Speculative Growth and the AI "Bubble"

Ricardo J. Caballero*

June 11, 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 boom 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 boom collapses. The technology can be real even when peak valuations are not sustained.

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

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

*MIT and NBER. Email: caball@mit.edu. I thank 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

The recent boom 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, a price boom can help create the fundamentals that later make part of the boom real. A high valuation finances capital accumulation and, by strengthening investment demand, raises interest rates during the buildout. If enough capital is installed, however, the economy may later support 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 boom leaves a legacy only if enough capital is installed before the correction arrives.

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, the boom raises funding demand and therefore the interest rate. The lower interest rate belongs to the eventual high-capital economy, after the installed capital has changed the distribution of income and the supply of saving.

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 boom collapses back to the low-capital path. Speculation therefore validates a destination, not a price: the technology can be real, peak valuations can correct, and the capital financed during the boom can still leave a permanent legacy.

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 buildout 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, following Straub (2019), so the required return falls with capitalist wealth. 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 episode gives this mechanism empirical relevance. Market capitalization has concentrated in AI-related firms while announced data-center, power, and computing investment points to a large real buildout (Fortune, 2025; Goldman Sachs Global Institute, 2025; Van Nieuwerburgh, 2026; McKinsey, 2025), and recent evidence suggests that AI news has moved long-term rates (Andrews and Farboodi, 2025). The model offers a way to organize that combination. High valuations can coexist with high investment and high transitional interest rates; a later correction in valuations need not imply that the technology was unreal or that the capital financed during the boom disappears. The relevant question is whether enough capital is installed before the correction. This is the same separation illustrated by earlier technology booms: the internet proved economically important even though peak dot-com valuations did not survive.

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. In CFH, a rational-expectations path can connect the low and high states. Here, rational pricing from the low-capital state remains on the low-capital path. The transition to the high-capital state requires temporary overvaluation, and the correction of that overvaluation is part of the mechanism.

The technological environment follows the task-based view that advanced AI changes which tasks reproducible capital can perform. The canonical formalization is Acemoglu and Restrepo (2018), in which automation expands the tasks done by capital, shifts income toward capital, and can sustain its marginal product. The same perspective underlies Restrepo (2025) on an AGI economy and Korinek and Suh (2024) on transition paths as tasks become automatable, and it informs broader assessments in Trammell and Korinek (2023), Acemoglu (2024), Aghion and Bunel (2024), Jones and Tonetti (2026), and Brynjolfsson et al. (2025a).

The funding feedback draws on work connecting inequality, saving, and low interest rates. The consumption rule follows Straub (2019), in which richer households have lower marginal propensities to consume, and relates to Mian, Straub, and Sufi (2021) 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 standard sentiment models, good news and recent price growth push prices above current-fundamental benchmarks (Barberis, Shleifer, and Vishny, 1998; Barberis et al., 2015; Greenwood and Shleifer, 2014). 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 emphasize situations in which expectations coordinate agents on one of several self-fulfilling equilibria (Rosenstein-Rodan, 1943; Murphy, Shleifer, and Vishny, 1989; Cooper and John, 1988; Azariadis and Drazen, 1990; Matsuyama, 1991). The mechanism here is different. The economy has multiple steady states, but from the low-capital state rational pricing selects the low-capital continuation. The high-capital state is self-sustaining once reached, yet it is not reached by a rational coordination switch. The transition instead requires a temporary overvaluation 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 deployment 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 in (2). 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 $K_{\text{sat}} = [1 + (\gamma + b)\bar{K}_\ell]/b$ 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,

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

The flat middle region is the interior deployment region. Its origin is the allocation

$$N = bK_c, \qquad K_\ell = \frac{bK - 1}{\gamma + b}, \tag{2}$$

under which each unit of deployment capital raises effective labor in step with conventional capital, fixing $K_c/N = 1/b$ and hence holding the marginal product of capital constant.

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

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

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

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,

$$\int_0^\infty e^{-\rho t}\left[\log c_t + \theta W_t\right] dt, \qquad \dot{W}_t = R_t W_t - c_t, \tag{4}$$

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

$$\frac{\dot{c}_t}{c_t} = R_t - \rho + \theta c_t. \tag{5}$$

Rearranging (5) defines the Euler-implied funding rate,

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

This is the interest-rate object used below. It is the return that prices intertemporal saving for capitalists and therefore the return that 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 (5) and the budget constraint. Setting $\dot{c} = 0$ in (5) gives $R = \rho - \theta c$. Setting $\dot{W} = 0$ in (4) gives $c = RW$. Solving these two equations for c and R gives

$$c^{ss}(W) = \frac{\rho W}{1 + \theta W}, \qquad R^{ss}(W) = \frac{\rho}{1 + \theta W}, \qquad \frac{dR^{ss}}{dW} = -\frac{\rho\theta}{(1 + \theta W)^2} < 0. \tag{7}$$

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:

$$\frac{\dot{K}_t}{K_t} = \psi \log q_t - \delta, \tag{8}$$

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

$$R_t^a = \frac{\dot{q}_t}{q_t} - \delta + \frac{r^K(K_t)}{q_t}, \qquad W_t = q_t K_t, \tag{9}$$

so q_t is both the wealth price and the investment price.

