NBER WORKING PAPER SERIES

SPECULATIVE GROWTH AND THE AI "BUBBLE"

Ricardo J. Caballero

Working Paper 34722
http://www.nber.org/papers/w34722

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2026 Revised July 2026

I thank Andrei Shleifer, Alp Simsek, Ludwig Straub, and Luohan Academy seminar participants for their comments, and Kei Uzui for research assistance. First draft: December 17, 2025. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

Speculative Growth and the AI "Bubble"
Ricardo J. Caballero
NBER Working Paper No. 34722
January 2026 Revised July 2026
JEL No. E21, E22, E24, E44, O33, O41

ABSTRACT

High valuations of AI-related firms are usually read in binary terms: either they reflect fundamentals or they are a bubble. This paper develops a third possibility: an optimistic valuation can have a permanent real legacy even when the valuation later corrects. AI capital expands productive capacity and shifts income toward high-saving capital owners. Over time, this wealth-saving feedback lowers the interest rate and can generate multiple steady states, including a self-sustaining high-capital economy. But rational pricing from the low-capital state does not take the economy there. The transition requires a temporary belief-supported valuation: investors perceive high returns, valuation rises, investment accelerates, and the interest rate rises during the transition. If enough capital has been installed before learning removes the wedge, the economy lands in the high-capital state, where the long-run interest rate is lower. If learning arrives too soon, the transition fails. The technology can be real, peak valuations can be unsustained, and the capital installed during the boom can remain. Workers receive higher wages at the high-capital destination despite a lower worker share, while capitalists finance the investment surge and are exposed to the correction in belief-supported prices.

Ricardo J. Caballero
Massachusetts Institute of Technology
Department of Economics
and NBER
caball@mit.edu

1 Introduction

> "While some part of the investment which was going on in the world at large was doubtless ill-judged and unfruitful, there can, I think, be no doubt that the world was enormously enriched by the constructions of the quinquennium from 1925 to 1929; its wealth increased in these five years by as much as in any other ten or twenty years of its history." —Keynes (1931)

The recent increase in AI-related valuations raises a familiar question: do high prices reflect fundamentals, or are they a bubble? The question takes the fundamentals as given and asks only whether the market has them right. But investment responds to valuation, and the capital installed during a boom changes the economy that later prices it. Once prices shape the fundamentals against which they are judged, the two answers are no longer exhaustive: a valuation can be unsustainable and still leave a permanent real legacy. This paper develops that third possibility. A temporary overvaluation raises investment; if enough capital is installed before the valuation corrects, the economy is left permanently with a larger capital stock, higher wages, and a lower interest rate. The mechanism is fragile: the legacy survives only if the correction arrives late enough. The paper therefore separates three objects that the bubble question runs together: whether the technology is productive, whether peak valuations are sustained, and whether the capital financed at those valuations remains.

AI provides the economic environment for the mechanism. AI capital performs tasks previously done by labor. As it accumulates, productive capacity expands and income shifts toward capital owners, who have a stronger saving motive. The resulting increase in saving capacity lowers the interest rate associated with a larger installed capital stock. This feedback can generate multiple steady states: a low-capital state and a high-capital state that is self-sustaining once enough capital has been installed.

The transition mechanism is separate from the destination. The model starts from the low-capital state; an ongoing AI episode can be interpreted as an economy already inside the transition. From that initial condition, rational pricing selects the low-capital path even though a high-capital steady state exists. A temporary overvaluation supplies the missing transition force. It raises perceived returns, pushes up valuation, accelerates investment, and moves capital toward the region where the high-capital economy can sustain itself. During this phase, investment demand raises the interest rate. The long-run effect is different: once installed capital has changed the distribution of income and the supply of saving, the high-capital steady state has a lower interest rate.

The overvaluation eventually corrects. The key question is whether the correction comes after enough capital has been installed. If it does, valuation can fall back to the high-capital rational path and the capital remains. If it comes too early, the perceived high-capital path disappears and the transition collapses back to the low-capital path.

The same distinction between the transition and the destination also shapes incidence. Workers benefit because AI deployment is accompanied by conventional capital deepening: workers operate with a larger conventional capital stock and wages rise even as the worker share falls. The capitalists who finance the investment boom carry the valuation exposure: they invest at prices supported by beliefs that later correct, and they compress their own consumption to finance it.

The model has three blocks. Investment follows a q-theory specification, so asset prices affect real accumulation. Capitalists have wealth-in-utility preferences. This generates a rising propensity to save with wealth, as in the non-homothetic structure in Straub (2019), and makes the interest cost of capital fall as capitalist wealth rises. Beliefs are Bayesian: investors estimate a persistent excess return from a noisy signal, use the posterior mean to value capital, and revise the posterior as evidence arrives. The baseline path is an optimistic-prior path that is not confirmed by subsequent data; equivalently, it can be read as an adverse ex-post realization under genuine uncertainty about a new technology. Valuation moves accumulation, wealth lowers required returns, and temporary beliefs supply a transition force that fades when the data do not confirm it. The contribution is to show that a temporary belief wedge on productive capital can have permanent real effects when the capital it finances moves the economy into a self-sustaining high-capital steady state.

The current AI cycle motivates this mechanism. Market capitalization has concentrated in AI-related firms, while announced data-center, power, and computing investment points to large real capital accumulation (Fortune, 2025; Goldman Sachs Global Institute, 2025; Van Nieuwerburgh, 2026; McKinsey, 2025).

Review of the literature. The closest antecedent is the speculative-growth mechanism of Caballero et al. (2006, henceforth CFH), in which asset values, funding conditions, and accumulation reinforce one another. This paper keeps the funding-feedback logic but changes both the technology and the source of optimism. AI provides the technological source: task substitution expands effective labor, shifts income toward capital owners, and lowers the eventual interest rate through the wealth-saving channel. Optimism comes instead from Bayesian learning: capitalists temporarily perceive an excess return to AI-related capital, and an optimistic posterior valuation can move capital across the separatrix before the posterior is revised down.

The technological environment follows the task-based view that advanced AI changes which tasks reproducible capital can perform. Acemoglu and Restrepo (2018) provide a canonical task-based formalization of automation and factor shares. Moll, Rachel, and Restrepo (2022) study how automation raises returns to wealth and increases income and wealth inequality. In the present model, automation also shifts income toward capital owners; because capitalists have a stronger wealth-saving motive, this redistribution lowers the destination real interest rate. Related work studies AGI and AI transition scenarios (Restrepo, 2025; Korinek and Suh, 2024) and broader macroeconomic implications of transformative AI (Trammell and Korinek, 2023; Acemoglu, 2024; Aghion and Bunel, 2024; Jones and Tonetti, 2026; Brynjolfsson et al., 2025).

The wealth-saving feedback draws on work connecting inequality, saving, and low interest rates. The wealth-in-utility specification reproduces the saving behavior generated by non-homothetic preferences in Straub (2019): richer households have lower marginal propensities to consume. It also relates to Mian, Straub, and Sufi (2020) on the saving glut of the rich. AI here provides an endogenous source of the concentration that drives the channel.

The transition trigger uses learning-driven valuation. Pástor and Veronesi (2009) show how learning about a new technology can generate high prices that later fall even when the technology is real. The additional force here is the wealth-saving feedback: learning-driven valuation affects capital accumulation, and installed capital can move the economy to a different steady state. The belief-to-investment link is related to evidence on expectations and investment in Gennaioli, Ma, and Shleifer (2016). Historical evidence from the late-1920s stock market provides a useful overvaluation benchmark: De Long and Shleifer (1991) use closed-end fund premia to measure investor sentiment. Because investment responds to valuation, a temporary belief wedge has real effects: it raises capital accumulation before valuation returns to rational pricing.

The paper is also related to rational-bubble models (Tirole, 1985; Martin and Ventura, 2012; Farhi and Tirole, 2012). These models study self-referential price components in overlapping-generations or liquidity-friction environments. The object here is different. The wedge is a perceived excess return on productive capital. When learning removes it, the price returns to the rational capital-claim valuation; the real legacy survives only if enough capital has already been installed to make the high-capital continuation available. Models of speculation with heterogeneous beliefs and resale motives (Scheinkman and Xiong, 2003; Simsek, 2013) study trade and pricing when investors disagree. In the representative-capitalist model here, all capitalists share the same belief, so the wedge reallocates resources intertemporally rather than across investors.

Finally, the paper is related more broadly to work on takeoff dynamics and multiple long-run outcomes. Big-push and coordination-failure models provide a benchmark for expectation-driven takeoff and multiple long-run outcomes (Rosenstein-Rodan, 1943; Murphy, Shleifer, and Vishny, 1989; Cooper and John, 1988; Azariadis and Drazen, 1990; Matsuyama, 1991). Here, multiplicity is in steady states rather than rational continuations from a fixed initial capital stock: the high-capital state is self-sustaining once reached, while the transition requires a belief-supported valuation that raises investment until the wealth-saving feedback generated by installed AI capital can sustain the high-capital outcome.

The remainder of the paper proceeds as follows. Section 2 builds the technology, saving, and valuation blocks. Section 3 develops the rational benchmark: AI deployment and the wealth-saving feedback generate multiple steady states, while the inherited capital stock selects a unique rational continuation. Section 4 studies the speculative transition: a temporary valuation wedge can move the economy toward the high-capital state, but the outcome depends on whether enough capital is installed before the wedge disappears. Section 5 concludes. The appendices collect the technology details, the wealth-in-utility derivation, the equilibrium system, the steady-state and local-geometry proofs, the Bayesian foundation for the belief wedge, the remaining transition proofs, and the illustrative numerical example behind the figures.

2 Model

The model has three blocks. The technology block describes how AI capital expands effective labor and changes factor shares. The saving block maps the resulting capitalist wealth into the interest rate. The investment block makes valuation matter for capital accumulation.

2.1 Technology

AI capital can perform worker tasks. A data center running trained models is capital on a firm's balance sheet, yet in production it performs tasks that would otherwise require labor. As AI capital accumulates, it expands the effective labor used with conventional capital. Because this AI labor is owned by capitalists, deployment shifts income toward capital owners. It also creates a range in which the marginal product of capital is constant: each additional unit of AI capital raises effective labor together with conventional capital, keeping the effective capital-labor ratio fixed.

Output uses conventional capital K_c and effective labor N,

$$Y = AK_c^{\alpha} N^{1-\alpha}, \qquad \alpha \in (0,1).$$

Raw labor supply is normalized to one. Total capital $K = K_c + K_\ell$ splits between conventional use and AI use. Each unit of AI capital produces AI labor at rate γ, so effective labor is $N = 1 + \gamma K_\ell$, and AI deployment faces a capacity constraint $\bar{K}_\ell$ reflecting limits on data, power and complementary infrastructure, or organizational absorption. Firms allocate capital optimally between the two uses. In the interior deployment region, the firm equates the marginal value of capital in conventional production to the marginal value of capital used to produce AI labor:

$$\alpha A K_c^{\alpha-1} N^{1-\alpha} = (1-\alpha) A K_c^{\alpha} N^{-\alpha} \gamma.$$

This condition gives $N = bK_c$, with $b \equiv (1-\alpha)\gamma/\alpha$. Thus, in the interior deployment region, AI labor and conventional capital expand together: deployment of AI capital is accompanied by conventional capital deepening rather than pure substitution away from the capital workers use. Combining the interior condition with $K = K_c + K_\ell$ and $N = 1 + \gamma K_\ell$ gives the allocation

$$N = bK_c, \qquad K_\ell = \frac{bK - 1}{\gamma + b}, \tag{1}$$

under which each unit of AI capital raises effective labor in step with conventional capital, fixing $K_c/N = 1/b$ and hence holding the marginal product of capital constant (details in Appendix A). The lower threshold $K_{\mathrm{AI}} = 1/b$ is where the interior allocation first activates AI deployment: $K_\ell = 0$ at the threshold and positive beyond it. The upper threshold $K_{\mathrm{sat}} = [1 + (\gamma + b)\bar{K}_\ell]/b$ is where it reaches the capacity constraint. Using $b = (1-\alpha)\gamma/\alpha$,

$$K_{\mathrm{sat}} - K_{\mathrm{AI}} = \frac{(\gamma + b)\bar{K}_\ell}{b} = \frac{\bar{K}_\ell}{1-\alpha} > 0.$$

The length of the interior deployment region is therefore the measure of deployable AI capacity scaled by $1/(1-\alpha)$. These conditions yield the three-region schedule for the marginal product of capital,

$$r^K(K) = \begin{cases} \alpha A K^{\alpha-1}, & K < K_{\mathrm{AI}}, \\ \alpha A b^{1-\alpha}, & K_{\mathrm{AI}} \le K < K_{\mathrm{sat}}, . \\ \alpha A (K - \bar{K}_\ell)^{\alpha-1}(1 + \gamma \bar{K}_\ell)^{1-\alpha}, & K \ge K_{\mathrm{sat}} \end{cases} \tag{2}$$

The dividend paid by the aggregate capital claim is average non-labor income per unit of installed capital,

$$d(K) \equiv \frac{Y(K) - w(K)}{K}. \tag{3}$$

For the Cobb–Douglas technology, this dividend schedule is

$$d(K) = \begin{cases} \alpha A K^{\alpha-1}, & K < K_{\mathrm{AI}}, \\ d_{\mathrm{flat}} \equiv \alpha A b^{1-\alpha}, & K_{\mathrm{AI}} \le K < K_{\mathrm{sat}}, \\ \dfrac{A(K-\bar{K}_\ell)^\alpha (1+\gamma\bar{K}_\ell)^{-\alpha}(\alpha+\gamma\bar{K}_\ell)}{K}, & K \ge K_{\mathrm{sat}}. \end{cases} \tag{4}$$

This dividend schedule is generated by technology: it is defined for every candidate capital stock K, before any steady-state condition is imposed. The flat middle region is the interior deployment region, generated by the allocation (1).

