Speculative Growth and the AI “Bubble”

Ricardo J. Caballero*

June 15, 2026

Link to the latest version

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: a price bubble can have a permanent real legacy. AI capital expands productive capacity and shifts income toward high-saving capital owners. Over time, this funding 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 overvaluation: investors perceive high returns, valuation rises, investment accelerates, and the transition itself raises the interest rate. The overvaluation must eventually correct; the question is whether it corrects too soon. If enough capital has been installed before the correction, the economy lands in the high-capital state. If not, the transition fails. The technology can be real even when peak valuations are not sustained. The incidence is asymmetric: workers receive higher wages at the high-capital destination despite a lower worker share, while capitalists finance the capital accumulation and bear the cost of the correction in belief-supported prices.

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

Keywords: AI, speculation, investment transition, multiple steady states, behavioral expectations, labor share, wealth-in-utility, saving glut, market capitalization.

---

[*] MIT and NBER. Email: caball@mit.edu. 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.

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? That binary framing misses an important possibility. When valuation affects investment, prices can shape the fundamentals they appear to anticipate. A high valuation strengthens investment demand and raises the interest rate during the transition. If enough capital is installed, the high-capital steady state later supports a larger capital stock at a lower interest rate. This paper develops a model in which a temporary overvaluation can create a permanent real legacy even though the valuation itself eventually corrects. The mechanism is fragile: the legacy survives only if the correction arrives after enough capital has been installed.

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 funding 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. Starting from the low-capital state, rational pricing keeps the economy on the low-capital path. The high-capital state exists, but rational dynamics do not take the economy there. 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 transition can leave a permanent real legacy even though the valuation that financed it later corrects: the technology can be real, peak valuations can fall, and the capital financed during the transition can remain.

The same distinction between the transition and the destination also shapes incidence. Workers benefit because AI deployment pulls in conventional capital alongside it: workers operate with a larger conventional capital stock and wages rise even as the worker share falls. The capitalists who finance the capital accumulation bear the valuation risk: 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 funding cost of capital fall as capitalist wealth rises. Beliefs follow Bayesian extrapolation: investors estimate a persistent excess return from a noisy signal and price capital as if the current estimate were permanent, revising the estimate as evidence arrives. Valuation moves accumulation, wealth lowers required returns, and temporary beliefs supply a transition force that fades when the data do not confirm it.

The current AI cycle gives this mechanism empirical relevance. 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). Large AI-related investment demand is one force behind the current high interest rate. Consistent with the destination force in this paper, Andrews and Farboodi (2025) find that news pointing to faster AI adoption lowers yields at medium-to-long maturities.

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 transition. 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. Extrapolative beliefs provide the transition force. Relative to CFH, the transition force here is behavioral: extrapolative valuations move capital across the separatrix before learning removes the wedge.

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. 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 funding 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 extrapolative valuation, in the spirit of sentiment and extrapolative-expectations models (Barberis, Shleifer, and Vishny, 1998; Barberis et al., 2015; Greenwood and Shleifer, 2014). Historical evidence from the late-1920s stock market also points to overvaluation: De Long and Shleifer (1991) use closed-end fund premia to measure investor sentiment and conclude that the S&P composite was priced substantially above fundamentals in late summer 1929. The ingredient here is Bayesian extrapolation: investors price the current perceived excess return as persistent while subsequent evidence updates the belief, which makes the impulse temporary by construction. Because investment responds to valuation, the temporary belief wedge has real effects: it raises capital accumulation before valuation returns to rational pricing.

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 funding 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, funding, and valuation blocks. Section 3 develops the rational benchmark: AI deployment and the funding 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 provide the numerical parameterization behind the figures, the Bayesian foundation for the extrapolative belief, technology details, the wealth-in-utility derivation, the equilibrium system, and the proofs.

2 Model

The model has three blocks. The technology block describes how AI capital expands effective labor and changes factor shares. The funding 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 diminishing returns are muted: 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, compute, or organization. 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.

With $\gamma > 0$ and $\bar{K}_\ell > 0$, 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 C). The lower threshold $K_{\text{AI}} = 1/b$ is where the interior allocation first activates AI deployment: $K_\ell = 0$ at the threshold and positive beyond it. The upper threshold $K_{\text{sat}} = [1 + (\gamma + b)\bar{K}_\ell]/b$ is where it reaches the capacity constraint. Since

$$K_{\text{sat}} - K_{\text{AI}} = \frac{(\gamma + b)\bar{K}_\ell}{b} > 0,$$

the interior deployment region has positive length. These conditions yield the three-region schedule for the marginal product of capital,

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

Let $r^K_{\text{flat}} \equiv \alpha A b^{1-\alpha}$ denote the flat-region marginal product; in the interior deployment region this also equals the capital-claim dividend, so $d_{\text{flat}} \equiv r^K_{\text{flat}}$. The flat middle region is the interior deployment region, generated by the allocation (1).

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

$$s_L(K) = \frac{w}{Y} = \frac{1 - \alpha}{N(K)} = \begin{cases} 1 - \alpha, & K < K_{\text{AI}}, \\ (1 - \alpha)\dfrac{\gamma + b}{b(\gamma K + 1)}, & K_{\text{AI}} \leq K < K_{\text{sat}}, \\ \dfrac{1 - \alpha}{1 + \gamma \bar{K}_\ell}, & K \geq K_{\text{sat}}. \end{cases} \tag{3}$$

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

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

$$d(K) \equiv \frac{Y(K) - w(K)}{K}. \tag{4}$$

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. In the pre-deployment and interior deployment regions, the deployment margin is adjustable and this 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]

2.2 Households and the funding 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,

$$\int_0^\infty e^{-\rho t}\left[\log c_t + \theta W_t\right] dt, \qquad \dot{W}_t = R_t W_t - c_t, \tag{5}$$

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

$$\frac{\dot{c}_t}{c_t} = R_t - \rho + \theta c_t. \tag{6}$$

Rearranging (6) defines the Euler-implied funding rate,

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

This is the interest-rate object used below. It is the return that prices capitalists' intertemporal saving and enters the Euler equation. Along a rational path, this return is the realized return on the capital claim, R_t^a. Along a behavioral path, this return is the perceived return R_t^p, which differs from R_t^a by the belief wedge defined below.

The steady-state funding schedule follows directly from (6) and the budget constraint. Setting $\dot{c} = 0$ in (6) gives $R = \rho - \theta c$. Setting $\dot{W} = 0$ in (5) gives $c = RW$. Solving these two equations for c and R gives

$$c^{ss}(W) = \frac{\rho W}{1 + \theta W}, \qquad R^{ss}(W) = \frac{\rho}{1 + \theta W}, \qquad \frac{dR^{ss}}{dW} = -\frac{\rho\theta}{(1 + \theta W)^2} < 0. \tag{8}$$

As capitalist wealth rises, desired saving rises and the steady-state interest rate falls. This downward-sloping funding 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.

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:

$$\frac{\dot{K}_t}{K_t} = \psi \log q_t - \delta, \tag{9}$$

---

[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.

