Consequently,

$$\frac{T^{CE}}{T^{SP}} > \frac{\kappa_\gamma(\psi^{SP})}{\kappa_\gamma(\psi^{CE})} > 1, \tag{189}$$

establishing (158).

Part (iv): Asymptotic approximation.

When $g_0 \gg a(\psi^j)/\kappa_\gamma(\psi^j)$, we have $x^j = a(\psi^j)/(\kappa_\gamma(\psi^j)g_0) \ll 1$. Using the expansion $\ln(1+x) = x - x^2/2 + O(x^3)$,

$$\frac{1}{a(\psi^j)} \ln\left(1 + \frac{a(\psi^j)}{\kappa_\gamma(\psi^j)g_0}\right) = \frac{1}{a(\psi^j)}\left(\frac{a(\psi^j)}{\kappa_\gamma(\psi^j)g_0} - \frac{1}{2}\left(\frac{a(\psi^j)}{\kappa_\gamma(\psi^j)g_0}\right)^2 + \cdots\right) = \frac{1}{\kappa_\gamma(\psi^j)g_0} - \frac{a(\psi^j)}{2\kappa_\gamma(\psi^j)^2 g_0^2} + \cdots \tag{190}$$

Thus $T^j = t_0 + 1/(\kappa_\gamma(\psi^j)g_0) + O(1/g_0^2)$, establishing (159). The ratio approximation (160) follows directly by substituting the leading-order terms. □