The flat payout is not a Cobb-Douglas knife edge. With any constant-returns technology $F(K_c, N)$, the interior deployment condition $F_{K_c} = \gamma F_N$ pins down the ratio K_c/N. Let ζ denote that ratio and let $w = F_N(\zeta, 1)$ be the wage per unit of raw labor. Euler's theorem gives $F(\zeta, 1) = \zeta F_{K_c}(\zeta, 1) + F_N(\zeta, 1) = w(1+\gamma\zeta)$. Since $K = N(\zeta + 1/\gamma) - 1/\gamma$, the capital-claim payout in the deployment region satisfies

$$d(K) = \frac{Y - w}{K} = \gamma w.$$

Thus the capital claim pays exactly the wage bill of the task labor it replaces. In the Cobb-Douglas case this identity reads $d_{\mathrm{flat}} = \gamma w_{\mathrm{flat}}$, where $w_{\mathrm{flat}} = (1-\alpha)Ab^{-\alpha}$.

The distributional object that matters is the *worker* share, not the Cobb-Douglas share paid to effective labor. The wage equals the marginal product of effective labor, $w = (1-\alpha)Y/N$; workers supply only the raw unit of labor, while AI-labor income accrues to capital owners. The worker share is therefore

$$s_L(K) = \frac{w}{Y} = \frac{1-\alpha}{N(K)} = \begin{cases} 1-\alpha, & K < K_{\mathrm{AI}}, \\ (1-\alpha)\dfrac{\gamma + b}{b(\gamma K + 1)}, & K_{\mathrm{AI}} \le K < K_{\mathrm{sat}}, \\ \dfrac{1-\alpha}{1+\gamma\bar{K}_\ell}, & K \ge K_{\mathrm{sat}}. \end{cases} \tag{5}$$

The Cobb-Douglas effective-labor share is constant; the worker share falls as AI labor expands and its income accrues to capital owners.

Capitalists own installed capital and the associated deployment capacity, so each unit of the aggregate capital claim receives a pro-rata claim on non-labor income. Before saturation, the deployment-capacity constraint is slack. Below K_{AI}, all non-labor income is the return to conventional capital; in the interior deployment region, the adjustable deployment margin pins down arbitrage across conventional and AI capital. In both regions, the capital-claim dividend coincides with the marginal product schedule, $d(K) = r^K(K)$. In the saturation region, deployment capacity is scarce and non-labor income includes an inframarginal capacity rent; the pro-rata capital claim pays this rent to capitalists.[1]

[1] The distinction is the standard fixed-factor distinction between a marginal product and average non-labor income. Below saturation, deployment capacity is adjustable and the arbitrage condition eliminates a separate capacity rent. Once the deployment-capacity constraint binds, the scarce deployment factor earns a quasi-rent.

2.2 Households and the interest-rate schedule

Workers supply labor, hold no assets, and consume their wage. Capitalists own the capital claim. Their saving behavior is summarized by wealth-in-utility preferences,

$$\int_0^\infty e^{-\rho t}\left[\log c_t + \theta W_t\right] dt, \qquad \dot{W}_t = R_t W_t - c_t, \tag{6}$$

where $\theta \geq 0$ governs the strength of the wealth-saving motive. The Euler equation (Appendix B) is

$$\frac{\dot{c}_t}{c_t} = R_t - \rho + \theta c_t. \tag{7}$$

The term θc_t is the wealth-saving force: utility from wealth raises the value of carrying resources forward, so desired saving rises with wealth. In steady state, this force appears as the downward-sloping interest-rate schedule in (9). Rearranging (7) defines the real interest rate,

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

This is the interest-rate object used below. Along a rational path, i_t equals the realized return on the capital claim, R_t^a.

The steady-state interest-rate schedule follows directly from (7) and the budget constraint. Setting $\dot{c} = 0$ in (7) gives $R = \rho - \theta c$. Setting $\dot{W} = 0$ in (6) gives $c = RW$. Solving these two equations for c and R gives

$$c^{ss}(W) = \frac{\rho W}{1 + \theta W}, \qquad R^{ss}(W) = \frac{\rho}{1 + \theta W}, \qquad \frac{dR^{ss}}{dW} = -\frac{\rho\theta}{(1 + \theta W)^2} < 0. \tag{9}$$

As capitalist wealth rises, desired saving rises and the steady-state interest rate falls. This downward-sloping interest-rate schedule is the key force behind the multiple-steady-state result. AI deployment matters because it shifts income toward capital owners and thereby strengthens the saving channel that lowers the eventual cost of holding a larger capital stock. This aggregate block captures the non-homothetic saving force emphasized by Straub (2019): higher wealth lowers the marginal propensity to consume and raises desired saving.

2.3 Investment, valuation, and beliefs

Let q_t be Tobin's q, the traded value of installed capital relative to replacement cost. Adjustment costs make the investment rate increasing in valuation:

$$\frac{\dot{K}_t}{K_t} = \psi \log q_t - \delta, \tag{10}$$

so $q_t = \bar{q} \equiv e^{\delta/\psi}$ is the valuation at which gross investment exactly covers depreciation. The capital claim pays the dividend flow $d(K_t)$ per unit of installed capital and has price q_t. Its realized return is the capital gain plus the payout yield, net of depreciation:

$$R_t^a = \frac{\dot{q}_t}{q_t} - \delta + \frac{d(K_t)}{q_t}, \qquad W_t = q_t K_t. \tag{11}$$

I summarize the belief-supported phase with a single object: a posterior belief x_t about an excess return on capital. The perceived return is

$$R_t^p = R_t^a + x_t. \tag{12}$$

This equation defines the return that enters the perceived-return Euler equation. Along a belief-supported path, the real interest rate equals the perceived return, $i_t = R_t^p = R_t^a + x_t$. At each date, capitalists value the capital claim using R_t^p, while actual consumption and resources are governed by R_t^a and by the resource constraint. The wedge changes valuation and intertemporal choices; it does not add output or relax feasibility. When realized returns fail to confirm the perceived excess return, the forecast error is what lowers x_t. The rational economy is the special case $x_t \equiv 0$.

Section 3 combines these blocks to characterize the steady states and the rational dynamics; Section 4 adds beliefs to study the transition.

3 Multiple Steady States and the Rational Benchmark

The rational benchmark studies how the technology and saving blocks can support more than one long-run allocation: a low-capital economy with little AI deployment and a high-capital economy in which enough AI capital has been installed to change production and saving. Both are rational steady states. The distinction between them is dynamic. Capital is predetermined, valuation is the jump variable, and the rational price must lie on an admissible continuation path from the inherited K. Starting from the low state, that continuation is the low-capital path. This section establishes both halves of the claim: first the multiplicity, then the selection.

3.1 Multiple steady states

A rational steady state solves the two stationary conditions of the rational dynamic system. The first condition is $\dot{K} = 0$: since investment responds to valuation through (10), stationarity of capital requires $q_t = \bar{q}$. Thus every steady state has the same valuation $q = \bar{q}$: the price at which gross investment exactly covers depreciation. Hence overvaluation is necessarily transitional: the high-capital steady state is a high-K object rather than a high-q object.

The second condition is stationarity of the capital price. The household block determines the real interest rate at which capitalists are willing to hold wealth $W = \bar{q}K$:

$$R^{ss}(\bar{q}K) = \frac{\rho}{1 + \theta\bar{q}K}.$$

At a steady state, $q = \bar{q}$. A unit of installed capital then has price $\bar{q}$, pays dividend $d(K)$, and has net return

$$-\delta + \frac{d(K)}{\bar{q}}.$$

Thus K is a steady state only if this return equals the required real interest rate:

$$-\delta + \frac{d(K)}{\bar{q}} = \frac{\rho}{1 + \theta\bar{q}K}.$$

Equivalently, define the carrying-cost schedule and the excess-payout gap by

$$\Phi(K) \equiv \bar{q}\left[\delta + \frac{\rho}{1+\theta\bar{q}K}\right], \qquad G(K) \equiv d(K) - \Phi(K). \tag{13}$$

Here $d(K)$ is what one unit of installed capital pays, while $\Phi(K)$ is what it must pay to be held at price $\bar{q}$. Steady states are the roots of the gap:

$$G(K) = 0, \qquad \text{equivalently} \qquad d(K) = \Phi(K).$$

The key feedback is the slope of the carrying-cost schedule. A larger K raises capitalist wealth $W = \bar{q}K$; with the wealth-saving motive in (6), higher wealth raises desired saving and lowers the interest rate $R^{ss}(\bar{q}K)$. AI deployment feeds this mechanism. When AI capital takes over tasks, more income accrues to capital owners. Since capital owners have the saving motive that generates the falling carrying-cost schedule, this redistribution raises aggregate saving capacity and lowers the interest cost of sustaining a larger installed capital stock.

Figure 1 plots the dividend schedule $d(K)$ against the carrying-cost schedule $\Phi(K)$. Their intersections are a low root, a middle threshold root, and a high root.

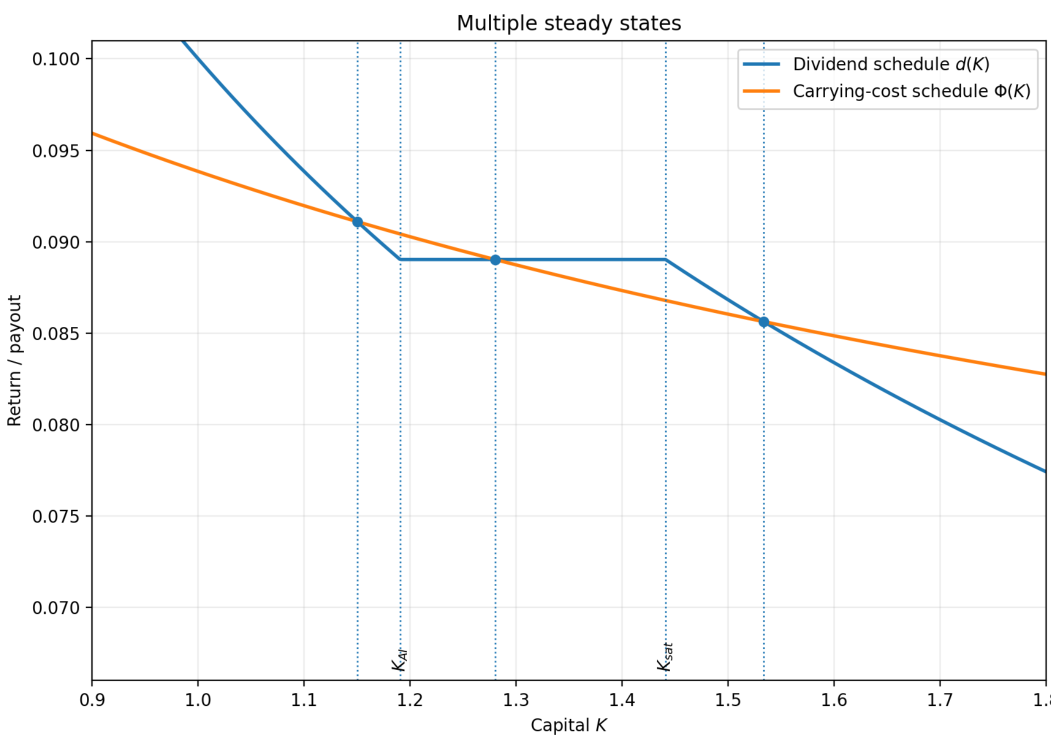

Figure 1: Multiple steady-state roots. The blue line is the dividend schedule $d(K)$ paid by the aggregate capital claim. The orange line is the carrying-cost schedule $\Phi(K) = \bar{q}[\delta+\rho/(1+\theta\bar{q}K)]$. Their intersections are a pre-deployment low root, a deployment-region middle threshold, and a post-saturation high root. All figures are illustrative; they are not a calibration or quantification.