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:

$$R_t^a = \frac{\dot{q}_t}{q_t} - \delta + \frac{d(K_t)}{q_t}, \qquad W_t = q_t K_t, \tag{10}$$

The behavioral phase is summarized by a single object: a posterior belief x_t about an excess return on capital. The perceived return is

$$R_t^p = R_t^a + x_t. \tag{11}$$

This equation defines the return that enters the behavioral Euler equation. Actual resources are governed by R_t^a; intertemporal choices, and therefore the Euler-implied funding rate i_t, are governed by R_t^p. 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 model has multiple steady states but a unique rational continuation from a given inherited capital stock. The technology and funding 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 (9), 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. Set $\dot{q} = 0$ in the return equation (10): the realized return on the claim reduces to its payout yield net of depreciation,

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

On the household side, stationary consumption in the Euler equation (6) and a stationary budget constraint deliver the funding schedule (8) evaluated at steady-state wealth $W = \bar{q}K$,

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

A steady state must satisfy both: the return the claim pays must be the return at which capitalists are content to hold their wealth constant. Equating the two expressions and multiplying through by $\bar{q}$ gives the steady-state condition

$$d(K) = \bar{q}\left[\delta + \frac{\rho}{1 + \theta\bar{q}K}\right], \qquad G(K) \equiv d(K) - \bar{q}\left[\delta + \frac{\rho}{1 + \theta\bar{q}K}\right]. \tag{12}$$

The first relation in (12) equates two functions of K. The dividend schedule $d(K)$ is the payoff on the aggregate capital claim. The funding curve is

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

the cost of carrying a unit of installed capital: depreciation plus the Euler-implied funding rate, both scaled by the replacement value $\bar{q}$. The gap $G(K)$ is the difference between these two functions. Steady states are the values of K at which $G(K) = 0$, or equivalently $d(K) = \Phi(K)$.

Figure 1 plots these two schedules against K. The blue line is $d(K)$. The orange line is $\Phi(K)$, which slopes downward because a larger installed capital stock raises capitalist wealth $W = \bar{q}K$, and the wealth-saving motive in (5) lowers the funding rate $R^{ss}(\bar{q}K)$. The three intersections are the low, middle, and high steady states.

The key feedback is the slope of the funding curve. A larger K raises capitalist wealth $W = \bar{q}K$; with the wealth-saving motive in (5), higher wealth raises desired saving and lowers the funding 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 funding curve, this redistribution raises aggregate saving capacity and lowers the funding cost of sustaining a larger installed capital stock.

The downward slope of the funding curve is the source of the multiplicity. If $\theta = 0$, the steady-state Euler equation requires $R = \rho$ at any level of wealth, and the funding curve 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 funding curve 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 captures, directly, the saving behavior generated by non-homothetic preferences in Straub (2019). AI deployment matters for the funding block precisely because it shifts income toward the agents with that behavior. The same parameter also governs the transition channel: Remark 2 shows that at $\theta = 0$ belief wedges are absorbed by the funding rate, with no effect on valuation or investment.

The flat payout region is the technological counterpart of this deployment phase. In that region, conventional capital and AI labor expand together, so the marginal product of capital remains high while the funding curve is falling. At low K, the economy can be stationary before AI deployment begins. As deployment begins, the expansion of AI labor shifts income toward capitalists and activates the saving feedback. In the deployment region, the capital payout remains high while the funding cost falls, so the gap $G(K)$ turns upward and can cross zero again. After saturation, diminishing returns reassert themselves; the high installed capital stock is sustained by the larger productive base and the lower funding rate created by the wealth-saving channel.

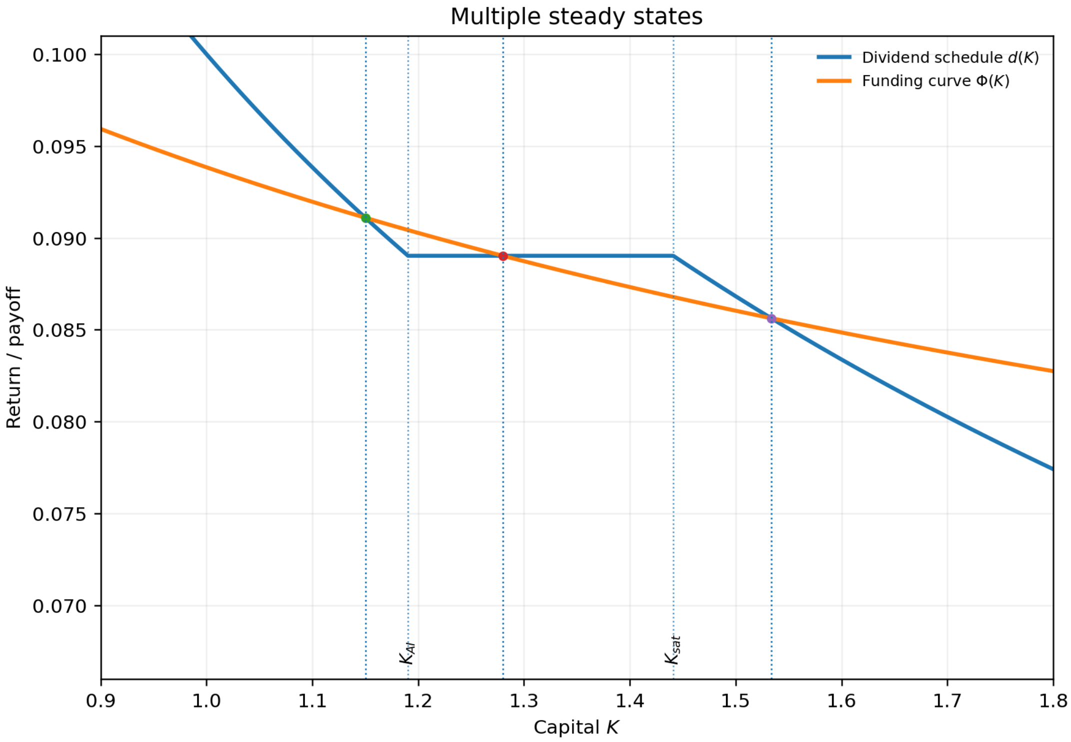

Figure 1: Multiple steady states. The blue line is the dividend schedule $d(K)$ paid by the aggregate capital claim. The orange line is the funding curve $\Phi(K) = \bar{q}[\delta + \rho/(1 + \theta\bar{q}K)]$. Their intersections are the steady states: a pre-deployment low state, a deployment-region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

The same primitives can therefore support a low-capital steady state, a middle steady state, and a high-capital steady state.

Proposition 1 (Multiple steady states)**.** *If the flat-region payout satisfies*

$$\bar{q}\left[\delta + \frac{\rho}{1 + \theta\bar{q}K_{\text{AI}}}\right] > d_{\text{flat}} > \bar{q}\left[\delta + \frac{\rho}{1 + \theta\bar{q}K_{\text{sat}}}\right], \tag{13}$$

then (12) *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,*

$$\left|d'(K)\right| > \frac{\bar{q}^2\rho\theta}{(1 + \theta\bar{q}K)^2} \qquad \text{for } K \in (0, K_{\text{AI}}) \cup (K_{\text{sat}}, \infty), \tag{14}$$

then the outer roots are unique and the economy has exactly three steady states, denoted $K^L < K^M < K^H$.

