

Online Appendix

The Upside of Being Replaceable:
AI, Market Structure, and Growth

Oliver Pardo

September 2026

[Main paper](#) | [Code and data](#)

Contents

S1 Additional competitive-to-monopoly simulation results	2
S2 Monopoly slows growth	13
S2.1 Numerical construction and checks	17
S3 Additional results with bounded AI efficiency	18
S4 Equilibria without an efficiency bound	21
S5 Numerical implementation and accuracy checks	30
S5.1 Computational coordinates and initial guesses	30
S5.2 Residuals and horizon checks	30

S1 Additional competitive-to-monopoly simulation results

I supplement the $\sigma = 1.5$ comparison in Section 6.4 of the main paper with the four-elasticity exercises at both research productivities: $\sigma = 0.9, 1.0, 1.1, 1.5$ and $\chi = 7.5, 1.5$. Each economy starts on a competitive labor-bottleneck balanced growth path. At date zero, an unexpected permanent grant of exclusive AI rights introduces monopoly pricing and incentives for RSI. Continued competition provides the benchmark from the same initial stocks. I use the existing simulated trajectories; no additional equilibria are computed for this presentation. Appendix B.1 of the main paper documents the initial conditions and numerical checks.

The level figures retain the longer horizons omitted from the main-text figure, with windows of years $-2-10$, $10-100$, and $10-500$. The growth, price, and distributional figures below show two pre-event years and years $0-10$ above, with years $10-500$ below. Thin horizontal lines show continued competition; black dotted long-run lines show the analytical monopoly limits at $\sigma = 1.5$. The benchmark lines also distinguish the competitive income shares from the monopoly limits: competitive AI revenue pays for inference, whereas monopoly revenue also finances research and developer profit. The date-zero level changes are not finite growth rates: open and filled dots distinguish the values on either side of the event, and growth curves describe the smooth trajectory on each side separately.

Table S1: Competitive initial capital and the impact of exclusive AI rights

σ	K_0	Change in output	Change in wage
0.9	6.235	-32.4%	-36.6%
1.0	6.278	-17.6%	-17.6%
1.1	6.324	-13.7%	-11.9%
1.5	6.535	-10.4%	-5.4%

Notes: The capital stocks are derived from the competitive BGP, not selected independently. All other parameters and initial stocks, except capital, are in Table 2 of the main paper. Changes compare levels immediately after and before the grant. They are identical for both values of χ , since K and B cannot jump. The ratio $K_0/Y_0 = 3.3$ applies before the grant, not after output falls.

Level comparisons across elasticities. Figures S2 and S3 preserve the level comparisons for all four elasticities. Research can produce level gains even where it does not permanently raise growth. With $\chi = 7.5$, output first regains its competitive counterpart around year 18.2 at $\sigma = 1.1$ and year 46.9 at $\sigma = 1$. With $\chi = 1.5$, those recoveries occur around years 87.3 and 212.3. At $\sigma = 0.9$, neither output nor wages regain their competitive counterparts within 500 years under either χ . The initial AI-service price in that case is about 131 times its competitive level, a large distortion implied by the illustrative parameters rather than an empirical estimate. These recovery dates compare contemporaneous levels; they do not rank welfare.

Higher research productivity. Figures S4–S6 report $\chi = 7.5$. Unlike the pure RSI-activation experiment, the grant immediately reduces the return on capital: output falls while capital cannot jump, so equation (11) in the main paper implies a lower interest rate. The Euler equation (5) then

gives lower consumption growth immediately after the grant, even where the level of consumption jumps upward. Subsequent RSI and capital adjustment determine recovery. Research absorbs resources before its gains are fully realized, and the distribution panels separate developer net profit from expenditure on inference and research.

Lower research productivity. Figures S7–S9 report $\chi = 1.5$. The impact pricing distortion is unchanged, but the initial research allocation and efficiency growth are smaller. At $\sigma = 1.5$, the eventual rise in growth and decline in labor’s share emerge later. At year 500, growth and returns still exceed their monopoly limits, so the final displayed values should not be read as those limits.

Replication. The repository’s [competitive-transition guide](#) records the parameters, the institutional experiment, and the replication commands. The simulation script is `scripts/simulate_competitive_to_monopoly.py`; the figure-generation script is `scripts/report_competitive_to_monopoly.py`. The guide links to the separate results folders for each research productivity, containing the original paths, competitive benchmarks, impact changes, and numerical audits. The main-text $\sigma = 1.5$ comparison and the AI-revenue figure below are generated by `scripts/report_competitive_main_comparison.py`, which reads these same audited paths without changing them. Figure windows only select dates from those stored paths; they do not change the simulated economy or the numerical solution horizon.

