Speculative-Growth and the AI "Bubble"

Ricardo J. Caballero*

May 13, 2026

Link to the latest version

Abstract

AI technology can generate speculative-growth equilibria. These equilibria are rational but fragile: elevated valuations support rapid capital accumulation, yet persist only as long as beliefs remain coordinated. The key force is a funding feedback—a saving-rate channel through which AI-driven redistribution toward capitalists lowers the required return. As AI capital performs tasks previously done by labor, the labor share falls and capitalist wealth rises. Because capitalists have a higher propensity to save, this wealth shift raises aggregate saving and helps validate the high valuations that support rapid investment. The task-based structure of AI also sustains the return to AI capital over a wider range of accumulation. Building on Caballero et al. (2006), the paper shows that these linked effects can produce multiple equilibria and a speculative transition toward a high-capital equilibrium. Interest rates remain elevated early in that transition, relative to their eventual low steady-state level, and fall sharply only later. At the same time, the equilibrium is fragile: a loss of confidence can precipitate a self-fulfilling crash and reversal. Higher AI productivity can make the high-capital outcome the unique equilibrium, while upward shifts in the funding schedule or faster obsolescence can eliminate the high-capital steady state entirely.

JEL Codes: E21, E22, E24, E44, O33, O41

Keywords: AI, speculation, investment boom, crash, multiple equilibria, labor share, non-homothetic preferences, saving rate, market capitalization, low interest rates, productivity, creative-destruction

*MIT and NBER. Email: caball@mit.edu. I thank Alp Simsek and Ludwig Straub for their comments, and Kei Uzui for excellent research assistance. First draft: December 17, 2025.

1 Introduction

High valuations of AI-related firms are often framed in binary terms: either they reflect fundamentals or they reflect a bubble. This paper develops a framework in which that distinction is too rigid. Along one equilibrium path, elevated valuations are consistent with fundamentals and raise Tobin's q, supporting rapid capital accumulation. Yet that same path is fragile, because it depends on coordinated beliefs about future profitability and investment.

The scale of the episode is already large on both sides of the asset-market and real-investment ledger. On the valuation side, estimates attributed to Goldman Sachs suggest that the market value of AI-related companies has risen by roughly $19 trillion since the public release of ChatGPT (Fortune, 2025).[1] On the investment side, Van Nieuwerburgh (2026) estimates that announced and planned U.S. data-center developments could imply roughly $8.2 trillion of spending on data centers, power infrastructure, and IT equipment over 2026–2032. Spread over the period, this would amount to about 2.8 percent of GDP per year, larger than the average annual investment shares associated with the railroad, electrification, interstate-highway, and telecom/fiber buildouts.[2] These numbers sharpen the central question: can exceptionally high valuations and exceptionally large investment be mutually validating along an equilibrium transition, or must one of them eventually give way?

The analysis builds on the speculative-growth framework of Caballero et al. (2006, henceforth CFH). In CFH, sustained high valuations require a *funding feedback*: as the economy moves toward the high-capital outcome, the interest rate must eventually decline enough to validate high asset prices. The contribution here is to show why AI technology can generate that feedback in a natural way, and to characterize the transition dynamics that follow. The key primitive is that AI capital performs labor tasks, and this generates two complementary effects:

1. **Distribution-side effect: lower labor share and stronger saving.** By performing tasks previously done by labor, AI capital shifts income away from labor and toward capital as AI is deployed. As wealth concentrates on capitalists, aggregate saving rises with wealth and the interest rate falls (Straub, 2019; Mian, Straub, and Sufi, 2021).

[1] A more conservative way to state the same concentration fact is that the largest technology firms have recently accounted for more than 30 percent of the market value of the S&P 500; see Goldman Sachs Global Institute (2025).

[2] Van Nieuwerburgh (2026) reports average annual investment shares of 2.4 percent of GDP for railroads, 1.1 percent for electrification, 1.6 percent for interstate highways, and 0.8 percent for telecom/fiber, compared with 2.8 percent for projected AI infrastructure. McKinsey (2025) provides a complementary global industry estimate of a $7 trillion data-center buildout.

This provides the funding feedback that can sustain elevated valuations at high levels of capital.

2. **Production-side effect: sustained returns during the buildout.** Because AI can substitute for labor across a broad range of tasks (e.g., Restrepo 2025), effective labor expands alongside AI capital, which keeps the capital-to-effective-labor ratio stable and attenuates the usual decline in the marginal product of capital. This supports high valuations over a wider range of accumulation and makes the high-capital outcome easier to sustain.

A third ingredient governs the speed and valuation intensity of the *transition*:

3. **Intermediate adjustment costs.** Adjustment costs must lie in an intermediate range: high enough that valuations can depart from replacement cost during the transition, yet low enough to permit rapid accumulation. When this condition holds, elevated valuations can sustain a transition toward the high-capital equilibrium.

When these forces are present, the economy can feature both a low-capital and a high-capital equilibrium. The distinction is not a permanently higher valuation ratio in the long run—in steady state, valuations return to replacement cost—but a different equilibrium transition. Under coordinated optimism, elevated valuations support rapid accumulation toward the high-capital outcome. Under pessimism, the same economy can reverse course and return to the low-capital equilibrium.

A useful implication of the model concerns the path of interest rates. The channel is inherently intertemporal: high valuations raise Tobin's q and support investment during the transition, while the expectation of lower interest rates in the high-capital equilibrium helps validate those valuations ex ante. The model therefore implies a two-stage transition. Early in the boom, resources are absorbed by the investment boom, so consumption adjusts with a lag; as consumption catches up, its growth remains elevated, which can keep the interest rate high for a time even as valuations rise. Later, as wealth shifts toward high-saving capitalists and aggregate saving increases, the funding feedback strengthens and the interest rate falls sharply. This sequencing is central to the paper's interpretation of speculative growth: high valuations and high investment can coexist with interest rates that are initially elevated relative to the eventual low-rate equilibrium.

Whether that destination exists at all depends on the parameter environment. Higher AI productivity strengthens both the distributional saving channel and the return to AI capital, and can make the high-capital outcome the unique equilibrium. Upward shifts in the funding schedule—for example from fiscal pressure or higher borrowing spreads—or faster obsolescence can eliminate the high-capital steady state entirely.

Related literature

The technological environment is motivated by recent work emphasizing that advanced AI should not be modeled only as a Hicks-neutral productivity shock. Instead, AI changes which tasks can be performed by reproducible capital rather than by human labor. The canonical task-based formalization is Acemoglu and Restrepo (2018), in which automation expands the set of tasks performed by capital, shifts income toward capital, and can sustain the marginal product of capital. This task-based perspective also underlies Restrepo (2025), who studies an AGI economy in which compute can perform economically valuable work, and Korinek and Suh (2024), who analyze transition paths as increasingly complex tasks become automatable. It is central to broader assessments of transformative AI and growth in Trammell and Korinek (2023), the benchmark discipline imposed by Acemoglu (2024), and the more optimistic growth accounting in Aghion and Bunel (2024). Jones and Tonetti (2026) emphasize that weak links and bottlenecks can tame the growth explosion even when many tasks are automated, while Brynjolfsson et al. (2025a) place transition dynamics, distribution, and institutional adaptation at the center of a broader research agenda for transformative AI. I share with this literature the premise that AI capital can perform labor tasks and that bottlenecks matter. The object here, however, is not to forecast the balanced-growth effect of AI. It is to use labor-like AI as the technological environment in which a speculative-growth equilibrium can arise.

The funding-feedback ingredient is motivated by work connecting inequality, saving, and low interest rates. The reduced-form consumption rule follows Straub (2019), in which richer households have lower marginal propensities to consume and aggregate saving rises with wealth concentration. The same logic is closely related to Mian, Straub, and Sufi (2021), who emphasize the macroeconomic importance of a saving glut of the rich for low interest rates. In this paper, AI provides an endogenous reason for such concentration: as AI substitutes for labor tasks, income shifts toward capital owners. The novelty is not the saving mechanism by itself, but its interaction with AI production and speculative-growth dynamics.

The transition mechanism uses adoption frictions. The speculative-growth transition requires adjustment costs in an intermediate range: high enough that valuations can depart from replacement cost, but low enough that accumulation can respond to those valuations. McElheran et al. (2025) provide useful empirical motivation for this ingredient by documenting J-curve dynamics in industrial AI adoption: short-run performance losses can precede medium-run productivity gains, consistent with complementary organizational investment and implementation frictions.

A final set of papers connects AI to asset prices and interest rates. This connection

has two subthemes. At the firm level, Babina et al. (2024) show that firms investing in AI experience faster growth in sales, employment, product innovation, and market valuations. In bond markets, Andrews and Farboodi (2025) study how news about AI progress is reflected in long-term interest rates. Björkegren (2025) provides related market evidence by separating closed/proprietary and open-weight AI model releases. This evidence is consistent with some of the implications of the model: AI news should affect both valuations and long rates, and the rate response can depend on how AI rents are distributed.

2 Model: AI Technology and the Funding Feedback

This section presents a model built around the AI labor primitive and the funding feedback described above, and includes adjustment costs, which shape transition dynamics. Derivations and supporting details are in the appendices: Appendix A derives the MPK schedule, Appendix B microfounds the consumption rule, and Appendix C collects the equilibrium dynamics.

2.1 Technology

In a standard neoclassical model, the return to capital falls steadily as capital accumulates. AI technology differs in an important respect. Following the task-based approach of Acemoglu and Restrepo (2018), as synthesized for AGI by Restrepo (2025), production involves many discrete tasks. Some tasks are performed by workers, others by machines. Traditional capital can only perform "machine tasks," but AI, like robotization, can also perform worker tasks.

