

Minimum Viable Scale

Extinction and Escape under Increasing Returns

Javier Milei

Demian Reidel

July 2026

Working Paper

Abstract

Aggregate increasing returns with bounded labor make absolute scale economically meaningful. This paper separates three notions of minimum viable scale: a physical floor, the least labor capacity at which capital can be maintained against depreciation; a marginal-return floor, the least capacity at which the return required by discounting and depreciation is attainable, necessary but not sufficient for an interior stationary allocation; and a global, value-based notion — above what capital does the planner select a high-scale path rather than middle behavior or exit? For the certified CES gross-complements family the static theory is exact: the sustainable set is an interval whose endpoints are geometric-mean conjugates about labor capacity, $\kappa \bar{K}_{\text{sus}} = \bar{L}^2$, and a productivity wedge moves both endpoints with elasticity $\sqrt{(m+1)/(m-1)}$, $m = 2A\bar{L}/\delta - 1$, which diverges at the physical floor. Wedges can collapse the economy: any productivity wedge condemns a band of previously viable capital stocks to certain extinction, the band widens with that elasticity, and beyond an explicit critical wedge, arbitrarily small near constant returns, no initial condition survives. Global selection is proved by a Bellman-barrier method: a first-passage argument reduces unrestricted selection to four finite families of inequalities, verified in exact rational arithmetic with directed rounding. The certificates prove high selection on $[\text{.098935}, 50]$ at demonstration primitives and on $[\text{.016}, 450]$ — within a factor eight of the extinction threshold — at annual discounting and depreciation, for every productivity in a certified interval. The certified economies are stylized proof-of-concept environments rather than calibrations; the archive records the finite objects from which every certified inequality is deterministically regenerated and checked.

Keywords. Minimum viable scale; increasing returns; comparative statics; Bellman barriers; high selection; finite certificates; computer-assisted proofs; closed-form scale boundaries.

JEL codes. C61, C62, C63, D24, E13, O41.

1 Introduction

Constant-returns growth models remove absolute scale. If labor, capital, and production can be scaled together, a small economy is a scaled-down large economy. Aggregate increasing returns with bounded labor destroy that invariance. Some capital stocks cannot maintain themselves. Some labor capacities can maintain capital physically but cannot generate the marginal return required by the Euler equation. Some economies can reach high scale, but only if residual middle paths and downward exits are globally dominated.

This is a certified-theory paper, not a calibrated macro paper. It distinguishes three notions of minimum viable scale. The first is physical: the least labor capacity at which some capital stock

can be maintained against depreciation. The second is marginal: the least capacity at which the return required by the Euler equation is attainable. The third is global: the set of initial capital stocks from which a high-scale path is value selected over middle stagnation or downward exit. The first two objects are static and admit closed-form characterizations in a CES gross-complements family; the third is a nonconcave dynamic selection problem, which we solve by constructing Bellman barriers that reduce unrestricted path comparisons to finite, verifier-checkable inequalities. The certified economies are proof-of-concept environments, showing that global selection can be proved without local Euler conditions or value-function approximation. We solve a planner problem because the object of interest is global selection under aggregate increasing returns; decentralization with external increasing returns is left for future work (Section 9).

This paper studies these distinctions in a deterministic one-sector planner problem. The starting point is a minimum-viable-scale diagnosis. Under homogeneous increasing returns,

$$F(K, L) = L^\theta y(K/L), \quad \theta > 1,$$

two primitive thresholds matter. The physical floor is

$$L_{\min}^{\text{avg}}(\delta) = \left(\frac{\delta}{\max_k y(k)/k} \right)^{1/(\theta-1)}, \quad (1.1)$$

the least labor capacity at which capital can be maintained against depreciation. The Euler floor is

$$L_{\min}^{\text{mg}}(D) = \left(\frac{D}{\max_k y'(k)} \right)^{1/(\theta-1)}, \quad D = \beta^{-1} - (1 - \delta), \quad (1.2)$$

the least labor capacity at which the marginal return required by discounting and depreciation is attainable — necessary for an interior stationary allocation, though not sufficient. These floors are distinct objects: the average-product peak governs maintenance, the marginal-product peak governs Euler feasibility, and the two peaks occur at different capital-labor ratios. Neither floor determines global selection.

The floors give a transparent theory of wedge amplification. If a policy wedge lowers usable productivity by a factor $1 - \tau_A$, both floors scale as $(1 - \tau_A)^{-1/(\theta-1)}$. Near constant returns the exponent is large: with $\theta = 1.05$, a one-percent productivity wedge raises both floors by 22.3 percent, and a five-percent wedge by 179 percent. With $\theta = 2$, the certified economy below, the elasticity is exactly one. We are explicit about this: the amplification channel is a near-constant-returns phenomenon, while the certified economy illustrates a second, distinct channel — proximity to a global selection threshold — which operates at any degree of increasing returns. Section 4 develops both channels and states precisely what a wedge-induced departure from the certified region does and does not imply. It also contains the model's collapse theorem (Proposition 4.1): every productivity wedge condemns a band of previously viable capital stocks to extinction, and a wedge beyond an explicit critical size, which vanishes near constant returns, collapses the economy from every initial condition.

For the certified family the static scale theory turns out to be exact (Section 3). The sustainable set is $[\kappa, \bar{K}_{\text{sus}}]$ with

$$\kappa = \bar{L}(m - \sqrt{m^2 - 1}), \quad \bar{K}_{\text{sus}} = \bar{L}(m + \sqrt{m^2 - 1}), \quad m = \frac{2A\bar{L}}{\delta} - 1;$$

nonempty exactly when labor capacity clears the physical floor. Its boundaries obey $\kappa \bar{K}_{\text{sus}} = \bar{L}^2$: extinction threshold and sustainability ceiling are reflections of one another about labor capacity.

A productivity wedge moves both with elasticity $E = \sqrt{(m+1)/(m-1)}$, and E diverges as capacity falls toward the floor. The value of the extinction region sits between explicit logarithmic bounds. Stationary policies in closed form come within 0.024 of the certified high values, and the steady state and the golden-rule point are roots of a one-parameter family of quartics in K/L (Proposition 3.3). All of this falls out of the quadratic the certificate evaluates. What resists closed form is the pair of families at the heart of the selection proof, the middle barrier and the exit ledgers; Section 9 quantifies why. The certificate itself is built twice: at the demonstration primitives $\beta = \delta = .7$, and again at annual discounting and depreciation, $\beta = .96$ and $\delta = .08$ (Section 7.4).

The main contribution is global and methodological. We prove high selection using Bellman barriers. A Bellman barrier is a bounded function B satisfying

$$B(K) \geq u(K, L, K') + \beta B(K')$$

on a restricted transition class; telescoping then bounds the value of every path in the class. A first-passage *splice* argument (Theorem 5.2) reduces the unrestricted selection problem to four ingredients: (i) a barrier B_M dominating paths that stay forever in the middle region; (ii) an upper-continuation barrier $\bar{U}$ dominating the unrestricted value after any downward exit; (iii) explicit high-admissible policies whose values exceed both B_M and every exit payoff; and (iv) forward-closure bounds ruling out upward escapes. Every path is then either dominated where it stays, or spliced at its first downward exit onto an explicit high policy that does strictly better.

The second methodological contribution is verification. Each of the four ingredients is a finite family of inequalities between explicitly recorded numbers, reduced from the continuum by an endpoint-ledger theorem for piecewise log-affine barriers (Theorem 6.1) and a monotone corner bound for the labor-reduced reward (Lemma 6.1). Every ledger entry is evaluated in exact rational arithmetic; logarithms are bounded by directed rounding with an explicit one-unit-in-the-last-place guard; reward suprema over labor are bounded by a monotone subdivision branch-and-bound that requires no root isolation. The result is a certificate in the sense of the computer-assisted-proof literature: a finite object whose validity can be re-checked line by line, independently of how it was found. Construction of the certificate (value iteration, policy search) is untrusted and plays no role in the proof.

The certified theorem is for the CES gross-complements family (Arrow et al., 1961)

$$F_A(K, L) = 4A \frac{K^2 L^2}{(K + L)^2}, \quad (1.3)$$

with fixed primitives

$$\beta = .7, \quad \delta = .7, \quad \eta = .7, \quad \bar{L} = .88.$$

For every productivity $A \in \mathcal{A} = [9.916, 10.04]$ we prove

$$V_A^*(K) = V_A^H(K) \quad \forall K \in [.098935, 50], \quad (1.4)$$

where V_A^* is the unrestricted value and V_A^H is the value over paths that reach and remain in $H = [2, 50]$, and we prove that every optimal path from $[.098935, 50]$ is of that form. The same primitives imply extinction below $\kappa = 0.018156974557\dots$, the smaller root of an explicit quadratic, and an upper sustainability bound at the larger root: capital above 42.66 cannot be maintained even at full labor capacity, so the upper part of H is transient. The model therefore delivers a scale taxonomy with one unclassified band $(\kappa, .1)$, which we take up in Section 9.

This paper contributes to three literatures. The nonconcave growth literature studies global dynamics and threshold states in models where local Euler equations do not characterize global optima (Skiba, 1978; Majumdar and Mitra, 1982; Dechert and Nishimura, 1983; Askenazy and Le Van, 1999; Kamihigashi and Roy, 2007). The increasing-returns, big-push, and poverty-trap literatures emphasize scale effects and threshold phenomena (Rosenstein-Rodan, 1943; Romer, 1986; Murphy et al., 1989; Azariadis and Drazen, 1990; Matsuyama, 1991; Krugman, 1991; Galor and Zeira, 1993; Benhabib and Farmer, 1994; Azariadis and Stachurski, 2005). Empirical evidence on poverty-trap thresholds is contested (Kraay and McKenzie, 2014; Balboni et al., 2022); the thresholds here are certified properties of a model rather than estimates. The monotone comparative-statics literature supplies lattice tools for selection (Milgrom and Shannon, 1994; Edlin and Shannon, 1998; Topkis, 1998; Quah, 2007); Appendix B records where increasing differences fail in this class, and the limits of that failure. Methodologically, the paper connects two strands rarely used in economic theory: barrier certificates for safety verification in control (Prajna et al., 2007; Ames et al., 2019; Lasserre et al., 2008) and rigorous computer-assisted proof (Tucker, 2002; Hales, 2005; Gonthier, 2008). Relative to numerical dynamic programming with a priori error bounds (Santos and Vigo-Aguiar, 1998; Santos, 2000) and to backward-error analyses of computed equilibria (Kubler and Schmedders, 2005), the trusted object here is not an approximation with an error bound but a finite list of certified inequalities; the construction stage does use value iteration, and nothing from that stage is trusted.

2 Model and minimum viable scale

The planner chooses $(C_t, L_t, K_{t+1})_{t \geq 0}$ to maximize

$$\sum_{t=0}^{\infty} \beta^t [\eta \log C_t + (1 - \eta) \log(1 - L_t)] \quad (2.1)$$

subject to

$$C_t + K_{t+1} = F(K_t, L_t) + qK_t, \quad q = 1 - \delta, \quad (2.2)$$

$$K_{t+1} > 0, \quad 0 \leq L_t \leq \bar{L} < 1, \quad C_t > 0. \quad (2.3)$$

Write $u(K, L, K') = \eta \log(F(K, L) + qK - K') + (1 - \eta) \log(1 - L)$ for the period payoff and

$$\Gamma(K) = \{(L, K') : 0 \leq L \leq \bar{L}, 0 < K' < F(K, L) + qK\}$$

for the transition correspondence. A *path* from K_0 is a sequence with $(L_t, K_{t+1}) \in \Gamma(K_t)$ for all t ; its value is $J(\pi) = \sum_t \beta^t u(K_t, L_t, K_{t+1}) \in [-\infty, \infty)$, well defined because period payoffs are bounded above on bounded state sets. The unrestricted value is $V^*(K) = \sup_{\pi} J(\pi)$. The planner formulation isolates the scale-selection problem from the decentralization issues raised by aggregate increasing returns.

For the general scale results, assume

$$F(K, L) = L^{\theta} y(K/L), \quad \theta > 1. \quad (2.4)$$

Assumption 2.1. $y : (0, \infty) \rightarrow (0, \infty)$ is twice continuously differentiable with

$$\lim_{k \rightarrow 0^+} \frac{y(k)}{k} = \lim_{k \rightarrow \infty} \frac{y(k)}{k} = 0, \quad \lim_{k \rightarrow 0^+} y'(k) = \lim_{k \rightarrow \infty} y'(k) = 0.$$

We extend F continuously to $L = 0$ by $F(K, 0) = 0$. Monotonicity of F in (K, L) is used only for the certified family, where it holds by inspection.

Under [Assumption 2.1](#) the average product $y(k)/k$ and the marginal product $y'(k)$ attain finite interior maxima

$$\bar{a} = \max_{k>0} \frac{y(k)}{k}, \quad \bar{\varphi} = \max_{k>0} y'(k).$$

The vanishing of y' at the origin is a *reverse Inada* condition: the marginal product of capital collapses at low capital-labor ratios. It is the gross-complements signature that makes low-scale traps possible, and it holds for the certified family (1.3).