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

$$R_t^p = R_t^a + x_t. \tag{10}$$

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 persistence of x_t matters for asset prices when investors value future capital returns. Section 4.2 specifies the anticipated-utility pricing rule and the learning process that govern that valuation. 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

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

stationarity of capital requires

$$q_t = \bar{q} \equiv e^{\delta/\psi}.$$

Thus every steady state has the same valuation $q = \bar{q}$: the price at which gross investment exactly covers depreciation. This simple observation does a lot of work later in the paper. The high-capital steady state is a high-K object, not a high-q object, so any valuation above $\bar{q}$ is necessarily a transitional phenomenon. There is no steady state for it to be the price of.

The second condition is stationarity of the capital price. Set $\dot{q} = 0$ in the return equation (9): the realized return on the claim reduces to its payout yield net of depreciation,

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

On the household side, stationary consumption in the Euler equation (5) and a stationary budget constraint deliver the funding schedule (7) 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

$$r^K(K) = \bar{q}\left[\delta + \frac{\rho}{1 + \theta\bar{q}K}\right], \qquad G(K) \equiv r^K(K) - \bar{q}\left[\delta + \frac{\rho}{1 + \theta\bar{q}K}\right]. \tag{11}$$

The first relation in (11) equates two functions of K. The MPK schedule is $r^K(K)$. The funding curve is

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

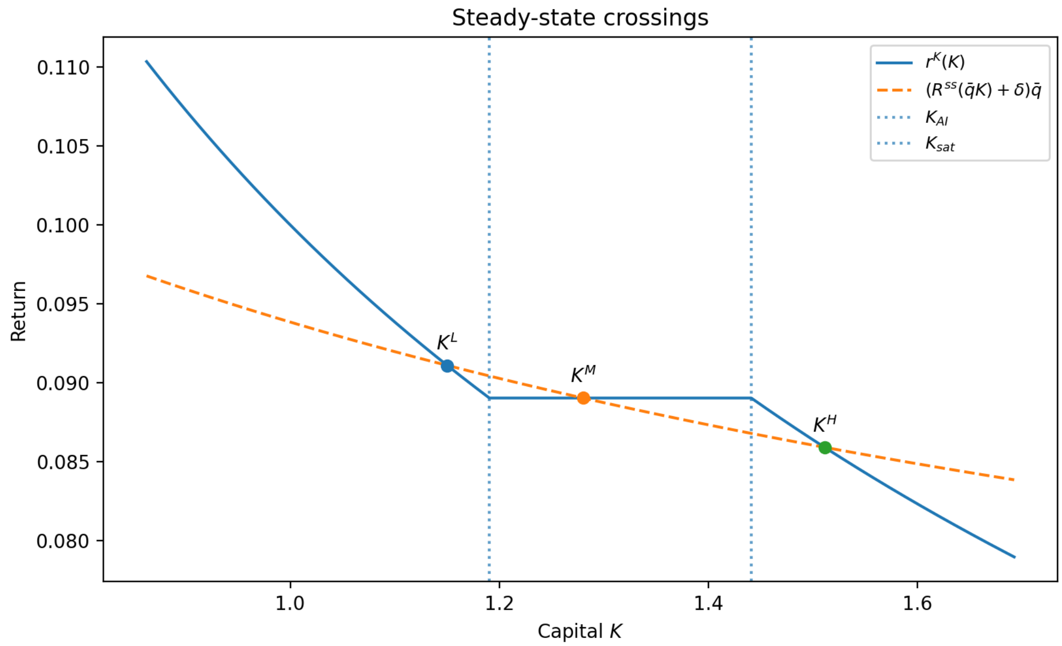

Figure 1: Multiple steady states. The blue line is the MPK schedule $r^K(K)$. The orange line is the funding curve $\Phi(K) = \bar{q}[\delta + \rho/(1 + \theta\bar{q}K)]$. Their intersections are the steady states: a pre-deployment low state, a deployment-region middle state, and a post-saturation high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

the 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 $r^K(K) = \Phi(K)$.

Figure 1 plots the two schedules against K. The blue line is the MPK schedule $r^K(K)$. It declines before AI deployment, is flat in the interior deployment region, and declines again after deployment saturates. The orange line is the funding curve $\Phi(K)$. It slopes downward because a larger installed capital stock raises capitalist wealth $W = \bar{q}K$, and the wealth-saving motive in (4) 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 (4), 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 r^K_{flat}, a horizontal line crosses the three-region MPK 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 is the saving behavior emphasized by Straub (2019), and AI deployment matters for the funding block precisely because it shifts income toward the agents with that behavior.