AI-industry revenue at $\sigma = 1.5$. Figure S1 supplements the main-text level comparison with AI-industry revenue as a share of each economy’s output. The share falls from 8.6% to 5.3% at the grant under either χ . The markup reduces AI use enough to lower this share, despite the higher AI-service price. As research improves efficiency and AI use expands, the share rises, reaching 37.4% with $\chi = 7.5$ and 12.9% with $\chi = 1.5$ at year 50. Revenue is not net profit: it also pays for inference and research compute.

AI-industry revenue | $\sigma = 1.5$





Figure S1: From competition to monopoly with RSI at $\sigma = 1.5$: AI-industry revenue, $p_X X/Y$, as a percentage of each economy's output. Panels cover years -2 – 10 and 10 – 50 . The gray dashed line denotes continued competition; blue solid and gold dash-dotted lines denote monopoly with $\chi = 7.5$ and $\chi = 1.5$, respectively.


Levels relative to continued competition | $\chi = 7.5$



Top: 100.0% = continued competition. Middle/bottom: multiples (log scale); 1x = competition. Open/filled dots: before/after exclusive rights. Growth rates exclude the date-zero level jumps.

Figure S2: From competition to monopoly with RSI, $\chi = 7.5$: levels relative to continued competition for all four elasticities. The rows cover years $-2-10$, $10-100$, and $10-500$. The upper row uses percentages; the middle and lower rows use multiples on logarithmic scales. The horizontal benchmark is 100.0% or 1x. Open and filled dots mark levels before and after the grant. Vertical scales differ across windows and between the two research-productivity exercises.

Levels relative to continued competition | $\chi = 1.5$



Top: 100.0% = continued competition. Middle/bottom: multiples (log scale); 1x = competition.
 Open/filled dots: before/after exclusive rights. Growth rates exclude the date-zero level jumps.

Figure S3: From competition to monopoly with RSI, $\chi = 1.5$: output per worker, wages, and consumption per person relative to continued competition for all four elasticities. Windows, normalizations, and line styles match Figure S2. Vertical scales are fitted to this exercise.

Output, AI services, and capital growth | $\chi = 7.5$




Figure S4: Competition to monopoly, $\chi = 7.5$: annual growth of output, AI services, and capital relative to effective labor. Zero denotes growth at $n + \gamma$. Growth rates exclude the discrete level changes at date zero.

Wages, interest rates, and AI prices | $\chi = 7.5$





Figure S5: Competition to monopoly, $\chi = 7.5$: annual wage growth, the net interest rate, and AI-service prices. Prices are measured in final-good units on a logarithmic scale. The wage level falls at the grant; panel A reports wage growth on either side, not that level change.

Income and compute expenditure shares | $\chi = 7.5$



Thin lines: continued competition. Black dotted long-run line: monopoly limit for $\sigma = 1.5$.
Open/filled dots: before/after exclusive rights. Growth rates exclude the date-zero level jumps.

Figure S6: Competition to monopoly, $\chi = 7.5$: labor income, developer net profit, inference expenditure, and research expenditure as shares of output. The first two rows cover years -2 – 10 ; the last two cover years 10 – 500 . Gross capital income remains 33.0% of output.

Output, AI services, and capital growth | $\chi = 1.5$




Figure S7: Competition to monopoly, $\chi = 1.5$: annual growth of output, AI services, and capital relative to effective labor. Windows and benchmarks match Figure S4.





Figure S8: Competition to monopoly, $\chi = 1.5$: annual wage growth, the net interest rate, and AI-service prices in final-good units. Windows, benchmarks, and the logarithmic price scale match Figure S5.

Income and compute expenditure shares | $\chi = 1.5$



Thin lines: continued competition. Black dotted long-run line: monopoly limit for $\sigma = 1.5$.
Open/filled dots: before/after exclusive rights. Growth rates exclude the date-zero level jumps.

Figure S9: Competition to monopoly, $\chi = 1.5$: labor income, developer net profit, inference expenditure, and research expenditure relative to output. Layout and benchmarks match Figure S6.

S2 Monopoly slows growth

I now ask whether monopoly can eliminate AI-dominated growth even while research makes AI more efficient. This is the opposite outcome to the exercise in Section 6.4 of the main paper: research gains fail to compensate for the restriction on AI use caused by monopoly pricing. Following Corollary 1 of the main paper, I choose $\sigma = 1.5$ and $\mathcal{B}^{\text{mc}}(1.5) < B_0 < \bar{B} < \mathcal{B}(1.5)$. Initial efficiency is high enough for AI-dominated growth under competition, but even the upper bound is insufficient under monopoly. As in the main-text experiment, the unexpected grant covers existing technology and future improvements. RSI is available throughout; the grant changes pricing and research incentives. The preceding simulations remain separate experiments with different initial states and efficiency bounds.