This distinction has implications for both factor shares and diminishing returns. As AI capital accumulates, it does not merely add machines alongside a fixed labor force. Instead, AI operates *as labor*, expanding effective labor (now comprising both humans and AI) that works alongside conventional capital. Because the additional effective labor is owned by capitalists, deployment shifts income toward capital owners. It also generates a constant-MPK region. In this "AI deployment" region, each additional unit of AI capital adds effective labor, keeping the effective capital-labor ratio constant. Since the MPK depends on this ratio, diminishing returns are forestalled.

The technology is embedded in a simple continuous-time model. Output is produced with capital and labor:

$$Y = AK_c^{\alpha} N^{1-\alpha},$$

where K_c is conventional capital, N is effective labor, and $\alpha \in (0, 1)$.

Capital can be used in two ways: as conventional capital K_c or as AI capital K_ℓ that substitutes for labor. Total capital is $K = K_c + K_\ell$. AI capital produces "AI labor" at a rate of γ per unit, so effective labor is $N = 1 + \gamma K_\ell$. However, AI deployment faces a capacity constraint $\bar{K}_\ell$—reflecting limits on data, compute, or organizational capacity (e.g., Jones and Tonetti 2026).

Firms allocate capital optimally between the two uses. As shown in Appendix A, this generates the three-region MPK schedule in Figure 1:

- **Region I** ($K < K_{\text{AI}}$): No AI deployment. Standard diminishing returns $r^K = \alpha A K^{\alpha-1}$.
- **Region II** ($K_{\text{AI}} \leq K < K_{\text{sat}}$): AI deployment phase. The MPK is constant at $r^K = \alpha A \left(\frac{(1-\alpha)\gamma}{\alpha}\right)^{1-\alpha}$.
- **Region III** ($K \geq K_{\text{sat}}$): AI saturated at $K_\ell = \bar{K}_\ell$. Diminishing returns resume $r^K = \alpha A (K - \bar{K}_\ell)^{\alpha-1}(1 + \gamma \bar{K}_\ell)^{1-\alpha}$.

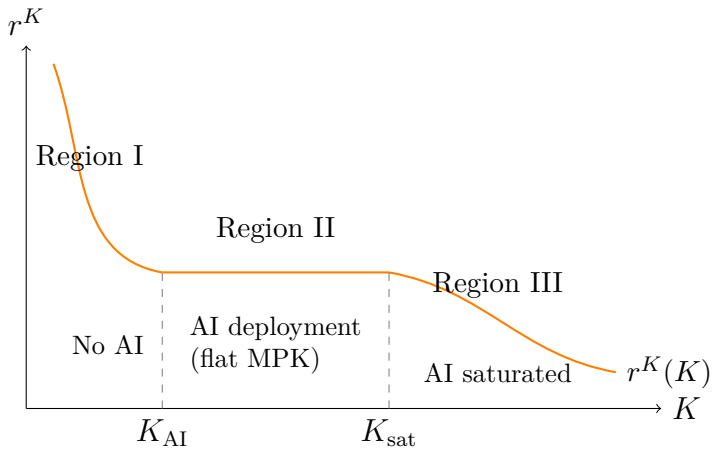

Figure 1: The marginal product of capital with AI technology. In Region II, AI deployment keeps the MPK flat as capital accumulates.

The thresholds K_{AI} and K_{sat} mark, respectively, the onset of AI deployment and the full utilization of AI capacity. Both are derived in Appendix A.

2.2 Households

The funding feedback is formalized through a two-group household structure: workers and capitalists.

Workers supply labor, earn wages, hold no assets, and consume their entire income: $c_w = w$. Worker consumption is therefore absorbed by the wage-bill component of output, while asset pricing and accumulation can be summarized by capitalist wealth.

Capitalists own all capital and have non-homothetic preferences. Following Straub (2019), their consumption is given by:

$$c = \kappa W^{\phi}, \qquad \kappa > 0, \quad 0 < \phi < 1, \tag{1}$$

where W is wealth and κ, ϕ are parameters. The key feature is $\phi < 1$: consumption rises less than proportionally with wealth. Equivalently, the saving rate rises with wealth.

This specification is a tractable approximation to optimal behavior under non-homothetic preferences. Appendix B provides microfoundations and shows how to choose κ and ϕ so that the isoelastic rule matches the exact steady-state consumption policy locally at a reference wealth level.

The declining funding schedule is the central implication of the household block. In steady state, the Euler condition gives $R = \rho - \theta c$ while the capitalist budget constraint gives $c = RW$. Combining the two,

$$c^{ss}(W) = \frac{\rho W}{1 + \theta W}, \qquad R^{ss}(W) = \frac{\rho}{1 + \theta W}, \qquad \frac{dR^{ss}(W)}{dW} = -\frac{\rho\theta}{(1 + \theta W)^2} < 0. \tag{2}$$

Thus, as capitalist wealth rises, desired saving rises and the required return falls.

2.3 Investment and Asset Pricing

Investment faces adjustment costs. Let q denote Tobin's q—the ratio of market value to replacement cost of capital. The investment rate responds to q, so that capital growth is driven by:

$$\frac{\dot{K}}{K} = \psi \ln q - \delta, \tag{3}$$

where $\psi > 0$ governs the responsiveness of investment to valuations and δ is the depreciation rate. When $q > 1$ (market value exceeds replacement cost), gross investment is positive; when $q < 1$, gross investment is negative (scrapping).

Asset pricing requires that the return on holding capital equals the required return, R:

$$\underbrace{\frac{\dot{q}}{q} - \delta}_{\text{capital gain}} + \underbrace{\frac{r^K(K)}{q}}_{\text{dividend yield}} = R. \tag{4}$$

Since there is no explicit risk in the model, R is the equilibrium interest rate. The funding feedback operates through the relation between capitalist wealth and the required return. The steady-state diagrams use the exact schedule in (2). For the transition dynamics I use the locally calibrated isoelastic rule in (1). Combining that rule with the Euler equation derived in Appendix C gives the transition return schedule

$$R(W) = \frac{\rho - \phi\kappa W^{\phi-1} - \theta\kappa W^{\phi}}{1-\phi}. \tag{5}$$

The $W^{\phi-1}$ term captures the consumption-growth component and can keep rates elevated early in the transition; the W^{ϕ} term captures the wealth-saving component and dominates near the high-capital steady state. Equation (5) is used in the phase diagrams and time paths, while (2) is used for steady-state comparisons.

3 Multiple Steady States

The AI production block and the funding feedback combine to produce multiple steady states. The production block shapes the return to AI capital as accumulation proceeds, while the funding feedback lowers the required return as capitalist wealth rises. This section characterizes these steady states; the next section asks whether and how the economy can transition between them.

At a steady state, investment exactly covers depreciation ($\dot{K} = 0$), which requires $\psi \ln q = \delta$, hence:

$$\bar{q} = e^{\delta/\psi}. \tag{6}$$

At this valuation, asset market clearing requires that the MPK equals the rental rate implied by the interest rate. Setting $\dot{q} = 0$ gives $r^K(K)/q = R^{ss} + \delta$, or equivalently $r^K(K) = [R^{ss} + \delta]q$. At a steady state where $q = \bar{q}$ and $W = \bar{q}K$, the condition is

$$r^K(K) = [R^{ss}(\bar{q}K) + \delta]\bar{q}. \tag{7}$$

The left side is the MPK; the right side is the rental rate implied by the interest rate.

Figure 2 plots both sides against K. The MPK follows the "down-flat-down" pattern from Figure 1. The rental rate schedule is strictly decreasing in K: higher capital means more wealth, which raises saving, lowers the steady-state interest rate $R^{ss}(\bar{q}K)$ through (2), and hence lowers the rental rate.

The three-crossing configuration has a simple interpretation. The flat portion of the AI-capital return schedule must lie below the funding cost at the beginning of the AI-deployment

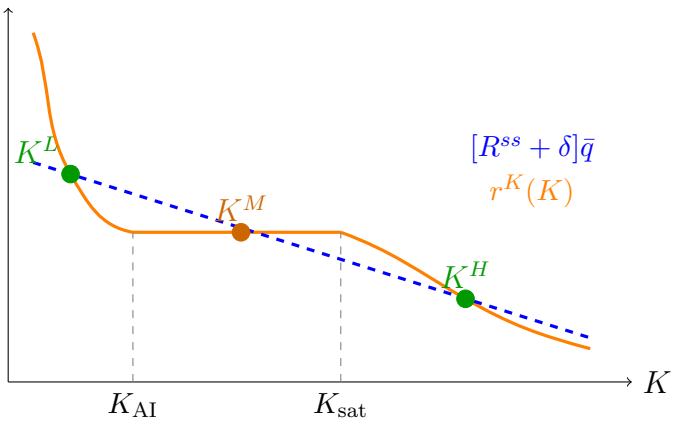

Figure 2: Multiple steady states. The economy admits K^L (low-capital, high interest rate) and K^H (high-capital, low interest rate) as stable equilibria. The middle intersection K^M is unstable.

region, but above it by the end of that region. The first inequality leaves a low-capital crossing before AI deployment; the second allows a high-capital crossing after AI saturation.

The curves cross three times, generating three steady states:

- K^L: Low-capital, no AI, high interest rate.
- K^M: Middle-capital, partial AI, intermediate interest rate. Unstable.
- K^H: High-capital, saturated AI, low interest rate.