The $\theta = 0$ counterfactual shows why this slope matters. If $\theta = 0$, the steady-state Euler equation requires $R = \rho$ at any level of wealth, and the carrying-cost schedule is horizontal at

$\Phi(K) = \bar{q}(\delta + \rho)$. Unless that level happens to coincide with d_{flat}, a horizontal line crosses the three-region dividend schedule exactly once, and the steady state is unique. Multiplicity therefore requires the carrying-cost schedule to fall as capital accumulates, which is what the wealth-saving motive delivers: wealthier capitalists are content with a lower return. This dependence of required returns on wealth directly captures the saving behavior generated by non-homothetic preferences in Straub (2019). AI deployment matters for the saving block precisely because it shifts income toward the agents with that behavior.

The deployment region is where the technological and saving forces align. Conventional capital and AI labor expand together, keeping the capital payout flat at d_{flat}, while the carrying-cost schedule falls with capitalist wealth. Hence $G(K)$ can turn upward in the deployment region and generate the middle crossing. After saturation, diminishing returns lower $d(K)$; the high root is sustained by the larger productive base and the lower carrying cost created by the wealth-saving channel. The same primitives can therefore support two outer destination candidates, separated by a middle threshold root.

Proposition 1 (Large-enough AI and multiple steady states)**.** *Suppose*

$$\Phi(K_{\text{AI}}) > d_{\text{flat}} > \bar{q}\delta. \tag{14}$$

Then there exists a deployment capacity $\bar{K}_\ell^* > 0$ such that, for every $\bar{K}_\ell > \bar{K}_\ell^*$, the steady-state equation $G(K) = 0$ has at least one low root in $(0, K_{\text{AI}})$, a unique middle root in $(K_{\text{AI}}, K_{\text{sat}})$, and at least one high root in (K_{sat}, ∞). If, in addition,

$$G'(K) < 0 \qquad \textit{for } K \in (0, K_{\text{AI}}) \cup (K_{\text{sat}}, \infty), \tag{15}$$

then the outer roots are unique, they are down-crossings of the dividend schedule and the carrying-cost schedule, and the economy has exactly three steady states, denoted $K^L < K^M < K^H$.

Condition (14) turns technology size into the sandwich restriction needed for multiplicity. The first inequality says that, at the point where AI deployment becomes marginally feasible, the economy is still too poor to carry the AI payout at the prevailing interest cost. Since $d_{\text{flat}} = \gamma w_{\text{flat}}$, the second says that the wage bill saved by a unit of AI capital exceeds the depreciation cost of installed capital at price $\bar{q}$. Since $K_{\text{sat}} = [1 + (\gamma + b)\bar{K}_\ell]/b$ rises with deployment capacity and $\Phi(K) \downarrow \bar{q}\delta$ as $K \to \infty$, sufficiently large deployment capacity makes $\Phi(K_{\text{sat}}) < d_{\text{flat}}$. Hence

$$\Phi(K_{\text{AI}}) > d_{\text{flat}} > \Phi(K_{\text{sat}}). \tag{16}$$

The fall in the carrying-cost schedule creates room for the middle crossing; the flat deployment region keeps the capital payout from falling while that wealth-saving feedback operates. After saturation, diminishing returns reassert themselves and the dividend schedule crosses the carrying-cost schedule from above. The outer monotonicity condition (15) is the same down-crossing condition that makes the outer steady states saddles in Lemma 1. The proof is in Appendix D.

The three roots have different economic roles. The low root K^L is the pre-deployment economy: capital is too scarce for AI deployment to have transformed production. The high root K^H is the post-deployment economy: enough capital has been installed that AI expands productive capacity and shifts income toward high-saving capital owners, lowering the interest rate consistent with the larger capital stock.[2] These two outer roots are the candidate long-run destinations.

The middle root K^M is not a candidate destination. It lies in the deployment region, where additional capital does not immediately run into diminishing returns and the wealth-saving feedback strengthens accumulation. The phase-diagram analysis below shows that K^M is the threshold separating the basins of attraction of the low and high steady states.

3.2 Rational inaccessibility

The previous subsection characterized the stationary roots of the rational economy. The present question is dynamic selection: starting from the low-capital state K^L, can rational pricing move the economy to the high-capital state K^H? The economic issue is simple. Since K_t is predetermined, an equilibrium path can choose only the initial price q_0. To converge to a long-run steady state, that price must place the economy on a stable branch. The question is therefore whether the high-capital stable branch passes through the inherited low capital stock K^L. The answer is no: under the rational geometry below, the high-capital branch is available only for inherited capital stocks $K \geq K^M$.

The argument has four steps. First, I write the rational economy as a two-dimensional system in the predetermined state K_t and the jump variable q_t. Second, I derive the resource-accounting identity for residual consumption. Third, the two nullclines classify the local geometry around the three stationary roots. Finally, the global branch structure shows that the high-capital saddle path is not available from K^L.

The rational dynamic system. The first law of motion is capital accumulation:

$$\dot{K}_t = K_t(\psi \log q_t - \delta).$$

Thus the $\dot{K} = 0$ locus is the horizontal line $q = \bar{q}$. Above that line, $q_t > \bar{q}$, investment exceeds depreciation and K_t rises. Below it, $q_t < \bar{q}$, depreciation exceeds investment and K_t falls.

Resource accounting and residual consumption. Before deriving the valuation equation, I close the resource side. Given the current state K_t, the current valuation q_t, and the investment law, the budget/resource constraint pins down residual feasible consumption. Start from capitalist wealth $W_t = q_t K_t$, so that $\dot{W}_t = \dot{q}_t K_t + q_t \dot{K}_t$. Using the return equation (11), total capital income is

$$R_t^a W_t = \left(\frac{\dot{q}_t}{q_t} - \delta + \frac{d(K_t)}{q_t}\right) q_t K_t = \dot{q}_t K_t - \delta q_t K_t + d(K_t) K_t.$$

Substituting this and $\dot{W}_t$ into the budget constraint $c_t = R_t^a W_t - \dot{W}_t$, the capital-gain term $\dot{q}_t K_t$ appears on both sides and cancels:

$$c_t = d(K_t)K_t - \delta q_t K_t - q_t \dot{K}_t = d(K_t)K_t - q_t\big(\dot{K}_t + \delta K_t\big).$$

[2] Related event-study evidence from the current AI cycle is consistent with this long-rate implication. Andrews and Farboodi (2025) find that long-maturity Treasury and TIPS yields fall around major AI model releases, which I interpret as news about faster progress at the AI frontier. In the model, such news maps naturally into a higher high-capital destination and, through the wealth-saving channel, a lower interest rate at that destination.

The cancellation has a direct economic implication: capital gains raise measured wealth and the measured return by exactly the same amount, so they free no resources for consumption. What remains is the dividend income of the capital claim, $d(K)K$, minus spending on gross investment, valued at q. Using $\dot{K}_t + \delta K_t = K_t \psi \log q_t$ from (10),

$$c_t = C(K_t, q_t) = K_t \big[d(K_t) - q_t \psi \log q_t \big]. \tag{17}$$

Admissibility imposes $C(K, q) > 0$. The identity $C(K, q)$ is the resource-side closure of the dynamic system: for each pair (K, q), the investment law determines absorption by capital accumulation, and the remaining resources are consumed. The Euler equation closes the valuation side next, through the law of motion for q_t. A higher q_t commits more output to investment and leaves less for capitalist consumption; this consumption compression during investment transitions returns as a central force in the incidence analysis of Section 4.4.

The valuation equation and nullclines. The price q_t moves to make the return on the capital claim consistent with capitalist intertemporal saving. Because residual consumption is given by $C(K_t, q_t)$, its growth along any path is

$$\dot{c}_t = C_K(K_t, q_t)\dot{K}_t + C_q(K_t, q_t)\dot{q}_t.$$

The Euler equation (7), written with the realized return (11), becomes

$$\underbrace{\frac{\dot{q}_t}{q_t} - \delta + \frac{d(K_t)}{q_t}}_{\text{return paid by the claim}} = \underbrace{\rho - \theta C(K_t, q_t) + \frac{C_K \dot{K}_t + C_q \dot{q}_t}{C(K_t, q_t)}}_{\text{return required by the Euler equation}}. \tag{18}$$

Every object in (18) is a known function of (K_t, q_t) except $\dot{q}_t$, which appears linearly on both sides. Solving for $\dot{q}_t$ yields the law of motion of the valuation; the explicit expression is recorded in Appendix C. Stacking the two laws of motion defines the rational vector field

$$(\dot{K}_t, \dot{q}_t) = F^0(K_t, q_t),$$

with first component $f(K, q) = K(\psi \log q - \delta)$ and second component $g(K, q)$ obtained from (18); it is convenient to write $g = Z/D$, where $Z(K, q)$ is the price-pressure numerator and $D(K, q)$ is the denominator of that solution. Appendix E shows that $D < 0$ for every admissible state, so the sign of valuation drift is the opposite of the sign of Z:

$$\operatorname{sign} \dot{q} = -\operatorname{sign} Z.$$

The local analysis below therefore reads off stability from the sign of Z and its derivatives.

The two nullclines determine the phase diagram. The horizontal nullcline $q = \bar{q}$ determines whether capital rises or falls. The $\dot{q} = 0$ locus determines whether valuation rises or falls. Setting $\dot{q}_t = 0$ in (18), the capital-gain term drops from the left side and the $C_q \dot{q}$ term from the right, leaving

$$-\delta + \frac{d(K_t)}{q_t} = \rho - \theta C(K_t, q_t) + \frac{C_K(K_t, q_t)}{C(K_t, q_t)} K_t(\psi \log q_t - \delta). \tag{19}$$

Equation (19) says that the current dividend return on the capital claim equals the return required by the Euler equation, including the consumption-growth correction generated by movement in K. Away from this locus, the Euler-pricing balance requires $\dot{q}_t \neq 0$, with the sign given by the direction field.

Local geometry and global branch selection. Figure 2 draws the two nullclines and the implied direction field. The horizontal line $q = \bar{q}$ determines the horizontal component: arrows point right above it and left below it. The $\dot{q} = 0$ curve determines the vertical component: arrows point up where $\dot{q} > 0$ and down where $\dot{q} < 0$. The steady states are the intersections of the two nullclines.

For the selection result, only two geometric facts matter. The middle root K^M is locally repelling, so it separates basins. The outer roots are saddles, so convergence to either long-run destination requires starting on the appropriate stable arm. On the horizontal nullcline $q = \bar{q}$, the sign of valuation drift is read directly from the steady-state gap:

$$\operatorname{sign} \dot{q}\Big|_{q=\bar{q}} = -\operatorname{sign} G(K). \tag{20}$$

Lemma 1 establishes these local facts rather than assuming them.