The flat MPK 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 marginal product 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. 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 MPK satisfies*

$$\bar{q}\left[\delta + \frac{\rho}{1 + \theta\bar{q}K_{\mathrm{AI}}}\right] > r^K_{\mathrm{flat}} > \bar{q}\left[\delta + \frac{\rho}{1 + \theta\bar{q}K_{\mathrm{sat}}}\right], \tag{12}$$

then (11) *has at least one low root in* $(0, K_{\mathrm{AI}})$*, a unique middle root in* $(K_{\mathrm{AI}}, K_{\mathrm{sat}})$*, and at least one high root in* $(K_{\mathrm{sat}}, \infty)$*. If, in addition,*

$$\left|r^{K\prime}(K)\right| > \frac{\bar{q}^2\rho\theta}{(1 + \theta\bar{q}K)^2} \qquad \text{for } K \in (0, K_{\mathrm{AI}}) \cup (K_{\mathrm{sat}}, \infty), \tag{13}$$

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

The inequalities in (12) have a direct graphical interpretation. At the start of AI deployment, the funding curve $\Phi(K_{\mathrm{AI}})$ lies above the flat MPK. By the time deployment saturates, the funding curve $\Phi(K_{\mathrm{sat}})$ has fallen below the same flat MPK. The fall in the funding curve is the force that creates room for the additional crossing; the flat MPK region keeps the MPK from falling while that funding feedback operates. After saturation, the MPK falls again and crosses the funding curve from above. The three crossings are the low, middle, and high steady states. Condition (13) is a mild 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

Starting from the low-capital steady state, capital is predetermined and valuation can jump. One might therefore expect a high valuation to anticipate the high-capital state and sustain the investment needed to reach it, as in models where expectations select among equilibria. The result of this subsection is that this selection is not available here. A high price today is a claim about returns along the entire continuation path, and that claim must be honored at every date by payouts and capital gains. The relevant question is therefore not whether the high-capital state is sustainable once reached—it is—but whether there exists an admissible price path from the inherited capital stock along which the claim is honored throughout.

The rational dynamics have two moving objects. 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 (9), total capital income is

$$R_t^a W_t = \Big(\frac{\dot{q}_t}{q_t} - \delta + \frac{r^K(K_t)}{q_t}\Big) q_t K_t = \dot{q}_t K_t - \delta q_t K_t + r^K(K_t) K_t.$$

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:

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

The cancellation has a clean economic meaning: capital gains raise measured wealth and the measured return by exactly the same amount, so they free no resources for consumption. What remains is output accruing to capital, $r^K K$, minus spending on gross investment, valued at q. Using $\dot{K}_t + \delta K_t = K_t \psi \log q_t$ from (8),

$$c_t = C(K_t, q_t) = K_t\big[r^K(K_t) - q_t \psi \log q_t\big], \tag{14}$$

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 booms returns as a central force in the incidence analysis of Section 4.3.

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 (5), written with the realized return (9), becomes

$$\underbrace{\frac{\dot{q}_t}{q_t} - \delta + \frac{r^K(K_t)}{q_t}}_{\text{return paid by the claim}} = \underbrace{\rho - \theta C(K_t, q_t) + \frac{C_K\dot{K}_t + C_q\dot{q}_t}{C(K_t, q_t)}}_{\text{return required by the Euler equation}}. \tag{15}$$

Every object in (15) 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 (15); it is convenient to write $g = Z/D$, where $Z(K, q)$ collects the numerator and $D(K, q)$ is the denominator of that solution. The denominator is negative on the relevant region; see Appendix G.

The $\dot{q} = 0$ locus follows directly from (15). 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

$$-\delta + \frac{r^K(K_t)}{q_t} = \rho - \theta C(K_t, q_t) + \frac{C_K(K_t, q_t)}{C(K_t, q_t)} K_t(\psi \log q_t - \delta). \tag{16}$$

Equation (16) is the $\dot{q} = 0$ curve, and each piece has an economic reading. The left side is the current payout 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 payout 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 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 MPK 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.