The flat region in the MPK schedule enables three crossings. In Region II, the MPK holds steady while the required rental rate continues to fall with K. This allows the curves to cross, separate, and cross again.

Appendix D provides the formal conditions for three steady states and proves that K^L and K^H are saddle-path stable while K^M is unstable.

Although q equals $\bar{q}$ at both stable steady states, total market capitalization $\bar{q}K$ is substantially higher at K^H. The high-capital equilibrium thus features not only a larger capital stock but also greater aggregate wealth.

The existence of multiple steady states is only the first step. The key question is whether asset prices can coordinate a transition from the low-capital to the high-capital outcome.

4 Speculative Growth and Fragility

Section 3 established that multiple steady states can exist. But can the economy transition from K^L to K^H? Whether such a transition exists depends on the magnitude of adjustment

costs, which I discuss next.

4.1 The Role of Adjustment Costs

Multiple steady states are necessary but not sufficient for a speculative-growth transition. As in CFH, the issue is not simply whether both a low-capital and a high-capital steady state exist. It is whether the phase diagram contains a transition path that starts near the low-capital steady state with an upward jump in q and then joins the saddle path leading to the high-capital steady state.

This is a condition on the global phase diagram. Capital is predetermined, while q can jump. A speculative-growth path therefore requires that, starting from the low-capital steady state, an upward jump in q place the economy on a trajectory that subsequently converges to the high-capital steady state. The transition is valuation-driven only if that initial jump is large enough to generate sustained investment, but not so large that the resulting accumulation path is inconsistent with convergence to the high-capital destination.

4.2 The Speculative Growth Path

Figure 3 plots the (K, q) phase diagram. The dashed horizontal line is the $\dot{K} = 0$ locus, which lies at the steady-state valuation $\bar{q}$ defined in (6). The non-monotonic orange curve is the $\dot{q} = 0$ locus. The green curve is the stable manifold of the high-capital steady state $(K^H, \bar{q})$; it describes the speculative-growth trajectory.

Capital cannot jump, but asset prices can. Starting from the low-capital steady state $(K^L, \bar{q})$, a speculative-growth episode proceeds as follows:

1. **Expectations shift.** Agents coordinate on optimistic beliefs.
2. **Valuations jump.** q rises discretely from $\bar{q}$ to $q_0 > \bar{q}$.
3. **Investment booms.** The rise in q makes investment profitable and capital starts to accumulate.
4. **AI is deployed.** As K crosses K_{AI}, firms deploy AI. By expanding effective labor, AI raises the capital share and concentrates wealth among capitalists, while also sustaining the return to AI capital over the deployment range.
5. **The interest rate eventually falls.** As capitalists become wealthier, their saving rate rises, strengthening the funding feedback and lowering the interest rate.

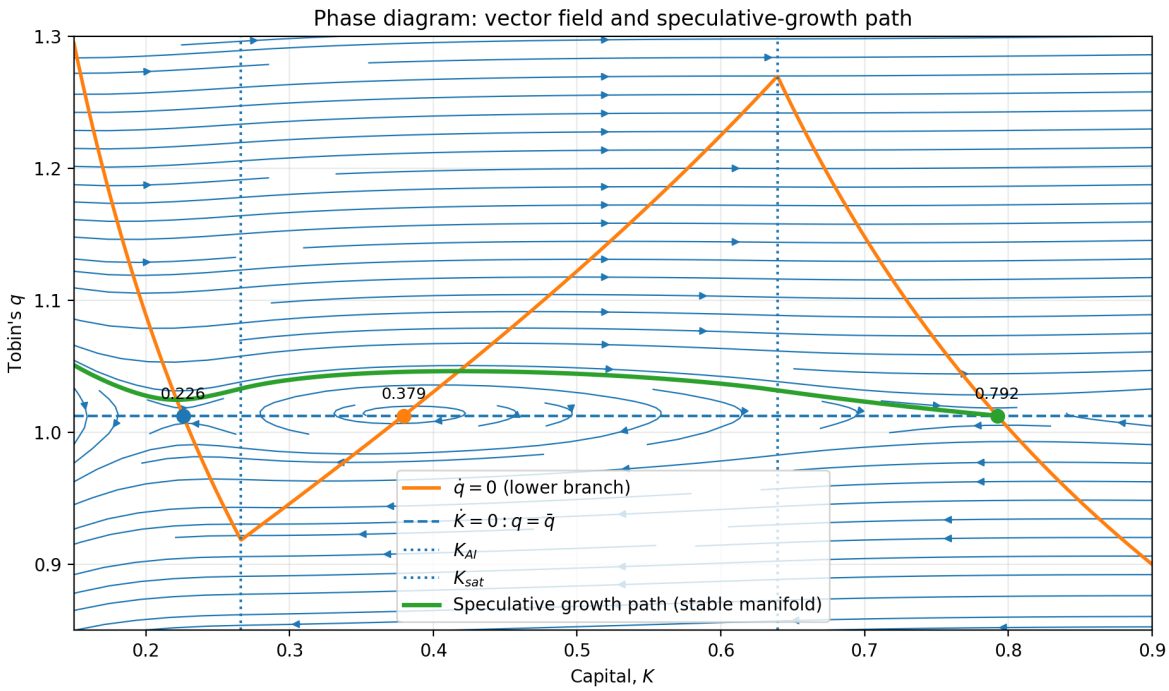

Figure 3: Phase diagram and speculative-growth path. Starting from $(K^L, \bar{q})$, optimistic beliefs trigger an upward jump in q. The economy then converges along the stable manifold (green) to $(K^H, \bar{q})$.

6. **Convergence.** The economy converges to $(K^H, \bar{q})$: capital reaches its high steady state and valuations eventually return to $\bar{q}$, now consistent with a lower interest rate.

The phase diagram also makes clear why elevated valuations are integral to the transition. At $(K^L, \bar{q})$ the economy is at rest; to induce capital accumulation one must have $q > \bar{q}$. Moreover, reaching K^H requires staying on the stable manifold, which lies above $\bar{q}$ throughout the transition. *High valuations are therefore not a symptom of irrational exuberance; they are the equilibrium force that makes the transition feasible.*

Remark: adjustment costs and transition geometry. The relevant object is geometric. Starting from $(K^L, \bar{q})$, the valuation jump must place the economy on a trajectory that reaches the stable arm converging to $(K^H, \bar{q})$. Adjustment costs shape both arms: when they are low, investment responds strongly to a valuation premium; when they are high, capital responds only slowly. A speculative-growth path obtains for an intermediate range in which the trajectory leaving the low-capital steady state meets the decreasing-q stable arm of the high-capital steady state. Appendix F formalizes this as a slope condition on the two arms.

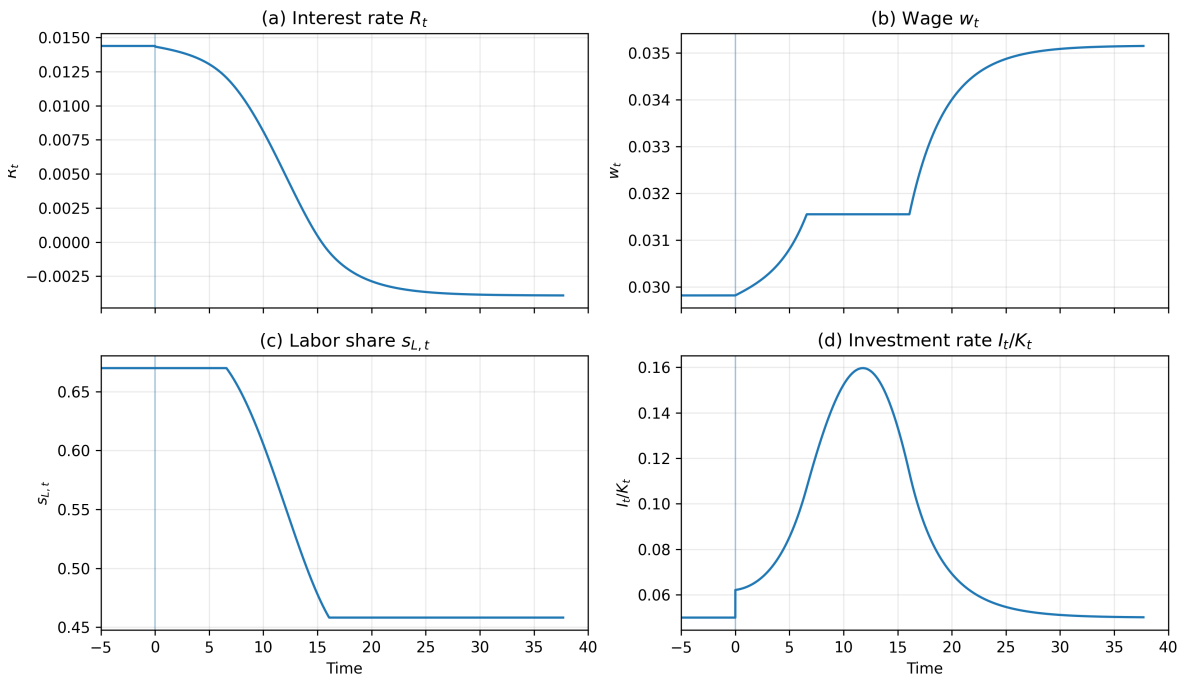

Figure 4: Time paths along the speculative-growth trajectory. The economy starts at the low-capital steady state and jumps onto the speculative-growth manifold. Panel (a): interest rate R_t; Panel (b): wage w_t; Panel (c): labor share $s_{L,t}$; Panel (d): investment rate I_t/K_t.