The inequalities in (13) have a direct graphical interpretation. At the start of AI deployment, the funding curve evaluated at K_{AI} lies above the flat payout. By the time deployment saturates, the funding curve evaluated at K_{sat} has fallen below the same flat payout. The fall in the funding curve is the force that creates room for the additional crossing; the flat deployment region keeps the capital payout from falling while that funding feedback operates. After saturation, diminishing returns reassert themselves and the dividend schedule crosses the funding curve from above. The three crossings are the low, middle, and high steady states. Condition (14) is a uniqueness requirement for the outer crossings; it says that diminishing returns dominate the slope of the funding curve away from the deployment region. It is the same force that later makes the outer steady states saddles, and it holds in the numerical example. The proof is in Appendix F.

It is useful to attach economic labels to the three states. The low state is a conventional economy: capital is too scarce for AI deployment to have transformed production. The middle state is the threshold region. Around it, additional capital does not immediately run into diminishing returns, and the funding feedback pushes in the same direction as accumulation. The high state 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. The high state is rational once reached, but it is not yet clear that rational prices can take the economy there.

This distinction is important. Capital is a state variable. The economy can support more than one rational long-run allocation, and the inherited capital stock determines which rational price paths are available. The next subsection shows that from K^L, rational pricing selects the low continuation and does not launch the economy toward K^H.

3.2 Rational inaccessibility

The steady-state analysis establishes that the rational economy can have both a low-capital and a high-capital steady state. This does not imply multiple rational equilibria from a given inherited capital stock. This section shows that, from K^L, rational pricing selects the low-capital continuation. The high-capital state is self-sustaining once reached, but it is not reached by rational pricing from the low-capital initial condition.

The rational system is a two-dimensional system in the predetermined state K_t and the jump variable q_t. The first is capital:

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

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.

The second moving object is the valuation q_t, and its law of motion comes from combining the budget constraint with the Euler equation. The first step is to show that consumption is a static function of the state (K_t, q_t). 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 (10), 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:

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

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 (9),

$$c_t = C(K_t, q_t) = K_t\big[d(K_t) - q_t \psi \log q_t\big], \tag{15}$$

which must be positive for admissibility. 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 price q_t then moves to make the return on the capital claim consistent with capitalist intertemporal saving. Because consumption is the function $C(K_t, q_t)$, its growth along any path is

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

and the Euler equation (6), written with the realized return (10), becomes

$$\underbrace{\frac{\dot{q}_t}{q_t} - \delta + \frac{d(K_t)}{q_t}}_{\text{return paid by the claim}} = \underbrace{\rho - \theta C(K_t, q_t) + \frac{C_K \dot{K}_t + C_q \dot{q}_t}{C(K_t, q_t)}}_{\text{return required by the Euler equation}}. \tag{16}$$

Every object in (16) is a known function of (K_t, q_t) except $\dot{q}_t$, which appears linearly on both sides. Solving for $\dot{q}_t$ therefore yields the law of motion of the valuation; the explicit expression is recorded in Appendix E. 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 (16); 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 G shows that $D < 0$ for every admissible state.

The $\dot{q} = 0$ locus follows directly from (16). Setting $\dot{q}_t = 0$, the capital-gain term drops from the left side and the $C_q \dot{q}$ term from the right, leaving

$$-\delta + \frac{d(K_t)}{q_t} = \rho - \theta C(K_t, q_t) + \frac{C_K(K_t, q_t)}{C(K_t, q_t)} K_t(\psi \log q_t - \delta). \tag{17}$$

Equation (17) characterizes the $\dot{q} = 0$ locus. The left side is the current dividend return on the capital claim, net of depreciation. The first two terms on the right, $\rho - \theta C$, are the funding rate that would be required if consumption were locally constant. The last term is the consumption-growth correction induced by capital accumulation: when $q > \bar{q}$, capital rises; this changes future consumption through C_K, and the asset price must incorporate that consumption-growth effect. Thus the $\dot{q} = 0$ curve collects the (K, q) pairs at which the current dividend return exactly equals the funding rate plus the transition adjustment. Away from that curve, q_t moves up or down to restore this Euler-pricing balance.

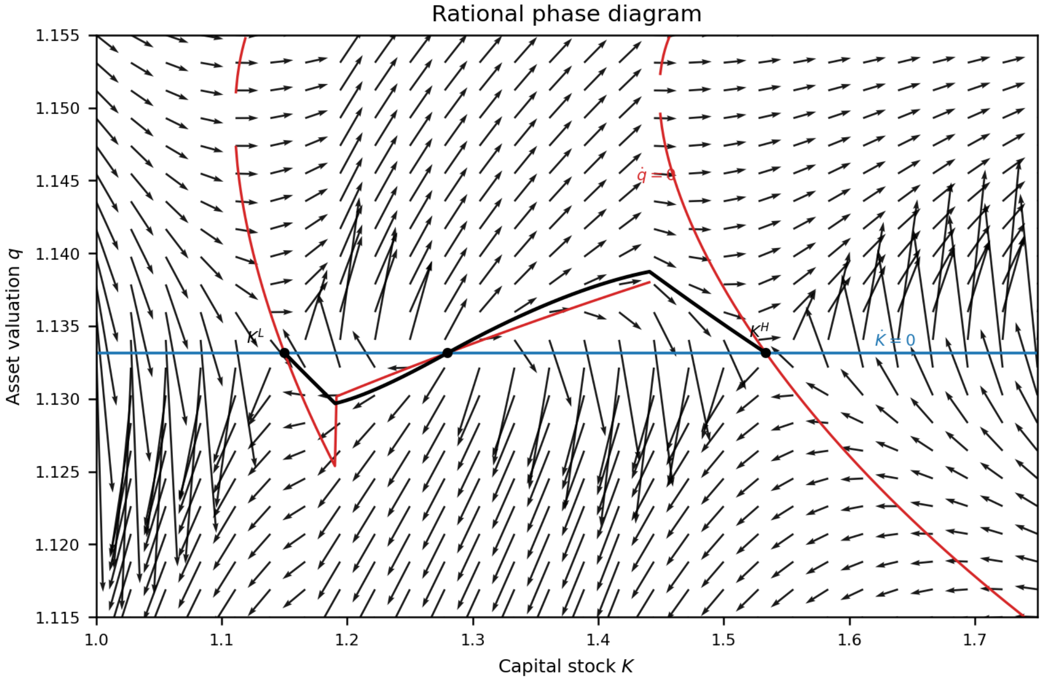

Figure 2: Rational phase diagram. The horizontal nullcline is $\dot{K} = 0$, equivalently $q = \bar{q}$. The second nullcline is $\dot{q} = 0$ from (17). Arrows point right above $q = \bar{q}$ and left below it; they point up where $\dot{q} > 0$ and down where $\dot{q} < 0$. The stable arms are the saddle paths of the low- and high-capital steady states, and the middle source separates the two basins.

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 $\dot{q} < 0$. Hence the regions separated by the two curves have the usual four motions: right-up, right-down, left-up, and left-down. The steady states are the intersections of the two nullclines. At the low and high intersections, the stable arms are saddle paths. At the middle intersection, the flat payout region and falling funding curve make the crossing locally repelling, so the middle steady state separates the basins of attraction. The next result establishes this geometry rather than assuming it.