Table S2 reports the new design. I retain the other parameters in Table 2 of the main paper and compare continued competition with monopoly under $\chi = 7.5$ and $\chi = 1.5$, all from the same initial stocks. Before granting monopoly rights, I let the competitive economy approach its AI-dominated growth path. I select the first whole year when output-per-worker growth, wage growth, and the interest rate are each within 0.1 percentage point of their long-run limits, and reset the date to zero. I inherit A_0 , N_0 , and K_0 jointly from that date, keeping B_0 and both values of χ unchanged. This is a near-long-run competitive allocation, not an exact steady state: $K/(AL)$ continues to grow. Just before the grant, output per worker grows at 5.8%, wages at 4.2%, and the interest rate is 9.8%. The capital-output ratio is 2.23. Section S2.1 below describes the additional competitive-transition solver and numerical checks.

Table S2: Parameters and initial stocks for the growth-reversal experiment

Object	Value	Choice
σ	1.5	Substitution
χ	7.5; 1.5	Same illustrative research productivities
$\bar{B}$	1,113.5	$0.90\mathcal{B}(1.5)$
B_0	1,002.2	$0.90\bar{B}$
A_0, N_0	6.6859; 1.7683	From the simulated competitive path
K_0	7.2907×10^9	Inherited from the same date
K_0/Y_{0-}	2.23	Competitive ratio before the grant

Notes: $\mathcal{B}^{\text{mc}}(1.5) = 829.0$ and $\mathcal{B}(1.5) = 1,237.3$. The upper bound and initial efficiency are chosen to satisfy the corollary's threshold inequalities. The three paths share the reported initial stocks; consumption and research are equilibrium choices. The initial stocks are obtained after 190 years of simulated competition; this duration is not a forecast of the grant date. Stocks retain their original units; these illustrative choices are not empirical estimates. Section S2.1 below describes the competitive path before the grant and its starting stocks.

Figure S10 shows annual output-per-worker growth, wage growth, and the net interest rate in panels A–C. The rows cover years -10 – 10 , 10 – 100 , and 100 – 500 ; negative dates show the common simulated competitive path before the grant. The dashed gray line denotes continued competition; the solid blue and dash-dotted gold lines denote monopoly with high and low research productivity. At the grant, output falls by 55.6%, wages by 33.5%, and AI services by 70.3%. These are instantaneous level changes at unchanged K and B , not annual growth rates in the figure. The

lower output–capital ratio reduces the interest rate from 9.8% to 1.6%.

After this level loss, high research productivity briefly raises output-per-worker growth to 13.4% as efficiency improves rapidly. With low research productivity, initial growth is only 0.1%. Neither response restores AI-dominated growth: by year 10, output per worker contracts under both monopoly paths. The inherited capital stock supported faster growth under competitive AI supply; restricting that supply lowers its return and induces capital decumulation. Consumption initially jumps upward by 141.6% and 132.2%, respectively, as lower compute use and investment release resources despite the fall in gross output. Consumption then falls: the interest rate initially lies below ρ , so the Euler equation implies negative consumption-per-person growth after the grant.

The longer windows show that the loss of AI-dominated growth is not confined to the initial repricing. Under continued competition, output-per-worker growth tends to 5.7%, wage growth to 4.1%, and the interest rate to 9.7%, as in equation (45) of the main paper. Under either monopoly path, the corresponding limits are 1.0%, 1.0%, and 5.0%, consistent with the labor-bottleneck regime in Proposition 1 of the main paper. Both monopoly economies are still adjusting at year 500: output-per-worker growth is about -0.1% , wage growth is 0.3% , and the interest rate is 3.9% . The large stock of capital inherited from the competitive expansion takes a long time to adjust to restricted AI supply. The monopoly limit therefore remains distinct from the last plotted date: monopoly removes permanently faster growth, not all long-run growth.

Figure S11 explains why efficiency gains do not prevent this outcome. Using the same lines and windows, panels A–C show $B/\bar{B}$, AI services relative to continued competition, and labor income as a share of output. In the first decade, high research productivity takes efficiency to 99.9% of its upper bound, compared with 97.1% under low research productivity. Both nevertheless supply fewer AI services than continued competition. In the longer run, efficiency approaches $\bar{B}$, an improvement of 11.1% over B_0 ; competition retains B_0 . Yet the monopoly markup prevents these improvements from sustaining the joint expansion of capital and AI services that occurs under competition. AI services under monopoly become negligible *relative to the competitive path*, not in absolute terms.