4.3 Time Paths Along the Speculative Growth Trajectory

Figure 4 shows the corresponding time paths along the speculative-growth trajectory. The figure includes a short pre-jump interval (so the initial discrete jump is visible) and then traces the post-jump dynamics on the speculative-growth stable manifold. The model implies a two-stage transition. First, elevated valuations raise Tobin's q and stimulate investment before the low-rate high-capital outcome has materialized. Later, the wealth-saving feedback strengthens and the interest rate falls. The key real-side channel is the investment response to valuations: the jump in q_t produces an immediate increase in the investment rate. Along the manifold, q_t may continue to rise for a time (even as K_t increases) because the prospect of lower future interest rates feeds back into asset prices. Eventually, once the decline in rates has done most of its work, q_t peaks and mean-reverts toward $\bar{q}$ as the economy approaches the high-capital steady state.

Interest Rate. The interest-rate path is central to the speculative-growth mechanism. The high-capital steady state is a low-interest-rate equilibrium, but the transition toward it need not begin with low interest rates. The reason is that the Euler equation contains two forces that move in opposite directions along the transition. With the locally calibrated

consumption rule in (1), evaluated at $W_t = q_t K_t$, the capitalist Euler equation is

$$R_t = \rho + \frac{\dot{c}_t}{c_t} - \theta c_t. \tag{8}$$

The term $\dot{c}_t/c_t$ is the consumption-growth force: when the boom raises future consumption growth, it raises the required return. The term $-\theta c_t$ is the wealth-saving force: as capitalist wealth rises and desired saving increases, it lowers the required return.

The decomposition can be written directly in terms of wealth. Since

$$\frac{\dot{c}_t}{c_t} = \phi \frac{\dot{W}_t}{W_t}$$

and aggregate capitalist wealth evolves according to

$$\dot{W}_t = R_t W_t - c_t,$$

we have

$$\frac{\dot{c}_t}{c_t} = \phi \left(R_t - \kappa W_t^{\phi - 1} \right).$$

Substituting into (8) recovers the transition return schedule in (5). Differentiating that schedule,

$$R'(W) = \phi \kappa W^{\phi - 2} \left[1 - \frac{\theta W}{1 - \phi} \right]. \tag{9}$$

Equation (9) clarifies how wealth affects the required return through the Euler equation. In the calibration underlying Figure 4, the relevant portion of $R(W)$ is downward sloping: higher capitalist wealth lowers the required return. The time path of the interest rate, however, depends not only on this slope but also on the speed at which wealth accumulates,

$$\dot{R}_t = R'(W_t) \dot{W}_t.$$

Early in the transition, wealth is still rising slowly, so the decline in the required return is muted. Elevated valuations raise investment and support future consumption growth, keeping the Euler-equation growth term elevated.

As AI deployment expands, the distributional force becomes stronger. AI lowers the labor share and raises capitalist wealth; because capitalists have a higher propensity to save, this increases aggregate saving and strengthens the downward pressure on the required return. Wealth then accumulates more rapidly, so the interest rate falls more sharply. As the economy approaches K^H and $q_t \to \bar{q}$, wealth stabilizes and R_t converges to its lower

high-capital steady-state level.

Wages and the labor share. Wages w_t and the labor share $s_{L,t}$ depend on the economy's effective labor N_t. As the trajectory enters the AI-deployment region, N_t rises because AI expands effective labor. Output increases while wages initially stagnate, compressing the labor share:

$$s_{L,t} = \frac{w_t L}{Y_t} = \frac{1-\alpha}{N_t}.$$

Once AI saturates (Region III), N_t stabilizes and the labor share converges to a permanently lower level, while wages resume rising with capital deepening.

Investment dynamics. The path of I_t/K_t follows the evolution of q_t along the speculative-growth manifold. The initial jump in q_t produces an immediate increase in investment. In Region I, investment continues to rise as valuations build in the expected transition toward the high-capital regime. In Region II, AI deployment both sustains the return to AI capital and strengthens the wealth-saving feedback, so valuations remain elevated despite ongoing capital deepening and investment remains high. Eventually, as q_t peaks and begins to mean-revert toward $\bar{q}$, the investment rate declines gradually and converges to its steady-state level.

4.4 Fragility: The Crash

The speculative-growth path is fragile. The same feedback that enables the boom also enables its reversal.

Consider an economy partway through the transition to the high-capital equilibrium: capital has accumulated to some $K > K^L$ and valuations remain elevated at $q > \bar{q}$. Suppose confidence weakens—due to negative news about AI capabilities, a financial shock, or a shift in sentiment.

If valuations decline—even absent any change in fundamentals—the economy can depart from the speculative-growth path. A sufficiently large decline places the economy on the only alternative equilibrium path—one converging to K^L rather than K^H.

Figure 5 illustrates this scenario. The red segment represents a crash: q drops discretely to the stable manifold associated with the low-capital steady state. The red path shows the subsequent dynamics: investment collapses, capital decumulates, and the economy returns to K^L.

The crash is self-fulfilling: a downward revision in beliefs lowers valuations today, which reduces investment and reverses capital accumulation. The weaker capital path then validates the pessimistic beliefs.

This analysis clarifies the sense in which AI valuations can be simultaneously "not a

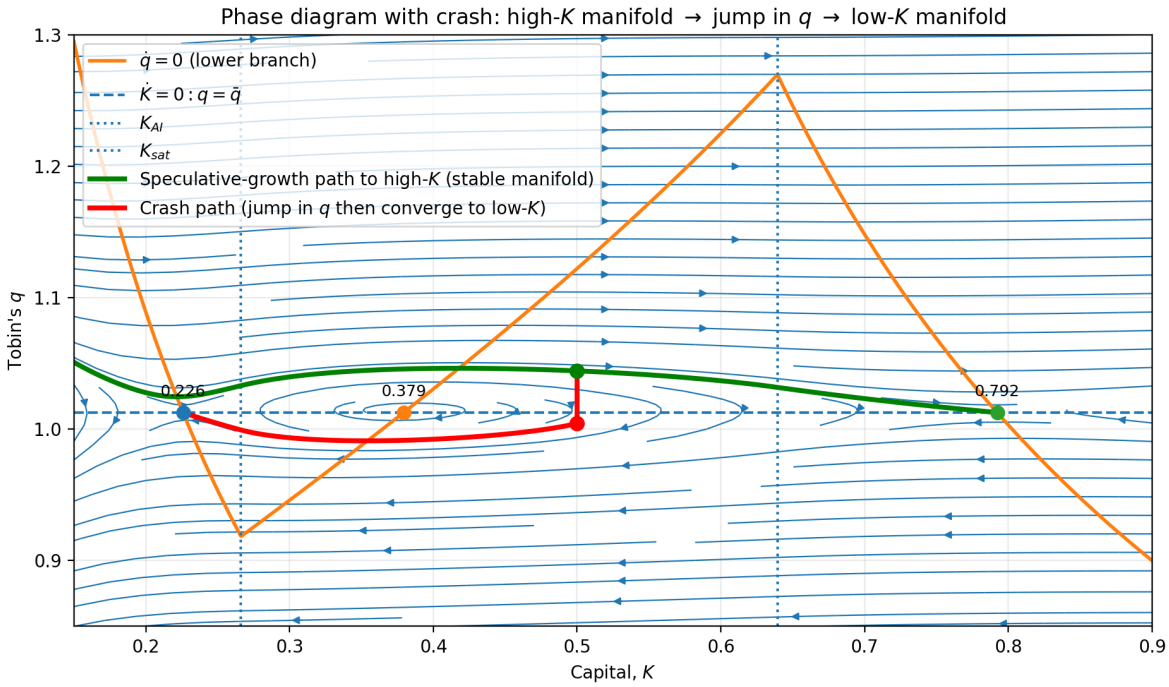

Figure 5: A crash along the speculative-growth path. A drop in valuations places the economy on a trajectory (red) leading back to K^L.

bubble" and fragile. They are not a bubble in the traditional sense because the growth and wealth they generate can ultimately validate those valuations. They are fragile because that validation requires sustained confidence throughout a potentially lengthy transition.

5 Comparative Statics

This section uses only the steady-state condition. It asks how changes in the AI technology and in the funding environment alter the number and location of steady states. The central point is simple. Shifts in the MPK schedule, generated here by changes in the AI labor-equivalence parameter γ, can eliminate either the low-capital or the high-capital steady state: sufficiently high γ removes K^L, leaving only K^H, while sufficiently low γ removes K^H. Upward shifts in the funding burden, generated here by fiscal pressure, higher borrowing spreads, or faster creative destruction, can eliminate the high-capital destination that validates the speculative-growth path.

Recall that steady states are crossings of the MPK schedule and the funding schedule:

$$r^K(K;\gamma) = \mathcal{F}(K;\delta), \qquad \mathcal{F}(K;\delta) \equiv \left[R^{ss}(\bar{q}K) + \delta\right]\bar{q}, \qquad \bar{q} = e^{\delta/\psi}. \tag{10}$$

Multiplicity requires the flat-MPK segment to lie below the funding curve at the start of AI deployment but above it by the time AI capacity is exhausted:

$$\mathcal{F}(K_{\text{AI}};\delta) > r^K_{\text{flat}}(\gamma) > \mathcal{F}(K_{\text{sat}};\delta). \tag{11}$$

5.1 Shifts in the MPK Curve: Changes in AI Labor Equivalence

A higher γ shifts the MPK schedule upward and to the left. The flat segment rises,

$$r^K_{\text{flat}}(\gamma) = \alpha A\Big(\tfrac{(1-\alpha)\gamma}{\alpha}\Big)^{1-\alpha},$$

and both deployment thresholds move left together, preserving the flat-region width:

$$K_{\text{AI}}(\gamma) = \frac{\alpha}{(1-\alpha)\gamma}, \qquad K_{\text{sat}}(\gamma) = K_{\text{AI}}(\gamma) + \frac{\bar{K}_\ell}{1-\alpha}. \tag{12}$$

The analysis below uses one regularity condition: once AI capacity is exhausted, the MPK falls at least as steeply as the funding schedule. Appendix G states this formally; Appendix H verifies it for the numerical example.