In a constant-returns world, capital scarcity is self-correcting: the fewer the machines, the higher their marginal product, and accumulation follows. Under gross complementarity the correction fails at low scale, because capital and labor are productive only together. Machines without operators produce little, and operators working with a near-empty capital stock add little more, so at low K the marginal product of capital is low precisely where capital is scarce. Scarcity reinforces itself instead of curing itself. The rest of the paper flows from that reversal.

Proposition 2.1 (Physical scale floor). *Let*

$$L_{\min}^{\text{avg}}(\delta) = \left(\frac{\delta}{\bar{a}} \right)^{1/(\theta-1)}.$$

If $\bar{L} < L_{\min}^{\text{avg}}(\delta)$, then every feasible capital path satisfies $K_t \rightarrow 0$. If $\bar{L} > L_{\min}^{\text{avg}}(\delta)$, there exists (K, L) with $L < \bar{L}$ such that $F(K, L) > \delta K$.

Proof. For any feasible K, L with $L \leq \bar{L}$,

$$\frac{F(K, L)}{K} = L^{\theta-1} \frac{y(K/L)}{K/L} \leq \bar{L}^{\theta-1} \bar{a}.$$

If $\bar{L} < L_{\min}^{\text{avg}}(\delta)$, then $F(K, L) \leq \mu K$ with $\mu = \bar{L}^{\theta-1} \bar{a} < \delta$. Since $C_t > 0$,

$$K_{t+1} < F(K_t, L_t) + qK_t \leq (\mu + q)K_t,$$

and $\mu + q < \delta + q = 1$, so $K_t \rightarrow 0$ geometrically. If $\bar{L} > L_{\min}^{\text{avg}}(\delta)$, choose k near the maximizer of $y(k)/k$ and $L < \bar{L}$ with $L^{\theta-1}y(k)/k > \delta$; then $K = kL$ satisfies $F(K, L) > \delta K$. $\square$

Proposition 2.2 (Marginal-return floor). *Let $D = \beta^{-1} - q$ and*

$$L_{\min}^{\text{mg}}(D) = \left(\frac{D}{\bar{\varphi}} \right)^{1/(\theta-1)}.$$

There exists a pair (K, L) with $L < \bar{L}$ and $F_K(K, L) = D$ if and only if $\bar{L} > L_{\min}^{\text{mg}}(D)$, and a pair with $L \leq \bar{L}$ if and only if $\bar{L} \geq L_{\min}^{\text{mg}}(D)$.

Proof. Since $F_K(K, L) = L^{\theta-1}y'(K/L)$, the largest marginal product attainable with labor at most $\bar{L}$ is $\bar{L}^{\theta-1}\bar{\varphi}$. It exceeds D strictly if and only if $\bar{L} > L_{\min}^{\text{mg}}(D)$, and weakly if and only if $\bar{L} \geq L_{\min}^{\text{mg}}(D)$. Under [Assumption 2.1](#), y' is continuous and vanishes at both ends, so every value in $(0, \bar{\varphi}]$ is attained; the result follows. $\square$

Remark 2.1 (Necessity, not sufficiency). The marginal-return floor is a necessary condition for an interior stationary allocation. Positive stationary consumption ($C = F - \delta K > 0$), the labor first-order condition ($\eta(1-L)F_L = (1-\eta)C$), and stationarity itself are additional requirements; for the certified family they are characterized jointly by the quartic of [Proposition 3.3](#).

Remark 2.2 (The two peaks differ). Let $\alpha(k) = k y'(k)/y(k)$ denote the capital elasticity of y . An interior maximizer of the average product solves $\frac{d}{dk} [y(k)/k] = 0$, that is, $\alpha(k) = 1$; an interior maximizer of the marginal product solves $y''(k) = 0$. These generally differ. For the certified family (1.3), $y(k) = 4Ak^2/(1+k)^2$, so $y(k)/k$ peaks at $k = 1$ with $\bar{a} = A$, while $y'(k) = 8Ak/(1+k)^3$ peaks at $k = 1/2$ with $\bar{\varphi} = 32A/27$. The ratio best suited to maintaining capital is not the ratio best suited to delivering the Euler return.

Example 2.1 (Floors of the certified economy). For (1.3), $\theta = 2$, so the floors are $L_{\min}^{\text{avg}} = \delta/A$ and $L_{\min}^{\text{mg}} = D/(32A/27) = 27D/(32A)$ with $D = \beta^{-1} - q = 79/70$. At $A = 10$: $L_{\min}^{\text{avg}} = 0.0700$ and $L_{\min}^{\text{mg}} = 2133/22400 \approx 0.0952$, both far below $\bar{L} = .88$. The certified economy sits well above both static floors; the content of the selection theorem in Section 7 is precisely that static viability does not settle which basin the economy selects.

3 The sustainable set in closed form

Everything static about the certified family (1.3) can be computed exactly, from the scale boundaries and their comparative statics to the value of the extinction region. Throughout this section $F_A(K, L) = 4AK^2L^2/(K+L)^2$ with $A > 0$, $\delta \in (0, 1)$, $q = 1 - \delta$, and $\bar{L} \in (0, 1)$; write

$$\phi_A(K) = qK + F_A(K, \bar{L}), \quad \lambda_A(K) = \frac{\phi_A(K)}{K} = q + \frac{4A\bar{L}^2K}{(K + \bar{L})^2}$$

for full-capacity resources and the maximal gross growth factor, and call K *maintainable* if $\phi_A(K) \geq K$.

Theorem 3.1 (Closed-form sustainable set). *Let $m = 2A\bar{L}/\delta - 1$. The maintainable set is nonempty if and only if $m \geq 1$, equivalently $\bar{L} \geq \delta/A = L_{\min}^{\text{avg}}(\delta)$, the physical floor of Proposition 2.1 at $\theta = 2$. For $m > 1$ it is the interval $[\kappa, \bar{K}_{\text{sus}}]$ with*

$$\kappa = \bar{L}(m - \sqrt{m^2 - 1}), \quad \bar{K}_{\text{sus}} = \bar{L}(m + \sqrt{m^2 - 1}), \quad (3.1)$$

and consequently

$$\kappa \bar{K}_{\text{sus}} = \bar{L}^2, \quad \log \frac{\bar{K}_{\text{sus}}}{\kappa} = 2 \operatorname{arccosh} m. \quad (3.2)$$

Proof. $\phi_A(K) \geq K$ if and only if $4A\bar{L}^2K \geq \delta(K + \bar{L})^2$. In $x = K/\bar{L}$ this reads $x^2 - 2mx + 1 \leq 0$ with m as displayed, whose solution set is the interval between the conjugate roots $x_{\pm} = m \mp \sqrt{m^2 - 1}$, nonempty if and only if $m \geq 1$; the product of the roots is 1, which gives (3.2). At $\theta = 2$ the average-product peak is $\bar{a} = A$ (Remark 2.2), so $m \geq 1$ if and only if $A\bar{L} \geq \delta$, i.e. $\bar{L} \geq L_{\min}^{\text{avg}}(\delta)$. Writing $m = \cosh t$ gives $x_{\pm} = e^{\pm t}$, hence the log-length $2t$. $\square$

Remark 3.1 (Duality of the scale boundaries). $\sqrt{\kappa \bar{K}_{\text{sus}}} = \bar{L}$ exactly: the extinction threshold and the sustainability ceiling are geometric-mean conjugates about labor capacity, and the sustainable set is log-symmetric around $K = \bar{L}$. At the certified primitives this is visible in (7.2), whose constant-to-leading-coefficient ratio is $0.54208/0.7 = .7744 = \bar{L}^2$; numerically, $0.018156974557 \dots \times 42.650277311 \dots = .7744$. At $m = 1$ the set degenerates to the single point $K = \bar{L}$; the existence of any sustainable scale is the physical-floor condition itself.

Proposition 3.1 (Exact wedge elasticities of the scale boundaries). *For $m > 1$ let $E = \sqrt{\frac{m+1}{m-1}} > 1$. Then*

$$\begin{aligned} \frac{\partial \log \kappa}{\partial \log A} &= -E, & \frac{\partial \log \bar{K}_{\text{sus}}}{\partial \log A} &= +E, & \frac{\partial \log \kappa}{\partial \log \delta} &= +E, & \frac{\partial \log \bar{K}_{\text{sus}}}{\partial \log \delta} &= -E, \\ \frac{\partial \log \kappa}{\partial \log \bar{L}} &= 1 - E < 0, & \frac{\partial \log \bar{K}_{\text{sus}}}{\partial \log \bar{L}} &= 1 + E. \end{aligned}$$

Proof. From (3.1), $d \log x_{\mp} / dm = \mp 1 / \sqrt{m^2 - 1}$. Since $m + 1 = 2A\bar{L} / \delta$,

$$\frac{\partial m}{\partial \log A} = \frac{\partial m}{\partial \log \bar{L}} = -\frac{\partial m}{\partial \log \delta} = m + 1,$$

so $\partial \log x_{\mp} / \partial \log A = \mp(m + 1) / \sqrt{m^2 - 1} = \mp E$; the $\bar{L}$ entries add $+1$ from the prefactor $\bar{L}$ in (3.1). $\square$

Remark 3.2 (Two channels, one formula). E is the global-boundary counterpart, at $\theta = 2$, of the floor-amplification exponent $1/(\theta - 1)$ of Section 4: both scale boundaries move against the economy under a productivity wedge with common magnitude E , and $E \sim \sqrt{2/(m - 1)} \rightarrow \infty$ as capacity approaches the physical floor, where the sustainable set collapses to the point $\bar{L}$. Far from the floor, E is close to one. The certified economy is far from the floor: at A_+ , $m = 24.24342857 \dots$ and $E = 1.042135 \dots$, and the two-point movements of the boundaries across $[A_-, A_+]$ have log elasticities ∓ 1.042415 , matching E at the geometric midpoint of $[A_-, A_+]$ to six decimals. Equivalently $E = \coth(\frac{1}{4} \log(\bar{K}_{\text{sus}}/\kappa))$: the wedge elasticity of both boundaries is a sufficient statistic of the observed log-width of the sustainable set. The comparative statics of the sustainable set reduce to the single parameter m .

Proposition 3.2 (The value of the extinction region). *Let $m > 1$. For every $K_0 \in (0, \kappa)$, $\lambda_A(K_0) < 1$ and*

$$\frac{\eta \log K_0}{1 - \beta} + \frac{\eta \log((1 - \beta)q)}{1 - \beta} + \frac{\eta \beta \log(\beta q)}{(1 - \beta)^2} \leq V_A^*(K_0) \leq \frac{\eta \log K_0}{1 - \beta} + \frac{\eta \log \lambda_A(K_0)}{(1 - \beta)^2}. \quad (3.3)$$

The leading coefficient $\eta/(1 - \beta)$ is exact on both sides. At the certified primitives the bracket width is below 14.2 on all of $(0, \kappa)$ and below 4.76 in the limit $K_0 \downarrow 0$.

Proof. Upper bound. On $(0, \bar{L}]$, λ_A is increasing, since $\frac{d}{dK} [K/(K + \bar{L})^2] = (\bar{L} - K)/(K + \bar{L})^3 > 0$, and $\lambda_A < 1$ on $(0, \kappa)$ by Theorem 3.1; note $\kappa < \bar{L}$ because $x_- < 1$. Along any feasible path from $K_0 \in (0, \kappa)$, $K_{t+1} < \phi_A(K_t) = \lambda_A(K_t)K_t$, so K_t is strictly decreasing, $\lambda_A(K_t) \leq \lambda_A(K_0)$, and by induction $K_t \leq \lambda_A(K_0)^t K_0$; moreover $C_t < \phi_A(K_t) \leq \lambda_A(K_0)^{t+1} K_0$. Since leisure utility is nonpositive and $\sum_{t \geq 0} \beta^t(t + 1) = (1 - \beta)^{-2}$,

$$J \leq \eta \sum_{t \geq 0} \beta^t [\log K_0 + (t + 1) \log \lambda_A(K_0)] = \frac{\eta \log K_0}{1 - \beta} + \frac{\eta \log \lambda_A(K_0)}{(1 - \beta)^2}.$$

Lower bound. The eat-capital policy $L_t = 0$, $K_{t+1} = \beta q K_t$ is feasible from any $K_0 > 0$, has zero leisure cost, and consumption $C_t = (1 - \beta)q(\beta q)^t K_0$; evaluating $\eta \sum_t \beta^t \log C_t$ gives the left side. (βq is the optimal linear eat-capital rate.) $\square$

Corollary 3.1 (All crash thresholds in closed form). *For $\lambda \in (q, 1)$ with $m_\lambda = 2A\bar{L}/(\lambda - q) - 1 > 1$,*

$$\phi_A(K) \leq \lambda K \quad \text{for all } K \in (0, \kappa_\lambda], \quad \kappa_\lambda = \bar{L}(m_\lambda - \sqrt{m_\lambda^2 - 1}),$$

and the threshold is exact: $\phi_A(K) > \lambda K$ immediately above κ_λ . Inverting $K = \kappa_\lambda$ gives $\lambda = \lambda_A(K)$, so (3.3) is the envelope of the sharpened fixed- λ crash bounds (the remark following Lemma 7.3).