Lemma 1 (Local rational geometry)**.** *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*

$$\operatorname{sign}\det DF^0(K^j, \bar{q}) = \operatorname{sign} Z_K(K^j, \bar{q}), \qquad Z_K = C\left[d'(K^j)\left(\tfrac{1}{\bar{q}} + \theta K^j\right) + \theta\left(d(K^j) - \bar{q}\delta\right)\right]. \tag{18}$$

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* $Z_K(K^M, \bar{q}) > 0$*; the trace is also positive there;*

(ii) the outer steady states K^L and K^H are saddles whenever

$$\left|d'(K^j)\right|\left(\frac{1}{\bar{q}} + \theta K^j\right) > \theta\left(d(K^j) - \bar{q}\delta\right), \qquad j \in \{L, H\}. \tag{19}$$

The economic content of the lemma is visible from the nullcline geometry. 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 Z to capital. The expression for Z_K in (18) then displays the two competing forces. The first term carries diminishing returns, $d' \leq 0$: more capital lowers the payout, which pushes the price down and pulls the economy back. The second term carries the funding 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, diminishing returns are active and condition (19) says they 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 funding feedback remains. The middle state is therefore a source for any admissible parameters—a property of the deployment technology, not of the example.

Remark 1 (Crossing direction)**.** *At any steady state, the determinant has the same sign as the slope of the steady-state gap G:*

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

Thus saddles are down-crossings, where the dividend schedule cuts the funding schedule from above, while sources are up-crossings. The middle steady state is a source because the flat payout region forces such an up-crossing.

The static diagram therefore already contains the local stability information used in the phase diagram. The proofs are in Appendix G. The one fact that cannot be settled locally—that the high-capital saddle path lies entirely on the far side of the source and does not extend back to K^L—is a global property of the branches, stated next and checked in the numerical example in Appendix G.

Condition 1 (Global branch separation)**.** *The stable arm of the high-capital saddle K^H is not defined at K^L: it lies on the high-capital side of the source K^M.*

Condition 1 makes the separatrix role of K^M global: at K^L, the admissible rational continuation lies on the low-capital branch. The high state is self-sustaining for capital stocks on its branch, but that branch begins only after enough capital has already been installed.

Proposition 2 (Rational inaccessibility)**.** *Under Lemma 1 and Condition 1, 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 dynamics.*

The rational benchmark identifies the transition problem. The economy has two saddle-stable long-run states supported by rational prices, each with its own saddle path: a low-capital state and a high-capital state. From K^L, however, the rational continuation lies on the low saddle arm. Because K is predetermined at K^L, an equilibrium price must place the economy on a convergent branch: any other initial valuation puts it on a path that eventually drives consumption to zero or violates admissibility, and so cannot be an equilibrium. The economy therefore cannot bootstrap itself with anticipated capital gains; the proof is in Appendix H. The high state is self-sustaining once reached, but the low-capital economy reaches it only if perceived returns change while capital accumulates. Section 4 supplies that force and characterizes its consequences and costs.

4 Speculative Transition, Fragility, and Incidence

The rational benchmark leaves a transition problem: the high-capital state is sustainable once reached, but rational pricing from the low-capital state does not move the economy there. This section studies a temporary perceived-return wedge that can bridge the gap. The wedge raises valuation and capital accumulation while actual resources remain governed by true returns; extrapolative beliefs provide a disciplined source for that wedge and determine whether the transition reaches the high rational arm or collapses to the low one.

The transition is a race between capital accumulation and learning. Capital K_t rises when the belief-supported valuation raises investment. The perceived-return wedge x_t falls when realized returns do not confirm the initial optimism. The required-wedge threshold $x^*(K_t)$ also falls with capital, because higher installed capital raises capitalist wealth and lowers the return needed to support the high-capital continuation. The transition succeeds if K_t reaches the rational high-capital domain $\mathcal{K}_0$, where no wedge is needed, before x_t falls below $x^*(K_t)$.

4.1 A constant wedge and the required-wedge frontier

This subsection holds the perceived-return wedge fixed. The question is: at a given inherited capital stock K, how large must the perceived-return wedge be for the high-capital continuation to exist?

Let $x \geq 0$ denote a constant perceived excess return on capital. Investors price capital using the perceived return $R^p = R^a + x$ from (11), while actual resources are governed by R^a. The wedge raises the private intertemporal return used to price capital. A unit of installed capital is therefore more valuable to investors, q rises, and higher q induces faster accumulation.

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 funding schedule (8) with ρ replaced by $\rho - x$: investors who perceive an extra return x are content to carry wealth at a true return lower by exactly the amount the belief supplies. 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

$$d(K) = \bar{q}\left[\delta + \frac{\rho - x}{1 + \theta\bar{q}K}\right]. \tag{20}$$

In the language of Figure 1, the wedge shifts the funding curve down by $\bar{q}x/(1 + \theta\bar{q}K)$.[2] The rational funding feedback lowers the funding curve as wealth accumulates; the wedge lowers it immediately, by belief rather than by wealth. For a fixed K, a large enough wedge can therefore make the high-capital continuation privately attractive even when rational pricing would not. For a fixed wedge, a higher K makes the high continuation easier to sustain, because the wealth that has already accumulated is doing part of the work.

Three objects organize the transition. $Q^H(K; x)$ is the high-capital saddle price when investors perceive a constant excess return x. $\mathcal{D}^H$ is the set of (K, x) pairs for which that continuation exists and leaves consumption positive. The required-wedge threshold $x^*(K)$ is the boundary of this set: the minimum wedge required at capital K. Finally, $\mathcal{K}_0$ is the subset of capital stocks for which $x^*(K) = 0$.

Formally, for the relevant range of x, the permanent-x economy has the same phase-diagram structure as Figure 2, with the $\dot{q} = 0$ locus shifted by the wedge; in particular, it has a high-capital saddle and a stable arm leading to it. Let $Q^H(K; x)$ denote that high-capital saddle branch. It is the price, at capital K, of the continuation that converges to the permanent-x high steady state. The rational high arm is the special case $Q^H(K; 0)$.

Two requirements define where this branch is available: the saddle branch $Q^H(K; x)$ exists, and the implied consumption $C(K, Q^H(K; x))$ is positive. The consumption restriction matters because, by (15), a valuation far above $\bar{q}$ commits current output to investment. Collect the pairs that satisfy both requirements in

$$\begin{aligned} \mathcal{D}^H &\equiv \{(K, x) : Q^H(K; x) \text{ exists and } C(K, Q^H(K; x)) > 0\}, \\ x^*(K) &\equiv \inf\{x \geq 0 : (K, x) \in \mathcal{D}^H\}. \end{aligned} \tag{21}$$

The frontier $x^*(K)$ is the required-wedge threshold. When K is close to K^L, fundamentals alone do not support the high-capital path, so the required wedge is positive. Below this threshold, the perceived high-capital continuation is not available at the inherited capital stock: the shifted phase diagram does not provide an admissible high-capital saddle branch with positive consumption. A wedge above the threshold shifts the perceived vector field enough to make that continuation privately supportable.

---

[2] At $\theta = 0$ the shift is maximal yet inoperative: the stationary schedule there implies $c = (\rho - x)W$ rather than the log policy $c = \rho W$, and the transversality condition selects the latter; see Remark 2.