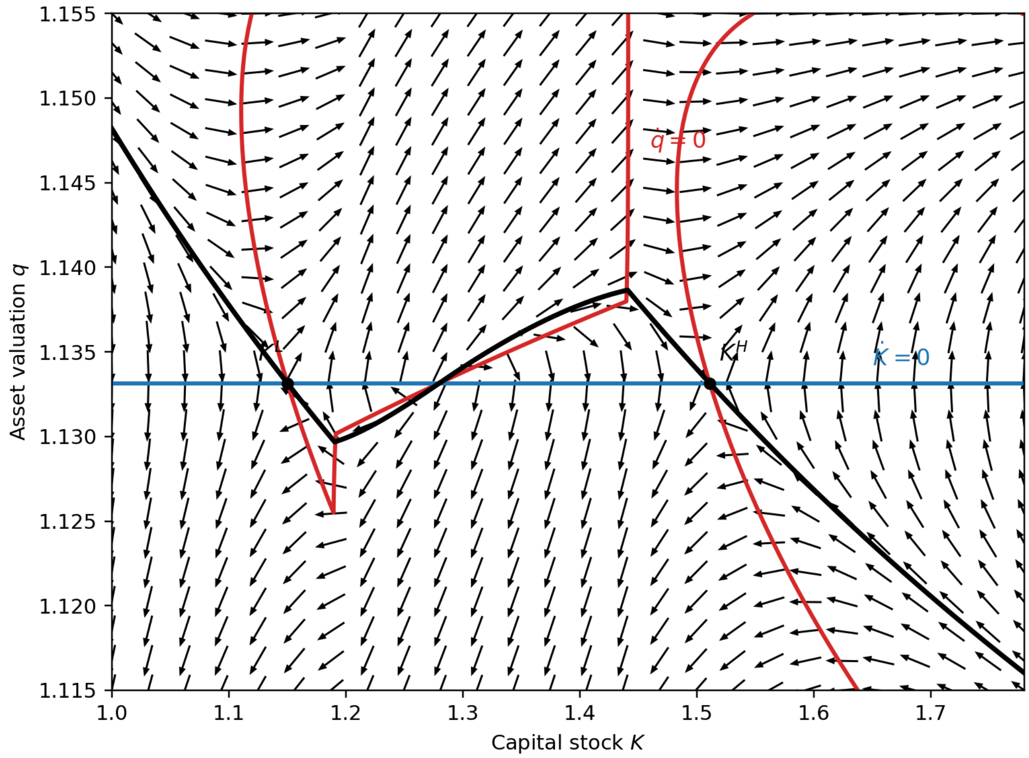

Figure 2: Rational phase diagram. The horizontal nullcline is $\dot{K} = 0$, equivalently $q = \bar{q}$. The second nullcline is $\dot{q} = 0$ from (16). Arrows point right above $q = \bar{q}$ and left below it; they point up where $\dot{q} > 0$ and down where $\dot{q} < 0$. The stable arms are the saddle paths of the low- and high-capital steady states, and the middle source separates the two basins.

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

$$\operatorname{sign} \det DF^0(K^j, \bar{q}) = \operatorname{sign} Z_K(K^j, \bar{q}), \qquad Z_K = C\left[r^{K\prime}(K^j)\left(\tfrac{1}{\bar{q}} + \theta K^j\right) + \theta\left(r^K(K^j) - \bar{q}\delta\right)\right]. \tag{17}$$

Consequently:

- *(i) the middle steady state K^M is a source for any admissible parameters. In the flat region, $r^{K\prime}(K^M) = 0$, 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*

$$\left|r^{K\prime}(K^j)\right|\left(\tfrac{1}{\bar{q}} + \theta K^j\right) > \theta\left(r^K(K^j) - \bar{q}\delta\right), \qquad j \in \{L, H\}. \tag{18}$$

The logic of the Lemma can be followed without the appendix. 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 (17) then displays the two competing forces. The first term carries diminishing returns, $r^{K\prime} \leq 0$: more capital lowers the payout, which pushes the price down and pulls the economy back. The second term carries the funding feedback, $\theta(r^K - \bar{q}\delta) > 0$: more capital raises wealth, lowers the required return, and pushes the price up. At the outer steady states, diminishing returns are active and condition (18) 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 MPK schedule cuts the funding schedule from above, while sources are up-crossings. The middle steady state is a source because the flat MPK region forces such an up-crossing.

The Remark makes the dynamics legible in the static diagram: Figure 1 already contains the stability information of Figure 2. 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 verified for the system 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 says that the high-capital rational continuation is not something the low-capital economy can simply choose by jumping to the right price. The high state is self-sustaining for capital stocks on its branch, but that branch begins only after enough capital has already been installed. This is the sense in which the high state is a destination, not a rational transition from the low state.

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 proposition is the hinge between the rational and speculative parts of the paper. The economy has two 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. What is missing is not sustainability of the high state, but a force that temporarily changes the return investors perceive while capital accumulates. Section 4 supplies that force and asks what it delivers and what it 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 the force that can bridge the gap: a temporary perceived-return wedge. 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 boom lands on the high rational arm or collapses to the low one.

The economic logic has three steps. First, a perceived excess return raises the private value of installed capital. Second, because investment responds to q, the higher valuation raises capital accumulation. Third, the required wedge falls as capital accumulates, because the economy moves closer to the region where the high-capital rational continuation is available without any wedge. Speculation can therefore create the conditions under which it is no longer needed. But this is a race: the belief must last long enough for capital to reach the rational high-capital domain.

4.1 A temporary perceived-return wedge makes the high state accessible

Let $x_t \geq 0$ denote a temporary perceived excess return on capital. Investors price capital using the perceived return

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

while actual resources are governed by R_t^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.

To see what the wedge does, start with its steady-state effect, which can be computed exactly as in Section 3.1. Suppose the wedge were permanent at level x. The stationary Euler condition now 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 (7) 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. Substituting the steady-state return $R^a = -\delta + r^K/\bar{q}$ and wealth $W = \bar{q}K$, the steady-state condition becomes

$$r^K(K) = \bar{q}\left[\delta + \frac{\rho - x}{1 + \theta\bar{q}K}\right]. \tag{19}$$

In the language of Figure 1, the wedge shifts the funding curve down by $\bar{q}x/(1 + \theta\bar{q}K)$. This is a useful way to think about the transition. 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.

The transition object, however, is not a steady state but a price path. 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: the price, at capital K, of the continuation that converges to the permanent-x high steady state. It is the natural generalization of the rational high arm, which is the special case $Q^H(K; 0)$. Two requirements define where this branch is available: the high-capital continuation has an admissible price, and the implied consumption $C(K, Q^H(K; x))$ is positive. The consumption restriction matters because, by (14), a valuation far above $\bar{q}$ commits current output to investment. Collect the pairs that satisfy both requirements in

$$\begin{aligned} \mathcal{D}^H &\equiv \{(K, x) : Q^H(K; x) \text{ exists and } C(K, Q^H(K; x)) > 0\}, \\ x^*(K) &\equiv \inf\{x \geq 0 : (K, x) \in \mathcal{D}^H\}. \end{aligned} \tag{20}$$

The frontier $x^*(K)$ is the minimum perceived-return wedge that keeps the high-capital continuation alive at capital K. When K is close to K^L, fundamentals alone do not support the high-capital path, so the required wedge is positive. As investment raises K, the funding feedback accumulates and the required wedge falls. The capital stocks at which no wedge is needed form the rational high-capital domain,

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

so τ_V is the first time at which the rational high-capital continuation is available at the inherited capital stock. 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.