Labor’s income share tends to 4.5% under monopoly rather than zero under competition. This larger limiting share coexists with slower wage growth. At year 500, the labor share is still only 2.1% and 2.2% under high and low research productivity, respectively, confirming that neither path has reached its labor-bottleneck limit.

Together with Section 6.4 of the main paper, this experiment illustrates opposite regime rankings. There, research allows the monopoly economy with $\sigma = 1.5$ to escape the labor bottleneck from a low initial efficiency. Here, initial efficiency already sustains AI-dominated growth under competition, and further efficiency gains do not compensate for the restriction on AI use caused by monopoly pricing. More efficient AI is therefore compatible with slower long-run growth.

The large initial adjustments also reflect the abrupt withdrawal of access to existing technology. I use this experiment to illustrate a possible growth reversal, not to predict the effects of an actual monopoly grant or rank household welfare.




Figure S10: Monopoly and the loss of AI-dominated growth: annual output-per-worker growth, wage growth, and the net interest rate. Rows show years $-10-10$, $10-100$, and $100-500$. The grant occurs at zero, after the common simulated competitive path. Rates describe the smooth paths in each scenario, not the instantaneous output and wage losses. Gray dotted lines in the bottom row indicate theoretical long-run limits. Vertical scales differ across windows.




Figure S11: More efficient AI but slower growth under monopoly: efficiency relative to its upper bound, AI services relative to continued competition, and labor's income share. Windows and line styles match Figure S10. Panel B uses percentages in the first row and multiples on logarithmic scales in the lower rows; its competitive benchmark is 100.0% or one. Vertical scales differ across windows.

S2.1 Numerical construction and checks

I construct the competitive prehistory with $B = B_0$ above the competitive threshold. At the start of this prehistory, I normalize $A = N = 1$ and set $K = 14,047.3$, the limiting monopoly capital–effective-labor ratio at the chosen upper bound. These seed values are illustrative; they are not the stocks at the monopoly grant. I advance the competitive path to year 190, the first whole year at which the three rates in (45) are each within 0.1 percentage point of their limits; all subsequent annual observations through year 1,000 also satisfy this criterion. I then set that date to zero and inherit A , N , and K without changing units. Re-solving the competitive equilibrium from those stocks reproduces the continued prehistory path to within 10^{-9} in log levels over the next 500 years. The figures include its last ten pre-grant years.

I solve the competitive resource constraint and Euler equation in (A25), with $M = 0$ and $B = B_0$, using the inherited capital stock and the terminal condition $C/K \rightarrow \rho - n$ from (A27). The numerical coordinates are $\log(C/K)$ and $\log(K/(AL)) - [\mathcal{R}^{\text{mc}}(B_0) - \rho - \gamma]t$. Removing the limiting trend avoids extremely large levels without imposing balanced growth on the transition. I use the `scipy.integrate.solve_bvp` routine described in Appendix B of the main paper. Extending its horizon from 1,200 to 1,800 years changes these coordinates by less than 10^{-9} over the displayed 500 years; independent finite-difference residuals are below 10^{-9} . The static supply condition, positive allocation, limiting growth rates, and household transversality diagnostics also pass the checks.

I solve the two monopoly paths using the continuation method and equilibrium checks in Appendix B of the main paper. Their final solution horizon is about 4,592 years, although the figures display only the first 500 years. Both pass the equation, horizon-extension, transversality, and developer-optimality checks, including denser checks over the first decade. The [replication guide for this experiment](#) records the inputs, data, audit reports, and commands.

S3 Additional results with bounded AI efficiency

Remark S1 (Initial conditions and the efficiency bound). Suppose $\sigma < 1$ and $\delta > 0$. For any $\bar{B} < \infty$, including $\bar{B} > \mathcal{B}(\sigma)$, there is a nonempty open set of initial states from which no equilibrium can satisfy $B_t \rightarrow \bar{B}$. With sufficiently low initial AI efficiency and capital, production cannot replace depreciating capital and leaves insufficient resources for research to raise efficiency enough to reverse the capital decline. This obstacle persists under $\eta < \min\{\alpha, 1/\sigma\}$. The argument does not establish the existence of an alternative equilibrium.

Proof of Remark S1. I bound the improvement in AI efficiency that the economy can finance from low initial efficiency and capital. The bound uses only the static monopoly allocation, the resource constraint, and the RSI technology; it therefore applies to any equilibrium from those initial states.