Proof. Replace δ by $\lambda - q$ in the proof of Theorem 3.1: $\phi_A(K) \leq \lambda K$ if and only if $4A\bar{L}^2K \leq (\lambda - q)(K + \bar{L})^2$, i.e. $x^2 - 2m_\lambda x + 1 \geq 0$ with $x = K/\bar{L}$, which holds on $(0, \kappa_\lambda]$, the region below the smaller root. Solving $\kappa_\lambda = K$ for λ returns $\lambda = q + 4A\bar{L}^2K/(K + \bar{L})^2 = \lambda_A(K)$. $\square$

3.1 Stationary points in radicals

Proposition 3.3 (One quartic for all interior stationary points). *Fix a required gross-of-depreciation return $R > 0$ and consider interior pairs (K, L) , $0 < L < 1$, satisfying*

$$F_{A,K}(K, L) = R, \quad \eta(1 - L)F_{A,L}(K, L) = (1 - \eta)(F_A(K, L) - \delta K).$$

Write $s = K/L$. Then

$$L = \frac{R(1 + s)^3}{8As}, \quad F_{A,L} = Rs^2, \quad F_A = \frac{R^2(1 + s)^4}{16A},$$

$$C = F_A - \delta K = \frac{R(1 + s)^3[R(1 + s) - 2\delta]}{16A},$$

and s solves the explicit quartic

$$16A\eta s^2 - 2\eta R s(1 + s)^3 - (1 - \eta)(1 + s)^3[R(1 + s) - 2\delta] = 0. \quad (3.4)$$

The quartic is solvable by radicals, so (K, L, C) are elementary in the primitives; in practice one evaluates the roots numerically and verifies them exactly. At $R = D = \beta^{-1} - q$, (3.4) characterizes interior Euler steady states; at $A = 10$ its unique admissible root ($0 < L < \bar{L}$, $C > 0$) is $s = 6.10734\dots$, giving

$$(K^*, L^*, C^*) = (5.06477\dots, 0.82929\dots, 16.7673\dots),$$

the point quoted for orientation in Appendix E. At $R = \delta$, (3.4) characterizes the golden-rule point maximizing stationary utility over maintainable stationary allocations; at A_- its admissible root is $s = 8.22125\dots$, giving $(K_g, L_g, C_g) = (6.91896\dots, 0.84159\dots, 17.4872\dots)$.

Proof. $F_{A,K} = 8AKL^3/(K + L)^3 = 8AsL/(1 + s)^3$, so $F_{A,K} = R$ gives $L = R(1 + s)^3/(8As)$. Since $F_{A,L}/F_{A,K} = (K/L)^2$, $F_{A,L} = Rs^2$. Substituting $L(s)$ into $F_A = 4As^2L^2/(1 + s)^2$ and $\delta K = \delta sL$ gives the displays; at $R = \delta$, $C = \delta^2(1 + s)^3(s - 1)/(16A)$, so interior consumption requires $s > 1$. Multiplying the labor condition by $16A/R$ and substituting yields (3.4). For the golden rule: stationarity of $u = \eta \log(F_A - \delta K) + (1 - \eta) \log(1 - L)$ in K is $\eta(F_{A,K} - \delta)/C = 0$, i.e. $R = \delta$, and stationarity in L is the displayed labor condition. $\square$

Table 1 collects the scale landmarks of both certified economies.

Table 1: Scale landmarks. Boundaries at $A_+ = 10.04$; the Euler steady state K^* at $A = 10$; the golden-rule point K_g at A_- (demonstration) and $A = 10$ (annual).

	κ	K^*	K_g	$\bar{K}_{\text{sus}}$
Demonstration ($\beta = \delta = .7$)	0.0182	5.06	6.92	42.65
Annual ($\beta = .96, \delta = .08$)	0.0020	17.98	22.61	386.99

Remark 3.3 (Best stationary floor). The golden-rule root is the interior stationary point of maximal stationary utility identified by [Proposition 3.3](#); it serves as a floor on $[K_g, \bar{K}]$ by [Lemma 6.2](#) (the certified floors below rely on no maximality claim): at A_- ,

$$\frac{\eta \log C_g + (1 - \eta) \log(1 - L_g)}{1 - \beta} = 4.83417 \dots,$$

against $4.78265 \dots$ for the convenient rational point $(5, \frac{5}{6})$ used in [Proposition 7.4](#). Both are closed form; the golden-rule value is exact in radicals, the rational one in logarithms of rationals.

Remark 3.4 (Beyond $\theta = 2$). Write the family as $F = A(2KL/(K + L))^\theta$, $\theta > 1$ (the certified case is $\theta = 2$). Maintainability at capacity reads $A2^\theta \bar{L}^{\theta-1} x^{\theta-1}/(1+x)^\theta \geq \delta$ in $x = K/\bar{L}$; the left side is single-peaked with maximizer $x = \theta - 1$, so the maintainable set is again an interval $[x_-, x_+]$ around $\theta - 1$, nonempty exactly when $\bar{L} \geq L_{\min}^{\text{avg}}(\delta)$ as in [Proposition 2.1](#). The boundaries are algebraic in the primitives only at $\theta = 2$, the exactly solvable member of the family, but their wedge elasticities are closed form *given* the boundary: total differentiation of the boundary condition gives

$$\frac{\partial \log x_*}{\partial \log A} = - \left[(\theta - 1) - \theta \frac{x_*}{1 + x_*} \right]^{-1},$$

negative at x_- and positive at x_+ , and reducing at $\theta = 2$ to $\mp E$ of [Proposition 3.1](#) via the identity $E = (1 + x_-)/(1 - x_-)$. The elasticity still diverges as the boundaries merge at $x = \theta - 1$: the floor divergence of [Remark 3.2](#) survives at every θ .

4 Comparative statics of the scale boundaries

The word *wedge* below is borrowed from the accounting literature ([Chari et al., 2007](#); [Restuccia and Rogerson, 2008](#); [Hsieh and Klenow, 2009](#)); here a wedge is a perturbation of the planner's primitives, and no equilibrium claim is made. Scale floors make one channel transparent. Suppose a productivity or usability wedge scales effective output by $1 - \tau_A \in (0, 1]$. Then

$$L_{\min}^{\text{avg}}(\delta; \tau_A) = L_{\min}^{\text{avg}}(\delta; 0) (1 - \tau_A)^{-1/(\theta-1)}. \quad (4.1)$$

If a return wedge raises the required marginal return by $\omega_R \geq 0$, the marginal floor becomes

$$L_{\min}^{\text{mg}}(D + \omega_R; \tau_A) = \left(\frac{D + \omega_R}{(1 - \tau_A)\bar{\varphi}} \right)^{1/(\theta-1)}. \quad (4.2)$$

A labor-capacity wedge that consumes a fraction τ_L of effective labor replaces $\bar{L}$ by $(1 - \tau_L)\bar{L}$. The economy crosses the physical floor when $(1 - \tau_L)\bar{L} < L_{\min}^{\text{avg}}(\delta; \tau_A)$ and the Euler floor when $(1 - \tau_L)\bar{L} < L_{\min}^{\text{mg}}(D + \omega_R; \tau_A)$. The elasticity of either floor with respect to usable productivity is

$$\frac{\partial \log L_{\min}}{\partial \log(1 - \tau_A)} = - \frac{1}{\theta - 1}. \quad (4.3)$$

Table 2: Floor amplification: percentage rise of both floors after a productivity wedge.

	$\theta = 1.05$	$\theta = 1.25$	$\theta = 1.5$	$\theta = 2$
$\tau_A = 1\%$	+22.3%	+4.1%	+2.0%	+1.0%
$\tau_A = 5\%$	+179.0%	+22.8%	+10.8%	+5.3%

Proposition 4.1 (Collapse from wedges). *Let a productivity wedge scale effective productivity to $(1 - \tau_A)A$, $\tau_A \in (0, 1)$, in the certified family.*

- (i) Every wedge dooms a band. *The extinction threshold rises strictly, $\kappa((1 - \tau_A)A) > \kappa(A)$, and from every $K_0 \in (\kappa(A), \kappa((1 - \tau_A)A)]$ — a state viable before the wedge — every feasible path under the wedge converges to zero.*
- (ii) Amplification. *The doomed band widens at the elasticity $E = \sqrt{(m + 1)/(m - 1)}$ of [Proposition 3.1](#), which diverges as capacity approaches the physical floor.*
- (iii) Total collapse at a critical wedge. *If $\tau_A > \tau^* = 1 - \delta/(A\bar{L})$, the sustainable set is empty and every feasible path from every $K_0 > 0$ converges to zero. For the general homogeneous family of [Section 2](#), $\tau^* = 1 - (L_{\min}^{\text{avg}}(\delta)/\bar{L})^{\theta-1}$, which vanishes as $\theta \downarrow 1$; near constant returns the total-collapse wedge is arbitrarily small, regardless of the capacity buffer.*

Proof. (i) κ has elasticity $-E < 0$ in A ([Proposition 3.1](#)), so it rises strictly when effective productivity falls; for $K_0 \leq \kappa((1 - \tau_A)A)$ the argument of [Lemma 7.1](#) applies verbatim at productivity $(1 - \tau_A)A$. (ii) is [Proposition 3.1](#). (iii) After the wedge, $m < 1$ if and only if $(1 - \tau_A)A\bar{L} < \delta$ ([Theorem 3.1](#)); then $\phi(K) < K$ for every $K > 0$, so $C_t > 0$ forces $K_{t+1} < K_t$ with limit ℓ satisfying $\ell \leq \phi(\ell)$, which fails for $\ell > 0$; hence $K_t \rightarrow 0$. The general- θ expression restates $\bar{L} < L_{\min}^{\text{avg}}(\delta)(1 - \tau_A)^{-1/(\theta-1)}$ in [Proposition 2.1](#) and (4.1). $\square$

This is the sense, and the only sense, in which the model licenses the statement that a wedge can destroy an economy: not by leaving the certified interval but by crossing the physical floor. Every productivity wedge, however small, condemns a band of previously viable capital stocks to certain extinction; the condemned band widens explosively as capacity nears the floor; and beyond the critical wedge τ^* no initial condition survives. The critical wedge can be small. At $\theta = 1.05$, an economy holding twice the required labor capacity collapses from every state under a 3.4 percent productivity wedge ($\tau^* = 1 - 2^{-.05} = .0341$). It can also be enormous: at the certified economies $\tau^* = .92$ and $.99$, which is what $m = 24$ and $m = 220$ mean. Within the model these are theorems. The model does not claim that any actual economy sits near its floor; locating m empirically is the step this paper does not take.

[Table 2](#) makes the quantitative content explicit, together with its limits. *Channel one — floor amplification — is a near-constant-returns phenomenon:* the exponent $1/(\theta - 1)$ is large only when θ is close to one. In the certified economy $\theta = 2$ and the elasticity is exactly one; a small wedge moves the floors by a comparably small amount, and both floors sit far below $\bar{L}$ in any case ([Example 2.1](#)).

Channel two — selection-threshold proximity — operates at any θ . Its static component is exact. A productivity wedge moves both sustainability boundaries with elasticity magnitude $E = \sqrt{(m + 1)/(m - 1)}$ ([Proposition 3.1](#)): $E = 3$ at $m = 1.25$, $\sqrt{3}$ at $m = 2$, $\sqrt{2}$ at $m = 3$, and 1.042 at the certified $m = 24.24$; the elasticity is dangerous only near the physical floor. In the certified economy, three scale objects respond to a productivity wedge. First, the extinction

threshold $\kappa(A)$ (the smaller root of the quadratic (7.2) below) rises as usable productivity falls: across the certified interval, a 1.235 percent drop in A (from 10.04 to 9.916) raises the extinction threshold by 1.304 percent (from 0.018157 to 0.018394); in logs this is an elasticity of -1.042415 , and the closed form of Proposition 3.1, evaluated at the geometric midpoint of $\mathcal{A}$, gives -1.042415 as well. Second, the sustainability ceiling (the larger root) falls, from 42.65 to 42.10, with two-point log elasticity $+1.042415$: the same magnitude with opposite sign, as Proposition 3.1 requires. Third, and most importantly, the certified high-selection region itself is proved only for $A \in [9.916, 10.04]$: a wedge that pushes usable productivity below 9.916 moves the economy outside the certificate.

One caveat deserves emphasis, because it is easy to overread threshold results. *Leaving the certified region is loss of guarantee, not proven collapse.* The certificate is a sufficient condition for high selection; nothing in this paper shows that high selection fails just below $A = 9.916$ or just below $K = .1$. What the wedge analysis establishes is a discontinuity in what can be *guaranteed*: inside the certified region, every optimal path reaches high scale; outside it, the planner's problem re-enters the zone where nonconcavity makes local diagnostics unreliable. For environments with increasing returns, this suggests a complementary risk notion: part of the cost of a wedge may be the loss of a certified high-scale guarantee rather than its direct static burden. The paper does not claim that every wedge lowers welfare. Three statements should be kept separate. *Static fact*: wedges move the closed-form boundaries with the elasticities of Proposition 3.1. *Certified fact*: if $(1 - \tau_A)A \in [A_-, A_+]$, the high-selection theorems apply verbatim; if $(1 - \tau_A)A < A_-$, they are silent. *Collapse fact*: wedges collapse economies exactly as Proposition 4.1 states, by raising the extinction threshold over previously viable states or by emptying the sustainable set outright; both require crossing the physical floor, not merely leaving the certified interval. *Interpretive claim*: losing certification is a form of model risk, not a collapse theorem.