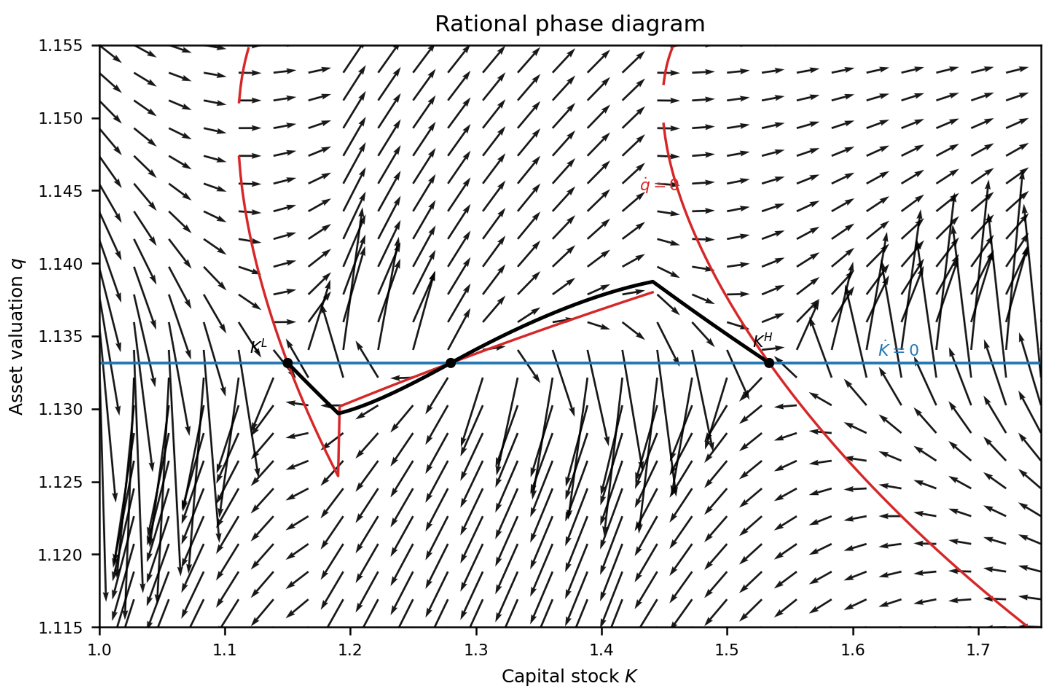

Figure 2: Rational phase diagram. The horizontal nullcline is $\dot{K} = 0$, equivalently $q = \bar{q}$. The second nullcline is $\dot{q} = 0$ from (19). Arrows point right above $q = \bar{q}$ and left below it; they point up where $\dot{q} > 0$ and down where $\dot{q} < 0$. The stable arms are the saddle paths of the low- and high-capital steady states, and the middle source separates the two basins.

Lemma 1 (Local rational geometry)**.** *Equation* (20) *gives the price-drift sign on the horizontal nullcline. At every steady state* $(K^j, \bar{q})$ *of* F^0*, the Jacobian satisfies* $f_K = 0$ *and* $f_q = K^j \psi / \bar{q} > 0$*, so that*

$$\operatorname{sign}\det DF^0(K^j, \bar{q}) = \operatorname{sign} Z_K(K^j, \bar{q}) = \operatorname{sign} G'(K^j), \tag{21}$$

with

$$Z_K = C(K^j, \bar{q})\left[d'(K^j)\left(\tfrac{1}{\bar{q}} + \theta K^j\right) + \theta\big(d(K^j) - \bar{q}\delta\big)\right]. \tag{22}$$

Consequently:

(i) the middle steady state K^M *is a source for any admissible parameters. In the flat region,* $d'(K^M) = 0$*, which forces* $G'(K^M) > 0$*; the trace is also positive there;*

(ii) any outer steady state K^j*,* $j \in \{L, H\}$*, is a saddle whenever it is a down-crossing,* $G'(K^j) < 0$*. In particular, the outer steady states are saddles under* (15).

The economic content of the lemma is that the same gap G has two readings. The level of $G(K)$ gives the direction of price drift on the horizontal nullcline: when the dividend schedule exceeds the carrying-cost schedule, $G(K) > 0$, valuation is pushed down; when $G(K) < 0$, valuation is pushed up. The slope of G at a root gives local stability. All three steady states sit on the same horizontal nullcline $q = \bar{q}$, so at a steady state the capital equation is flat in K: $f_K = \psi \log \bar{q} - \delta = 0$. With one diagonal entry of the Jacobian equal to zero, the determinant collapses to a single product, and its sign is the sign of Z_K—the response of the price-pressure numerator Z to capital. Hence G read at a point gives the vertical direction field, while G read as a slope gives the saddle/source classification.

The expression for Z_K in (22) displays the two competing forces. The first term carries diminishing returns, $d' \le 0$: more capital lowers the payout, which pushes the price down and pulls the economy back. The second term carries the wealth-saving feedback, $\theta(d - \bar{q}\delta) > 0$: more capital raises wealth, lowers the required return, and pushes the price up. At the outer steady states, the down-crossing condition $G'(K^j) < 0$ says that diminishing returns dominate; the steady states are saddles. At the middle steady state, the economy is in the flat region, the diminishing-returns term vanishes identically, and only the destabilizing wealth-saving feedback remains. The middle state is therefore a source for any admissible parameters—a property of the deployment technology, not of the example. In this sense, the outer roots are down-crossings, while the middle root is the up-crossing separating the basins.

The local geometry does not yet settle the transition question. The remaining issue is global admissibility: whether the high-capital saddle branch is available at the inherited low capital stock with positive capitalist consumption.

Let $Q^H(K; 0)$ denote the rational high-capital saddle price wherever that branch exists. Define the rational high-capital domain as

$$\mathcal{K}_0 \equiv [K^M, \infty). \tag{23}$$

This closed set includes the threshold K^M by convention. For $K > K^M$, it is the set of inherited capital stocks from which the high-capital continuation is available without any belief wedge; at $K = K^M$, the high branch reaches the middle source only as a limiting endpoint.

Lemma 2 (Rational high-capital domain)**.** *Under the phase-diagram geometry of Lemma 1, assume that the admissible region of F^0 contains no closed orbits and that the backward high arm remains admissible until it reaches the middle source. Then the high-capital continuation is available for $K > K^M$, with K^M as the limiting threshold endpoint; I write the closed high-capital domain as $\mathcal{K}_0 = [K^M, \infty)$. The low-capital stable arm is available for inherited capital stocks $K \in (K^L, K^M)$.*

The lemma turns branch separation into a consequence of the maintained rational geometry. The high state is self-sustaining for capital stocks strictly above the threshold, while the low steady state satisfies $K^L < K^M$ and is outside this domain. The same geometry gives the low rational continuation for capital stocks between the low root and the middle threshold. The illustrative numerical example satisfies these global properties.

Proposition 2 (Rational inaccessibility)**.** *Under the local geometry of Lemma 1 and the maintained global conditions of Lemma 2, there is no admissible rational equilibrium path with $K(0) = K^L$ that converges to K^H. The high-capital steady state is self-sustaining once reached, but it is not reached from the low-capital state by rational asset pricing.*

Thus the rational benchmark identifies the transition problem. The economic content is that rational asset pricing cannot bootstrap the economy into the high-capital steady state: the high-capital price is consistent with abundant capital, but not with the low capital stock inherited at K^L. Starting from K^L, rational pricing selects the low-capital continuation. Any other initial valuation eventually drives consumption to zero or violates transversality, and so cannot be an equilibrium. To reach K^H, the economy needs a temporary force that raises valuation while capital accumulates. Section 4 introduces that force as a perceived-return wedge.

4 Speculative Transition, Fragility, and Incidence

Section 3.2 showed that rational pricing from K^L selects the low-capital continuation. The high-capital continuation is available only once inherited capital has crossed into the rational high-capital domain $\mathcal{K}_0 = [K^M, \infty)$. This section asks how a temporary belief-supported valuation can move the economy from K^L into that domain.

The argument separates feasibility from speed. Section 4.1 fixes a perceived excess return and asks whether it is large enough, at a given capital stock K, to make the high-capital continuation available. This defines the required-wedge frontier. Section 4.2 lets the wedge fade according to Bayesian learning and asks whether induced investment raises capital into $\mathcal{K}_0$ before learning removes the perceived high-capital branch. The remaining subsections describe the landing after success or failure and the distributional incidence of the transition.

4.1 Feasibility: a constant wedge and the required-wedge frontier

Begin with a fixed perceived excess return $x \geq 0$. Investors price capital using the perceived return

$$R^p = R^a + x.$$

A larger perceived return raises the price of installed capital, and a higher price raises investment through the q-theory law.

This is the same pricing channel as in the rational phase diagram, with one change. In the valuation equation, the perceived return is $R^p = R^a + x$, so a positive x lowers the true return investors require from a given dividend stream. In the (K, q) phase diagram, the wedge shifts the $\dot{q} = 0$ locus relative to the rational vector field in Figure 2.

The steady-state effect of the wedge can be computed exactly as in Section 3.1. Suppose the wedge were permanent at level x. The stationary Euler condition holds at the perceived return:

$$R^a + x = \rho - \theta c.$$

The budget constraint is unchanged, because the wedge changes what investors believe, not what resources exist:

$$c = R^a W.$$

Solving the two equations for R^a gives

$$R^a = \frac{\rho - x}{1 + \theta W},$$

the interest-rate schedule (9) with ρ replaced by $\rho - x$: investors who perceive an extra return x are content to carry wealth at a lower true return, with pass-through scaled by the same wealth term that governs the interest-rate schedule. For pricing purposes, a permanent perceived excess return therefore acts like a reduction in the required return on the capital claim. Substituting the steady-state return $R^a = -\delta + d/\bar{q}$ and wealth $W = \bar{q}K$, the steady-state condition becomes

$$d(K) = \bar{q}\left[\delta + \frac{\rho - x}{1 + \theta \bar{q} K}\right]. \tag{24}$$

In the language of Figure 1, the wedge shifts the carrying-cost schedule down by $\bar{q}x/(1 + \theta\bar{q}K)$. Write the perceived steady-state gap as

$$G_x(K) \equiv d(K) - \bar{q}\left[\delta + \frac{\rho - x}{1 + \theta \bar{q} K}\right],$$

so $G_0(K) = G(K)$. The wedge lowers the carrying-cost schedule directly. With $\theta > 0$, capital accumulation then lowers the carrying-cost schedule further through the wealth-saving channel. This is the handoff mechanism: a temporary belief can raise valuation long enough for capital accumulation to move the economy into the rational high-capital domain, where the destination no longer requires the wedge.

The steady-state calculation shows how a permanent wedge shifts the perceived phase diagram. The dynamic object needed for the transition is the associated high-capital branch. For each fixed wedge x, let $Q^H(K; x)$ denote the price on the perceived high-capital saddle branch, when that branch exists, and let $K^H(x)$ denote the associated perceived high-capital root. The pair (K, x) is feasible for a high-capital continuation if this branch exists and leaves consumption positive. Formally,

$$\begin{aligned} \mathcal{D}^H &\equiv \{(K, x) : Q^H(K; x) \text{ exists and } C(K, Q^H(K; x)) > 0\}, \\ x^*(K) &\equiv \inf\{x \geq 0 : (K, x) \in \mathcal{D}^H\}. \end{aligned} \tag{25}$$

The function $x^*(K)$ is the required-wedge frontier. It is the minimum perceived excess return needed, at capital K, to make the high-capital continuation available. For launch from K^L, I maintain that the admissible launch beliefs $\{x \geq 0 : (K^L, x) \in \mathcal{D}^H\}$ have finite supremum $\bar{x} < \infty$.

Under the maintained global geometry, applied also to the perceived system for each fixed x, the required-wedge frontier admits a closed form. The middle steady state satisfies

$$K^M = \frac{1}{\theta \bar{q}} \left[\frac{\bar{q}\rho}{d_{\text{flat}} - \bar{q}\delta} - 1 \right], \tag{26}$$

and the required wedge is

$$x^*(K) = \theta\big(d_{\text{flat}} - \bar{q}\delta\big)\big(K^M - \max\{K, K_{\text{AI}}\}\big) \qquad \text{for } K \leq K^M. \tag{27}$$

Thus $x^*(K)$ is flat for $K < K_{\text{AI}}$, then declines linearly to zero at K^M. In the deployment region, the frontier has a simple interpretation: $x^*(K)$ is the wedge that makes the inherited capital stock the perceived middle root, $G_x(K) = 0$. Below K_{AI}, the analogous root-collision condition binds at the deployment threshold K_{AI}, which is why the frontier is flat. Once deployment begins, the launch point itself must clear the perceived middle root, and each additional unit of capital lowers the required wedge through the wealth it adds. The frontier captures the handoff from belief to capital: before deployment the belief must be large enough on its own to open the high branch, and once deployment begins each additional unit of installed capital lowers the belief still needed to keep that branch feasible.