The threshold falls with capital through the same wealth-saving feedback that generated the high steady state. As investment raises K, wealth $W = \bar{q}K$ rises and the funding schedule requires a lower true return. The belief wedge therefore has less work to do. The direct shift from a given wedge is scaled by $1 + \theta W$, but the relevant object is the threshold $x^*(K)$: it falls because accumulated wealth itself increasingly supports the high-capital continuation.

The capital stocks at which no wedge is needed form the rational high-capital domain,

$$\mathcal{K}_0 \equiv \{K : (K, 0) \in \mathcal{D}^H\}, \qquad \tau_V \equiv \inf\{t : K_t \in \mathcal{K}_0\}.$$

Thus τ_V is the first time at which the rational high-capital continuation is available at the inherited capital stock. In the numerical example, $x^*(K)$ is monotone over the relevant range, so $\mathcal{K}_0$ is a threshold region: once capital crosses the lower boundary of $\mathcal{K}_0$, no belief wedge is needed to remain on the high-capital rational continuation.

4.2 Learning, declining beliefs, and the race to $\mathcal{K}_0$

The previous subsection defined the static threshold $x^*(K)$. This subsection introduces the belief path x_t and compares it with that threshold. A wedge path is *feasible* if $(K_t, x_t) \in \mathcal{D}^H$ along the induced path, so that the perceived high-capital branch exists and consumption stays positive throughout. The transition succeeds when

$$x_t \geq x^*(K_t) \quad \text{for all } t < \tau_V, \qquad K_{\tau_V} \in \mathcal{K}_0.$$

The first condition keeps the perceived high-capital continuation available before τ_V. The second says that capital reaches the rational high-capital domain before the belief wedge is no longer sufficient.

The pointwise comparison between x_t and the static threshold $x^*(K_t)$ follows from permanent-belief pricing. At date t, investors price the capital claim as if the current posterior mean x_t were permanent. The relevant perceived branch is therefore the constant-wedge branch $Q^H(K_t; x_t)$, evaluated at the current state. This is why the static frontier from Section 4.1 governs feasibility along the dynamic belief path.

Figure 3 plots the race in the (K, x) plane. The horizontal axis is the capital stock and the vertical axis is the perceived-return wedge. The solid frontier is $x^*(K)$, the minimum wedge required to keep the high-capital continuation feasible at capital K. The shaded region above the frontier is the feasible high-continuation region. The successful path starts at (K^L, x_0), moves right as capital accumulates and down as learning lowers the belief, and reaches the boundary $K_0^{\min} = K^M$ before crossing the frontier. The failed path also starts at (K^L, x_0), but learning lowers the belief fast enough that the path crosses the frontier at τ_*. At that point the perceived high-capital branch is no longer available.

Proposition 3 (Temporary perceived-return wedge)**.** *Under Lemma 1 and Condition 1, let $K_0 = K^L$. If a feasible wedge path $\{x_t\}$ has $\tau_V < \infty$ along the induced path, then the wedge implements a transition to the high-capital rational continuation. At and after τ_V, the wedge is no longer required to keep the high-capital continuation available.*

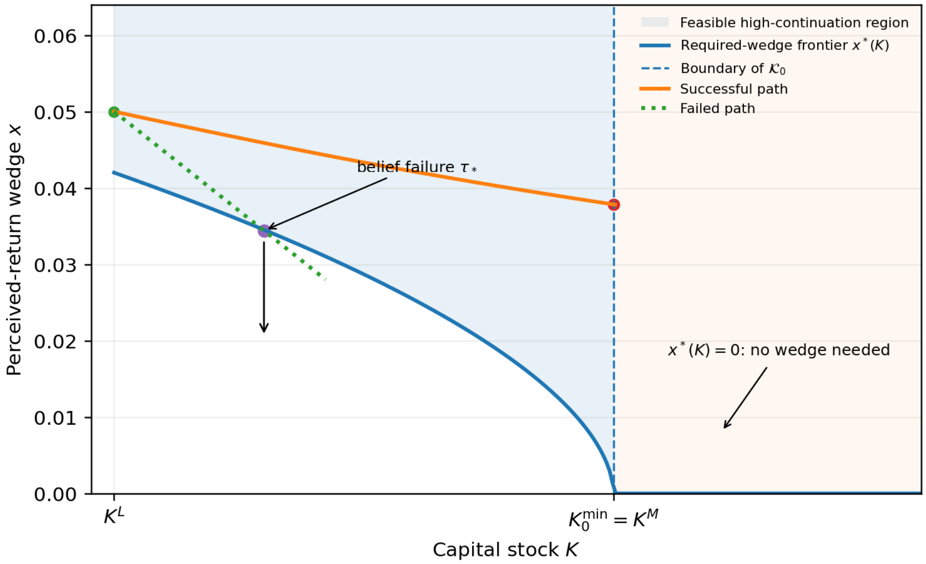

Figure 3: Required-wedge frontier and belief paths. The figure is a schematic guide to the objects in Section 4. The launch belief $x_0 = 0.05$ is taken from the numerical example; the boundary of the rational high-capital domain is $K_0^{\min} = K^M \simeq 1.28$. The pre-entry frontier and failed path are drawn to illustrate the feasibility condition $x_t \geq x^*(K_t)$. The frontier $x^*(K)$ is the minimum perceived-return wedge that keeps the high-capital continuation available at inherited capital K. The shaded region is the feasible high-continuation region. The successful path reaches $\mathcal{K}_0$ before the belief falls below the frontier, while the failed path crosses the frontier first and exits the perceived high-capital branch.

The proposition formalizes the race described at the start of the section. 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 H.

Remark 2 (Belief pass-through requires $\theta > 0$)**.** *The wealth term plays two roles. It makes the required return fall with capitalist wealth, generating the downward-sloping funding schedule; and it gives beliefs traction on valuation and investment. The knife-edge case is* $\theta = 0$*. There the Euler equation and the transversality condition imply the log policy rule* $c_t = \rho W_t = \rho q_t K_t$ *along any deterministic perceived-return path: income and substitution effects cancel, so the consumption-wealth ratio is invariant to beliefs. Combining this rule with the budget constraint* (15) *gives*

$$\rho q_t = d(K_t) - q_t \psi \log q_t.$$

The left side of the equivalent form $q_t(\rho + \psi \log q_t) = d(K_t)$ *has derivative* $\rho + \psi(1 + \log q_t) \geq \psi > 0$ *wherever it is nonnegative, so given* K_t *the valuation is pinned uniquely:* q_t *ceases to be a forward-looking variable, and belief wedges are absorbed one-for-one by the funding rate,*

$i_t = R_t^a + x_t$, with no effect on valuation or investment. For $\theta > 0$, the consumption-wealth ratio falls with the perceived return, so optimism raises desired saving and can finance the capital accumulation. The same term that tilts the funding schedule and creates the high-capital destination gives beliefs traction during the transition—the same role the high intertemporal elasticity of substitution plays in CFH.