For $\sigma < 1$, the CES aggregator (8) implies $Z \leq \omega_X^{\sigma/(\sigma-1)} X$. Also, $e_X \geq \alpha$ in (15), so (31) gives $X \leq B(1-\alpha)^2 Y$. Substituting both inequalities into production (7) and rearranging yields the following bound at every date:

$$\frac{Y_t}{K_t} \leq \left[(1-\alpha)^2 B_t \omega_X^{\sigma/(\sigma-1)} \right]^{(1-\alpha)/\alpha} = \frac{\mathcal{R}(B_t) + \delta}{\alpha}.$$

The right-hand side tends to zero as $B_t \downarrow 0$. Since $\delta > 0$, I can choose an efficiency level $b < \bar{B}$ such that

$$d \equiv \delta - \frac{\mathcal{R}(b) + \delta}{\alpha} > 0.$$

Choose $B_0 < b/2$ and K_0 small enough that

$$\chi b^\eta K_0^\eta \left(\frac{1-\eta}{d\eta} \right)^{1-\eta} < \frac{b}{2}.$$

These strict inequalities define a nonempty open set of initial states, with no further restriction on A_0 or N_0 .

As long as $B_t \leq b$, the output bound and resource constraint (26), with $C_t, U_t \geq 0$, imply $\dot{K}_t \leq -dK_t - M_t$. Multiplying by e^{dt} and integrating gives

$$e^{dT} K_T + \int_0^T e^{dt} M_t dt \leq K_0.$$

Since capital is nonnegative, Hölder's inequality then bounds the cumulative research input to the RSI technology:

$$\begin{aligned} \int_0^T M_t^\eta dt &= \int_0^T (e^{dt} M_t)^\eta e^{-d\eta t} dt \\ &\leq \left(\int_0^T e^{dt} M_t dt \right)^\eta \left(\int_0^T e^{-d\eta t/(1-\eta)} dt \right)^{1-\eta} \\ &\leq K_0^\eta \left(\frac{1-\eta}{d\eta} \right)^{1-\eta}. \end{aligned}$$

By the RSI technology (17), $\dot{B}_t \leq \chi b^\eta M_t^\eta$ while $B_t \leq b$. Hence

$$B_T - B_0 \leq \chi b^\eta K_0^\eta \left(\frac{1-\eta}{d\eta} \right)^{1-\eta} < \frac{b}{2}.$$

If T were the first date at which $B_T = b$, this inequality would contradict $b - B_0 > b/2$. Thus efficiency never reaches b , and the same bound applies at every horizon. Since B_t is nondecreasing, its limit exists and is strictly below $b < \bar{B}$. The argument holds for every admissible η and χ , including $\eta < \min\{\alpha, 1/\sigma\}$. $\square$

Proposition S1 (An AI bottleneck under complementarity). *Suppose $\sigma < 1$ and $\bar{B} < \mathcal{B}(\sigma)$. In any equilibrium, long-run average output-per-worker growth is bounded strictly below the rate of labor-augmenting technological progress:*

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log \left(\frac{Y_t/N_t}{Y_0/N_0} \right) \leq \mathcal{R}(\bar{B}) - \rho < \gamma.$$

AI production services become scarce relative to effective labor, and labor's income share vanishes:

$$\frac{X}{AL} \rightarrow 0, \quad s_X \rightarrow 1, \quad \frac{wL}{Y} \rightarrow 0.$$

Proof of Proposition S1. Consider any equilibrium with $0 < \sigma < 1$ and $0 < B_0 < \bar{B} < \mathcal{B}(\sigma)$. The increasing share map in equation (A9), established in Step 1 of the proof of Proposition 1, and $B_t \leq \bar{B}$ imply

$$r_t \leq \mathcal{R}(B_t) \leq \mathcal{R}(\bar{B}) < \rho + \gamma, \quad \frac{Y_t}{K_t} \leq \frac{\mathcal{R}(\bar{B}) + \delta}{\alpha}. \quad (\text{S3.1})$$

The present value of the developer's net profit from any date onward is nonnegative: the developer can stop research and earn nonnegative operating profit. Wage income is also nonnegative. Let $D_{t,s} = \exp(-\int_t^s r_u du)$. Integrating the household budget (4), using the household transversality condition (6) and the Euler equation (5), gives

$$\frac{C_t}{\rho - n} = K_t + \int_t^\infty D_{t,s}(w_s N_s + \Pi_s) ds.$$

The integral is nonnegative by the preceding argument, yielding the first inequality below. Integrating the Euler equation with the return bound gives the second:

$$C_t \geq (\rho - n)K_t, \quad \frac{C_t}{N_t} \leq \frac{C_0}{N_0} e^{[\mathcal{R}(\bar{B}) - \rho]t}. \quad (\text{S3.2})$$