The constant-wedge analysis establishes feasibility. The frontier $x^*(K)$ governs the size of the wedge needed at each capital stock; the learning-speed problem asks whether a fading wedge remains above that frontier long enough.

Proposition 3 (Constant-wedge implementation)**.** *Under the maintained local and global geometry, a constant admissible wedge $x > x^*(K^L)$ implements a transition from $K(0) = K^L$ to the rational high-capital domain in finite time.*

The proposition formalizes feasibility. The wedge raises the perceived return, lifts q, and accelerates accumulation before fundamentals alone justify the high-capital path. If the induced accumulation reaches $\mathcal{K}_0$, the high branch is self-sustaining and the wedge is redundant. At that point, the high-capital rational continuation is available at the inherited capital stock; the elevated valuation used during the transition is no longer required. The proof is in Appendix G.

4.2 Speed: fading beliefs and the learning limit

The constant-wedge experiment identifies the feasibility frontier. I now let the wedge fade according to Bayesian learning. Appendix F derives the filtering equations; this subsection uses their deterministic certainty-equivalent path. Let x_t denote the posterior mean of the perceived excess return. Asset pricing uses persistent-belief pricing: at each date, investors price capital as if the current posterior mean were permanent, so the belief enters the pricing equation as

$$R^p_t = R^a_t + x_t.$$

Thus x_t is the perceived-return wedge in Proposition 3.

The appendix defines h_t as the Kalman learning gain. A higher h_t means that new observations receive more weight in updating the posterior. Along the deterministic path studied below, realized returns do not confirm the optimistic belief, so the posterior mean declines according to

$$\dot{x}_t = -h_t x_t.$$

The learning gain measures how responsive beliefs are to new evidence. As investors observe more realizations of AI-related returns, adoption outcomes, and productivity effects, posterior uncertainty about the excess return falls, so each new observation receives less weight. The learning gain therefore declines:

$$\dot{h}_t = -h_t^2.$$

Solving these two equations gives

$$x_t = \frac{x_0}{1 + h_0 t}, \qquad h_t = \frac{h_0}{1 + h_0 t}. \tag{28}$$

Both objects decline over time. The belief never reaches zero in finite time, but it fades steadily as confirming evidence fails to arrive.

The learning phase is built from three equations. The price is the perceived high-branch value $q_t = Q^H(K_t; x_t)$; capital follows the investment law (10); and beliefs follow the law summarized in (28). Together these give the temporary-equilibrium system

$$q_t = Q^H(K_t; x_t), \qquad \dot{K}_t = K_t\big[\psi \log Q^H(K_t; x_t) - \delta\big], \qquad \dot{x}_t = -h_t x_t, \qquad \dot{h}_t = -h_t^2, \tag{29}$$

which governs the economy while the perceived high-capital branch exists, with consumption still given by (17).

The transition succeeds if the belief-supported high-capital continuation remains available until capital enters the rational high-capital domain $\mathcal{K}_0$. There are therefore two stopping times. The first is the arrival time,

$$\tau_V \equiv \inf\{t : K_t \in \mathcal{K}_0\},$$

the first date at which accumulated capital is high enough for the rational high-capital continuation to be available without the belief wedge. The second is the exit time,

$$\tau_* \equiv \inf\{t : (K_t, x_t) \notin \mathcal{D}^H\},$$

the first date at which the fading belief is no longer large enough to support the perceived high-capital branch at the current capital stock.

Success requires arrival before exit, $\tau_V < \tau_*$. In that case, the belief has lasted long enough to move the economy into the rational high-capital domain. If instead $\tau_* < \tau_V$, the path fails: the belief fades before capital has crossed the threshold. For the decaying-belief paths considered here, failure occurs when the belief falls below the required-wedge frontier,

$$x_t < x^*(K_t).$$

At the frontier, $q = \bar{q}$; since $\dot{x}_t < 0$, touching it means exit. Fragility is the race: capital must enter $\mathcal{K}_0$ before evidence erodes the belief financing the move.

Fix an admissible launch belief x_0. A higher learning gain h_0 makes x_t fall faster; a lower h_0 makes the belief more persistent. Define

$$h^*(x_0) \equiv \sup\{h_0 : \tau_V < \tau_*\}.$$

The frontier $x^*(K)$ governs feasibility, while $h^*(x_0)$ governs speed.

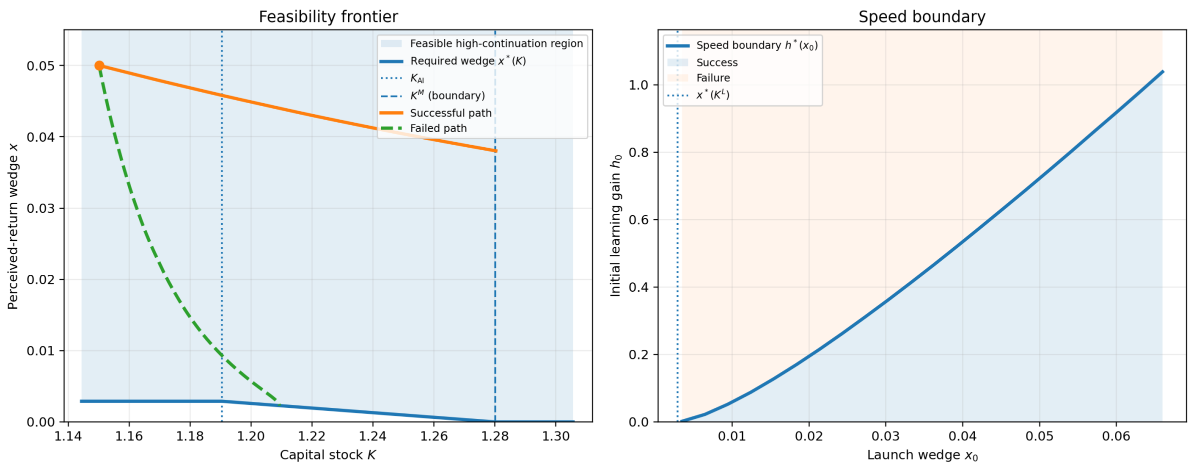

Figure 3: Feasibility and speed. The left panel plots the required-wedge frontier $x^*(K)$, which is flat before AI deployment and then declines linearly to zero at K^M, the boundary of $\mathcal{K}_0$. The successful and failed paths use the same launch belief but different learning gains; both paths are generated from the transition system. The right panel plots the corresponding speed boundary $h^*(x_0)$: below the boundary, learning is slow enough for capital to reach $\mathcal{K}_0$; above it, the belief fades too quickly.

Proposition 4 (Learning speed limit)**.** *There exists $\bar{h} < \infty$ such that, if $h_0 > \bar{h}$, no admissible launch belief implements the transition from $K(0) = K^L$ to $\mathcal{K}_0$. For every admissible $x_0 > x^*(K^L)$, $h^*(x_0) > 0$: sufficiently slow learning implements the transition, while learning above $h^*(x_0)$ causes exit from the perceived high-capital domain before the high-capital rational continuation is reached.*

The proposition gives the fragility result. The belief need not be correct forever, but it must last long enough for capital accumulation to make the high-capital rational continuation available without the wedge. Feasibility alone is not enough: a belief can lie above the static frontier while the induced valuation is only slightly above $\bar{q}$, so capital moves too slowly. Feasibility fixes how large the belief must be; the speed limit fixes how durable it must be, and a belief can satisfy the first without satisfying the second. The speed boundary in Figure 3 records this second constraint.

4.3 Success, failure, and interest rates along the transition

The landing depends on whether capital reaches $\mathcal{K}_0$ before the belief wedge falls below the required threshold. If learning removes the wedge first, the perceived high-capital continuation is no longer available and valuation jumps to the low rational arm at the inherited capital stock.

Corollary 1 (Correction and fragility)**.** *Along any belief-supported path governed by* (29)*, if $\tau_V < \tau_*$, valuation converges to the rational high-capital arm $Q^H(K_t; 0)$ and the economy converges to K^H. If $\tau_* < \tau_V$, the perceived high-capital branch ceases to exist; since K_t is predetermined, valuation jumps to the low rational arm and the economy returns to K^L.*

The asymmetry follows from the availability of rational continuations at the inherited capital stock. In a successful transition, learning removes the wedge only after $K_t \in \mathcal{K}_0$. The rational high arm $Q^H(K_t; 0)$ is then available, so valuation can move onto that arm as the wedge disappears. In a failed transition, the perceived high branch disappears before K_t reaches $\mathcal{K}_0$. Since capital is predetermined, the adjustment must occur through the jump variable q_t; the only admissible rational continuation is then the low arm.

A failed transition does not erase the capital already installed, but it changes its continuation value. The economy lands on the low saddle arm at a capital stock above K^L and converges back to K^L from above. In a successful transition, the economy enters $\mathcal{K}_0$, the wedge becomes redundant, and valuation converges along the high rational arm. Figures 4 and 5 illustrate these dynamics: the first focuses on early valuation, rates, and accumulation, while the second shows arrival in $\mathcal{K}_0$ and continuation on the high arm.

Remark 1 (Interest rates along the transition)**.** *Along the belief-supported transition, the relevant interest-rate object is the real interest rate $i_t = R_t^p = R_t^a + x_t$. The same belief wedge that raises valuation and investment demand therefore raises the interest rate during the transition. The destination rate is different: in any steady state $i^j = R^{ss}(\bar{q}K^j)$, so $K^H > K^L$ implies $i^H < i^L$. A successful transition can therefore feature a temporarily high interest rate and a lower steady-state interest rate. Since $R_t^a = i_t - x_t$, realized returns can be low while the interest rate is high. The gap is the capitalist-side valuation exposure characterized below.*

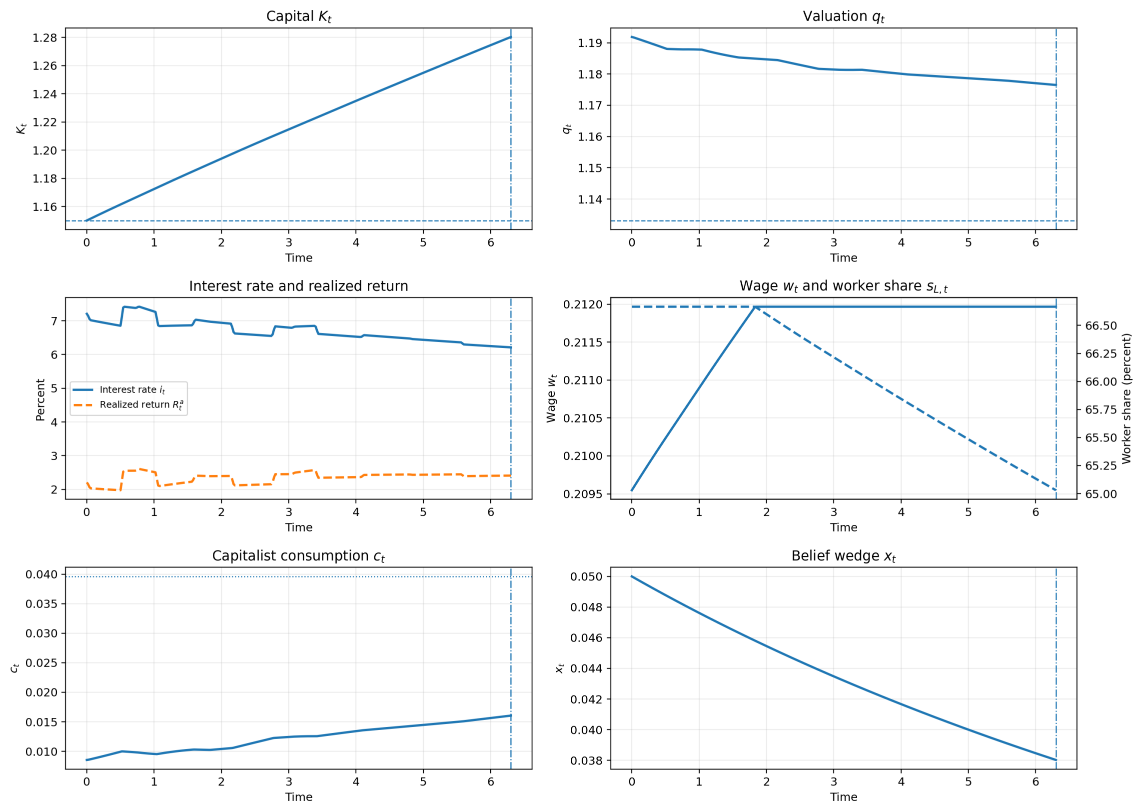

Figure 4: Successful transition: early dynamics. The window focuses on the first years of the transition, when most of the qualitative action occurs. Optimism raises q on impact and triggers rapid capital accumulation. The rate panel reports the real interest rate $i_t = R^p_t$ and the realized return R^a_t during the perceived-branch phase; the interest rate rises while the realized return remains low. The valuation panel reports q_t; the dashed horizontal line is $\bar{q}$, the value at which $\dot{K} = 0$. The vertical line marks the entry into $\mathcal{K}_0$, at which point the belief wedge is no longer needed for feasibility. The worker share is flat before AI deployment begins and then declines once AI capital substitutes for labor tasks, while the wage never falls, as in Lemma 3.