The belief model supplies the wedge as a persistent perceived return. Let x_t denote the posterior mean of that perceived excess return. Asset pricing uses anticipated-utility pricing: at each date, investors price capital as if the current posterior mean were permanent, so the belief enters the pricing equation as

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

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

The posterior follows a standard Gaussian filter. The latent perceived excess return is permanent, and h_t is the posterior precision: a higher h_t means that new observations receive more weight in updating the belief. Along the deterministic path studied below, realized returns do not confirm the optimistic belief. The forecast error is therefore negative relative to the current posterior mean, so the posterior mean declines at rate $h_t x_t$:

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

The precision also evolves. With a permanent latent component, each observation reduces posterior uncertainty, but the marginal value of additional observations falls as precision rises. This gives the Riccati equation

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

Solving the two equations gives

$$x_t = \frac{x_0}{1 + h_0 t}, \qquad h_t = \frac{h_0}{1 + h_0 t}. \tag{22}$$

Both objects decline over time. The belief never reaches zero in finite time, but it fades steadily as confirming evidence fails to arrive. Appendix B derives (22) and explains the permanent-belief pricing convention.

The extrapolative 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 (9); and beliefs follow the law summarized in (22). Together these give the temporary-equilibrium system

$$q_t = Q^H(K_t; x_t), \qquad \dot{K}_t = K_t\big[\psi \log Q^H(K_t; x_t) - \delta\big], \qquad \dot{x}_t = -h_t x_t, \qquad \dot{h}_t = -h_t^2, \tag{23}$$

which governs the economy while the perceived high-capital branch exists. Actual consumption and resources remain governed by (15); the wedge changes the return used in valuation, not the resource constraint.

The transition succeeds if the belief path stays above the required-wedge threshold until the threshold reaches zero. Let

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

be the first exit time from the perceived high-capital domain. Along the decaying-belief paths studied here, the binding exit is the lower edge, $x_t < x^*(K_t)$.

Fix x_0. A higher learning gain h_0 makes x_t fall faster; a lower h_0 makes the belief more persistent. Both curves in the race move downward: learning lowers x_t, while capital accumulation lowers the threshold $x^*(K_t)$. A low gain keeps the belief curve above the threshold curve long enough for capital to enter $\mathcal{K}_0$; a high gain makes the belief curve cross the threshold first. Define

$$h^*(x_0) \equiv \sup\{h_0 : \text{the induced path reaches } \mathcal{K}_0 \text{ while feasible}\}.$$

Proposition 4 (Exuberance as a private wedge)**.** *Under Lemma 1 and Condition 1, if there exist $x_0 > 0$ and $h_0 > 0$ such that the solution of* (23) *from $K_0 = K^L$ satisfies $\tau_V < \tau_*$, then the economy admits a transition to the high-capital rational continuation. Along the behavioral phase the belief x_t is the private intertemporal wedge in Proposition 3. Once K_t reaches the high-capital domain, the continuation is self-sustaining without the belief wedge.*

The proposition gives the fragility result. Persistent optimism can finance the transition, but optimism that is removed too quickly cannot. The belief need not be correct forever. It only has to last until capital accumulation makes the high-capital rational continuation available without the wedge.

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 behavioral path governed by* (23)*, elevated valuation is temporary: any convergent continuation to a rational steady state must return to rational saddle geometry. If the path reaches $\mathcal{K}_0$ before belief failure, then under the continuous learning process* (22) *valuation can correct continuously toward the rational high-capital arm, $Q^H(K_t; 0)$, and the economy converges to K^H. If belief failure occurs before $K_t \in \mathcal{K}_0$, 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.

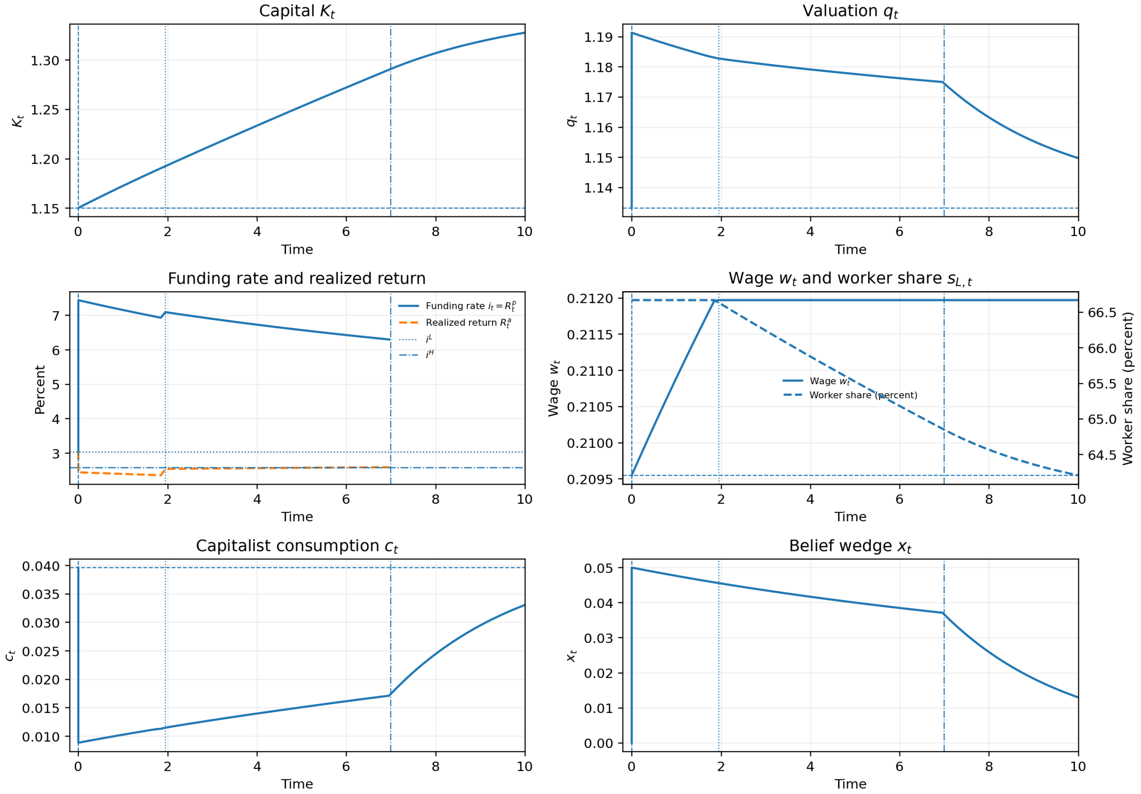

Figure 4: Successful transition: early dynamics. The window focuses on the first years of the transition, when most of the qualitative action occurs. Optimism raises q on impact and triggers rapid capital accumulation. The rate panel reports the Euler-implied funding rate $i_t = R_t^p$ and the realized return R_t^a during the perceived-branch phase; the funding 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 $\dot{K} = 0$. Once capital enters $\mathcal{K}_0$, valuation corrects smoothly toward the rational high arm rather than crashing to the low arm. 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 2.