Together with the output-capital bound in (S3.1), these consumption inequalities yield

$$\frac{Y_t}{N_t} \leq \frac{\mathcal{R}(\bar{B}) + \delta}{\alpha(\rho - n)} \frac{C_0}{N_0} e^{[\mathcal{R}(\bar{B}) - \rho]t}, \quad \limsup_{t \rightarrow \infty} \frac{1}{t} \log \left(\frac{Y_t/N_t}{Y_0/N_0} \right) \leq \mathcal{R}(\bar{B}) - \rho < \gamma. \quad (\text{S3.3})$$

Consequently $Y/(AL) \rightarrow 0$. Since equation (31) gives $X/Y \leq \bar{B}(1 - \alpha)^2$ under complementarity, $X/(AL) \rightarrow 0$: AI production services become scarce relative to effective labor. Equation (10) then implies $s_X \rightarrow 1$ and hence $wL/Y \rightarrow 0$. The growth bound concerns the long-run average, not convergence of the instantaneous growth rate. Neither convergence of C/Y nor $B_t \rightarrow \bar{B}$ is assumed. The growth bound and vanishing labor share are necessary properties of any equilibrium in this region, not a proof of equilibrium existence. $\square$

S4 Equilibria without an efficiency bound

Proof of Proposition 5. Dropping the positive AI term from the CES sum gives an upper bound for Z , because the outer exponent $\sigma/(\sigma - 1)$ is negative. This gives the production bounds in (C2) and (C3). Define $k = K/(AL)$. The resource constraint (28a), nonnegative consumption and compute expenditure, and $L = N$ imply

$$\dot{k} \leq \omega_L^{\sigma(1-\alpha)/(\sigma-1)} k^\alpha - (\delta + n + \gamma)k. \quad (\text{S4.1})$$

Since $\delta + n + \gamma > 0$ and $0 < \alpha < 1$, the right-hand side is negative whenever k exceeds a finite positive level. In particular,

$$k_t \leq \max \left\{ k_0, \left[\frac{\omega_L^{\sigma(1-\alpha)/(\sigma-1)}}{\delta + n + \gamma} \right]^{1/(1-\alpha)} \right\}.$$

Equation (C3) then bounds $Y/(AL)$ uniformly. Since $A_t = A_0 e^{\gamma t}$ and $L = N$, Y/N is bounded above by a constant times $e^{\gamma t}$. Taking logarithms, dividing by t , and taking the upper limit gives (C4). $\square$

Proof of Proposition 6. I first construct and verify the reference BGP, denoted by a superscript $*$, and then establish convergence from nearby initial stocks.

Step 1. Construction of the reference BGP. At unit elasticity, $s_X = \omega_X$ and $e_X = 1 - (1 - \alpha)\omega_X$. The AI-demand and monopoly conditions, equations (13) and (19), give

$$\frac{p_X X}{Y} = (1 - \alpha)\omega_X, \quad \frac{U}{Y} = (1 - \alpha)^2 \omega_X^2. \quad (\text{S4.2})$$

Substituting $X = BU = (1 - \alpha)^2 \omega_X^2 BY$ into production yields

$$Y^{1-(1-\alpha)\omega_X} = K^\alpha (AL)^{(1-\alpha)\omega_L} [(1 - \alpha)^2 \omega_X^2 B]^{(1-\alpha)\omega_X}. \quad (\text{S4.3})$$

On an exact BGP with positive constant K/Y , C/Y , and M/Y , log differentiation of this equation and the RSI technology (C1) gives the growth restrictions (C6) and (C7). Define

$$\Delta \equiv (1 - \alpha)(1 - \eta - \omega_X) > 0.$$

The unique solution to the two growth restrictions has

$$g_Y^* = \frac{(1 - \eta)(1 - \omega_X)}{1 - \eta - \omega_X} (n + \gamma), \quad (\text{S4.4})$$

and g_B^* as in (C12). The equality $g_K = g_C = g_Y^*$ follows from constant capital and consumption ratios; equations (12) and (5) then give (C9) and (C10). Since $n + \gamma > 0$, the two growth restrictions are inconsistent if $\Delta = 0$ and imply $g_B < 0$ if $\Delta < 0$, contrary to the RSI technology. Thus

$\eta + \omega_X \geq 1$ excludes this positive finite-rate BGP, not every possible equilibrium trajectory.