Proposition 3 (Temporary perceived-return wedge)**.** *Under Lemma 1 and Condition 1, let $K_0 = K^L$. If a feasible wedge path $\{x_t\}$ satisfies $x_t \geq x^*(K_t)$ for all $t < \tau_V$, with $\tau_V < \infty$ 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.*

The proposition formalizes the transition. The wedge raises the perceived return, lifts q, and accelerates accumulation before fundamentals alone justify the high-capital path. The transition succeeds if this induced accumulation reaches $\mathcal{K}_0$. Once it does, the high branch is self-sustaining and the wedge is redundant. The object validated by the transition is the high-capital destination, not the elevated valuation used to get there. The proof is in Appendix H.

The proposition deliberately treats the wedge as an instrument: anything that raises the perceived return temporarily will do. The discipline comes from asking where such a wedge comes from and why it would be temporary. That is the role of beliefs.

4.2 Exuberance as a private wedge, and its fragility

It is useful to describe the sequence of events that the formal apparatus disciplines. Starting from the low-capital steady state, investors come to perceive a persistent excess return on capital. They infer it from a noisy signal and price claims as if the current estimate were permanent. The higher perceived return raises valuation, and the higher valuation raises investment. From that point, two processes operate simultaneously. Along the paths studied here, realized returns do not confirm the perceived excess return, so the estimate is revised downward at a rate governed by learning. At the same time, installed capital shifts income toward capital owners and lowers the return required to sustain a larger capital stock, so the wedge needed to keep the high-capital continuation available falls as capital accumulates. The transition succeeds if the perceived wedge remains above the required wedge until the required wedge reaches zero; it fails if the falling belief crosses the falling frontier while the frontier is still positive. In the formal apparatus, the perceived wedge is the belief x_t, the required wedge is the frontier $x^*(K_t)$ of Section 4.1, and the date at which the required wedge reaches zero is the entry time τ_V into $\mathcal{K}_0$.

Extrapolative beliefs supply the wedge endogenously. Investors interpret recent high returns or favorable news as evidence of a persistent excess return on capital. The posterior is updated as evidence arrives, while at each date prices are computed by treating the current posterior mean as if it were permanent. This anticipated-utility rule is what makes the belief matter for valuation: an investor who expects the excess return to persist capitalizes it into the price of the claim. The current posterior mean x_t thereby becomes exactly the perceived-return wedge of Proposition 3, and the behavioral Euler equation is

$$\frac{\dot{c}_t}{c_t} = R_t^a + x_t - \rho + \theta c_t. \tag{21}$$

When $x = 0$, pricing is rational and the high branch is available exactly on $\mathcal{K}_0$. When $x > 0$, investors can price the high-capital continuation before fundamentals alone support it.

The belief is disciplined by learning, and the discipline can be derived in two lines. Investors observe a return signal $dY_t = \mu\, dt + \sigma\, dB_t$ with a Gaussian prior on the latent excess return μ, and update by Kalman filtering with gain $h_t = P_t/\sigma^2$, where P_t is the posterior variance. As evidence accumulates the posterior variance falls, $\dot{P}_t = -P_t^2/\sigma^2$, so the gain obeys $\dot{h}_t = -h_t^2$: each new observation is less informative relative to what is already known. Integrating gives $h_t = h_0/(1 + h_0 t)$. Along the path on which realized returns carry no excess return—the relevant adverse path, in which the optimism is never confirmed—the filter innovation is $dY_t - x_t\, dt = -x_t\, dt$, so the posterior mean decays at the gain rate, $\dot{x}_t = -h_t x_t$. Dividing by x_t and integrating,

$$\dot{x}_t = -h_t x_t, \qquad \dot{h}_t = -h_t^2, \qquad \text{so} \qquad x_t = \frac{x_0}{1 + h_0 t}, \quad h_t = \frac{h_0}{1 + h_0 t}. \tag{22}$$

The belief is a hyperbola: it never reaches zero in finite time, but it fades steadily as confirming evidence fails to arrive. This decay is what makes the mechanism economically interesting. If beliefs never faded, a large enough wedge would trivially hold the economy on the perceived high branch forever. If beliefs faded immediately, there would be no real buildout. The model studies the intermediate case: optimism is temporary, but it may last long enough to install capital.

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 then follows the investment law (8); and beliefs follow the learning law (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. 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, $x_t < x^*(K_t)$. After τ_V, pricing continues under the same rule (23); what changes at τ_V is availability: $Q^H(K_t; 0)$ now exists, so the high-capital rational continuation is self-sustaining without the wedge.