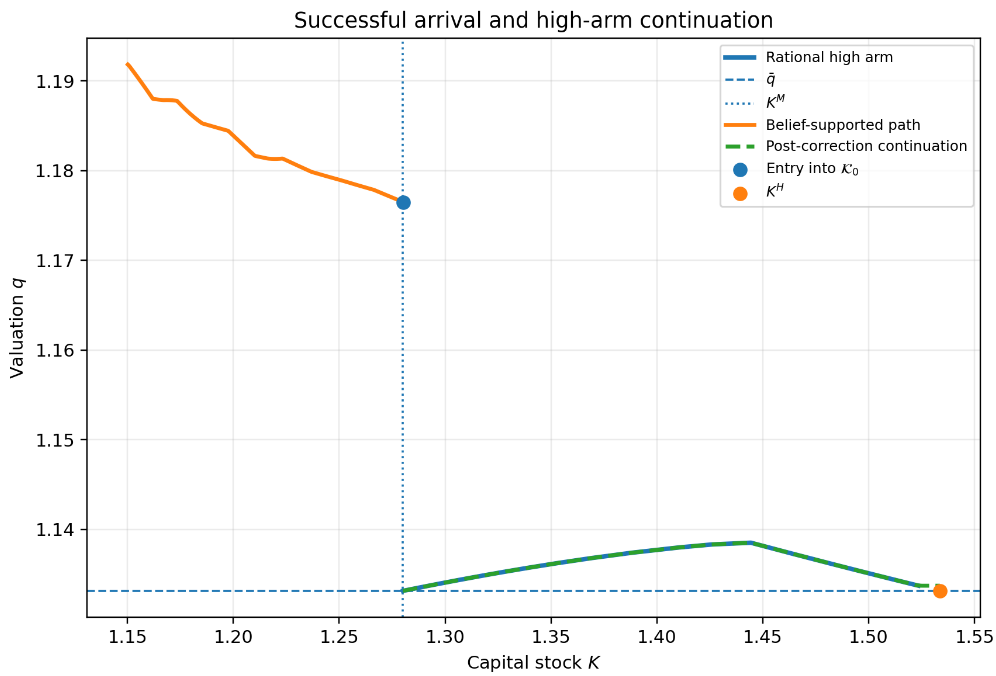

Figure 5: Successful arrival and high-arm continuation. The belief-supported path reaches $\mathcal{K}_0$, after which the wedge is no longer required. The price corrects to the rational high arm rather than to the low arm, and the subsequent continuation converges to K^H as the wedge fades.

4.4 Distributional incidence

The incidence of the transition has two dimensions: destination versus transition, and workers versus the capitalists who finance the investment boom.

The destination is straightforward on the capitalist side. Steady-state wealth is $W = \bar{q}K$, so substituting into the consumption rule (9) gives capitalist consumption

$$c_C(K) = \frac{\rho\bar{q}K}{1+\theta\bar{q}K},$$

which is increasing in K: the high-capital steady state is a wealthier, higher-consumption state for capital owners.

Workers face a share-versus-income distinction, and the distinction turns on a single object: the ratio of conventional capital to effective labor. The wage is the marginal product of effective labor times the single raw unit each worker supplies,

$$w = (1-\alpha)A\left(\frac{K_c}{N}\right)^{\alpha},$$

so the wage tracks K_c/N, and the technology of Section 2 pins down this ratio region by region. Before deployment, $K_c = K$ and $N = 1$, so $K_c/N = K$ and the wage rises with capital: ordinary capital deepening. During deployment, the interior allocation (1) fixes $N = bK_c$, so $K_c/N = 1/b$ and the wage is flat: AI labor expands exactly in step with conventional capital, and the dilution of workers by AI labor exactly offsets the capital deepening. After saturation, $N = 1 + \gamma\bar{K}_\ell$ is fixed while $K_c = K - \bar{K}_\ell$ keeps rising, so the wage rises again. The three expressions coincide at the thresholds—$K_{\mathrm{AI}} = 1/b$ at the first, and $(K_{\mathrm{sat}} - \bar{K}_\ell)/(1+\gamma\bar{K}_\ell) = 1/b$ at the second—so the wage is continuous and weakly increasing in K: rising, flat, rising. The worker share moves differently. From (5), $s_L = (1-\alpha)/N$: constant before deployment, strictly falling while AI labor expands, constant again after saturation. The share falls; the wage never does.

This resolves the apparent tension between AI substitution and rising wages. AI capital substitutes for raw labor tasks, which lowers the worker share as effective labor N expands through AI labor. But wages depend on conventional capital per unit of effective labor, K_c/N. The deployment technology never lets this ratio fall, and at the high-capital destination it is strictly higher than at the start.

At the destination, the comparison compounds these movements with the output expansion. The wage rises between steady states exactly when output grows by more than effective labor,

$$w(K^H) > w(K^L) \iff \frac{Y(K^H)}{Y(K^L)} > \frac{N(K^H)}{N(K^L)}, \tag{30}$$

and in the maintained configuration of Proposition 1 this condition is automatic: K^L lies in the first rising region and K^H in the second, so $w(K^H) > w(K_{\mathrm{sat}}) = w(K_{\mathrm{AI}}) > w(K^L)$.

Lemma 3 (Worker wage monotonicity)**.** *In the technology of Section 2, the worker wage $w(K) = (1-\alpha)Y(K)/N(K)$ is weakly increasing in K, and is strictly higher at the high-capital destination: $w(K^H) > w(K^L)$.*

The monotone wage converts immediately into a statement about paths, because capital never falls below K^L along any belief-supported path. During the transition, $q_t > \bar{q}$ and capital rises. On a successful transition, capital continues to rise while the belief persists, and the economy converges to K^H as the wedge fades. After a crash, the economy lands on the low saddle arm at a capital stock above K^L and converges back to K^L from above, never crossing it. Wages therefore never fall below their initial level, in success or in failure.

Corollary 2 (Worker wage incidence)**.** *Along successful and failed belief-supported paths satisfying Proposition 4 and Corollary 1, the worker wage flow does not fall below its low-capital benchmark* $w(K^L)$*. On a successful transition, workers reach a high-capital destination with higher wages even though the worker share is lower.*

Within the model, then, workers are protected on the downside: a failed transition returns wages to their initial level, and a successful transition raises wages permanently. The proof of both results is in Appendix G.

The capitalist side is different, because capitalists hold the claims priced by the belief-supported valuation and finance the capital accumulation. During the transition, capitalists appear wealthier: q is high and $W = qK$ is marked at the belief-supported price. But two forces work against them at the same time. First, the realized return on the claim is low precisely when the interest rate is high—the gap between the two is the perceived excess return x_t of Remark 1, a return that is priced but never paid. Second, the investment surge itself is financed out of their consumption: by (17), a high q commits output to investment, so the transition is a low-consumption, high-investment phase for its financiers.

The transition has an asymmetric income-and-return profile. Workers receive the wage upside of successful deployment without falling below the low-capital wage path. Capitalists' realized return combines three forces: the destination gain from owning claims in the high-capital economy, the consumption compression during the investment bridge, and the price correction as belief-supported prices converge to rational prices after the wedge disappears.

The learning dynamics also sharpen the capitalist-side cost. A constant wedge just above $x^*(K^L)$ can remain above the weakly decreasing frontier by construction. A fading belief cannot rely on that logic: while $K_t < K_{\mathrm{AI}}$, the frontier is flat but x_t is already falling, so the launch belief must start with extra slack. That extra slack means a higher initial valuation and stronger consumption compression. The same learning dynamics that make optimism temporary also make it costly: the transition must arrive with enough excess belief to survive the subsequent evidence. The numerical parameterization used for the figures is reported in Appendix H.

5 Conclusion

A temporary overvaluation can leave a real legacy when valuation affects investment. In the model, a belief wedge raises valuation and compresses current capitalist consumption relative to investment; the real interest rate rises during the transition even though realized capital returns can be low. The installed AI capital then changes the economy it enters: it substitutes for labor tasks, shifts income toward capital owners, deepens the pool of saving, and lowers the interest rate consistent with a larger capital stock.

This feedback creates a high-capital steady state that is self-sustaining once reached. Rational pricing from the low-capital state remains on the low-capital path, so the transition relies on a temporary overvaluation. The mechanism is fragile because the overvaluation must correct. If enough capital has been installed before the correction, valuation lands on the high rational arm and the capital remains. If learning removes the perceived high branch too soon, valuation crashes to the low arm and the transition collapses.

The distributional implication follows from the same logic. Workers benefit from the installed capital stock through higher wages at the high-capital destination. Capitalists finance the belief-supported valuation; their realized return reflects the high-capital destination gain, the compressed consumption, and the price correction required to get there. The broader lesson is that an asset-price boom should be judged not only by whether current prices are justified by current fundamentals, but also by whether the boom finances the capital accumulation that can make the high-valuation outcome self-sustaining.

A Technology Details

Firms choose the split $K = K_c + K_\ell$, $0 \le K_\ell \le \bar{K}_\ell$, to maximize $Y = AK_c^\alpha N^{1-\alpha}$ with $N = 1 + \gamma K_\ell$. The interior first-order condition equates the marginal value of a unit of capital in its two uses,

$$\alpha A K_c^{\alpha-1} N^{1-\alpha} = (1-\alpha) A K_c^\alpha N^{-\alpha} \gamma,$$

which gives $N = bK_c$ with $b = (1-\alpha)\gamma/\alpha$, hence the allocation (1). Substituting yields the constant marginal product $r^K = \alpha A b^{1-\alpha}$ on the interior region. Below $K_{\rm AI} = 1/b$ the constraint $K_\ell \ge 0$ binds, all capital is conventional, and $r^K = \alpha A K^{\alpha-1}$. Above $K_{\rm sat} = [1 + (\gamma + b)\bar{K}_\ell]/b$ the capacity constraint $K_\ell = \bar{K}_\ell$ binds, $N = 1 + \gamma \bar{K}_\ell$ is fixed, and $r^K = \alpha A (K - \bar{K}_\ell)^{\alpha-1}(1 + \gamma \bar{K}_\ell)^{1-\alpha}$. This is the schedule (2), which is continuous at both thresholds. The wage $w = (1-\alpha)Y/N$ and the worker share (5) follow directly.