5 Bellman barriers and high selection

Let S be a set of admissible transitions, $S(K) \subseteq \Gamma(K)$, on a state set X . For a function B on X define the restricted Bellman operator

$$(\mathcal{T}_S B)(K) = \sup\{u(K, L, K') + \beta B(K') : (L, K') \in S(K)\}. \quad (5.1)$$

Call a path S -admissible if all its transitions lie in S , and let V_S denote the supremum of J over S -admissible paths.

Theorem 5.1 (Bellman barrier). *Suppose payoffs are bounded above on S and B is bounded on X with $B \geq \mathcal{T}_S B$. Then $V_S \leq B$ on X .*

Proof. For any S -admissible path, $B(K_t) \geq u(K_t, L_t, K_{t+1}) + \beta B(K_{t+1})$ for every t . Iterating to horizon T ,

$$B(K_0) \geq \sum_{t=0}^{T-1} \beta^t u(K_t, L_t, K_{t+1}) + \beta^T B(K_T) \geq \sum_{t=0}^{T-1} \beta^t u_t + \beta^T \inf B.$$

Since B is bounded, $\beta^T \inf B \rightarrow 0$, and since payoffs are bounded above the partial sums converge in $[-\infty, \infty)$ to J . Hence $J \leq B(K_0)$; take suprema. $\square$

Remark 5.1 (On a converse). When payoffs are bounded and the feasible correspondence is compact-valued and continuous, V_S is the unique bounded fixed point of $\mathcal{T}_S$ (Stokey et al., 1989), so $V_S \leq U$ holds if and only if some bounded $B \leq U$ satisfies $B \geq \mathcal{T}_S B$ (take $B = V_S$). Log utility is unbounded below, so we do not rely on any such converse: the certificate only ever uses the direct implication of Theorem 5.1, which survives payoffs equal to $-\infty$ on infeasible transitions.

5.1 The first-passage bridge

Fix $0 < K_{\text{lo}} < K_{\text{hi}} < \bar{K}$ and let

$$\mathcal{M} = [K_{\text{lo}}, K_{\text{hi}}], \quad \mathcal{M}^\circ = [K_{\text{lo}}, K_{\text{hi}}), \quad H = [K_{\text{hi}}, \bar{K}].$$

A path is *high-admissible* if there is a finite τ with $K_t \in H$ for all $t \geq \tau$; let $V^H(K)$ be the supremum of J over high-admissible paths from K (with $V^H = -\infty$ if none exists). Since high-admissible paths are feasible, $V^* \geq V^H$ always. Four objects carry the argument:

Object	Economic role	Mathematical role
B_M	bounds middle stagnation	residual Bellman barrier on $\mathcal{M}^\circ$
$\bar{U}$	bounds post-exit continuation	upper barrier for V^* below K_{hi}
W^M	value of explicit escape policy	beats B_M and every exit payoff
W^H	value of explicit high policy	beats every exit payoff from H

Theorem 5.2 (First-passage bridge). *Suppose there exist functions $B_M : \mathcal{M} \rightarrow \mathbb{R}$ bounded, $W^M : \mathcal{M}^\circ \rightarrow \mathbb{R}$, $W^H : H \rightarrow \mathbb{R}$, and $\bar{U} : (0, K_{\text{hi}}) \rightarrow \mathbb{R}$ such that:*

- (D1) Forward closure: *every feasible successor from $K \in \mathcal{M} \cup H$ satisfies $K' < \bar{K}$.*
- (D2) Residual barrier and sandwich: *$B_M(K) \geq u(K, L, K') + \beta B_M(K')$ for every $K \in \mathcal{M}^\circ$ and $(L, K') \in \Gamma(K)$ with $K' \in \mathcal{M}^\circ$; and $B_M \leq W^M$ on $\mathcal{M}^\circ$.*
- (D3) Explicit high policies: *for every $K \in \mathcal{M}^\circ$ there is a high-admissible path from K with value at least $W^M(K)$, and for every $K \in H$ a high-admissible path with value at least $W^H(K)$.*
- (D4) Downward-exit domination: *$\bar{U} \geq V^*$ pointwise on $(0, K_{\text{hi}})$, and*

$$\begin{aligned} W^M(K) &\geq u(K, L, K') + \beta \bar{U}(K') && \text{for all } K \in \mathcal{M}^\circ, (L, K') \in \Gamma(K), K' < K_{\text{lo}}, \\ W^H(K) &\geq u(K, L, K') + \beta \bar{U}(K') && \text{for all } K \in H, (L, K') \in \Gamma(K), K' < K_{\text{hi}}. \end{aligned}$$

Then $V^* = V^H$ on $\mathcal{M} \cup H$.

Proof. Only $V^* \leq V^H$ requires proof. Fix $K_0 \in \mathcal{M} \cup H$ and a feasible path $\pi = (K_t, L_t)$. Define the first downward-exit time

$$\sigma = \inf\{t \geq 0 : (K_t \in \mathcal{M}^\circ \text{ and } K_{t+1} < K_{\text{lo}}) \text{ or } (K_t \in H \text{ and } K_{t+1} < K_{\text{hi}})\}.$$

Before σ the path remains in $\mathcal{M}^\circ \cup H$: from $\mathcal{M}^\circ$ it can only move within $\mathcal{M}^\circ$, into H (using (D1), $K' < \bar{K}$, so $K' \geq K_{\text{hi}}$ lands in H), or exit downward (a σ -event); from H it can only stay in H (again by (D1)) or exit downward.

Case 1: $\sigma = \infty$ and π never enters H . Then every transition of π stays in $\mathcal{M}^\circ$. By (D2) and Theorem 5.1 applied to the residual class, $J(\pi) \leq B_M(K_0) \leq W^M(K_0)$, and by (D3) $W^M(K_0) \leq V^H(K_0)$.

Case 2: $\sigma = \infty$ and π enters H at some finite time. After entry, every transition stays in H (a departure below K_{hi} would be a σ -event, above is impossible). Hence π is high-admissible and $J(\pi) \leq V^H(K_0)$.

Case 3: $\sigma < \infty$. Write $X = M$ if $K_\sigma \in \mathcal{M}^\circ$ and $X = H$ if $K_\sigma \in H$. The continuation of π from $K_{\sigma+1}$ is a feasible path, so its value is at most $V^*(K_{\sigma+1}) \leq \bar{U}(K_{\sigma+1})$, where $K_{\sigma+1}$ lies in the domain of $\bar{U}$ by the definition of σ . Then, by (D4),

$$J(\pi) \leq \sum_{t < \sigma} \beta^t u_t + \beta^\sigma \left[u_\sigma + \beta \bar{U}(K_{\sigma+1}) \right] \leq \sum_{t < \sigma} \beta^t u_t + \beta^\sigma W^X(K_\sigma).$$

By (D3) there is a high-admissible path ρ from K_σ with $J(\rho) \geq W^X(K_\sigma)$. Let π'' follow π up to time σ and then follow ρ . Then π'' is feasible and high-admissible (it eventually enters and stays in H along ρ), and

$$J(\pi) \leq \sum_{t < \sigma} \beta^t u_t + \beta^\sigma J(\rho) = J(\pi'') \leq V^H(K_0).$$

The initial segment is common to π and π'' , so no bound on it is needed.

The three cases are exhaustive, so $J(\pi) \leq V^H(K_0)$ for every feasible π ; take suprema. $\square$

The splice in Case 3 is the load-bearing step, and it is worth stating what makes it work: the exit-domination inequalities in (D4) compare the *exit payoff plus its best unrestricted continuation* with the value of an explicit high policy *at the exit state*, so the deviation is punished exactly where it occurs; discounting transports the comparison back to time zero without any bound on the pre-exit segment.

Corollary 5.1 (Optimal paths select high scale). *Assume, in addition to (D1)–(D4), that the residual-barrier inequality in (D2) and the exit inequalities in (D4) hold with a margin $\varepsilon > 0$. Then every feasible path from $K \in \mathcal{M} \cup H$ that is not high-admissible satisfies $J(\pi) < V^H(K)$. In particular every optimal path from $\mathcal{M} \cup H$, if one exists, reaches H in finite time and remains in H .*

Proof. In Case 1 the margin gives $J(\pi) \leq B_M(K) - \varepsilon/(1 - \beta) < V^H(K)$ (iterate the strict barrier inequality). In Case 3 the margin gives $J(\pi) \leq J(\pi'') - \beta^\sigma \varepsilon < V^H(K)$. Case 2 paths are high-admissible. Optimal paths exist (Proposition C.1); the corollary's bounds do not depend on existence. $\square$

6 Finite certificates

For the certified economy, define the labor-reduced reward

$$R_A(K, K') = \sup_{0 \leq L \leq \bar{L}} u_A(K, L, K'), \quad (6.1)$$

with $R_A = -\infty$ when no feasible labor exists.

Lemma 6.1 (Monotone reward). *For the CES family F_A , $R_A(K, K')$ is nondecreasing in A and K and nonincreasing in K' on the feasible domain. Consequently, for any box $[K_1, K_2] \times [K'_1, K'_2]$,*

$$\sup\{R_A(K, K') : K \in [K_1, K_2], K' \in [K'_1, K'_2]\} \leq R_A(K_2, K'_1).$$

Proof. For fixed L , the payoff is nondecreasing in A and K (output rises, consumption rises) and nonincreasing in K' (consumption falls); suprema preserve these monotonicities. The corner bound is immediate. $\square$

Lemma 6.2 (Monotone extension of policy values). *Fix a node x and a policy from x : an action (L_0, K'_1) feasible at x followed by any continuation path from K'_1 , with total value w . Then from every $K \geq x$ the same action and continuation are feasible and have value at least w .*

Proof. $C = F(K, L_0) + qK - K'_1$ is increasing in K , so feasibility and the period payoff only improve; the continuation from K'_1 is unchanged. $\square$

6.1 Endpoint ledgers

Let $x_0 < \dots < x_n$ be nodes and let B be *log-affine on cells*: affine in $\log K$ on each $[x_i, x_{i+1}]$. Let a *successor partition* $\mathcal{J}$ consist of groups $J = [x_{j(J)}, x_{\bar{j}(J)}]$ of consecutive cells. A *reward ledger* assigns to each current cell i and group J a number Q_{iJ} with

$$Q_{iJ} \geq \sup\{R(K, K') : K \in [x_i, x_{i+1}], K' \in J, (K, K') \text{ feasible}\},$$

for instance $Q_{iJ} \geq R(x_{i+1}, x_{j(J)})$ by Lemma 6.1.

Theorem 6.1 (Endpoint-ledger certificate). *Suppose that for every current cell i and every group J containing a feasible successor,*

$$\min\{B(x_i), B(x_{i+1})\} \geq Q_{iJ} + \beta \max_{x_m \in J} B(x_m) + \varepsilon \quad (6.2)$$

for some $\varepsilon > 0$. Then $B(K) \geq R(K, K') + \beta B(K') + \varepsilon$ for every feasible transition with K in a covered cell and K' in a covered group.

Proof. B is affine in $\log K$ on each cell, so on the current cell $B(K) \geq \min\{B(x_i), B(x_{i+1})\}$, and on the group J the piecewise-affine function B attains its maximum at a node of J , so $B(K') \leq \max_{x_m \in J} B(x_m)$. Combine with $R(K, K') \leq Q_{iJ}$ and (6.2). $\square$

Grouping successor cells coarsens the ledger soundly (the successor maximum is taken over more nodes) and is used away from binding regions; near binding regions groups are single cells.

6.2 Exact reward evaluation

With $\eta = .7$, the interior labor first-order condition for $u_A(K, \cdot, K')$ is

$$7(1 - L)F_L(K, L) - 3(F_A(K, L) + qK - K') = 0. \quad (6.3)$$

Let $Y = K' - qK$. Since $F_L = 8AK^3L/(K + L)^3$, multiplying by $(K + L)^3$ gives the cubic

$$(3Y - 12AK^2)L^3 + (9YK - 68AK^3)L^2 + (56AK^3 + 9YK^2)L + 3YK^3 = 0. \quad (6.4)$$

Theorem 6.2 (Exact reward reduction). *For fixed (A, K, K') , the maximizers of $u_A(K, L, K')$ over $[0, \bar{L}]$ lie in the finite set consisting of $L = 0$ and $L = \bar{L}$ (when feasible) and the real roots of (6.4) in $(0, \bar{L})$ with positive consumption.*

Proof. The payoff is continuously differentiable on the feasible open interval and $-\infty$ at infeasible labor. Interior maximizers satisfy the first-order condition, which is equivalent to (6.4) after clearing the positive denominator $(K + L)^3$; boundary candidates complete the list. $\square$

Theorem 6.2 is the structural reason a finite certificate exists: the reward is determined by finitely many algebraic candidates. For *verification* we use an even more robust route that requires no root isolation. Consumption $C(L) = F_A(K, L) + qK - K'$ is increasing in L , so on any labor subinterval $[\ell_1, \ell_2]$,

$$u_A(K, L, K') \leq \eta \overline{\log} C(\ell_2) + (1 - \eta) \overline{\log}(1 - \ell_1), \quad (6.5)$$

where $\overline{\log}$ denotes an upper bound on the logarithm computed by directed rounding. A branch-and-bound over labor subintervals, with C evaluated in exact rational arithmetic and logarithms bounded with a one-ulp guard, yields a rigorous upper bound on $R_A(K, K')$ at any rational point, to any prescribed tightness. Policy values are bounded below symmetrically (exact C , downward-rounded logarithms, exact geometric sums along eventually-cyclic policy paths). [Appendix D](#) records the full protocol.