The race has two state variables and a moving frontier. Capital accumulation moves K_t toward $\mathcal{K}_0$ and, in doing so, lowers the frontier $x^*(K_t)$. Learning moves x_t toward zero. The boom succeeds if the falling belief stays above the falling frontier until capital enters $\mathcal{K}_0$. Fixing x_0, define the critical learning gain

$$h^*(x_0) \equiv \sup\{h_0 > 0 : \tau_V(x_0, h_0) < \tau_*(x_0, h_0)\}. \tag{24}$$

The transition succeeds for $h_0 < h^*(x_0)$ and fails for $h_0 > h^*(x_0)$. The economic message is that belief persistence matters as much as belief size: a modest but stubborn optimism can succeed where a large but quickly corrected one fails.

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 is the formal version of speculative growth. Optimism raises valuation; valuation raises investment; investment raises capital; and capital accumulation reduces the optimism needed to sustain the high-capital path. A temporary mispricing can therefore produce a permanent real legacy. But the proposition is deliberately conditional: the belief must be large enough and persistent enough to survive until accumulation has crossed into $\mathcal{K}_0$. The proof is in Appendix H.

What happens to the elevated valuation itself? Because every steady state has $q = \bar{q}$, the answer is built into the model: the elevated valuation cannot last. Whatever the outcome of the race, the price must eventually return to a rational arm; the only question is which arm is available when it does. To measure the distance to the high arm, define the rational-arm gap

$$v_t^H \equiv Q^H(K_t; x_t) - Q^H(K_t; 0) \tag{25}$$

whenever $K_t \in \mathcal{K}_0$; before the rational high arm is available, the relevant comparison is to the low rational continuation selected from the inherited capital stock. In a successful transition, the correction occurs after the economy has installed enough capital for the high branch to be rationally available: the residual belief fades, v_t^H closes, and the capital remains. In a failed transition, learning removes the perceived high branch while the capital stock is still outside $\mathcal{K}_0$; the price cannot land on the high arm, so it falls to the low arm and the boom collapses. In both cases, valuation returns to a rational arm—to the saddle geometry of Figure 2.

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 corollary separates the two endings by the form of the same event. A price correction is unavoidable because no steady state has $q \neq \bar{q}$; success and failure differ in the rational arm on which the correction lands, and in how it lands. Successful speculation corrects smoothly: past τ_V the rational branch is available, and the residual belief merely fades along it. Failed speculation crashes: the branch the price was riding ceases to exist while capital cannot move, so the price must jump. This is why the paper separates the reality of the technology from the sustainability of peak valuations. Peak prices correct in both cases; what differs is whether enough capital has been installed for the correction to leave a high-capital legacy. The proof is in Appendix H.

The transition also has a sharp interest-rate signature, which connects the model to the current episode.

Remark 2 (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 buildout. 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 boom can therefore have a high-rate transition and a low-rate destination. 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.*

Figures 3 and 4 show the two horizons of a successful transition. The early window isolates the buildout: optimism lifts q on impact, capital accumulation accelerates, the funding rate jumps while realized returns stay low, and capitalist consumption is compressed to finance the investment surge. The late window shows the approach to the high-capital destination after valuation has rejoined the high rational arm: q_t settles near $\bar{q}$ while capital completes its climb, the wage sits above its initial level, and the worker share sits below it. The next subsection explains those last two facts.

Figure 3 reads the early window through the equations above; the vertical lines mark the launch, the start of AI deployment, and entry into $\mathcal{K}_0$. The capital panel reflects (8): with $q_t > \bar{q}$, capital rises from K^L. The valuation panel shows that q_t is transitional: it jumps above $\bar{q}$ and later falls because no steady state has elevated q. The rate panel uses Remark 2: the funding rate is $i_t = R_t^a + x_t$, so the gap between the funding rate and the realized return is exactly the perceived excess return. The wage and share panels use the technology of Section 2 and Lemma 2: wages do not fall, while the worker share declines once AI labor expands N. The consumption panel follows (14): the jump in q commits output to investment, so capitalist consumption drops on impact. The belief panel follows (22). After capital enters $\mathcal{K}_0$, the figures allow the residual belief to decay faster than (22) alone would imply; this display convention shortens the plotting window and does not affect the path before entry.