When γ is high enough that

$$r^K_{\text{flat}}(\gamma) \geq \mathcal{F}(K_{\text{AI}}(\gamma);\delta), \tag{13}$$

the MPK is already above the funding schedule at the start of AI deployment. Since the funding schedule slopes down, the two curves do not cross on the flat segment, and K^L disappears. The key is that K_{AI} moves left with γ: once it passes below the former low-capital crossing, that crossing exits Region I and ceases to be a steady state. The high-capital crossing survives because the post-saturation MPK rises with γ.

When γ is low enough that

$$r^K_{\text{flat}}(\gamma) \leq \mathcal{F}(K_{\text{sat}}(\gamma);\delta), \tag{14}$$

the MPK exits the flat region already below the funding schedule. The regularity condition then rules out a crossing further along, so K^H disappears.

Figure 6 illustrates the implication. Starting from a multiplicity region, a higher γ both raises the flat MPK and moves K_{AI} leftward; jointly, these effects eliminate the low-capital steady state. A sufficiently low γ lowers the flat segment enough that the high-capital equilibrium disappears.

These comparative statics sharpen the interpretation of the AI boom. If realized AI

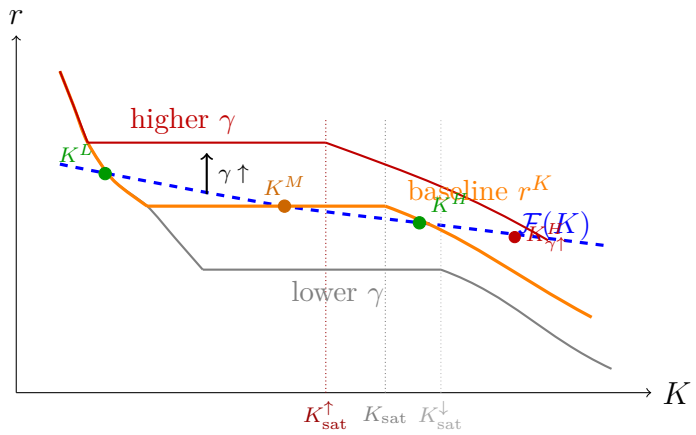

Figure 6: Comparative statics for AI labor equivalence. A higher γ shifts the flat-MPK segment up and left. For sufficiently high γ, the low-capital steady state disappears; for sufficiently low γ, the high-capital steady state disappears.

productivity is high enough, the economy no longer has a meaningful low-capital steady state: high capital accumulation is the only outcome. If AI productivity disappoints, the high-capital steady state may cease to exist and optimistic valuations cannot be validated.

5.2 Shifts in the Funding Curve: Fiscal Deficits, Borrowing Spreads, and Creative Destruction

Now hold the MPK schedule fixed and shift the funding schedule. Three forces are especially relevant.

First, fiscal deficits absorb saving that would otherwise help fund the high-capital equilibrium. I represent this by a funding wedge $\chi \geq 0$ added to R^{ss}:

$$R^{ss}(W;\chi) = \frac{\rho}{1+\theta W} + \chi, \qquad \mathcal{F}(K;\delta,\chi) = \big[R^{ss}(\bar{q}K;\chi) + \delta\big]\bar{q}, \qquad \frac{\partial \mathcal{F}}{\partial \chi} = \bar{q} > 0. \tag{15}$$

A larger fiscal claim on saving therefore shifts the funding curve up.

Second, the borrowing cost of the firms financing the AI buildout may rise as the scale of investment presses against internal funds and balance-sheet capacity. In reduced form, this raises the required return on AI capital by adding a borrowing-spread wedge $\zeta \geq 0$ to the funding cost:

$$R^{ss}(W;\chi,\zeta) = \frac{\rho}{1+\theta W} + \chi + \zeta, \qquad \mathcal{F}(K;\delta,\chi,\zeta) = \big[R^{ss}(\bar{q}K;\chi,\zeta) + \delta\big]\bar{q}, \qquad \frac{\partial \mathcal{F}}{\partial \zeta} = \bar{q} > 0. \tag{16}$$

The spread ζ should be interpreted broadly as any premium required to finance the AI

buildout once internal funds and balance-sheet capacity become scarce.[3] Like the fiscal wedge χ, a higher ζ shifts the funding schedule upward. It therefore makes the high-capital steady state harder to sustain and can eliminate the speculative-growth destination.

Third, creative destruction raises the effective depreciation burden on installed AI capital. I denote the obsolescence hazard by λ^{CD}. In steady state it acts like an increase in depreciation:

$$\delta^{\text{eff}} = \delta + \lambda^{CD}, \qquad \bar{q}^{\,\text{eff}} = e^{\delta^{\text{eff}}/\psi}, \qquad \mathcal{F}(K; \delta^{\text{eff}}, \chi, \zeta) = \left[R^{ss}(\bar{q}^{\,\text{eff}} K; \chi, \zeta) + \delta^{\text{eff}}\right] \bar{q}^{\,\text{eff}}. \quad (17)$$

The direct effect is to raise the required rental rate: capital must now cover ordinary depreciation plus the probability that a vintage becomes obsolete. All three forces shift the funding curve up.

The high-capital steady state disappears once

$$r^K_{\text{flat}}(\gamma) \leq \mathcal{F}(K_{\text{sat}}(\gamma); \delta^{\text{eff}}, \chi, \zeta). \quad (18)$$

At the boundary of AI saturation, the MPK is then already weakly below the funding schedule. Under the post-saturation regularity condition, the gap remains weakly negative thereafter, so there is no post-saturation crossing. Fiscal deficits and borrowing spreads operate through additive wedges to R^{ss}; creative destruction operates through δ^{eff} and the induced change in $\bar{q}^{\,\text{eff}}$. All three raise the funding burden and can eliminate the high-capital destination that validates the speculative-growth path.

Figure 7 summarizes these forces. An upward shift in $\mathcal{F}$ eliminates the high-capital destination that validates the speculative-growth path.

6 Conclusion

This paper develops a model of speculative growth in which elevated asset valuations and rapid capital accumulation can reinforce one another. In the AI context, this channel can produce multiple equilibria, including a high-capital equilibrium sustained by optimistic valuations and a fragile transition path vulnerable to self-fulfilling reversals.

[3]Recent market commentary is consistent with this reduced-form wedge. BIS (2026) document that hyperscalers' gross corporate bond issuance topped $100 billion in 2025 and that CDS spreads rose as AI-related borrowing expanded. MUFG (2025) reports that the "big five" hyperscalers' capital expenditures are forecast to exceed $600 billion in 2026, with roughly $450 billion tied directly to AI infrastructure, and that capital expenditures are increasingly outpacing free cash flow. Van Nieuwerburgh (2026) describes the Hyperion–Meta transaction as an example in which off-balance-sheet AI infrastructure debt carried a yield at least 100 basis points above comparable Meta unsecured debt. Chicago Fed (2026) emphasizes the associated bank-exposure tail risk.

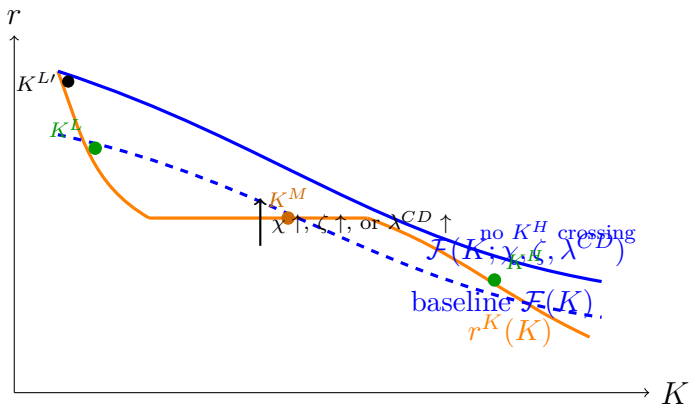

Figure 7: Comparative statics for the funding schedule. Fiscal deficits and borrowing spreads raise the required return by adding wedges to R^{ss}. Creative destruction raises effective depreciation, $\delta^{\text{eff}} = \delta + \lambda^{CD}$. All three raise the funding burden and can eliminate the high-capital equilibrium; the shifted schedule has no post-saturation crossing.

The model gives a qualified answer to the bubble question posed at the outset. Along the speculative-growth path, high AI valuations are not a bubble in the usual sense of prices detached from fundamentals: they support investment, raise future capital, and can be validated by the funding feedback generated by AI adoption. But they are fragile fundamentals. The same valuations are sustained only if beliefs coordinate the transition toward the high-capital destination and if the funding schedule remains low enough to validate that destination.

Three ingredients do the work. On the distribution side, AI deployment reduces the labor share and shifts income and wealth toward capitalists, whose saving rises with wealth. The resulting increase in aggregate saving lowers the interest rate in the high-capital equilibrium and helps validate elevated valuations. On the production side, the same task-based structure sustains the return to AI capital over a wider range of accumulation, supporting the investment path that carries the economy toward the high-capital outcome. On the transition side, intermediate adjustment costs allow valuations to depart from replacement cost while still generating enough accumulation to reach the high-capital steady state.