Remark 6.1 (The reward in radicals). The cubic (6.4) is solvable in closed form, so by [Theorem 6.2](#) the labor-reduced reward $R_A(K, K')$ is, at every point of its domain, the maximum of at most five explicit closed-form candidates: two endpoints and up to three Cardano–Viète roots. The certificate deliberately uses neither fact — monotone subdivision requires no root isolation and no radical arithmetic — but the closed form makes R_A an elementary function in the classical sense, available wherever an analytical expression is preferred.

7 A verifier-checked high-selection theorem

The certified family is (1.3) with

$$\mathcal{A} = [A_-, A_+] = [9.916, 10.04], \quad \beta = .7, \quad q = .3, \quad \eta = .7, \quad \bar{L} = .88, \quad (7.1)$$

and regions

$$\mathcal{M} = [.098935, 2], \quad H = [2, 50], \quad \bar{K} = 50.$$

7.1 Exact scale boundaries

Recall from [Section 3](#) the full-capacity resources $\phi_A(K) = qK + F_A(K, \bar{L})$. Capital is maintainable at K if and only if $\phi_A(K) \geq K$, which reduces to the quadratic equation

$$0.7 K^2 - (4A\bar{L}^2 - 1.4\bar{L})K + 0.7\bar{L}^2 = 0. \quad (7.2)$$

At $A_+ = 10.04$ the coefficients are exact decimals, $0.7K^2 - 29.867904K + 0.54208 = 0$, with roots

$$\kappa = 0.018156974557\dots, \quad \bar{K}_{\text{sus}} = 42.650277311\dots \quad (7.3)$$

At $A_- = 9.916$ the roots are $0.018393721\dots$ and $42.101322850\dots$. These are the evaluations of the closed forms (3.1), with $m = 2A\bar{L}/\delta - 1 = 24.24342857\dots$ at A_+ ; in particular $\kappa \bar{K}_{\text{sus}} = \bar{L}^2 = .7744$ exactly, the coefficient ratio $0.54208/0.7$ of (7.2) ([Theorem 3.1](#)). Since ϕ_A is increasing in A , maintainability at any $A \in \mathcal{A}$ implies maintainability at A_+ . Two consequences:

Lemma 7.1 (Extinction below κ). *For every $A \leq A_+$ and every feasible path with $K_0 \in (0, \kappa]$, K_t is strictly decreasing and $K_t \rightarrow 0$.*

Proof. On $(0, \kappa]$, $\phi_A(K) \leq \phi_{A_+}(K) \leq K$, with equality only at $K = \kappa$, $A = A_+$. Since $C_t > 0$, $K_{t+1} < \phi_A(K_t) \leq K_t$. The decreasing sequence converges to some $\ell \geq 0$ with $\ell \leq \phi_{A_+}(\ell)$ by continuity; on $(0, \kappa)$ this fails, so $\ell = 0$. $\square$

Remark 7.1 (Upper sustainability bound). The larger root is easy to overlook: for $K > \bar{K}_{\text{sus}}$, $\phi_A(K) < K$ even at full labor, so capital above 42.66 (at A_+ ; above 42.11 at A_-) cannot be maintained and the band $(\bar{K}_{\text{sus}}, 50]$ of H is transient. High selection is a statement about reaching and remaining in H , not about sustaining its upper edge. The physically sustainable set at A_+ is exactly $[\kappa, \bar{K}_{\text{sus}}]$; see [Figure 2](#).

Lemma 7.2 (Forward closure). *For every $A \in \mathcal{A}$: (i) every feasible successor from $K \in (0, 50]$ satisfies $K' < 45.04 < 50$; (ii) every feasible successor from $K \in (0, 60]$ satisfies $K' < 48.21 < 60$.*

Proof. $\phi_A \leq \phi_{A_+}$ and ϕ_{A_+} is increasing (both terms are), so successors from $(0, 50]$ are below $\phi_{A_+}(50) = 1138490/25281 = 45.0334\dots$ and from $(0, 60]$ below $\phi_{A_+}(60) = 48.2073\dots$; both are exact rational evaluations. $\square$

Lemma 7.3 (Low-capital crash bound). *Suppose $qK + F_A(K, \bar{L}) \leq \lambda K$ for all $K \leq \kappa_\lambda$, with $\lambda < 1$. Then for every $K_0 \leq \kappa_\lambda$, every feasible path satisfies $K_t \leq \lambda^t K_0$ and $C_t \leq \lambda^{t+1} K_0$, and*

$$V_A^*(K_0) \leq \frac{\eta}{1-\beta} \log K_0 + \frac{\eta\beta}{(1-\beta)^2} \log \lambda.$$

Proof. From the resource constraint, $C_t + K_{t+1} \leq \lambda K_t$, so $C_t \leq \lambda K_t$ and $K_{t+1} < \lambda K_t$; induction gives $K_t \leq \lambda^t K_0$ and hence $C_t \leq \lambda^{t+1} K_0 \leq \lambda^t K_0$. Leisure utility is nonpositive, so

$$V_A^*(K_0) \leq \eta \sum_{t \geq 0} \beta^t (\log K_0 + t \log \lambda) = \frac{\eta}{1-\beta} \log K_0 + \frac{\eta\beta}{(1-\beta)^2} \log \lambda. \quad \square$$

Remark 7.2. Using the sharper $C_t \leq \lambda^{t+1} K_0$ instead of $\lambda^t K_0$ improves the constant to $\eta/(1-\beta)^2$; the certificate conservatively uses the displayed weaker constant, which only strengthens its inequalities. The K -dependent optimal choice of λ yields the closed-form envelope of [Proposition 3.2](#), which is strictly tighter (-17.78 against -16.14 at $K = .005$); the certificate keeps the fixed- λ form.

Lemma 7.4 ($\lambda = \frac{1}{2}$ below the crash edge). *At A_+ (hence for every $A \in \mathcal{A}$), $qK + F_A(K, \bar{L}) \leq \frac{1}{2}K$ for all $K \leq 0.005$. Indeed $F_A(K, \bar{L})/K$ is increasing in K on $(0, \bar{L})$, and the exact evaluation at the edge gives $\phi_{A_+}(0.005) = 0.0024926\dots \leq 0.0025$; the condition in fact holds up to $K = 0.0050372\dots$, the smaller root of $0.2K^2 - 30.747904K + 0.15488 = 0$.*

Proof. $F_A(K, \bar{L})/K = 4A\bar{L}^2K/(K+\bar{L})^2$ and $\frac{d}{dK}[K/(K+\bar{L})^2] = (\bar{L}-K)/(K+\bar{L})^3 > 0$ for $K < \bar{L}$. Hence the constraint $F_A(K, \bar{L}) \leq (\frac{1}{2} - q)K = 0.2K$ on $(0, 0.005]$ follows from the exact endpoint evaluation; clearing denominators gives the displayed quadratic for the exact threshold. $\square$

Remark 7.3. [Lemma 7.4](#) is the $\lambda = \frac{1}{2}$ evaluation of the closed form of [Corollary 3.1](#): $\kappa_{1/2} = \bar{L}(m_{1/2} - \sqrt{m_{1/2}^2 - 1})$ with $m_{1/2} = 2A_+\bar{L}/(\frac{1}{2} - q) - 1 = 87.352$, giving $0.0050372\dots$ as displayed.

Consequently the *crash continuation bound*

$$\bar{U}_A(K) = \frac{\eta}{1-\beta} \log K + \frac{\eta\beta}{(1-\beta)^2} \log \frac{1}{2} \quad (0 < K < 0.005) \quad (7.4)$$

satisfies $\bar{U}_A \geq V_A^*$ on $(0, 0.005)$ for every $A \in \mathcal{A}$, and its value at the edge, $-16.13654\dots$, bounds it from above on that band.

7.2 The certificate

Definition 7.1 (Certificate $\mathfrak{C}$). The certificate is the tuple $\mathfrak{C} = (\mathcal{G}, B_M, U^+, \pi^M, \pi^H, \Gamma_\varepsilon)$ where:

- (i) $\mathcal{G}$: three node grids with exact decimal nodes — $\mathcal{G}_M$ on $[\text{.098935}, 2]$ (1,050 nodes), $\mathcal{G}_U$ on $[\text{.005}, 60]$ (781 nodes, containing .005, .098935, 2, 50, and 60 exactly), $\mathcal{G}_H$ on $[2, 50]$ (220 nodes);
- (ii) B_M : node values of a piecewise log-affine *middle barrier* on $\mathcal{G}_M$;
- (iii) U^+ : node values of a piecewise log-affine *upper-continuation barrier* on $\mathcal{G}_U$;
- (iv) π^M : for each $\mathcal{G}_M$ node, a target node (a strictly higher $\mathcal{G}_M$ node or a $\mathcal{G}_H$ node) and a rational labor choice — an acyclic *ladder* into H ;
- (v) π^H : for each $\mathcal{G}_H$ node, a target $\mathcal{G}_H$ node and a rational labor choice — an eventually-cyclic policy inside H ;
- (vi) Γ_ε : the verified ledger margins of [Table 3](#).

The *policy values* W^M, W^H are the exact values of π^M, π^H at A_- , computed with downward rounding (cycle values by exact geometric sums), extended to cells as step functions from the left node via [Lemma 6.2](#). The extended upper barrier is

$$\bar{U}_A(K) = \begin{cases} U^+(K), & K \in [\text{.005}, 60], \\ (7.4), & 0 < K < \text{.005}. \end{cases}$$

Productivity monotonicity lets one certificate cover the whole interval $\mathcal{A}$: reward upper bounds ([Lemma 6.1](#)) are verified at A_+ and policy lower bounds at A_- ; for $A \in \mathcal{A}$ the former only fall and the latter only rise.

Proposition 7.1 (Upper continuation). *The certificate verifies, at A_+ and with margin $\varepsilon_U > 0$: for every current cell i of $\mathcal{G}_U$ and successor group J ,*

$$\min\{U_i^+, U_{i+1}^+\} \geq Q_{iJ} + \beta \max_J U^+ + \varepsilon_U,$$

and the crash handoff row

$$\min\{U_i^+, U_{i+1}^+\} \geq \bar{R}_{A_+}(x_{i+1}, 0) + \beta \bar{U}(0.005^-) + \varepsilon_U,$$

where $\bar{R}(x, 0)$ bounds the reward of any transition to $K' < 0.005$ (consumption at most $F + qK$) and $\bar{U}(0.005^-)$ bounds (7.4) on $(0, .005)$. Then $\bar{U}_A \geq V_A^*$ on $(0, 60]$ for every $A \in \mathcal{A}$.

Proof. Fix $A \in \mathcal{A}$ and $K_0 \in [\text{.005}, 60]$; successors remain below 60 by [Lemma 7.2\(ii\)](#). For any feasible path, let τ be the first entry into $(0, .005)$. Until τ , [Theorem 6.1](#) (rows at A_+ majorize rows at A) gives the barrier inequality at every step, including — via the handoff row — the step that crosses the edge; telescoping to τ (or to infinity, using boundedness of U^+) and bounding the continuation after τ by [Lemmas 7.3](#) and [7.4](#) yields $J \leq U^+(K_0)$. On $(0, .005)$, $\bar{U}_A \geq V_A^*$ is [Lemma 7.3](#) directly. $\square$

The crash handoff row closes a genuine logical gap: without it, the barrier property of U^+ would be silent about feasible transitions that jump below the barrier's domain, and the telescoping in [Proposition 7.1](#) would not go through. Earlier drafts of this analysis omitted it.

Proposition 7.2 (Middle residual domination). *The certificate verifies, at A_+ with margin $\varepsilon_M > 0$, the endpoint-ledger rows of B_M over residual transitions ($K, K' \in \mathcal{M}^\circ$), and, at A_- with margin $\varepsilon_S > 0$, the sandwich*

$$W_i^M \geq \max\{B_{M,i}, B_{M,i+1}\} + \varepsilon_S \quad \text{for every cell of } \mathcal{G}_M.$$

Then for every $A \in \mathcal{A}$: every path remaining forever in $\mathcal{M}^\circ$ has value at most $B_M \leq W^M \leq V_A^H$.

Proof. [Theorem 6.1](#) and [Theorem 5.1](#) give domination by B_M for every $A \leq A_+$. The ladder π^M strictly ascends $\mathcal{G}_M$ or enters $\mathcal{G}_H$, so from every node it reaches H in finitely many steps and then follows π^H inside H ; its exact value at A_- is W^M , and period payoffs increase in A , so $W^M \leq V_A^H$ for every $A \geq A_-$, using [Lemma 6.2](#) to pass from nodes to cells. $\square$

Proposition 7.3 (Downward exits are dominated). *The certificate verifies, with margin $\varepsilon_X > 0$, for every $\mathcal{G}_M$ cell i and every $\mathcal{G}_U$ cell j with $x_{j+1} \leq K_{lo} = .098935$:*

$$W_i^M \geq \bar{R}_{A_+}(x_{i+1}, x_j) + \beta \max\{U_j^+, U_{j+1}^+\} + \varepsilon_X,$$

together with a crash row per cell (successors below .005), and the analogous families from $\mathcal{G}_H$ cells over $\mathcal{G}_U$ cells with $x_{j+1} \leq 2$. Then condition (D4) of [Theorem 5.2](#) holds for every $A \in \mathcal{A}$ with $\bar{U} = \bar{U}_A$.