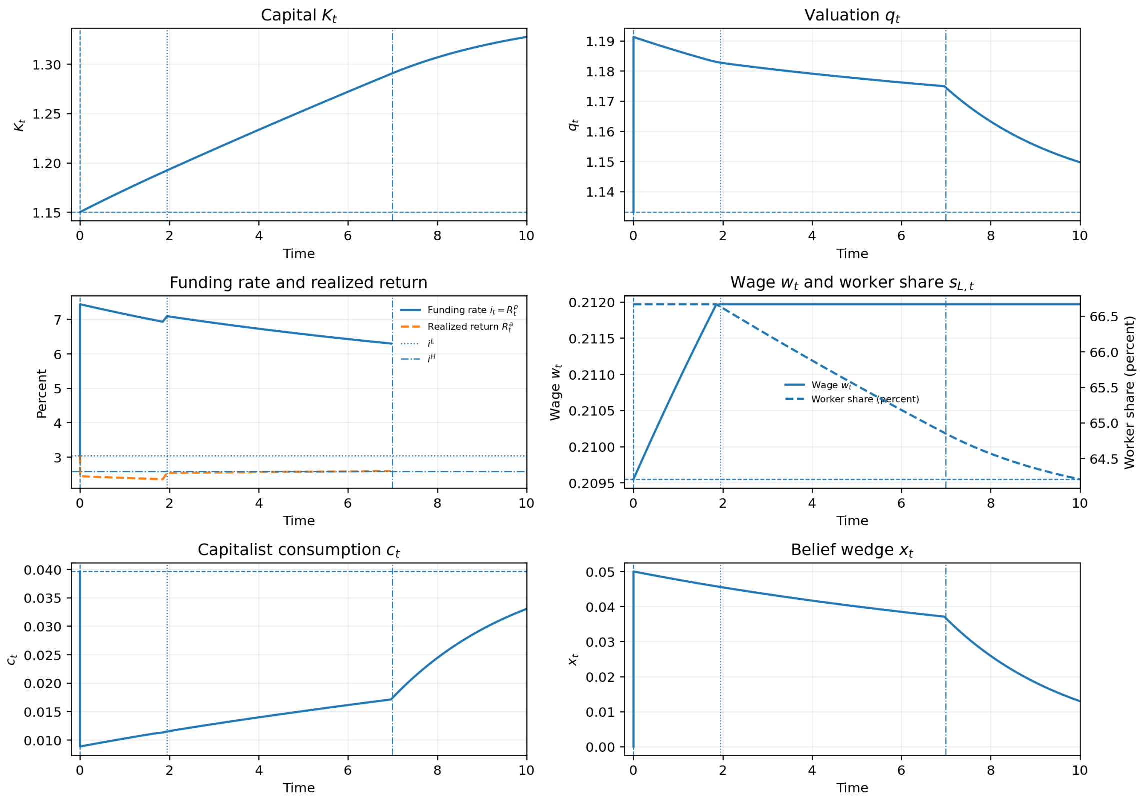

Figure 3: Successful transition: early dynamics. The window focuses on the first years of the transition, when most of the qualitative action occurs. Optimism raises q on impact and triggers rapid capital accumulation. The rate panel reports the Euler-implied funding rate $i_t = R_t^p$ and the realized 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.

4.3 Distributional incidence

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

The destination is straightforward on the capitalist side. Steady-state wealth is $W = \bar{q}K$, so substituting into the consumption rule (7) 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-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

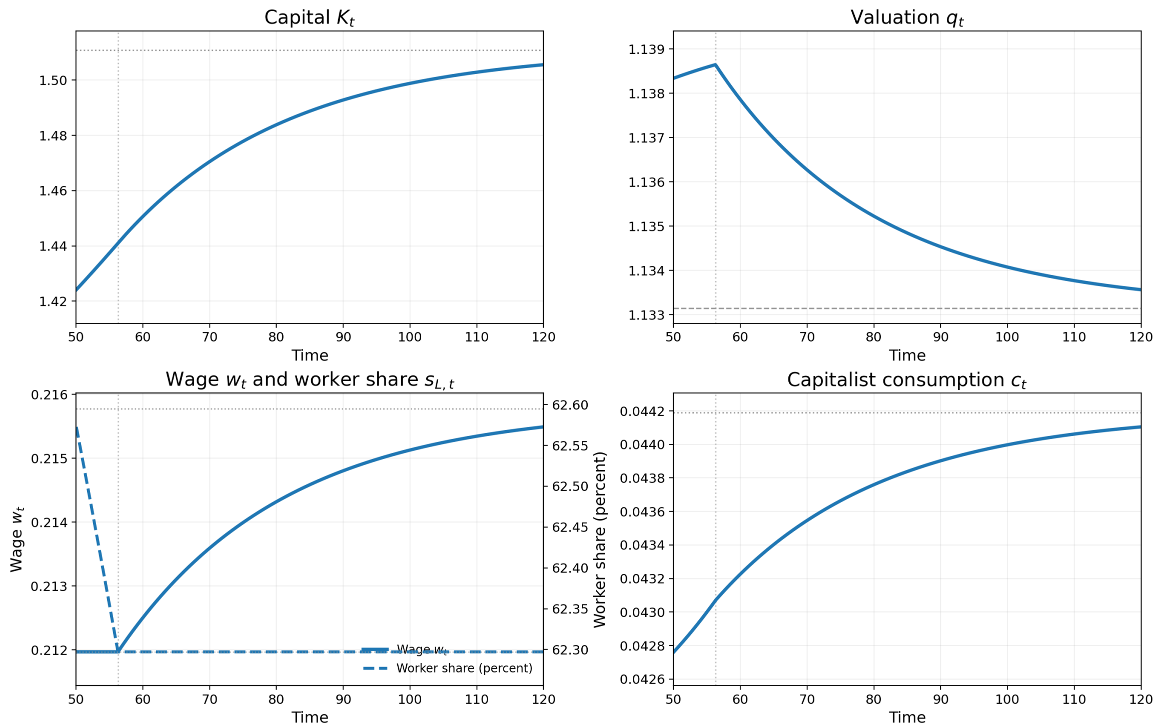

Figure 4: Successful transition: late dynamics. The terminal window illustrates the qualitative approach to the rational high-capital continuation after valuation has rejoined the high arm. Capital and consumption move toward their high-state levels, q_t remains close to $\bar{q}$ while capital keeps moving, and the wage is above its initial level while the worker share is below it.