The model also implies a distinctive transition path for interest rates. Elevated valuations raise Tobin's q and stimulate investment during the transition, even though the lower interest rates that help support those valuations materialize only later. Early in the boom, resources are absorbed by the investment boom, so consumption adjusts with a lag; as consumption catches up, its growth remains elevated, which can keep the interest rate high for a time even as valuations rise. Interest rates fall sharply only later, as wealth accumulation and higher saving dominate near the high-capital equilibrium.

These results suggest that AI-driven valuation booms may be both more persistent and more fragile than standard narratives imply. They can be sustained by forward-looking beliefs about future productivity and future low interest rates, yet they remain vulnerable to shifts in confidence that can trigger self-fulfilling crashes and reversals.

To be clear, this is a possibility argument. My goal is to isolate a coherent force that could rationalize the joint behavior of valuations and investment, not to provide conclusive evidence. Nor do I mean to imply that the valuations and investment rates we are currently observing are fully consistent with a rational expectations model. The point is that narratives can matter when they coordinate beliefs around a feasible equilibrium path. Narratives that align with such a path can persist, not because agents understand the underlying model, but because the equilibrium itself sustains beliefs that happen to point in the right direction.

References

Acemoglu, Daron. 2024. "The Simple Macroeconomics of AI." *NBER Working Paper No. 32487.*

Acemoglu, Daron, and Pascual Restrepo. 2018. "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review* 108(6): 1488–1542.

Aghion, Philippe, and Simon Bunel. 2024. "AI and Growth: Where Do We Stand?" *FRBSF Policy Note*, June.

Andrews, Isaiah, and Maryam Farboodi. 2025. "Do Markets Believe in Transformative AI?" *NBER Working Paper No. 34243.*

Björkegren, Daniel. 2025. "Market Beliefs about Open vs. Closed AI." *arXiv:2512.14969.*

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. 2024. "Artificial Intelligence, Firm Growth, and Product Innovation." *Journal of Financial Economics* 151: 103745.

Brynjolfsson, Erik, Anton Korinek, and Ajay Agrawal. 2025. "A Research Agenda for the Economics of Transformative AI." *NBER Working Paper No. 34256.*

Korinek, Anton, and Donghyun Suh. 2024. "Scenarios for the Transition to AGI." *NBER Working Paper No. 32255.*

McElheran, Kristina, Mu-Jeung Yang, Zachary Kroff, and Erik Brynjolfsson. 2025. "The Rise of Industrial AI in America: Microfoundations of the Productivity J-Curve(s)." *Center for Economic Studies Working Paper No. 25-27.*

Mian, Atif R., Ludwig Straub, and Amir Sufi. 2021. "The Saving Glut of the Rich." *Quarterly Journal of Economics* 136(4): 2243–2307.

Trammell, Philip, and Anton Korinek. 2023. "Economic Growth under Transformative AI." *NBER Working Paper No. 31815.*

Caballero, Ricardo J., Emmanuel Farhi, and Mohamad L. Hammour. 2006. "Speculative Growth: Hints from the U.S. Economy." *American Economic Review* 96(4): 1159–1192.

Jones, Charles I., and Christopher Tonetti. 2026. "Past Automation and Future A.I.: How Weak Links Tame the Growth Explosion." *Stanford GSB Working Paper*, January.

Restrepo, Pascual. 2025. "We Won't be Missed: Work and Growth in the AGI World." *Yale Working Paper*, October.

Straub, Ludwig. 2019. "Consumption, Savings, and the Distribution of Permanent Income." *Harvard Working Paper*, June.

Eren, Egemen, Ingomar Krohn, and Karamfil Todorov. 2026. "Financing the AI Infrastructure Boom: On- and Off-Balance Sheet Borrowing." *BIS Quarterly Review*, March.

Cohen, Greg, Cooper Killen, and Simon Lau. 2026. "Tail Risk for Banks Posed by Investments in Generative Artificial Intelligence." *Chicago Fed Insights*, February.

Fortune. 2025. "Goldman Says the Stock Market Has Already Priced in the AI Boom." November 17.

Goldman Sachs Global Institute. 2025. "A Generational Infrastructure Buildout Might Hinge on AI Agents." Goldman Sachs, December.

McKinsey & Company. 2025. "The $7 Trillion Data Center Build-Out: How Industrials Can Capture Their Share."

MUFG. 2025. "AI Chart Weekly: Financing the AI Supercycle." December 19.

Van Nieuwerburgh, Stijn. 2026. "Financing the AI Buildout." Draft, March 19.

A Technology Details

A.1 Setup

Total capital $K = K_c + K_\ell$ is divided between conventional capital $K_c \geq 0$ and AI capital $K_\ell \geq 0$. Effective labor is

$$N = 1 + \gamma \min\{K_\ell, \bar{K}_\ell\},$$

where $\gamma > 0$ is the labor equivalence of AI and $\bar{K}_\ell$ is the AI capacity constraint.

Production is Cobb-Douglas: $Y = AK_c^\alpha N^{1-\alpha}$.

A.2 Optimal Allocation

Given K, firms choose K_c and K_ℓ to maximize output. At an interior solution:

$$\frac{\partial Y}{\partial K_c} = \frac{\partial Y}{\partial K_\ell} \quad \Rightarrow \quad \frac{\alpha}{K_c} = \frac{(1-\alpha)\gamma}{N}.$$

Define $b \equiv (1-\alpha)\gamma/\alpha$. The optimality condition becomes $N = bK_c$, yielding:

$$K_c = \frac{\gamma K + 1}{\gamma + b}, \qquad K_\ell = \frac{bK - 1}{\gamma + b}.$$

This interior solution is valid when $K_\ell \in [0, \bar{K}_\ell]$, i.e., when $K \in [K_{\text{AI}}, K_{\text{sat}}]$ where:

$$K_{\text{AI}} = \frac{1}{b} = \frac{\alpha}{(1-\alpha)\gamma}, \qquad K_{\text{sat}} = \frac{1 + (\gamma + b)\bar{K}_\ell}{b}.$$

A.3 The MPK Schedule

The marginal product of capital is:

$$r^K(K) = \begin{cases} \alpha A K^{\alpha-1}, & K < K_{\text{AI}}, \\ \alpha A b^{1-\alpha} \equiv r^K_{\text{flat}}, & K_{\text{AI}} \leq K < K_{\text{sat}}, \\ \alpha A (K - \bar{K}_\ell)^{\alpha-1}(1 + \gamma \bar{K}_\ell)^{1-\alpha}, & K \geq K_{\text{sat}}. \end{cases}$$

In Region II, the MPK is constant because as K rises, K_ℓ rises proportionally, keeping the capital-to-effective-labor ratio $K_c/N = 1/b$ constant.

B Microfoundations for the Consumption Rule

B.1 Preferences

Capitalists maximize:

$$\int_0^\infty e^{-\rho t} \left[\ln c_t + \theta W_t\right] dt,$$

subject to $\dot{W}_t = R_t W_t - c_t$. The term θW_t captures direct utility from wealth (wealth as a "luxury good").

B.2 Euler Equation

The Hamiltonian is $H = \ln c + \theta W + \mu(RW - c)$. First-order conditions:

$$\frac{\partial H}{\partial c} = 0 \quad \Rightarrow \quad \frac{1}{c} = \mu,$$

$$\dot{\mu} = \rho\mu - \frac{\partial H}{\partial W} = \rho\mu - \theta - \mu R.$$

Combining: $\dot{\mu}/\mu = \rho - \theta c - R$. Since $\mu = 1/c$, we have $\dot{\mu}/\mu = -\dot{c}/c$, yielding the Euler equation stated in (8).

B.3 Steady-State Consumption

At a steady state, $\dot{c} = \dot{W} = 0$, so $R^{ss} = \rho - \theta c^{ss}$ and $c^{ss} = R^{ss} W$. Solving these two equations gives the steady-state consumption and required-return schedules reported in (2).

B.4 The Isoelastic Approximation

The consumption-wealth elasticity at a steady state is:

$$\phi(W) \equiv \frac{d \ln c^{ss}}{d \ln W} = \frac{1}{1 + \theta W} \in (0, 1).$$

I approximate the optimal policy with $c = \kappa W^\phi$, calibrating ϕ and κ at a reference wealth W^*:

$$\phi = \frac{1}{1 + \theta W^*}, \qquad \kappa = \rho\phi(W^*)^{1-\phi}.$$

Equivalently, calibrating to match the exact steady-state policy at W^*,

$$\kappa = \frac{c^{ss}(W^*)}{(W^*)^\phi} = \frac{\rho W^*/(1 + \theta W^*)}{(W^*)^\phi} = \rho\,\phi\,(W^*)^{1-\phi},$$

where the last equality uses $1 + \theta W^* = 1/\phi$.

This approximation is exact at W^* and first-order accurate nearby. Taking $W^* = \bar{q}K^L$ (the low steady state) yields an upper bound on ϕ for the transition, since $\phi(W)$ is decreasing. The transition dynamics below use this locally calibrated isoelastic rule rather than the global exact steady-state schedule. This preserves the consumption-growth term in the Euler equation and is the source of the interest-rate dynamics shown in Figure 4.

C Equilibrium Dynamics

C.1 Laws of Motion

Capital accumulates according to (3). Multiplying by K gives the level form used in the numerical system.

C.2 Required Return

The transition system uses the locally calibrated isoelastic consumption rule from (1). The Euler-equation steps in Section 4 imply the transition return schedule (5). At the reference wealth level W^* used to calibrate κ and ϕ, this schedule coincides with the exact steady-state return in (2). Away from W^*, the transition system uses the isoelastic approximation.