Proof. Any feasible exit transition from $K \in \mathcal{M}^\circ$ with $K' < .098935$ lies in some box covered by the family ($\mathcal{G}_U$ contains .098935 and .005 as exact nodes; successors below .005 are covered by the crash rows), so [Lemma 6.1](#), the step-function property of W^M ([Lemma 6.2](#)), and [Proposition 7.1](#) give $W^M(K) \geq R_A(K, K') + \beta \bar{U}_A(K')$. Identically for exits from H below 2. $\square$

Theorem 7.1 (High selection on a productivity interval, verifier-checked). *For every $A \in [9.916, 10.04]$,*

$$V_A^*(K) = V_A^H(K) \quad \text{for every } K \in [.098935, 50],$$

and every optimal path from $[.098935, 50]$ reaches $H = [2, 50]$ in finite time and remains there. Moreover every feasible path from $K \in (0, \kappa]$ converges to zero, where κ is given in (7.3).

Proof. Extinction is [Lemma 7.1](#). For high selection, apply [Theorem 5.2](#) with $K_{lo} = .098935$, $K_{hi} = 2$, $\bar{K} = 50$: (D1) is [Lemma 7.2\(i\)](#); (D2) and (D3) are [Proposition 7.2](#); (D4) is [Proposition 7.3](#) with [Proposition 7.1](#). All ledger families carry strictly positive margins ([Table 3](#)), so [Corollary 5.1](#) applies. $\square$

Remark 7.4 (Stagnation is never optimal). The theorem's economic content is sharper than its statement. The middle region contains states that are physically sustainable forever: from any $K \in \mathcal{M}$ the economy could hover below the ceiling indefinitely, consuming a positive stream. High selection says the planner never chooses this, at any certified productivity, from any certified initial stock; the consumption sacrificed on the climb is always repaid. Within the certified region, sustained middle-scale stagnation is not a trap an optimizing planner can fall into.

Table 3: Rigorous certificate margins (exact arithmetic, directed rounding). Sandwich, closure, and crash-edge margins are exact family minima (exact rational comparisons). Ledger and exit margins are certified lower bounds on the family minimum: every row is verified in exact arithmetic against an adaptive floor ($4 \cdot 10^{-3}$, tightened to $5 \cdot 10^{-4}$ on near-binding rows). Positivity of every entry certifies the theorem.

Family	Content	Rows	Min. margin
B_M ledger	residual barrier on $\mathcal{M}^\circ$, at A_+	181,833	$\geq 5.0 \cdot 10^{-4}$
U^+ ledger	upper continuation on $[\cdot 005, 60]$, at A_+ , incl. crash rows	167,583	$\geq 9.9 \cdot 10^{-4}$
Sandwich	$W^M \geq \max$ -endpoint $B_M + \varepsilon_S$, at A_-	1,049	0.0104632
$\mathcal{M}$ exits	W^M vs. exits below .098935, incl. crash rows	2,098	$\geq 4.0 \cdot 10^{-3}$
H exits	W^H vs. exits below 2, incl. crash rows	11,998	$\geq 4.0 \cdot 10^{-3}$
Closure	$\phi_{A_+}(50) < 50$, $\phi_{A_+}(60) < 60$ (exact rational)	2	4.9665, 11.79
Crash edge	$\phi_{A_+}(\cdot 005) \leq \cdot 0025$ (exact rational)	1	$7.3 \cdot 10^{-6}$




Figure 1: The certified scale taxonomy at the primitives (7.1). Extinction below $\kappa \approx 0.0182$ (Lemma 7.1); certified high selection on $[\cdot 098935, 50]$ (Theorem 7.1); the band $(\kappa, .1)$ is not classified by the present certificate. $\bar{K}_{\text{sus}} \approx 42.65$ is the upper sustainability root (Remark 7.1). On this log axis κ and $\bar{K}_{\text{sus}}$ are mirror images about $K = \bar{L}$: $\kappa \bar{K}_{\text{sus}} = \bar{L}^2$ (Theorem 3.1).

7.3 What the certificate consists of, in one paragraph

Every inequality behind Theorem 7.1 is one of: (a) an exact rational comparison (closure, crash edge, quadratic roots); (b) an endpoint-ledger row comparing recorded barrier node values against a rigorous corner bound on the reward, computed by monotone subdivision with directed-rounding logarithms at exact rational arguments; (c) an exact re-evaluation of a recorded policy graph with downward rounding. The counts and worst margins are in Table 3; ledger and exit entries are certified lower bounds on family minima, while the sandwich, closure, and crash margins are exact. The certificate data and the verifier are described in Appendix D; a fully worked ledger row appears in Appendix E. No value function is ever computed in the trusted path, and no floating-point comparison enters any verified inequality.

7.4 The certificate at annual discounting and depreciation

The primitives $\beta = \delta = .7$ compress the economics into a small state space. To show that the method survives realistic time preferences, we certify the same family a second time at

$$\beta = \frac{24}{25} = .96, \quad \delta = \frac{2}{25} = .08, \quad \eta = .7, \quad \bar{L} = .88, \quad A \in [9.916, 10.04].$$

Only β and δ are annual; the technology is unchanged, so this economy has realistic impatience and depreciation alongside degree-two returns to scale and a capital share near .09. It is a stress test of the method rather than a calibration to moments. The closed forms of Section 3 locate it




Figure 2: Maximal gross return $\phi_A(K)/K = (qK + F_A(K, \bar{L}))/K$. Capital is maintainable where the curve exceeds one: between the extinction root κ and the sustainability root $\bar{K}_{\text{sus}}$ of (7.2). The two curves bound the certified productivity interval. The crossings with 1 are log-symmetric about $\bar{L}$ (Remark 3.1).

Table 4: The annual-calibration certificate ($\beta = .96$, $\delta = .08$; grids 1,959 nodes on $[.0004, 520]$, 579 on $[.016, 1.5]$, 318 on $[1.5, 450]$). Conventions as in Table 3.

Family	Content	Rows	Min. margin
B_M ledger	residual barrier on $ [.016, 1.5)$, at A_+	46,476	$\geq 2.0 \cdot 10^{-4}$
U^+ ledger	upper continuation on $ [.0004, 520]$, incl. crash rows	637,011	$\geq 2.0 \cdot 10^{-4}$
Sandwich	$W^M \geq \max\text{-endpoint } B_M + \varepsilon_S$, at A_-	578	7.8589
$\mathcal{M}$ exits	W^M vs. exits below .016, incl. crash rows	1,234	$\geq 4.0 \cdot 10^{-3}$
H exits	W^H vs. exits below 1.5, incl. crash rows	26,815	$\geq 4.0 \cdot 10^{-3}$
Closure	$\phi_{A_+}(450) < 450$, $\phi_{A_+}(520) < 520$ (exact rational)	2	5.0213, 10.60
Crash edge	$\phi_{A_+}(.0004) \leq .94 \cdot .0004$ (exact rational)	1	$1.5 \cdot 10^{-6}$

immediately: $m = 2A_+\bar{L}/\delta - 1 = 219.88$, so the sustainable set is $[\kappa, \bar{K}_{\text{sus}}]$ with $\kappa = 0.0020011\dots$ and $\bar{K}_{\text{sus}} = 386.9868\dots$ (the roots of $0.08K^2 - 30.959104K + 0.061952$, with $\kappa\bar{K}_{\text{sus}} = \bar{L}^2$ as always), the boundary elasticity is $E = 1.00456$, and the quartic (3.4) puts the interior steady state at $(K^*, L^*, C^*) = (17.985, 0.8256, 23.487)$ at $A = 10$.

Theorem 7.2 (High selection at annual β and δ , verifier-checked). *At these primitives, for every $A \in [9.916, 10.04]$,*

$$V_A^*(K) = V_A^H(K) \quad \text{for every } K \in [.016, 450],$$

where V_A^H is the value over paths that reach and remain in $H = [1.5, 450]$, and every optimal path from $[.016, 450]$ reaches H in finite time and remains there. Every feasible path from $K \in (0, \kappa]$ converges to zero.

Three features of the annual certificate are worth recording, because they refine the scaling estimate of Section 9. First, the cost of patience is not spread evenly over the state space; it concentrates on the *production-saturation plateau* $[30, \bar{K}_{\text{sus}}]$, where output is flat, the value function is nearly flat, and near-optimal over-accumulation paths transmit barrier slack from coarse top

cells down the entire climb chain. At the density of the $\beta = .7$ grid the upper barrier exceeded the value function by 0.78 at the exit boundary; tripling the plateau density cut the excess to 0.15, inside the exit-domination budget. Second, patience starves the bottom: optimal consumption falls to $O(10^{-3})$ below $K \approx .1$ and to $O(10^{-4})$ below .01, reward slopes scale like η/C , and grid cost follows; a dedicated 550-node band on $[\text{.006}, \text{.05}]$ carries the certificate down to $K_{\text{lo}} = .016 = 8.0 \kappa$, and going further is a matter of node counts, with the limit recorded in [Conjecture 9.1](#). Third, the exit-domination room shrinks as K_{hi} rises toward the steady state (0.65 at $K_{\text{hi}} = 1$, 0.45 at 1.5, 0.14 at 4, measured at construction), which fixes $H = [1.5, 450]$. The certified sandwich margin, by contrast, is enormous (7.86 against 0.0104 at $\beta = .7$): patience makes hovering below the ceiling catastrophic relative to escape, so the middle family is the easy one at an annual calibration. In total the annual certificate verifies 711,536 ledger and exit rows against 363,512 at the demonstration primitives: moving β from .7 to .96 and δ from .7 to .08 roughly doubled the rows at matched domain ambition, close to the $[(1 - \beta)\varepsilon]^{-1/2}$ estimate, with the composition shift described above. Matching factor shares would mean re-parameterizing the CES family, which [Section 3](#) makes routine and which is the natural companion certificate; the technology is held fixed here to isolate the effect of (β, δ) on certificate cost.

7.5 Closed-form floors for the high value

The certified policy values have closed-form counterparts within 0.024.

Proposition 7.4 (Stationary high floors). *Let (K_s, L_s) satisfy $K_s \in H$, $0 \leq L_s \leq \bar{L}$, and $C_s = F_A(K_s, L_s) - \delta K_s > 0$. Then*

$$V_A^H(K) \geq \frac{\eta \log C_s + (1 - \eta) \log(1 - L_s)}{1 - \beta} \quad \text{for every } K \in [K_s, \bar{K}].$$

At $(K_s, L_s) = (5, \frac{5}{6})$, $F_A(5, \frac{5}{6}) = 100A/49$ exactly, so for every $A \geq A_-$ and every $K \in [5, 50]$,

$$V_A^H(K) \geq \frac{7 \log(\frac{100A}{49} - \frac{7}{2}) + 3 \log \frac{1}{6}}{3} = 4.782654 \dots \text{ at } A_-.$$

For $K \in [2, 5)$, the one-step policy with labor $\frac{4}{5}$ and target $K' = 5$ is feasible, $F_A(2, \frac{4}{5}) = 64A/49$ exactly, and

$$V_A^H(K) \geq \frac{7 \log(\frac{64A}{49} - \frac{22}{5}) + 3 \log \frac{1}{5}}{10} + \beta \frac{7 \log(\frac{100A}{49} - \frac{7}{2}) + 3 \log \frac{1}{6}}{3} = 4.367302 \dots \text{ at } A_-.$$

Proof. From any $K \geq K_s$, the action (L_s, K_s) is feasible with consumption $F_A(K, L_s) + qK - K_s \geq C_s > 0$ ([Lemma 6.2](#)); remaining at (K_s, L_s) forever is then feasible with per-period consumption C_s , and the path never leaves H , so its value is at least $u_s/(1 - \beta)$ with $u_s = \eta \log C_s + (1 - \eta) \log(1 - L_s)$, the first period being weakly better. The evaluations are exact rational identities: $F_A(5, \frac{5}{6}) = 4A \cdot 25 \cdot \frac{25}{36} / (\frac{35}{6})^2 = 100A/49$ and $F_A(2, \frac{4}{5}) = 4A \cdot 4 \cdot \frac{16}{25} / (\frac{14}{5})^2 = 64A/49$, with first-step consumption $\frac{64A}{49} + \frac{3}{5} - 5 = \frac{64A}{49} - \frac{22}{5} > 0$ at A_- ; for $K \in (2, 5)$ the first-step consumption only rises ([Lemma 6.2](#)). Payoffs are nondecreasing in A , so the A_- evaluations bound the whole interval $\mathcal{A}$. $\square$

Remark 7.5 (Distance to the certificate). At A_- the certified value at the bottom of H is $W^H(2) = 4.396665 \dots$ ([Appendix E](#)), and the two-line closed form of [Proposition 7.4](#) gives 4.367302 $\dots$. Choosing the bridge labor by the closed-form cubic root of (6.4) at $(K, K') = (2, 5)$, which is $L^* = 0.829625 \dots$ at A_- , improves the floor to 4.373288 $\dots$, a gap of 0.0234; the golden-rule

point of [Proposition 3.3](#) raises the stationary floor to 4.83417... on [6.919, 50]. Wherever a certified exit margin exceeds such a gap, the closed-form floor can stand in for the recorded policy π^H in (D4) without recomputation. The middle ladder is different: hand-built two-step ladders from $K = .14$ lose about 0.3 against the certified middle-ladder values ($W^M(.14) \approx .3938$), and this is why π^M stays in the certificate.