B Wealth-in-Utility and the Consumption Rule

Capitalists maximize (6) subject to $\dot{W} = RW - c$. The current-value Hamiltonian is $\mathcal{H} = \log c + \theta W + \lambda(RW - c)$, with first-order condition $1/c = \lambda$ and costate equation $\dot{\lambda} = \rho\lambda - \theta - \lambda R$. Eliminating λ gives the Euler equation (7). In steady state $\dot{c} = \dot{W} = 0$, so $R = \rho - \theta c$ and $c = RW$; solving gives $c^{ss}(W) = \rho W/(1 + \theta W)$ and $R^{ss}(W) = \rho/(1 + \theta W)$, with $dR^{ss}/dW < 0$ as in (9).

C Equilibrium Dynamics

The consumption flow implied by the budget and $W = qK$ is $C(K, q) = K[d(K) - q\psi \log q]$, admissible when $C > 0$. Writing $\dot{c} = C_K \dot{K} + C_q \dot{q}$ with $\dot{K} = K(\psi \log q - \delta)$ and using the Euler and asset-pricing equations gives the rational valuation law. Setting $\dot{q} = 0$ yields the nullcline condition (19). This equation and (10) define F^0. For a fixed perceived excess return x, persistent-belief pricing adds x to the return term in the rational valuation law, replacing $-\delta + d/q - \rho + \theta C$ by $-\delta + d/q + x - \rho + \theta C$. The high-capital branch of the resulting system is $Q^H(K; x)$, and $x = 0$ recovers the rational system.

D Proof of Proposition 1

The steady-state gap is

$$G(K) = d(K) - \Phi(K), \qquad \Phi(K) = \bar{q}\left[\delta + \frac{\rho}{1 + \theta \bar{q} K}\right].$$

Since $\Phi(K) \downarrow \bar{q}\delta$ as $K \to \infty$, the condition $d_{\rm flat} > \bar{q}\delta$ implies that there exists $\bar{K}_\ell^* > 0$ such that, for every $\bar{K}_\ell > \bar{K}_\ell^*$,

$$\Phi(K_{\rm sat}) < d_{\rm flat}.$$

Together with $\Phi(K_{\mathrm{AI}}) > d_{\mathrm{flat}}$, this gives the sandwich (16). The identity $K_{\mathrm{sat}} = K_{\mathrm{AI}} + \bar{K}_{\ell}/(1-\alpha)$ also implies $K_{\mathrm{sat}} > \bar{K}_{\ell}/(1-\alpha)$, so the saturation-region dividend schedule is strictly decreasing above K_{sat}. The outer down-crossing condition (15) is still needed because $G'(K)$ also contains the wealth-saving term from the carrying-cost schedule.

The existence of the three roots follows from continuity. As $K \to 0^+$, $d(K) \to \infty$ while $\Phi(K)$ is bounded, so $G(K) > 0$. At K_{AI}, the left inequality in (16) gives $G(K_{\mathrm{AI}}) < 0$, so a low root exists in $(0, K_{\mathrm{AI}})$. On the flat region, $d(K) = d_{\mathrm{flat}}$ and Φ is strictly decreasing, so G is strictly increasing. The right inequality in (16) gives $G(K_{\mathrm{sat}}) > 0$, so the middle root exists and is unique. Beyond saturation, $d(K) \to 0$ while $\Phi(K) \to \bar{q}\delta > 0$, so $G(K) \to -\bar{q}\delta < 0$; with $G(K_{\mathrm{sat}}) > 0$, a high root exists. If (15) holds, G is strictly decreasing in both outer regions, so each outer root is unique. $\square$

E Local Geometry: Proof of Lemma 1

Write $F^0 = (f, g)$ with $f(K,q) = K(\psi \log q - \delta)$ and $g = Z/D$, where

$$Z(K,q) = C(K,q)\Big[-\delta + \tfrac{d(K)}{q} - \rho + \theta C(K,q)\Big] - C_K(K,q)\,K(\psi \log q - \delta), \qquad D = C_q - \tfrac{C}{q}.$$

Since $C(K,q) = K[d(K) - q\psi \log q]$, direct computation gives

$$D = C_q - \frac{C}{q} = -K\left(\psi + \frac{d(K)}{q}\right) < 0$$

for every admissible (K,q). On the horizontal nullcline $q = \bar{q} = e^{\delta/\psi}$, the capital-growth term vanishes and

$$Z(K,\bar{q}) = C(K,\bar{q})\Big[-\delta + \tfrac{d(K)}{\bar{q}} - \rho + \theta C(K,\bar{q})\Big] = C(K,\bar{q})\frac{1+\theta\bar{q}K}{\bar{q}}\,G(K).$$

Since $C > 0$ and $D < 0$ on the admissible region, this gives (20). At a steady state $q = \bar{q}$, so $f_K = \psi \log \bar{q} - \delta = 0$ and $f_q = K\psi/\bar{q} > 0$. Since $Z(K^j, \bar{q}) = 0$, $g_K = Z_K/D$ there, and

$$\det DF^0 = f_K g_q - f_q g_K = -f_q \frac{Z_K}{D}.$$

At a steady state, the denominator becomes

$$D = -K\Big[\psi + \delta + R^{ss}(\bar{q}K)\Big] < 0.$$

Hence $\operatorname{sign}\det DF^0 = \operatorname{sign} Z_K$.

To evaluate Z_K, let $\mathcal{B} = -\delta + d/q - \rho + \theta C$. At a steady state $\mathcal{B} = 0$ and $\psi \log q - \delta = 0$, so differentiating $Z = C\mathcal{B} - C_K K(\psi \log q - \delta)$ in K at fixed $q = \bar{q}$ leaves only $Z_K = C\mathcal{B}_K$. With $\mathcal{B}_K = d'/\bar{q} + \theta C_K$ and $C_K = (d - \bar{q}\delta) + Kd'$,

$$Z_K(K,\bar{q}) = C(K,\bar{q})\Big[d'\big(\tfrac{1}{\bar{q}} + \theta K\big) + \theta(d - \bar{q}\delta)\Big],$$

which is (22). Using the steady-state condition, $d - \bar{q}\delta = \bar{q}\rho/(1 + \theta\bar{q}K)$, this can be written as

$$Z_K(K, \bar{q}) = \frac{C(K, \bar{q})(1 + \theta\bar{q}K)}{\bar{q}} G'(K).$$

Thus $\operatorname{sign}\det DF^0(K, \bar{q}) = \operatorname{sign} G'(K)$. At K^M, $d' = 0$ and $d - \bar{q}\delta > 0$, so $G'(K^M) > 0$, $Z_K > 0$, and $\det DF^0 > 0$; the trace $g_q = Z_q/D$ is positive there (in the flat region $Z_q = -K(d - \bar{q}\delta)[d/\bar{q}^2 + \theta K(\psi + \delta) + \psi/\bar{q}] < 0$ with $d - \bar{q}\delta > 0$, and $D < 0$), so K^M is a source. At an outer steady state, $G'(K^j) < 0$ implies $\det DF^0 < 0$, hence one stable and one unstable eigenvalue: the steady state is a saddle. The stable branches for the illustrative numerical example are shown in Figure 2. □

F Bayesian Learning and Persistent-Belief Pricing

This appendix derives the belief process used in Section 4.2. Investors observe a noisy signal about a persistent excess return component on AI-related capital,

$$dS_t = \mu\, dt + \sigma\, dB_t,$$

where μ is constant and unknown. Learning enters the aggregate dynamics through the posterior x_t, while the signal process itself evolves independently of (K, q). The prior is Gaussian,

$$\mu \sim \mathcal{N}(x_0, P_0),$$

with posterior mean and variance

$$x_t \equiv \mathbb{E}_t[\mu], \qquad P_t \equiv \mathrm{Var}_t(\mu).$$

The posterior mean x_t is the perceived excess return that enters asset pricing. The posterior variance determines the learning gain,

$$h_t \equiv \frac{P_t}{\sigma^2}.$$

A larger h_t means that new observations receive more weight in updating the posterior.

The continuous-time Kalman filter gives

$$dx_t = h_t(dS_t - x_t\, dt), \qquad \dot{P}_t = -\frac{P_t^2}{\sigma^2}.$$

Equivalently,

$$\dot{h}_t = -h_t^2, \qquad h_t = \frac{h_0}{1 + h_0 t}, \qquad h_0 \equiv \frac{P_0}{\sigma^2}.$$

Thus learning is faster when the prior is more uncertain relative to signal noise. The posterior mean has the explicit representation

$$x_t = \frac{x_0 + h_0 S_t}{1 + h_0 t}.$$

The deterministic transition studied in the text follows the noise-free path under the true value $\mu = 0$. Along this certainty-equivalent path, evidence does not confirm the initial excess-return belief and $S_t \equiv 0$, so

$$\dot{x}_t = -h_t x_t, \qquad x_t = \frac{x_0}{1 + h_0 t}.$$

The belief therefore fades continuously and never reaches zero in finite time. The parameter h_0 is the initial speed of learning: for a fixed launch belief x_0, a larger h_0 makes the perceived excess return decay faster. These closed forms are the laws reported in (28).

The pricing convention is persistent-belief pricing. At each date, capitalists price capital as if the current posterior mean x_t were a persistent excess return. Matching (12), the perceived return is

$$R_t^p = R_t^a + x_t.$$

As evidence arrives, the posterior mean x_t is revised according to the filter above. Hence valuation is computed from the current persistent belief, while the belief itself evolves over calendar time. This is the belief process used in the transition system: while the perceived high-capital branch exists, the price is $q_t = Q^H(K_t; x_t)$, capital follows the q-theory accumulation law, and x_t follows the learning law derived here.

G Proofs for Sections 3–4

Lemma 2. In the maintained phase diagram, K^M is the source that separates the low and high basins. Under the global proviso in the statement of the lemma, the stable manifold of the high saddle, followed backward, cannot cycle inside the admissible region and cannot exit admissibility before reaching the middle source. It also cannot cross the low stable arm, by uniqueness of solutions. The only remaining backward limit is therefore K^M. Hence the rational high-capital continuation is available exactly for inherited capital stocks on the high side of the source, with K^M as the limiting threshold endpoint. With the closed-domain convention, $\mathcal{K}_0 = [K^M, \infty)$. The low stable arm is available for inherited capital stocks in (K^L, K^M). The same backward-limit argument applies on the low side, and admissibility is automatic along the low arm: there $q \leq \bar{q}$ and $K \leq K_{\text{sat}}$, so $q\psi \log q \leq \bar{q}\delta < d_{\text{flat}} \leq d(K)$, which gives $C(K, q) > 0$. □

Closed-form required-wedge frontier. In the flat deployment region, a permanent perceived wedge changes the steady-state condition to

$$d_{\text{flat}} = \bar{q}\left[\delta + \frac{\rho - x}{1 + \theta\bar{q}K}\right].$$

Solving for the perceived middle root gives

$$K^M(x) = \frac{1}{\theta\bar{q}}\left[\frac{\bar{q}(\rho - x)}{d_{\text{flat}} - \bar{q}\delta} - 1\right] = K^M - \frac{x}{\theta(d_{\text{flat}} - \bar{q}\delta)}.$$

By the perceived-system analogue of Lemma 2—under the same global proviso, applied to the perceived system for each fixed x, as verified in the illustrative numerical example—the high-capital perceived branch is available at state values $K > K^M(x)$, with $K^M(x)$ as the threshold endpoint, as long as the perceived middle root remains in the deployment region. Hence, for $K \in [K_{\text{AI}}, K^M]$, the minimum wedge is the value of x that sets the threshold at K, so $K = K^M(x)$:

$$x^*(K) = \theta(d_{\text{flat}} - \bar{q}\delta)(K^M - K).$$

For $K < K_{\text{AI}}$, the binding obstruction is at the deployment threshold: the high branch first becomes available only when the perceived low and middle roots annihilate at K_{AI}. This gives the plateau $x^*(K) = x^*(K_{\text{AI}})$ below K_{AI}. Combining the two cases yields (27). On the declining segment $K \in [K_{\text{AI}}, K^M]$, the lower frontier has endpoint value $q = \bar{q}$, so $C(K, \bar{q}) = K[d(K) - \bar{q}\delta] > 0$ because $d(K) \geq d_{\text{flat}} > \bar{q}\delta$. On the plateau segment, the maintained admissible perceived high arm supplies the positive-consumption branch, as verified in the illustrative numerical example. The upper admissibility edge used in the learning-speed proof below is where the consumption constraint binds.