C.3 Asset Pricing

The return condition is (4), with $R = R(qK)$ along the transition. Rearranging gives

$$\dot{q} = [R(qK) + \delta]q - r^K(K). \tag{19}$$

C.4 Phase Diagram Loci

The $\dot{K} = 0$ locus is $q = \bar{q}$, with $\bar{q}$ defined in (6).

The $\dot{q} = 0$ locus is given by setting the right-hand side of (19) equal to zero; it varies with the three-region MPK structure.

C.5 Output, Wages, and Labor Share

Output in each region:

$$Y = \begin{cases} AK^{\alpha}, & K < K_{\text{AI}}, \\ AK_c^{\alpha}N^{1-\alpha} \text{ with } K_c = \frac{\gamma K+1}{\gamma+b}, N = bK_c, & K_{\text{AI}} \leq K < K_{\text{sat}}, \\ A(K-\bar{K}_\ell)^{\alpha}(1+\gamma\bar{K}_\ell)^{1-\alpha}, & K \geq K_{\text{sat}}. \end{cases}$$

The wage equals the marginal product of human labor:

$$w = (1-\alpha)\frac{Y}{N}.$$

The labor share (human labor's share of output) is:

$$s_L = \frac{wL}{Y} = \frac{1-\alpha}{N},$$

where $L = 1$ is human labor supply. In Region II, N rises with K, so s_L falls.

D Proof of Three Steady States

Define $\Delta(K) \equiv r^K(K) - [R^{ss}(K) + \delta]\bar{q}$. A steady state exists where $\Delta(K) = 0$.

Here $R^{ss}(K)$ is shorthand for the required return $R^{ss}(W)$ evaluated at steady-state wealth $W = \bar{q}K$, i.e., $R^{ss}(K) \equiv R^{ss}(\bar{q}K)$.

Region I: $\Delta(K) \to +\infty$ as $K \to 0$ (since $r^K \to \infty$). If the multiplicity condition holds, $\Delta(K_{\text{AI}}) < 0$. By continuity, $\exists K^L \in (0, K_{\text{AI}})$ with $\Delta(K^L) = 0$.

Region II: r^K is constant while $R^{ss}(K)$ is decreasing, so Δ is increasing. The multiplicity condition implies $\Delta(K_{\text{AI}}) < 0$ and $\Delta(K_{\text{sat}}) > 0$. By continuity, $\exists K^M \in (K_{\text{AI}}, K_{\text{sat}})$ with $\Delta(K^M) = 0$.

Region III: At K_{sat}, $\Delta > 0$. As $K \to \infty$, $r^K \to 0$ while $[R^{ss} + \delta]\bar{q} \to \delta\bar{q} > 0$, so $\Delta < 0$. By continuity, $\exists K^H \in (K_{\text{sat}}, \infty)$ with $\Delta(K^H) = 0$.

Multiplicity condition:

$$[R^{ss}(K_{\text{AI}}) + \delta]\bar{q} > r^K_{\text{flat}} > [R^{ss}(K_{\text{sat}}) + \delta]\bar{q}. \tag{20}$$

E Local Stability

The linearized system around a steady state $(K^*, \bar{q})$ is based on the return schedule $R(W)$ used in the transition dynamics:

$$\begin{pmatrix} \dot{K} \\ \dot{q} \end{pmatrix} = J \begin{pmatrix} K - K^* \\ q - \bar{q} \end{pmatrix}.$$

The Jacobian elements are:

$$\begin{aligned}
J_{11} &= \left.\frac{\partial \dot{K}}{\partial K}\right|_{ss} = 0, \\
J_{12} &= \left.\frac{\partial \dot{K}}{\partial q}\right|_{ss} = \frac{\psi K^*}{\bar{q}} > 0, \\
J_{21} &= \left.\frac{\partial \dot{q}}{\partial K}\right|_{ss} = -(r^K)'(K^*) + \bar{q}^2 R'(W^*), \\
J_{22} &= \left.\frac{\partial \dot{q}}{\partial q}\right|_{ss} = R(W^*) + \delta + \bar{q} R'(W^*) K^*.
\end{aligned}$$

When the local funding schedule slopes downward at the steady state, as in the baseline calibration, the trace is:

$$\mathrm{tr}(J) = J_{22} = R(W^*) + \delta + \bar{q} R'(W^*) K^* > 0$$

for reasonable parameters.

The determinant is:

$$\det(J) = -J_{12} \cdot J_{21} = -\frac{\psi K^*}{\bar{q}} \left[-(r^K)'(K^*) + \bar{q}^2 R'(W^*) \right].$$

The sign of $\det(J)$ depends on $(r^K)'(K^*)$:

- At K^L and K^H: $(r^K)'(K^*) < 0$ (diminishing returns). Since $R'(W^*) < 0$, the bracketed term is positive provided $-(r^K)'(K^*) > \bar{q}^2 |R'(W^*)|$. Under the baseline calibration this condition holds at K^L and K^H, hence $\det(J) < 0$. These are saddle points.
- At K^M: $(r^K)'(K^M) = 0$ (flat region). The bracketed term is negative, so $\det(J) > 0$. With $\mathrm{tr}(J) > 0$, K^M is an unstable node.

Since K is predetermined and q is a jump variable, saddle-path stability at K^L and K^H means these are locally stable steady states, while K^M is unstable. With one predetermined

variable and one jump variable, a saddle point—one stable and one unstable eigenvalue—implies a unique convergent path.

F Intermediate Adjustment Costs and the Speculative-Growth Path

This appendix states the phase-diagram condition behind the adjustment-cost discussion in Section 4. The issue is whether the path that leaves the low-capital steady state at elevated valuations can connect to the stable arm of the high-capital steady state. The condition is global, but its economic content is visible in the local geometry of the two stable steady states.

F.1 Setup and definitions

Consider the dynamical system (Appendix C):

$$\dot{K} = (\psi \ln q - \delta)\, K, \tag{21}$$

$$\dot{q} = F(K, q) \equiv [R(qK) + \delta]q - r^K(K). \tag{22}$$

For a given ψ, let $W^s_\psi(K^H)$ denote the one-dimensional stable manifold of the high-capital saddle steady state $(K^H(\psi), \bar{q}(\psi))$, where $\bar{q}(\psi) = e^{\delta/\psi}$. Let $W^u_\psi(K^L)$ denote the upward branch of the unstable manifold of the low-capital saddle steady state $(K^L(\psi), \bar{q}(\psi))$.

Since capital is predetermined and q is a jump variable, a speculative-growth episode starting from the low-capital steady state requires an upward jump in q that places the economy on the branch that reaches the stable manifold of K^H.

Definition 1 (Reach-back at elevated valuations)**.** *We say that the high-capital stable arm reaches back to the low-capital steady state if*

$$W^s_\psi(K^H) \cap \{(K, q) : K = K^L,\ q > \bar{q}(\psi)\} \neq \emptyset.$$

Equivalently, there exists an elevated valuation $q_0 > \bar{q}(\psi)$ such that the solution starting from (K^L, q_0) converges to $(K^H, \bar{q}(\psi))$.

F.2 Local geometry

At either stable steady state $j \in \{L, H\}$, write the Jacobian as

$$J_j = \begin{pmatrix} 0 & a_j \\ b_j & c_j \end{pmatrix}, \qquad a_j = \frac{\psi K^j}{\bar{q}} > 0,$$

where

$$b_j = -(r^K)'(K^j) + \bar{q}^2 R'(\bar{q}K^j), \qquad c_j = R(\bar{q}K^j) + \delta + \bar{q}R'(\bar{q}K^j)K^j.$$

Let $\lambda_j^s < 0 < \lambda_j^u$ denote the stable and unstable eigenvalues. The slope of the phase-diagram arm associated with eigenvalue λ is

$$\frac{dq}{dK} = \frac{\lambda}{a_j}. \tag{23}$$

Thus the upward branch leaving the low-capital steady state has local slope

$$\left.\frac{dq}{dK}\right|_{W^u(K^L)} = \frac{\lambda_L^u}{a_L} > 0, \tag{24}$$

and the branch entering the high-capital steady state from below has local slope

$$\left.\frac{dq}{dK}\right|_{W^s(K^H)} = \frac{\lambda_H^s}{a_H} < 0. \tag{25}$$

These slopes summarize the role of adjustment costs. When adjustment costs are very low (ψ is large), investment reacts strongly to a given valuation premium. The stable arm of the high-capital steady state is then locally flat in (K, q) space, so the decreasing-q branch near K^H lies only modestly above $\bar{q}$. The increasing-q trajectory coming from the low-capital region can therefore arrive above the stable arm and miss the high-capital saddle path.

When adjustment costs are very high (ψ is small), a valuation premium generates little capital accumulation. The upward branch leaving the low-capital steady state is then steep: q rises substantially before K moves much. This again tends to place the trajectory above the decreasing-q branch that enters the high-capital steady state.

For intermediate adjustment costs, these two local forces can be balanced. The increasing-q trajectory out of the low-capital region can meet the decreasing-q stable arm of the high-capital steady state. When the intersection is transverse, the connection is robust to small parameter perturbations by the Stable Manifold Theorem and the Implicit Function Theorem. This is the formal version of the transition-geometry condition discussed in Section 4.

G Proofs for Section 5

This appendix collects the formal arguments behind the comparative statics in Section 5.