Remark 3 (Interest rates along the transition)**.** *Along the behavioral transition, the relevant interest-rate object is the Euler-implied funding rate* $i_t = R_t^p = R_t^a + x_t$*. The same belief wedge that raises valuation and investment demand therefore raises the funding 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 funding rate is high. The gap is the capitalist-side valuation risk characterized below.*

The frontier in Figure 3 is the (K, x) representation of the same vector-field shift: the belief wedge temporarily moves the rational $\dot{q} = 0$ locus and the associated high saddle branch, while learning gradually removes that shift as capital moves to the right. Figures 4 and 5 then report the corresponding time paths. The early window shows the valuation-driven investment phase,

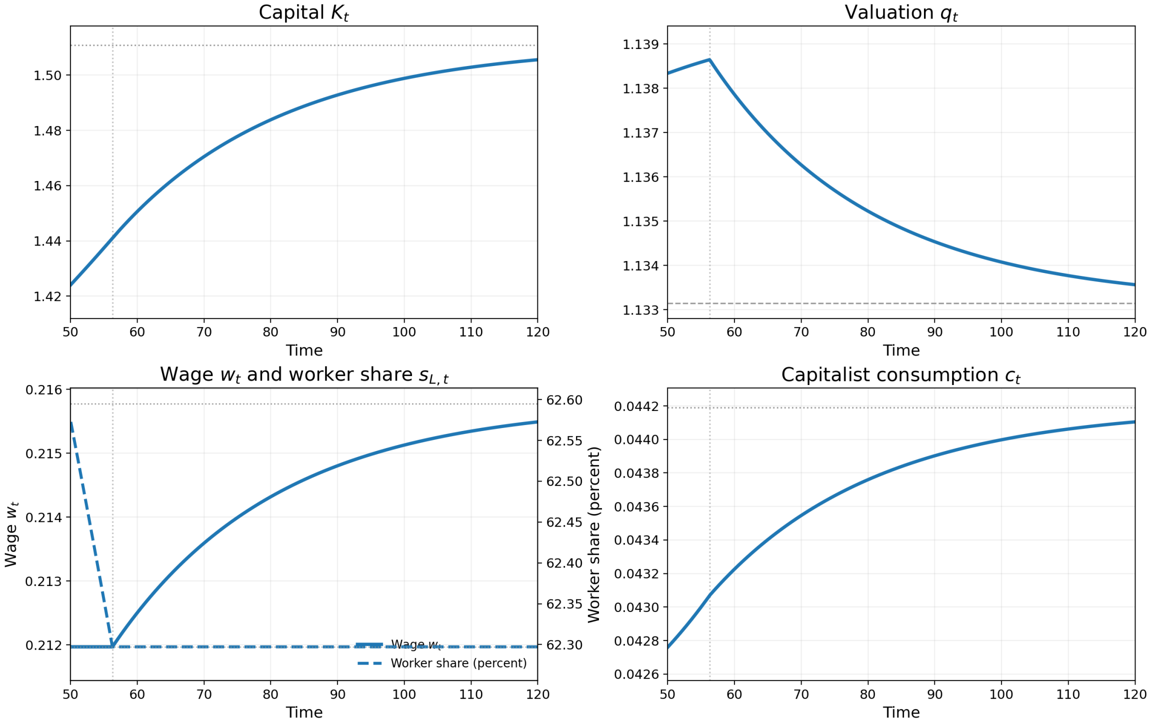

Figure 5: Successful transition: late dynamics. The terminal window illustrates the qualitative approach to the rational high-capital continuation after valuation has rejoined the high arm. Capital and consumption move toward their high-state levels, q_t remains close to $\bar{q}$ while capital keeps moving, and the wage is above its initial level while the worker share is below it.

the temporary increase in the interest rate, and the decline in the belief wedge. The late window shows the subsequent convergence to the high-capital steady state after the wedge has become redundant. The captions describe the individual panels.

4.4 Distributional incidence

The incidence of the transition separates cleanly along two lines: destination versus transition, and workers versus the capitalists who finance the capital accumulation.

The destination is straightforward on the capitalist side. Steady-state wealth is $W = \bar{q}K$, so substituting into the consumption rule (8) 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—$K_{\text{AI}} = 1/b$ at the first, and $(K_{\text{sat}} - \bar{K}_\ell)/(1 + \gamma\bar{K}_\ell) = 1/b$ at the second—so the wage is continuous and weakly increasing in K: rising, flat, rising. The worker share moves differently. From (3), $s_L = (1-\alpha)/N$: 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 the transition also raises the conventional capital stock K_c. Since the wage is a marginal product, it depends on conventional capital per unit of effective labor, K_c/N. The technology keeps this ratio from falling: it rises before deployment, remains constant during interior deployment because AI labor and conventional capital expand in lockstep, and rises again after AI deployment saturates. Thus workers receive a lower share of output, but the conventional capital they work with does not decline; 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,

$$w(K^H) > w(K^L) \iff \frac{Y(K^H)}{Y(K^L)} > \frac{N(K^H)}{N(K^L)}, \tag{24}$$

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 $w(K^H) > w(K_{\text{sat}}) = w(K_{\text{AI}}) > w(K^L)$.

Lemma 2 (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 behavioral path. During the transition, $q_t > \bar{q}$ and capital rises. After a successful correction, capital continues toward K^H. 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 behavioral 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 H.

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 funding rate is high—the gap between the two is the perceived excess return x_t of Remark 3, a return that is priced but never paid. Second, the investment surge itself is financed out of their consumption: by (15), a high q commits output to investment, so the transition is a low-consumption, high-investment phase for its financiers.

To record these effects, let c_t^B, W_t^B, and w_t^B denote realized capitalist consumption, capitalist wealth, and the worker wage along a behavioral path, evaluated at true returns, and define the payoff changes relative to remaining at the low-capital state,

$$\Delta V_C^B = \int_0^\infty e^{-\rho t}\left[\log c_t^B + \theta W_t^B - (\log c_C^L + \theta W^L)\right]dt, \tag{25}$$

$$\Delta V_W^B = \int_0^\infty e^{-\rho t}\left[\log w_t^B - \log w^L\right]dt. \tag{26}$$

In this notation, Corollary 2 implies $\Delta V_W^B > 0$ on any behavioral path that leaves K^L for a positive interval. The capitalist term ΔV_C^B records the transition cost evaluated at true returns: the gains from the eventual high-capital destination, weighed against the consumption compressed during the transition and the valuation losses as the belief-supported price corrects. During the speculative phase, capitalists perceive a high return and choose an aggressive investment path; under true returns, that valuation support is temporary, and the later correction turns the perceived return into an ex-post loss for sufficiently strong optimism. The limiting case makes the sign sharp. As the initial belief approaches the admissibility edge—the largest initial belief on the perceived high branch consistent with positive launch consumption—launch consumption approaches zero and ΔV_C^B becomes arbitrarily negative.

Finally, extrapolation overshoots. Because learning makes the belief fade, a successful initial belief must start above the frontier it eventually needs to cross: a fading x_t can satisfy $x_t \geq x^*(K_t)$ all the way to $\mathcal{K}_0$ only by starting with room to spare. A constant wedge that merely tracked the frontier would require a smaller initial distortion, a smaller launch valuation, and a milder consumption compression. The same learning dynamics that make optimism temporary also make it costly: a fading belief must start above the frontier it later needs, so it overshoots relative to the constant wedge that would just clear the frontier. The numerical parameterization used for the figures is reported in Appendix A.

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 Euler-implied funding 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 and bear the risk that the prices supporting the capital accumulation correct. The model therefore separates three objects: the productivity of the technology, the sustainability of peak valuations, and the permanence of the capital installed during the transition.