Define the candidate ratios $k^* = K/Y$, $u^* = U/Y$, $i^* = (\dot{K} + \delta K)/Y$, $m^* = M/Y$, and $c^* = C/Y$ by

$$k^* = \frac{\alpha}{r^* + \delta}, \quad u^* = (1 - \alpha)^2 \omega_X^2, \quad (\text{S4.5})$$

$$i^* = (g_Y^* + \delta)k^*, \quad (\text{S4.6})$$

$$m^* = \frac{u^* \eta g_B^*}{\rho - n + \eta g_Y^*}, \quad (\text{S4.7})$$

$$c^* = 1 - u^* - i^* - m^*. \quad (\text{S4.8})$$

Here k^* is a capital–output ratio, not capital per unit of effective labor, and i^* is gross physical investment as a share of output. The research share comes from the developer’s first-order and costate conditions; the consumption share comes from market clearing. These ratios are endogenous, not fixed saving or research rules.

All candidate ratios are positive. The denominator of the research share m^* , $\rho - n + \eta g_Y^*$, is positive. Moreover,

$$i^* = \frac{\alpha(g_Y^* + \delta)}{\rho - n + g_Y^* + \delta} < \alpha, \quad m^* < \frac{u^* \eta}{1 - \eta}.$$

Since $\Delta = (1 - \alpha)(1 - \eta - \omega_X) > 0$, $\omega_X < 1 - \eta$ and $\omega_X < 1$. Consequently,

$$u^* + m^* < \frac{u^*}{1 - \eta} < (1 - \alpha)\omega_X < 1 - \alpha, \quad c^* = 1 - u^* - i^* - m^* > 0.$$

Thus positive consumption follows from the stated conditions; it is not an additional assumption.

For given $A_0, N_0 > 0$, choose the reference initial stocks as

$$B_0^* = \left[A_0 N_0 (k^*)^{\alpha/[(1-\alpha)\omega_L]} (u^*)^{\omega_X/\omega_L} m^* \left(\frac{\chi}{g_B^*} \right)^{1/\eta} \right]^{\eta\omega_L/(1-\eta-\omega_X)}, \quad (\text{S4.9})$$

$$Y_0^* = A_0 N_0 (k^*)^{\alpha/[(1-\alpha)\omega_L]} (u^* B_0^*)^{\omega_X/\omega_L}, \quad K_0^* = k^* Y_0^*. \quad (\text{S4.10})$$

The outer exponent in the expression for B_0^* is positive. These choices solve simultaneously the initial production equation and $g_B^* B_0^* = \chi (B_0^* m^* Y_0^*)^\eta$. Set $U_0^* = u^* Y_0^*$, $M_0^* = m^* Y_0^*$, $C_0^* = c^* Y_0^*$, $X_0^* = B_0^* U_0^*$, and

$$q_0^* = \frac{M_0^*}{\eta g_B^* B_0^*}. \quad (\text{S4.11})$$

Let $Y_t^* = Y_0^* e^{g_Y^* t}$, $B_t^* = B_0^* e^{g_B^* t}$, $K_t^* = k^* Y_t^*$, $C_t^* = c^* Y_t^*$, $U_t^* = u^* Y_t^*$, $M_t^* = m^* Y_t^*$, and $q_t^* = q_0^* e^{(g_Y^* - g_B^*) t}$. Use the exogenous A_t, N_t and set $L_t = N_t$ and $X_t^* = B_t^* U_t^*$. The static equilibrium equations determine all remaining quantities and prices.

The two growth identities (C6) and (C7) ensure that the production equation and the RSI technology hold at every date once they hold initially. Equation (S4.11) gives $M = \eta q \dot{B}$, which is the research first-order condition (27k) with $\psi = 1$. The definitions of k^* , r^* , and c^* verify the

capital-return, Euler, and resource equations. The costate equation reduces to

$$g_Y^* - g_B^* = r^* - \frac{u^* \eta g_B^*}{m^*} - \eta g_B^*.$$

This holds by equation (S4.7) and $(1 - \eta)g_B^* = \eta g_Y^*$. All allocations are positive and finite at every finite date.

Both transversality conditions are explicit:

$$e^{-\rho t} \frac{N_t K_t^*}{C_t^*} = \frac{N_0 k^*}{c^*} e^{-(\rho-n)t} \longrightarrow 0, \quad (\text{S4.12})$$

$$e^{-r^* t} q_t^* B_t^* = q_0^* B_0^* e^{-(\rho-n)t} \longrightarrow 0. \quad (\text{S4.13})$$

Discounted developer profit is a constant multiple of $e^{-(\rho-n)t}$. Discounted household utility is $N_0 e^{-(\rho-n)t} [\log(C_0^*/N_0) + (g_Y^* - n)t]$. Both improper objective integrals are finite.