Assumption 1 (Post-saturation monotone gap)**.** *For the funding schedule under consideration, the post-saturation gap between the MPK and the funding schedule is weakly decreasing. In the notation used for the shifted funding schedules—with the baseline case obtained by setting $\delta^{\text{eff}} = \delta$ and $\chi = 0$—this condition is*

$$\frac{d}{dK}\Big[r^K(K) - \mathcal{F}(K; \delta^{\text{eff}}, \chi)\Big] \leq 0, \qquad \text{for all } K \geq K_{\text{sat}}(\gamma). \tag{26}$$

A sufficient primitive condition is

$$\alpha(1-\alpha)A(K - \bar{K}_\ell)^{\alpha-2}(1 + \gamma\bar{K}_\ell)^{1-\alpha} \geq \frac{\rho\theta(\bar{q}^{\text{eff}})^2}{\big(1 + \theta\bar{q}^{\text{eff}}K\big)^2}, \qquad \text{for all } K \geq K_{\text{sat}}(\gamma). \tag{27}$$

Indeed, in Region III,

$$r^{K\prime}(K) = -\alpha(1-\alpha)A(K - \bar{K}_\ell)^{\alpha-2}(1 + \gamma\bar{K}_\ell)^{1-\alpha},$$

whereas

$$\mathcal{F}_K(K; \delta^{\text{eff}}, \chi, \zeta) = -\frac{\rho\theta(\bar{q}^{\text{eff}})^2}{\big(1 + \theta\bar{q}^{\text{eff}}K\big)^2}.$$

Thus condition (27) *is exactly $r^{K\prime}(K) \leq \mathcal{F}_K(K; \delta^{\text{eff}}, \chi, \zeta)$ throughout Region III, which implies condition* (26)*.*

Proposition 1 (MPK shifts and disappearance of steady states)**.** *Fix the funding schedule $\mathcal{F}$. Under Assumption 1, if condition* (13) *holds, then the low- and middle-capital steady states disappear, while the high-capital steady state remains. If instead condition* (14) *holds, then the high-capital steady state disappears.*

Proof. In the multiplicity region, the low crossing occurs on the pre-AI branch, the middle crossing occurs on the flat segment, and the high crossing occurs after AI saturation. On the pre-AI branch, the MPK is $\alpha AK^{\alpha-1}$, which is independent of γ. Hence the pre-AI crossing that would define K^L is fixed by the equation $\alpha AK^{\alpha-1} = \mathcal{F}(K; \delta)$ as long as that crossing lies to the left of $K_{\text{AI}}(\gamma)$. As γ rises, however, $K_{\text{AI}}(\gamma)$ moves left. Once $K_{\text{AI}}(\gamma)$ falls below the former low-capital crossing, that crossing no longer belongs to Region I and the pre-AI equation is no longer the operative steady-state condition there. To the left of the new $K_{\text{AI}}(\gamma)$, the pre-AI MPK lies strictly above $\mathcal{F}(K; \delta)$ because the sole Region I crossing was

at the former $K^L > K_{\text{AI}}(\gamma)$; hence no new Region I steady state replaces the one that was absorbed into the AI-deployment region.

In Region II, the relevant condition is $r^K_{\text{flat}}(\gamma) = \mathcal{F}(K;\delta)$. Condition (13) implies $r^K_{\text{flat}}(\gamma) \geq \mathcal{F}(K_{\text{AI}}(\gamma);\delta)$. Since $\mathcal{F}$ is decreasing, $r^K_{\text{flat}}(\gamma) > \mathcal{F}(K;\delta)$ for every $K > K_{\text{AI}}(\gamma)$, so there is no crossing on the flat segment. Thus both the low- and middle-capital steady states disappear. The high-capital crossing remains: in Region III,

$$r^K(K;\gamma) = \alpha A(K - \bar{K}_\ell)^{\alpha-1}(1 + \gamma\bar{K}_\ell)^{1-\alpha},$$

which is increasing in γ for each post-saturation K. Moreover, at the entry to Region III, condition (13) and the monotonicity of $\mathcal{F}$ imply $r^K_{\text{flat}}(\gamma) > \mathcal{F}(K_{\text{sat}}(\gamma);\delta)$, while $r^K(K;\gamma) \to 0$ and $\mathcal{F}(K;\delta) \to \delta\bar{q} > 0$ as $K \to \infty$. By continuity, there is a post-saturation crossing. Hence only the high-capital steady state remains.

Conversely, if condition (14) holds, then $r^K_{\text{flat}}(\gamma) \leq \mathcal{F}(K_{\text{sat}}(\gamma);\delta)$, so the post-saturation MPK enters Region III weakly below the funding schedule. By Assumption 1, the MPK–funding gap is weakly decreasing on the post-saturation branch. Hence no crossing exists in Region III and the high-capital steady state disappears. Strict versions of the inequalities give strict disappearance; equality corresponds to the threshold case. □

Proposition 2 (Funding shifts and disappearance of the high-capital steady state)**.** *Fix the MPK schedule. An increase in the fiscal wedge χ, in the borrowing-spread wedge ζ, or in the creative-destruction hazard λ^{CD} raises the funding burden facing the high-capital equilibrium. Under Assumption 1, if condition* (18) *holds, then the high-capital steady state disappears.*

Proof. The high-capital steady state requires the MPK schedule to cross the funding schedule after AI saturation. A necessary threshold for that crossing is that, at the end of the flat-MPK region, the return on capital still exceeds the funding cost. Condition (18) violates this threshold. Under Assumption 1, the gap $r^K(K) - \mathcal{F}(K;\delta^{\text{eff}},\chi,\zeta)$ is weakly decreasing throughout Region III. Since the gap is weakly negative at K_{sat} under condition (18), it remains weakly negative thereafter. Hence no crossing occurs on the post-saturation branch, and the high-capital destination is absent. □

H Numerical Example

H.1 Parameter Values

The figures use:

$$A = 0.0729, \quad \alpha = 0.33, \quad \gamma = 1.85, \quad \bar{K}_\ell = 0.25, \quad \rho = 0.08, \quad \theta = 20, \quad \delta = 0.05, \quad \psi = 3.0.$$

H.2 Derived Quantities

From the technology block:

$$b = \frac{(1-\alpha)\gamma}{\alpha} = \frac{0.67 \times 1.85}{0.33} \approx 3.76.$$

The boundaries of the flat-MPK region:

$$K_{\text{AI}} = \frac{1}{b} = 0.266, \qquad K_{\text{sat}} = \frac{1 + (\gamma + b)\bar{K}_\ell}{b} = 0.641.$$

The flat-region MPK:

$$r^K_{\text{flat}} = \alpha A b^{1-\alpha} = 0.33 \times 0.0729 \times 3.76^{0.67} \approx 0.058.$$

The steady-state valuation from (6) is

$$\bar{q} = e^{0.05/3.0} \approx 1.0168.$$

H.3 Computing Steady States

For the numerical transition system, the required return is evaluated using the locally calibrated isoelastic schedule in (5), with $W = \bar{q}K$.[4] Steady states solve

$$r^K(K) = [R(\bar{q}K) + \delta]\bar{q}.$$

For the baseline calibration, the three solutions are:

$$K^L = 0.224, \qquad K^M = 0.383, \qquad K^H = 0.792.$$

[4]The transition consumption rule is a local isoelastic approximation calibrated around $W^* = \bar{q}K^L$. It matches the exact steady-state consumption rule locally but need not match it exactly far from the calibration point. In the baseline simulation, this accounts for the slightly negative terminal value of the transition return schedule, even though the exact high-capital steady-state rate remains positive.

H.4 Verifying the Multiplicity Condition

The multiplicity condition (20) requires:

$$[R(K_{\text{AI}}) + \delta]\bar{q} > r^K_{\text{flat}} > [R(K_{\text{sat}}) + \delta]\bar{q}.$$

At the boundaries, using the same isoelastic transition return:

$$[R(K_{\text{AI}}) + \delta]\bar{q} \approx 0.063,$$
$$[R(K_{\text{sat}}) + \delta]\bar{q} \approx 0.050.$$

Since $0.063 > 0.058 > 0.050$, the multiplicity condition is satisfied.

H.5 Verifying the Monotone-Gap Condition

The sufficient Region III condition (27) requires the absolute slope of the post-saturation MPK schedule to be at least as large as the absolute slope of the funding schedule. At the baseline, $\delta^{\text{eff}} = \delta$ and $\chi = 0$, so $\bar{q}^{\text{eff}} = \bar{q}$, and condition (27) reduces to

$$\alpha(1-\alpha)A(K - \bar{K}_\ell)^{\alpha-2}(1 + \gamma\bar{K}_\ell)^{1-\alpha} \geq \frac{\rho\theta\bar{q}^2}{(1+\theta\bar{q}K)^2}, \qquad K \geq K_{\text{sat}}.$$

At the baseline value $K_{\text{sat}} = 0.641$, the two sides are

$$\alpha(1-\alpha)A(K_{\text{sat}} - \bar{K}_\ell)^{\alpha-2}(1 + \gamma\bar{K}_\ell)^{1-\alpha} \approx 0.0998,$$
$$\frac{\rho\theta\bar{q}^2}{(1+\theta\bar{q}K_{\text{sat}})^2} \approx 0.0084.$$

The inequality therefore holds by a wide margin at the entry to Region III. In the baseline calibration, the ratio of the left-hand side to the right-hand side is minimized in the interior of Region III at approximately $K = 1.76$ (verified by numerical evaluation over $K \in [K_{\text{sat}}, 5]$), where it remains about 8.55. Thus condition (27) holds throughout the relevant post-saturation branch for the baseline calibration.