A Illustrative Parameterization

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 $K_{\mathrm{sat}} = [1 + (\gamma + b)\bar{K}_\ell]/b$. The transition figures use $x_0 = 0.05$ and $h_0 = 0.05$ in the learning law (22). After the path enters $\mathcal{K}_0$, the plotted residual belief is accelerated so the late convergence to the high-capital arm appears within the figure window.

B Bayesian Extrapolation and Permanent-Belief Pricing

Investors observe a noisy signal of a latent persistent excess return μ on capital, $dY_t = \mu\, dt + \sigma\, dB_t$, and hold a Gaussian prior $\mu \sim \mathcal{N}(x_0, P_0)$. Kalman filtering gives the posterior mean $x_t = \mathbb{E}_t[\mu]$ and variance P_t with

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

so $\dot{h}_t = -h_t^2$. Along the deterministic path on which the realized signal carries no excess return, $dY_t = 0$, the filter innovation is $dY_t - x_t\, dt = -x_t\, dt$. This is the relevant adverse path for fragility, and it gives $\dot{x}_t = -h_t x_t$, hence the closed forms in (22). Pricing is by anticipated utility: at each date investors treat the current posterior mean x_t as the permanent excess return and price accordingly, while updating x_t over time. This is the belief that enters the perceived return (11).

C 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 $K_{\mathrm{AI}} = 1/b$ the constraint $K_\ell \ge 0$ binds, all capital is conventional, and $r^K = \alpha A K^{\alpha-1}$. Above $K_{\mathrm{sat}} = [1 + (\gamma + b)\bar{K}_\ell]/b$ 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 (3) follow directly.

D Wealth-in-Utility and the Consumption Rule

Capitalists maximize (5) 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 (6). 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 (8).

E 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 (17). This equation and (9) define F^0. For a fixed perceived excess return x, anticipated-utility 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.

F Proof of Proposition 1

$G(K)$ in (12) is continuous, with d continuous across regions. As $K \to 0^+$, $d(K) \to \infty$ while the funding term is bounded, so $G > 0$; just past K_{AI}, $d = d_{\text{flat}}$ and the left inequality of (13) gives $G(K_{\text{AI}}) < 0$, so a root exists in $(0, K_{\text{AI}})$. On the flat region d is constant while the funding term decreases, so G is strictly increasing; the right inequality gives $G(K_{\text{sat}}) > 0$, so a root $K^M \in (K_{\text{AI}}, K_{\text{sat}})$ exists and is unique there. Beyond saturation $d(K) \to 0$ while the funding term tends to $\delta\bar{q} > 0$, so $G \to -\delta\bar{q} < 0$; with $G(K_{\text{sat}}) > 0$, a root exists in (K_{sat}, ∞). Finally,

$$G'(K) = d'(K) + \frac{\bar{q}^2 \rho\theta}{(1 + \theta\bar{q}K)^2}.$$

Condition (14) makes $G'(K) < 0$ throughout the two outer regions, so each outer root is unique. □

G 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

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

At a steady state $q = \bar{q} = e^{\delta/\psi}$, 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

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

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). At a steady state this becomes

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

Hence $\operatorname{sign}\det DF^0 = \operatorname{sign} Z_K$, which is (18).

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

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

which is (18). Using the steady-state condition, $d - \bar{q}\delta = \bar{q}\rho/(1 + \theta\bar{q}K)$, this can be written as

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

Thus $\operatorname{sign}\det DF^0(K, \bar{q}) = \operatorname{sign} G'(K)$, as stated in Remark 1. At K^M, $d' = 0$ and $d - \bar{q}\delta > 0$, so $Z_K > 0$ and $\det DF^0 > 0$; the trace $g_q = Z_q/D$ is positive there (in the flat region $Z_q = -KS[d/\bar{q}^2 + \theta K(\psi + \delta) + \psi/\bar{q}] < 0$ with $S = d - \bar{q}\delta > 0$, and $D < 0$), so K^M is a source. At K^L and K^H, $d' < 0$, so $\det DF^0 < 0$ (a saddle) iff (19) holds. Condition 1 is verified for the example by computing the stable branches shown in Figure 2. □

H Proofs for Sections 3–4

Proposition 2. We restrict rational equilibria to admissible price paths: $C(K_t, q_t) > 0$ along the path and the capitalist transversality condition holds. In the rational two-dimensional system, an admissible continuation from a given inherited K must therefore lie on a convergent stable branch. Off-branch continuations either leave the positive-consumption region or contain an explosive valuation component inconsistent with transversality. By Lemma 1, K^L and K^H are saddles with one-dimensional stable arms and K^M is a source. By Condition 1, the stable arm of K^H is not defined at K^L. Hence the only stable branch through $K = K^L$ is that of K^L itself, and the unique admissible rational path stays at K^L. No rational equilibrium from K^L converges to K^H, while the high-capital arm is self-sustaining for any K_0 on it. □

Proposition 3. Feasibility gives $(K_t, x_t) \in \mathcal{D}^H$ along the induced path, so $Q^H(K_t; x_t)$ exists with $C > 0$ and, by the definition of the frontier in (21), $x_t \ge x^*(K_t)$: the path remains on the high-capital perceived branch. The hypothesis $\tau_V < \infty$ means that K_t reaches $\mathcal{K}_0$. At that point the rational high-capital branch is defined at the inherited capital stock and is self-sustaining by Condition 1. Hence the wedge can be removed without sending the economy back to the low rational branch. □

Proposition 4. The behavioral system (23) is the perceived-wedge system of Proposition 3. If $\tau_V < \tau_*$, then $(K_t, x_t) \in \mathcal{D}^H$ for all $t < \tau_V$, so the high-capital perceived branch exists with $C > 0$ throughout and K_t reaches $\mathcal{K}_0$ at τ_V. Along the decaying-belief paths considered in the text, this domain condition is equivalently the lower-frontier inequality $x_t \ge x^*(K_t)$. For $t \ge \tau_V$, the rational high-capital branch is available and self-sustaining. □

Corollary 1. The behavioral price is $Q^H(K_t;x_t)$ while the perceived high-capital branch exists. Since all rational steady states satisfy $q = \bar{q}$, any convergent continuation to a rational steady state must return to rational saddle geometry. 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 (22); hence $Q^H(K_t;x_t)$ can approach $Q^H(K_t;0)$ continuously. 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, valuation jumps to the low-capital stable arm and the economy converges to K^L. □

Lemma 2 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. Along the behavioral phase, capital rises from K^L. If the transition succeeds, the path reaches $K^H > K^L$. If it fails, the post-crash path is the stable arm of K^L starting from capital above K^L; by uniqueness of solutions it cannot cross K^L in finite time and converges to K^L from above. Thus $K_t \geq K^L$ along both successful and failed paths, and worker wages do not fall below $w(K^L)$. □

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.

Barberis, Nicholas, Andrei Shleifer, and Robert Vishny. 1998. “A Model of Investor Sentiment.” *Journal of Financial Economics* 49(3): 307–343.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer. 2015. “X-CAPM: An Extrapolative Capital Asset Pricing Model.” *Journal of Financial Economics* 115(1): 1–24.

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.

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.

Greenwood, Robin, and Andrei Shleifer. 2014. “Expectations of Returns and Expected Returns.” *Review of Financial Studies* 27(3): 714–746.

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*.

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*.

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

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.

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

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.