8 Active geometry

The high-selection theorem is proved by dominance, not by comparative statics, and no certified theorem uses monotone comparative statics. This section is diagnostic only; it records the geometry that makes monotone selection natural in this class, and [Appendix B](#) records why it cannot replace the barrier argument.

On active transitions $qK < K' < qK + F(K, \bar{L})$, write $Y = K' - qK$. The labor first-order condition can be written

$$G_{K,Y}(L) = \frac{(1-L)F_L(K, L)}{F(K, L) - Y} = \frac{1-\eta}{\eta}.$$

For the certified family, interior solutions are roots of the cubic (6.4), so there are at most three; on the subregion where $G_{K,Y}$ is strictly decreasing in L the active labor choice is unique. At an interior optimum satisfying the second-order condition, the sign of $R_{KK'}$ is governed by

$$S = (F_K + q)\{F_L - (1-L)F_{LL}\} + (1-L)F_L F_{KL}.$$

Gross complementarity makes the second term positive, and a large consumption weight makes the first nonnegative, which is why active high-scale policies are robust in the certified economy. None of this enters the proof of [Theorem 7.1](#).

9 Discussion

The paper delivers a minimum-scale taxonomy with an explicit audit trail. Below κ , capital cannot maintain itself and every feasible path collapses ([Lemma 7.1](#)). On [.098935, 50], for every productivity in the certified interval, high selection is global and strict: residual middle paths, downward exits, and unrestricted continuations are all dominated by explicit high policies, with verifier-checked positive margins, and every optimal path joins the high region in finite time ([Theorem 7.1](#)). Above $\bar{K}_{\text{sus}} \approx 42.65$, capital is transient even inside H ([Remark 7.1](#)). The band $(\kappa, .1)$ is unclassified; at A_- it is (0.01839, .1). That band is not hidden — it is the boundary of the present certificate, and [Figure 1](#) displays it as such. What the untrusted stage sees is recorded as a conjecture.

Conjecture 9.1 (Selection frontier). *For each economy there is a threshold $K_f(A) \in (\kappa, 1.2\kappa)$ such that high selection holds on $(K_f, \bar{K}]$ and fails on (κ, K_f) . Node-restricted value iteration on the certificate grids places $K_f \approx 1.03\kappa$ at the demonstration primitives and $K_f \approx 1.16\kappa$ at annual β and δ .*

The certified regions reach 5.4κ and 8.0κ (each relative to κ at A_+). The remaining gap has a known price (per-decade node counts scale like inverse consumption, which vanishes toward κ) and a known limit: at K_f itself the sandwich and exit rooms vanish by definition, so no positive-margin certificate exists there. A certified two-sided bracket of K_f , with a bound on its productivity

elasticity, is the natural next target; it would turn the frontier from a conjecture into this paper’s central theorem.

Read together, the certificates and the conjecture carry a sharp implication. If the conjecture is right, value-based traps occupy a sliver. States from which decline is optimal lie between κ and roughly 1.2κ , while escape is optimal everywhere above (certified down to 5.4κ and 8.0κ , conjectured down to the sliver’s edge). In this economy “trapped but viable” is a nearly empty category: an economy is either close enough to extinction that decline is efficient, or escape pays. Poverty persistence at intermediate scale must then come from what the model omits (frictions, expectations, institutions, decentralized externalities), rather than from optimal dynamics. The sliver offers one model-based reading of the empirical ambivalence about poverty traps (Kraay and McKenzie, 2014; Balboni et al., 2022).

The economic mechanism is not local stability. The high branch is selected because every alternative path class is value dominated: a residual middle path is bounded by a middle Bellman barrier; a downward exit is bounded, at the exit state, by an unrestricted continuation barrier; explicit high policies beat both. This is why the theorem is global, and why it can tolerate the nonconcavities that defeat Euler-equation arguments and, in part of the parameter space, lattice arguments (Appendix B).

The wedge analysis separates two channels. Floor amplification (Table 2) is quantitatively dramatic only near constant returns. Selection-threshold proximity operates at any curvature and is the channel the certified economy quantifies: a productivity wedge moves the extinction threshold, shrinks the sustainable set, and — beyond a known, certified boundary — forfeits the guarantee of high selection. We stress once more that leaving the certificate is loss of guarantee, not proven collapse; we offer the certified guarantee as a complement to estimates of the threshold itself: it is the one object a planner near a threshold cannot re-derive from local diagnostics.

Why the certificate survives the closed forms. Everything static in this paper is now exact: boundaries, elasticities, crash thresholds, the value bracket on the extinction region, the stationary points, the high-value floors. Two objects are not: the middle barrier and the exit ledgers. The obstruction is quantitative. The binding sandwich margin at $K = .099$ is 0.0105 , a per-period slack budget of $(1 - \beta) \times 0.0105 \approx 3.1 \times 10^{-3}$. The natural analytical middle barrier is a one-step resource bound such as

$$B(K) = \eta \log \phi_{A+}(K) + \frac{\beta\eta}{1-\beta} \log \phi_{A+}(2),$$

valid on $\mathcal{M}^\circ$ by the same telescoping (ϕ_A is increasing, so $\phi_A(K') \leq \phi_A(2)$ for $K' < 2$), but worth 3.75 at $K = .099$ against a required -0.79 — a gap more than four hundred times the budget — because it lets a path consume peak resources and keep peak capital at once, period after period. Untangling that double count to 5×10^{-3} per period across a two-decade state box is what the 181,833-row residual ledger is for. A closed form that tight would be a closed-form solution of the nonconcave dynamic program itself; none is known even for the concave benchmark away from Cobb–Douglas with full depreciation, and the reverse-Inada condition that creates the phenomenon rules Cobb–Douglas out by design (Assumption 2.1). The certificate is confined to the two families where closed forms fail.

Calibration and certificate cost. We chose $\beta = \delta = .7$ to keep the certificate small, and for no other reason: strong discounting shortens the horizon over which barrier slack accumulates, and rapid depreciation compresses the sustainable set into two decades of state space. The closed forms of Section 3 do not care — β enters only through D , the crash envelope (3.3), and the quartic — so realism is a question about the certified families alone. A family with margin ε

has per-period budget $(1 - \beta)\varepsilon$, and the interpolation error of a log-affine barrier falls with the square of node density, so the required nodes per state decade should grow like $[(1 - \beta)\varepsilon]^{-1/2}$: roughly 2.7 per decade for $\beta = .96$. This estimate has now been tested. The annual certificate of [Section 7.4](#) roughly doubled the verified rows (711,536 against 363,512) at a $1.4\times$ node count, in line with the exponent, but with the cost concentrated where the estimate did not look — on the production-saturation plateau, where near-flat values let coarse cells leak slack into every climb chain, and in the low-capital starvation zone, where patience drives consumption and therefore grid spacing toward zero. The refined statement is that $\beta \rightarrow 1$ is affordable wherever consumption is bounded away from zero and values have curvature; the certified statements are [Theorems 7.1](#) and [7.2](#).

Several extensions are natural. First, the unclassified band invites a finer low-scale certificate; nothing in the method caps the region except construction effort at small consumption scales, where reward slopes grow like $\eta(F_K + q)/C$. Second, the productivity interval can be widened by strengthening the binding families (the sandwich at $K = .14$ and the H -exit rows near $K' = 2$); the margins in [Table 3](#) show exactly how much room each family has. Third, other primitives $(\eta, \beta, \bar{L})$ can be intervalized by the same two-endpoint monotonicity device wherever the reward is monotone in the primitive. The quartic (3.4) makes the steady state’s comparative statics available in radicals, and the implicit-elasticity formula of [Remark 3.4](#) extends the wedge theory to every $\theta > 1$. Fourth, stochastic shocks call for robust or probabilistic barriers, for which the control-theoretic literature offers templates ([Prajna et al., 2007](#)). Finally, the planner problem was deliberately solved first, and its limits should be stated plainly: with aggregate returns of degree two, marginal-product factor payments would exhaust twice output, so no competitive decentralization exists. The natural equilibrium counterpart has constant-returns firms and a Romer-type external effect reproducing F in the aggregate ([Romer, 1986](#); [Benhabib and Farmer, 1994](#)); wedges would then sit on prices, and the gap between planner and equilibrium thresholds would itself be big-push content ([Murphy et al., 1989](#); [Krugman, 1991](#)). Certifying that economy is the right sequel.

A CES tangent majorants

For the general CES family

$$F(K, L) = A[\omega K^\rho + (1 - \omega)L^\rho]^{\theta/\rho}, \quad \rho < 0,$$

weighted AM–GM gives, for every $p \in (0, 1)$,

$$F(K, L) \leq A \left[\left(\frac{\omega}{p} \right)^p \left(\frac{1 - \omega}{1 - p} \right)^{1-p} \right]^{\theta/\rho} K^{\theta p} L^{\theta(1-p)}.$$

These Cobb–Douglas tangent envelopes are useful for *discovering* barrier shapes. The certificate of [Section 7](#) does not use them: it bounds the exact reward directly.

B Sharpness: why lattice methods do not apply

The natural alternative route to high selection is monotone comparative statics: if the labor-reduced reward $R(K, K')$ had increasing differences, Topkis-style arguments would order policies and basins. This appendix records that increasing differences *fail* on active transitions in part of the parameter space of this class; the barrier route requires no such structure anywhere.

Let $F_0(K, L) = BL^\theta$ and write $z = qK - K'$; active means $z < 0$. At an interior labor optimum the first-order condition gives

$$z = BL^{\theta-1}\{\chi\theta - (\chi\theta + 1)L\}, \quad \chi = \frac{\eta}{1 - \eta}.$$

For $\eta < 1/2$, choose θ large enough that $\chi\theta/(\chi\theta + 1) < (\theta - 1)/\theta$ and take L between these bounds: the lower bound makes $z < 0$ (active), and the upper bound produces $R_{KK'} < 0$, i.e., strictly *decreasing* differences, at that transition. Since R depends on (K, K') only through z for pure-labor technologies, the sign computation is one-dimensional. The construction requires $\eta < 1/2$, and at $\theta = 2$ it requires $\eta < 1/3$; it therefore establishes the failure mode for the class, not for the certified economies, which have $\eta = .7$. High-precision scans at $\eta = .7$ find no violation of increasing differences on the certified state box (violations appear only near $K \approx 100$, outside forward closure, through a passive-transition mechanism), so whether a lattice argument could deliver a selection theorem for the certified economies is open. What the appendix establishes is narrower: supermodularity is not a property of the class, so a lattice route cannot be generic. The certified theorems are in any case obtained from dominance, which asks for no supermodularity anywhere.

C Existence of optimal paths

Proposition C.1. *Let $X = [0, \bar{B}]$ be compact with $K_0 \in X$ and $\phi_A(\bar{B}) \leq \bar{B}$, so that every feasible path from K_0 remains in X . Then an optimal path exists: $V^*(K_0) = J(\pi^*)$ for some feasible π^* .*

Proof. If $V^*(K_0) = -\infty$ there is nothing to prove, so assume $V^*(K_0) > -\infty$ (on the certified regions this holds because explicit high policies have finite value). Let π^n be a maximizing sequence. Each π^n lives in the compact product $\prod_{t \geq 0} (X \times [0, \bar{L}])$, so by diagonal extraction we may assume $(K_t^n, L_t^n) \rightarrow (K_t^*, L_t^*)$ for every t , with $K_0^* = K_0$. The closed constraints $0 \leq L \leq \bar{L}$ and $0 \leq K_{t+1} \leq F_A(K_t, L_t) + qK_t$ pass to the limit. The period payoff, extended by $u = -\infty$ where $C \leq 0$, is continuous into $[-\infty, \bar{u}]$ with $\bar{u} = \sup_{K \in X} \eta \log \phi_A(K) < \infty$. Applying Fatou's lemma to the nonnegative sequence $\bar{u} - u_t^n$ against the finite measure $\sum_t \beta^t \delta_t$,

$$\limsup_n \sum_t \beta^t u_t^n \leq \sum_t \beta^t \limsup_n u_t^n \leq \sum_t \beta^t u_t^* = J(\pi^*),$$

so $J(\pi^*) \geq V^*(K_0)$. In particular $J(\pi^*) > -\infty$, which forces $C_t^* > 0$ and $K_{t+1}^* > 0$ for every t ; hence π^* is feasible in the original open-constraint sense and optimal. See [Le Van and Morhaim \(2002\)](#) for a general treatment of dynamic programming with unbounded returns. $\square$

D Certificate and verification protocol

The certificate package accompanying the paper contains the grids, barrier node values, policy graphs, and margins of [Definition 7.1](#), together with two programs.