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 (2) 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,

$$w(K^H) > w(K^L) \iff \frac{Y(K^H)}{Y(K^L)} > \frac{N(K^H)}{N(K^L)}, \tag{26}$$

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 boom, $q_t > \bar{q}$ 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 pre-boom 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, the worker side of the episode is protected on the downside: a failed boom leaves workers where they started, and a successful one leaves them better paid. The statement concerns the wage flow workers consume in the model. 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 buildout. During the boom, 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 2, a return that is priced but never paid. Second, the investment surge itself is financed out of their consumption: by (14), a high q commits output to investment, so the boom 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,

$$\Delta V_C^B = \int_0^\infty e^{-\rho t} \left[\log c_t^B + \theta W_t^B - (\log c_C^L + \theta W^L) \right] dt, \tag{27}$$

$$\Delta V_W^B = \int_0^\infty e^{-\rho t} \left[\log w_t^B - \log w^L \right] dt. \tag{28}$$

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 buildout 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 aggressive 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. This is a sign statement about aggressive overvaluation, not a calibrated magnitude.

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 learning dynamics that make optimism temporary are therefore also what make it a blunt instrument: exuberance can finance the crossing, but at a cost a measured wedge would avoid. The numerical parameterization used for the figures is reported in Appendix A.

5 Conclusion

A speculative boom 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 buildout 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 boom 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 buildout correct. The model therefore separates three objects that are often conflated in discussions of technology booms: the reality of the technology, the sustainability of peak valuations, and the permanence of the capital installed during the boom.

The model is deliberately spare: one sector, a representative capitalist, and a single belief process. These choices isolate the mechanism; broader ownership, sectoral heterogeneity, and heterogeneous beliefs would refine incidence without changing the separation between transition, correction, and installed capital.

The framework organizes the observable features of the current episode without requiring a quantitative stand. In the model, concentrated market capitalization, a large investment program, and elevated long-term rates can arise together from a belief-supported valuation financing a buildout. The model adds two qualitative implications: the worker share declines as deployment proceeds while wages do not fall, and the low interest rate associated with abundant capital belongs to the destination, not to the boom. A later correction in valuations is consistent with the mechanism rather than evidence against it; what determines the legacy is the capital installed by the time the correction arrives.

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 associated paths are generated by the system in the text using the branch and learning equations described in the text.

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 (10).

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 (2). 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 (1), 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 (4) 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 (5). 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 (7).

E Equilibrium Dynamics

The consumption flow implied by the budget and $W = qK$ is $C(K,q) = K[r^K(K) - q\psi \log q]$, 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 (16). This equation and (8) 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 $-\delta + r^K/q - \rho + \theta C$ by $-\delta + r^K/q + x - \rho + \theta C$. 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 (11) is continuous, with r^K continuous across regions. As $K \to 0^+$, $r^K \to \infty$ while the funding term is bounded, so $G > 0$; just past K_{AI}, $r^K = r^K_{\text{flat}}$ and the left inequality of (12) gives $G(K_{\text{AI}}) < 0$, so a root exists in $(0, K_{\text{AI}})$. On the flat region r^K 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 $r^K \to 0$ 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,

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

Condition (13) 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

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

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

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

Direct computation gives $C_q = -K\psi(\log q + 1)$, so at $\bar{q}$,

$$D = -K\Big[(\psi + \delta) + \tfrac{r^K(K) - \bar{q}\delta}{\bar{q}}\Big] < 0,$$

using $r^K = \bar{q}[\delta + \rho/(1 + \theta\bar{q}K)] > \bar{q}\delta$ at every root. Hence $\operatorname{sign}\det DF^0 = \operatorname{sign} Z_K$, which is (17).

To evaluate Z_K, let $B = -\delta + r^K/q - \rho + \theta C$. 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 $B_K = r^{K\prime}/\bar{q} + \theta C_K$ and $C_K = (r^K - \bar{q}\delta) + Kr^{K\prime}$,

$$Z_K = C\left[r^{K\prime}\left(\tfrac{1}{\bar{q}} + \theta K\right) + \theta(r^K - \bar{q}\delta)\right],$$

which is (17). Using the steady-state condition, $r^K - \bar{q}\delta = \bar{q}\rho/(1 + \theta\bar{q}K)$, 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, $r^{K\prime} = 0$ and $r^K - \bar{q}\delta > 0$, so $Z_K > 0$ and $\det DF^0 > 0$; the trace $g_q = Z_q/D$ is positive there (in the flat region $Z_q = -KS[r^K/\bar{q}^2 + \theta K(\psi + \delta) + \psi/\bar{q}] < 0$ with $S = r^K - \bar{q}\delta > 0$, and $D < 0$), so K^M is a source. At K^L and K^H, $r^{K\prime} < 0$, so $\det DF^0 < 0$ (a saddle) iff (18) 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. With K predetermined and q free, an admissible rational equilibrium from $K_0 = K^L$ must lie on a convergent (stable) branch. 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 $C > 0$. The lower-bound condition $x_t \geq x^*(K_t)$ ensures that the path remains on the high-capital perceived branch until τ_V. On the induced path, 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 $x_t \geq x^*(K_t)$. For $t \geq \tau_V$, 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.

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.

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. 2021. "The Saving Glut of the Rich." *Quarterly Journal of Economics* 136(4): 2243–2307.

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.