Proposition 2. I restrict rational equilibria to admissible price paths: $C(K_t, q_t) > 0$ along the path and the capitalist transversality condition holds. The maintained no-closed-orbits condition rules out bounded non-convergent admissible continuations; positivity and transversality rule out paths that leave the admissible region or diverge. Hence an admissible continuation from a given inherited K must lie on a convergent stable branch. By Lemma 2, $K^L < K^M$ implies $K^L \notin \mathcal{K}_0$. Hence the rational high-capital branch is not available at the low inherited capital stock. The only admissible rational branch through $K = K^L$ is the low-capital branch itself, and the unique admissible rational path stays at K^L. □

Proposition 3. A constant admissible wedge $x > x^*(K^L)$ places (K^L, x) in $\mathcal{D}^H$, so the perceived high-capital branch $Q^H(K; x)$ is available at the launch point. The maintained perceived-system branch geometry gives an admissible branch over the compact traverse from K^L to the rational high-capital domain. Let G_x denote the perceived steady-state gap defined in Section 4.1. Since $x^*(K^L) = x^*(K_{\text{AI}})$, the strict inequality $x > x^*(K^L)$ implies $G_x(K_{\text{AI}}) > 0$: the perceived low and middle roots have annihilated at the deployment threshold. On the outer regions $K \in (0, K_{\text{AI}}) \cup (K_{\text{sat}}, \infty)$,

$$G_x'(K) = G'(K) - \frac{\theta\bar{q}^2 x}{(1 + \theta\bar{q}K)^2} < 0,$$

where the inequality uses the maintained outer down-crossing condition (15) and $x \geq 0$. On the flat deployment region, $d(K) = d_{\text{flat}}$ and $\Phi_x(K)$ is strictly decreasing in K, so $G_x(K) = d_{\text{flat}} - \Phi_x(K)$ is strictly increasing. Hence there is no additional crossing below the high root $K^H(x)$; $G_x(K) > 0$ for all $K < K^H(x)$. The sign identity in Lemma 1 extends to the perceived system after replacing G by G_x: at $q = \bar{q}$,

$$Z^x(K, \bar{q}) = C(K, \bar{q})\frac{1 + \theta\bar{q}K}{\bar{q}}G_x(K),$$

so every crossing of $q = \bar{q}$ below the perceived high root is downward. A point on the high stable arm below $\bar{q}$ would have $\dot{K} < 0$ and could not converge to the perceived high saddle; since crossings are one-way, it could not re-cross above $q = \bar{q}$. Hence $Q^H(K; x) > \bar{q}$ on the relevant part of the branch. Since the interval from K^L to the rational high-capital domain is compact and the branch is continuous, $\psi \log Q^H(K; x) - \delta$ is bounded below by a positive number there. Capital therefore reaches $\mathcal{K}_0$ in finite time. □

Proposition 4. Admissibility bounds the valuation above along admissible paths. Before the path reaches $\mathcal{K}_0$, Proposition 3 gives $Q^H(K_t; x_t) > \bar{q}$, so capital is increasing and $K_t \geq K^L$. The dividend schedule is weakly decreasing globally by the technology in Appendix A: it decreases before deployment, is flat in the deployment region, and decreases after saturation. Hence $d(K_t) \leq d(K^L)$. Since $C(K, q) > 0$ implies $q\psi \log q < d(K)$, there is a finite $\bar{Q}$ such that $q_t \leq \bar{Q}$. Hence capital growth is bounded above by

$$\frac{\dot{K}}{K} = \psi \log q - \delta \leq \psi \log \bar{Q} - \delta \equiv \bar{g}.$$

Since the function $q \mapsto q\psi \log q$ is increasing for $q \geq 1$, and $d(K^L) > \bar{q}\delta = \bar{q}\psi \log \bar{q}$, the solution $\bar{Q}$ to $\bar{Q}\psi \log \bar{Q} = d(K^L)$ satisfies $\bar{Q} > \bar{q}$, so $\bar{g} > 0$. Crossing from K^L to K^M therefore takes at least

$$T_{\min} = \frac{1}{\bar{g}} \log \frac{K^M}{K^L} > 0.$$

The growth bound implies that at time $T_{\min}/2$,

$$K_{T_{\min}/2} \leq K^L e^{\bar{g} T_{\min}/2} = (K^L K^M)^{1/2} \equiv \widehat{K} < K^M.$$

Since the required-wedge frontier is weakly decreasing and strictly positive below K^M, define

$$\varepsilon \equiv x^*(\widehat{K}) = \theta(d_{\text{flat}} - \bar{q}\delta)\big(K^M - \max\{\widehat{K}, K_{\text{AI}}\}\big) > 0.$$

Thus $x^*(K_{T_{\min}/2}) \geq \varepsilon$. A successful path is still in the feasible set at $T_{\min}/2$, so it must satisfy

$$\frac{x_0}{1 + h_0 T_{\min}/2} \geq x^*(K_{T_{\min}/2}) \geq \varepsilon.$$

By the maintained launch-admissibility restriction, $x_0 \leq \bar{x}$. Since $x_0 \leq \bar{x}$, this inequality fails for all

$$h_0 > \frac{2}{T_{\min}} \left(\frac{\bar{x}}{\varepsilon} - 1 \right) \equiv \bar{h} < \infty.$$

Thus sufficiently fast learning prevents any admissible launch belief from completing the transition. The positive lower bound $h^*(x_0) > 0$ for each admissible $x_0 > x^*(K^L)$ follows from Proposition 3 and continuity: the constant-wedge path succeeds, so small enough learning gains also succeed. □

Corollary 1. The belief-supported price is $Q^H(K_t; x_t)$ while the perceived high-capital branch exists. If $\tau_V < \tau_*$, the path reaches $\mathcal{K}_0$ while the perceived high branch exists. The branch $Q^H(K; x)$ is continuous in x on this domain, and x_t follows the continuous learning law (28); hence valuation moves onto the rational high arm rather than to the low arm. The perceived high root $K^H(x)$ is increasing in x: on the saturation region, G_x rises pointwise with x and is strictly decreasing in K at the high root. As $x_t \downarrow 0$, $K^H(x_t) \downarrow K^H$. The high branch pulls capital toward the current perceived high root—above $\bar{q}$ when $K_t < K^H(x_t)$, below $\bar{q}$ when $K_t > K^H(x_t)$. Therefore the successful path converges to K^H. If $\tau_* < \tau_V$, then $(K_{\tau_*}, x_{\tau_*}) \notin \mathcal{D}^H$, so $Q^H(K_{\tau_*}; x_{\tau_*})$ does not exist. With K predetermined and $K_{\tau_*} \in (K^L, K^M)$, valuation jumps to the low-capital stable arm identified in Lemma 2, and the economy converges to K^L. □

Lemma 3 and Corollary 2. The wage is $w(K) = (1-\alpha)A(K_c/N)^\alpha$. Before AI deployment, $K_c = K$ and $N = 1$, so $w(K) = (1-\alpha)AK^\alpha$, which is strictly increasing. In the deployment region, the first-order condition implies $N = bK_c$, so $K_c/N = 1/b$ and $w(K)$ is constant. After saturation, $N = 1 + \gamma\bar{K}_\ell$ is fixed and $K_c = K - \bar{K}_\ell$, so $w(K)$ is strictly increasing. The formulas coincide at the two thresholds, hence $w(K)$ is weakly increasing everywhere.

It remains to check failed paths. By the sign identity in Lemma 1, $G(K) < 0$ on (K^L, K^M) implies that the vector field points strictly upward across $q = \bar{q}$. The low stable arm identified in Lemma 2 and converging to $(K^L, \bar{q})$ from the right therefore lies below $q = \bar{q}$ on this interval; if it lay above, then $\dot{K} > 0$ and it would move away from K^L. Hence $\dot{K} < 0$ on the post-crash low arm and convergence to K^L is monotone from above. Since $K_t \geq K^L$ on both successful and failed paths and $w(K)$ is weakly increasing, worker wages never fall below $w(K^L)$. □

H Illustrative Numerical Example

The figures use one admissible parameterization of the model:

$$\alpha = \frac{1}{3}, \quad A = 0.3, \quad \rho = 0.07, \quad \delta = 0.05, \quad \theta = 1, \quad \gamma = 0.42, \quad \bar{K}_\ell = 0.167, \quad \psi = 0.4.$$

These values imply $\bar{q} = e^{\delta/\psi}$, $b = (1-\alpha)\gamma/\alpha$, $K_{\mathrm{AI}} = 1/b$, and $K_{\mathrm{sat}} = K_{\mathrm{AI}} + \bar{K}_\ell/(1-\alpha)$. The parameterization satisfies the global conditions used in the phase-diagram arguments: the sandwich condition is $\Phi(K_{\mathrm{AI}}) = 0.0904 > d_{\mathrm{flat}} = 0.0890 > \Phi(K_{\mathrm{sat}}) = 0.0868$, $G' < 0$ on both outer regions, and the three roots are $K^L = 1.1502$, $K^M = 1.2801$, and $K^H = 1.5336$. The required launch wedge is $x^*(K^L) = 0.0029$. Numerical integration of the saddle arms gives a backward high arm that remains admissible until the middle source; the perceived-system analogue holds over the wedge range used in the transition figures.

The transition figures use a common launch belief $x_0 = 0.05$. The successful path uses $h_0^{\mathrm{succ}} = 0.05$, while the failed path uses $h_0^{\mathrm{fail}} = 1$. These values are chosen only to display the qualitative cases of success and failure; the figures are illustrative and are not calibrated to match empirical magnitudes.

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*.

Azariadis, Costas, and Allan Drazen. 1990. “Threshold Externalities in Economic Development.” *Quarterly Journal of Economics* 105(2): 501–526.

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

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.

Cooper, Russell, and Andrew John. 1988. “Coordinating Coordination Failures in Keynesian Models.” *Quarterly Journal of Economics* 103(3): 441–463.

De Long, J. Bradford, and Andrei Shleifer. 1991. “The Stock Market Bubble of 1929: Evidence from Closed-End Mutual Funds.” *Journal of Economic History* 51(3): 675–700.

Farhi, Emmanuel, and Jean Tirole. 2012. “Bubbly Liquidity.” *Review of Economic Studies* 79(2): 678–706.

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

Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer. 2016. “Expectations and Investment.” *NBER Macroeconomics Annual* 30(1): 379–431.

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

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.

Keynes, John Maynard. 1931. “An Economic Analysis of Unemployment.” In Quincy Wright, ed., *Unemployment as a World-Problem: Lectures on the Harris Foundation*. Chicago: University of Chicago Press.

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

Martin, Alberto, and Jaume Ventura. 2012. “Economic Growth with Bubbles.” *American Economic Review* 102(6): 3033–3058.

Matsuyama, Kiminori. 1991. “Increasing Returns, Industrialization, and Indeterminacy of Equilibrium.” *Quarterly Journal of Economics* 106(2): 617–650.

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

Mian, Atif R., Ludwig Straub, and Amir Sufi. 2020. “The Saving Glut of the Rich.” *NBER Working Paper No. 26941*.

Moll, Benjamin, Lukasz Rachel, and Pascual Restrepo. 2022. “Uneven Growth: Automation’s Impact on Income and Wealth Inequality.” *Econometrica* 90(6): 2645–2683.

Murphy, Kevin M., Andrei Shleifer, and Robert W. Vishny. 1989. “Industrialization and the Big Push.” *Journal of Political Economy* 97(5): 1003–1026.

Pástor, Luboš, and Pietro Veronesi. 2009. “Technological Revolutions and Stock Prices.” *American Economic Review* 99(4): 1451–1483.

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

Rosenstein-Rodan, Paul N. 1943. “Problems of Industrialisation of Eastern and South-Eastern Europe.” *Economic Journal* 53(210–211): 202–211.

Scheinkman, José A., and Wei Xiong. 2003. “Overconfidence and Speculative Bubbles.” *Journal of Political Economy* 111(6): 1183–1220.

Simsek, Alp. 2013. “Belief Disagreements and Collateral Constraints.” *Econometrica* 81(1): 1–53.

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

Tirole, Jean. 1985. “Asset Bubbles and Overlapping Generations.” *Econometrica* 53(6): 1499–1528.

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

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