Construction (untrusted). Grids are log-spaced with extra density where reward slopes are large (low $\mathcal{M}$, successor band near $K' = 2$). Candidate barriers are least fixed points of the finite ledger operator, started from node-restricted value iteration; candidate policies are node-restricted value-iteration policies (π^H , cyclic) and an ascending node dynamic program (π^M). Nothing from this stage is trusted; it merely proposes the certificate.

Verification (trusted). A single program re-establishes every inequality:

- (i) all grid nodes, barrier values, labors are exact decimal rationals; consumption is evaluated in exact rational arithmetic;
- (ii) logarithms are bounded via correctly rounded `decimal` evaluations at 15 significant digits, with the argument first rounded outward and the result then moved one unit in the last place outward (`next_plus/next_minus`); all other arithmetic in bound position is exact rational;
- (iii) reward corner bounds $\bar{R}_{A+}(x_{i+1}, x_j)$ use branch-and-bound over labor subintervals with the monotone bound (6.5); the returned value is a valid upper bound at any stopping point; ledger and exit rows are certified against adaptive floors ($4 \cdot 10^{-3}$, tightened to $5 \cdot 10^{-4}$ on near-binding rows), and any floating-point computation is used only to *schedule* refinement, never inside a verified inequality;
- (iv) feasibility of every box is decided exactly at the most-feasible corner, independently of the construction stage;
- (v) policy values are re-derived from the recorded graphs: cycle values by exact geometric sums in rational arithmetic, then backward unrolling; every policy transition’s consumption is verified positive exactly;
- (vi) closure, crash-edge, and quadratic-root statements are exact rational comparisons.

The verified margins are those of Table 3. The mathematical object is the finite list of inequalities; the programs are a reproduction device.

Artifact. The supplementary package `MVS_certificate_package` distributed with this paper contains the complete pipeline and the verified certificate: `certlib.py` (trusted primitives: exact rational consumption, directed-rounding logarithms with a one-ulp guard), `rowengine.py` (trusted row certification; float-guided partition search proposes a partition, every box of which is then certified exactly), `construct.py` (untrusted construction: grids, barriers by least fixed points of the ledger operator, node-restricted value-iteration policies), `verify3.py` (the resumable trusted verifier), `certificate.json` (grids, barrier node values, policy graphs), and `margins.json` (the verified counts and margins reproduced in Table 3). The archive contains the finite data needed to regenerate every certified inequality: node grids, barrier node values, policy graphs, and verification margins. The construction scripts are untrusted and are included only for reproducibility of search. The trusted path re-derives policy values from the recorded graphs, evaluates consumption in exact rational arithmetic, bounds logarithms by outward-rounded decimal intervals, and certifies each row by exact box inequalities. The archive is self-contained (all paths relative); `make verify-smoke` re-certifies the exact families together with both full sandwiches and a sample of rows in about a minute, and `make verify` reproduces both margins files end-to-end (roughly ten and fifteen minutes on a single core; tested under Python 3.10.12 with libmpdec 2.5.1). An independent verifier in a second language, `verify_independent.js` (Node.js, big-integer rationals, logarithm bounds from the arctanh series with explicit remainder terms — no dependence on libmpdec, CPython, or floating point), re-checks the exact rational families and both sandwiches in full, with sampled ledger rows re-certified from scratch; `make verify-independent` runs it in about a minute. The theorem-to-file mapping and hashes are:

Theorem 7.1	<code>certificate2.json / margins2.json / verify5.py</code> <code>SHA-256(certificate2.json) = bbec8188c5ae04f7fbc9462936a0831d</code> <code>b586906110fd1c6c5fde34e6bd879d8f</code> <code>SHA-256(margins2.json) = 342d88ac0ea31bc1b987f101a3f00a3c</code> <code>930c20932ebb024d0b7478853504a930</code>
Theorem 7.2	<code>calibrated/certificate_cal3.json / margins_cal3.json / verify6_cal.py</code> <code>SHA-256(certificate_cal3.json) = 900fdd5e5d9fb809465521c7b92cbc24</code> <code>a95a7d8cf8d1013dc2eea1dd1135a3bf</code> <code>SHA-256(margins_cal3.json) = 0e017269330a0232471c4ded8b2d0997</code> <code>f733b22f8e4953308bb8ea9b713ab051</code>

The margins files are the outputs of the end-to-end verification runs, with two informational keys renamed afterward for clarity; the archived drivers are the executed drivers with the package-root constant made relative, and the trusted modules are unmodified. Earlier certificates are retained in the archive for the audit trail. The subdirectory `calibrated/` contains the annual-calibration files (`certlib_cal.py`, `rowengine_cal.py`, constructors and verifiers for both the $[.1, 450]$ and the extended $[.016, 450]$ certificates): 711,536 rows verified under the identical trusted protocol. Both operative certificates were verified end-to-end under the shipped code: `verify5.py` on `certificate2.json` (all six families, including the U^+ and H -exit families) and `verify6_cal.py` on `certificate_cal3.json`; `margins2.json` and `margins_cal3.json` are the outputs of those runs.

Provenance. The trusted computing base is CPython’s `fractions` and `decimal` modules (libmpdec’s correctly rounded `ln`) together with the one-ulp guard described above; SHA-256 prefixes of the operative certificates are recorded in the package README. “Verifier-checked” throughout means checked by the standalone program described above over exact rational data; the theorems are not formalized in a mechanized proof system, and formalizing the bridge theorem and the ledger conditions in Lean or Isabelle is a natural further step. Conflict-of-interest and funding disclosures will accompany journal submission. An auditor can re-check any single row from the printed data with a hand computation of (6.5) at a handful of labor points.

E A worked ledger row and a worked policy

A ledger row. Take the B_M family at $A_+ = 10.04$ and the box $K \in [1.1 - h, 1.1]$, $K' \in [1.2, 1.2 + h']$. The corner reward bound needs $R_{A_+}(1.1, 1.2)$. Here $Y = K' - qK = 0.87$ and the cubic (6.4) has exact coefficients

$$(c_3, c_2, c_1, c_0) = (-143.1708, -900.08732, 757.81574, 3.47391),$$

with a unique root in $(0, \bar{L})$ at $L^* = 0.75610473\dots$, consumption $C = 7.19376813\dots$, and

$$R_{A_+}(1.1, 1.2) = 0.7 \log C + 0.3 \log(1 - L^*) = 0.95794566\dots$$

The verifier’s branch-and-bound returns the upper bound $0.958210\dots$ (within its $4 \cdot 10^{-4}$ tolerance) without computing L^* : it only evaluates (6.5) on subintervals. The ledger row then compares recorded node values of B_M against this bound plus β times the successor-cell maximum of B_M , plus ε_M .

A policy. At $A_- = 9.916$, resources at the bottom of $\mathcal{M}$ are $\phi_{A_-}(.098935) = 0.3434\dots$, so no one-step jump reaches H ; the recorded ladder ascends and enters $\mathcal{G}_H$ within at most three steps from every node of $\mathcal{G}_M$, after which π^H is followed. The exact re-evaluated values include

$$W^M(.098935) = -0.794210\dots, \quad W^H(2) = 4.396665\dots,$$

and the certified sandwich margin at the binding cell ($K = .098935$) is $0.0104632 \dots$ (Table 3). For orientation, the unrestricted steady state at $A = 10$ sits at $(K^*, L^*) \approx (5.06, 0.829)$ with consumption ≈ 16.8 , comfortably inside H .

References

- Ames, Aaron D., Samuel Coogan, Magnus Egerstedt, Gennaro Notomista, Koushil Sreenath, and Paulo Tabuada. 2019. Control barrier functions: Theory and applications. In *Proceedings of the 18th European Control Conference*, 3420–3431.
- Arrow, Kenneth J., Hollis B. Chenery, Bagicha S. Minhas, and Robert M. Solow. 1961. Capital-labor substitution and economic efficiency. *Review of Economics and Statistics* 43(3): 225–250.
- Askenazy, Philippe, and Cuong Le Van. 1999. A model of optimal growth strategy. *Journal of Economic Theory* 85(1): 24–51.
- Azariadis, Costas, and Allan Drazen. 1990. Threshold externalities in economic development. *Quarterly Journal of Economics* 105(2): 501–526.
- Azariadis, Costas, and John Stachurski. 2005. Poverty traps. In *Handbook of Economic Growth*, vol. 1A, ed. Philippe Aghion and Steven N. Durlauf, 295–384. Amsterdam: Elsevier.
- Balboni, Clare, Oriana Bandiera, Robin Burgess, Maitreesh Ghatak, and Anton Heil. 2022. Why do people stay poor? *Quarterly Journal of Economics* 137(2): 785–844.
- Benhabib, Jess, and Roger E. A. Farmer. 1994. Indeterminacy and increasing returns. *Journal of Economic Theory* 63(1): 19–41.
- Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan. 2007. Business cycle accounting. *Econometrica* 75(3): 781–836.
- Dechert, W. Davis, and Kazuo Nishimura. 1983. A complete characterization of optimal growth paths in an aggregated model with a non-concave production function. *Journal of Economic Theory* 31(2): 332–354.
- Edlin, Aaron S., and Chris Shannon. 1998. Strict monotonicity in comparative statics. *Journal of Economic Theory* 81(1): 201–219.
- Galor, Oded, and Joseph Zeira. 1993. Income distribution and macroeconomics. *Review of Economic Studies* 60(1): 35–52.
- Gonthier, Georges. 2008. Formal proof—the four-color theorem. *Notices of the American Mathematical Society* 55(11): 1382–1393.
- Hales, Thomas C. 2005. A proof of the Kepler conjecture. *Annals of Mathematics* 162(3): 1065–1185.
- Hsieh, Chang-Tai, and Peter J. Klenow. 2009. Misallocation and manufacturing TFP in China and India. *Quarterly Journal of Economics* 124(4): 1403–1448.
- Kamihigashi, Takashi, and Santanu Roy. 2007. A nonsmooth, nonconvex model of optimal growth. *Journal of Economic Theory* 132(1): 435–460.

- Kraay, Aart, and David McKenzie. 2014. Do poverty traps exist? Assessing the evidence. *Journal of Economic Perspectives* 28(3): 127–148.
- Krugman, Paul. 1991. History versus expectations. *Quarterly Journal of Economics* 106(2): 651–667.
- Kubler, Felix, and Karl Schmedders. 2005. Approximate versus exact equilibria in dynamic economies. *Econometrica* 73(4): 1205–1235.
- Lasserre, Jean B., Didier Henrion, Christophe Prieur, and Emmanuel Trélat. 2008. Nonlinear optimal control via occupation measures and LMI-relaxations. *SIAM Journal on Control and Optimization* 47(4): 1643–1666.
- Le Van, Cuong, and Lisa Morhaim. 2002. Optimal growth models with bounded or unbounded returns: A unifying approach. *Journal of Economic Theory* 105(1): 158–187.
- Majumdar, Mukul, and Tapan Mitra. 1982. Intertemporal allocation with a non-convex technology: The aggregative framework. *Journal of Economic Theory* 27(1): 101–136.
- Matsuyama, Kiminori. 1991. Increasing returns, industrialization, and indeterminacy of equilibrium. *Quarterly Journal of Economics* 106(2): 617–650.
- Milgrom, Paul, and Chris Shannon. 1994. Monotone comparative statics. *Econometrica* 62(1): 157–180.
- Murphy, Kevin M., Andrei Shleifer, and Robert W. Vishny. 1989. Industrialization and the big push. *Journal of Political Economy* 97(5): 1003–1026.
- Prajna, Stephen, Ali Jadbabaie, and George J. Pappas. 2007. A framework for worst-case and stochastic safety verification using barrier certificates. *IEEE Transactions on Automatic Control* 52(8): 1415–1428.
- Quah, John K.-H. 2007. The comparative statics of constrained optimization problems. *Econometrica* 75(2): 401–431.
- Restuccia, Diego, and Richard Rogerson. 2008. Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic Dynamics* 11(4): 707–720.
- Romer, Paul M. 1986. Increasing returns and long-run growth. *Journal of Political Economy* 94(5): 1002–1037.
- Rosenstein-Rodan, Paul N. 1943. Problems of industrialisation of Eastern and South-Eastern Europe. *Economic Journal* 53(210/211): 202–211.
- Santos, Manuel S. 2000. Accuracy of numerical solutions using the Euler equation residuals. *Econometrica* 68(6): 1377–1402.
- Santos, Manuel S., and Jesús Vigo-Aguiar. 1998. Analysis of a numerical dynamic programming algorithm applied to economic models. *Econometrica* 66(2): 409–426.
- Skiba, A. K. 1978. Optimal growth with a convex-concave production function. *Econometrica* 46(3): 527–539.

- Stokey, Nancy L., Robert E. Lucas, Jr., and Edward C. Prescott. 1989. *Recursive Methods in Dynamic Economics*. Harvard University Press.
- Topkis, Donald M. 1998. *Supermodularity and Complementarity*. Princeton University Press.
- Tucker, Warwick. 2002. A rigorous ODE solver and Smale's 14th problem. *Foundations of Computational Mathematics* 2(1): 53–117.