Step 2. Global optimality.

The constructed path satisfies the necessary conditions, but these do not yet rule out profitable deviations. For the developer, holding the paths of K, A, L fixed, operating profit is AI revenue less inference expenditure. Substituting the static shares (S4.2) and then eliminating output using (S4.3) gives

$$\begin{aligned} \pi_t(B) &= (1 - \alpha)\omega_X [1 - (1 - \alpha)\omega_X] [K_t^\alpha (A_t L_t)^{(1-\alpha)\omega_L}]^{1/[1-(1-\alpha)\omega_X]} \\ &\quad \times [(1 - \alpha)^2 \omega_X^2 B]^{(1-\alpha)\omega_X/[1-(1-\alpha)\omega_X]}. \end{aligned} \quad (\text{S4.14})$$

To verify developer optimality by concavity, define the invertible state transformation and its current-value costate by

$$\nu = \frac{1 - \eta}{\eta}, \quad S = B^\nu, \quad p = \frac{q}{\nu B^{\nu-1}}. \quad (\text{S4.15})$$

The exponent ν is chosen so that the transformed research technology has constant returns in S and M :

$$\dot{S} = \nu \chi S^{1-\eta} M^\eta,$$

which is jointly concave in S, M for every $0 < \eta < 1$. At given K, A, L , operating profit in equation (S4.14) is a positive time-dependent coefficient times

$$S^{\frac{(1-\alpha)\omega_X}{\nu[1-(1-\alpha)\omega_X]}}.$$

Since $\eta + \omega_X < 1$ implies $(1 - \alpha)\omega_X < 1 - \eta$, this exponent lies strictly between zero and one. Operating profit is therefore concave in S . With $p > 0$, the transformed Hamiltonian $\pi_t(S^{1/\nu}) - M + p_t \nu \chi S^{1-\eta} M^\eta$ is jointly concave in S, M . The research condition (22) and costate equation (23), with $\psi = 1$, give the transformed problem's control first-order condition and costate equation by the chain rule.

Write $D_t = \exp(-\int_0^t r_s ds)$ and $J_T = \int_0^T D_t[\pi_t(B_t) - M_t] dt$. For any admissible alternative with the same initial B , concavity, the research condition, and the costate equation give, after integration by parts,

$$\tilde{J}_T - J_T \leq -D_T p_T (\tilde{S}_T - S_T) \leq D_T p_T S_T = \frac{D_T q_T B_T}{\nu} \longrightarrow 0. \quad (\text{S4.16})$$

The last limit follows from the developer TVC. Since the transformation preserves feasible research paths and their payoffs, no admissible alternative has a greater limiting payoff, or a greater upper limit of truncated payoff, than this finite-valued candidate. The same verification applies to nearby candidates satisfying the uncapped equations and the developer TVC; it does not require exact balanced growth.

Household optimality follows from log-utility concavity, the linear budget, the Euler equation, and the TVC, by the same integration-by-parts argument as in Step 4 of the proof of Proposition 1.

Competitive final-good firms face a concave constant-returns technology, and Lemma 1 establishes global static monopoly optimality. The constructed path satisfies Definition 1 with the uncapped RSI technology (C1).

Step 3. Local stability and selection of initial consumption and research. The initial stocks are predetermined. I must select initial consumption and the shadow price so that the resulting path converges to the BGP; the research condition then determines initial research expenditure. For the given exogenous paths of population and labor-augmenting technology, compare each variable with its value on the reference BGP at the same date. Define log deviations

$$\xi_K = \log \frac{K}{K^*}, \quad \xi_B = \log \frac{B}{B^*}, \quad \xi_C = \log \frac{C}{C^*}, \quad \xi_q = \log \frac{q}{q^*}. \quad (\text{S4.17})$$

Time subscripts are suppressed: the reference quantities in the denominators follow the BGP rather than remain fixed at their initial values. Substitution into equations (28a)–(28d), with $\psi = 1$ and $\psi' = 0$, gives a smooth autonomous system $\dot{\xi} = f(\xi)$ with $f(0) = 0$, where $\xi = (\xi_K, \xi_B, \xi_C, \xi_q)'$. Let $J^* = Df(0)$ denote its Jacobian. No constant growth rates or allocation shares are imposed along these transitions.

I next establish two properties needed for this selection. First, the linearized system has exactly two independent directions of convergence and no eigenvalues with zero real part. Writing E^s for the real invariant subspace associated with the convergent directions, this means

$$\text{spec}(J^*) \cap i\mathbb{R} = \emptyset, \quad \dim E^s = 2. \quad (\text{S4.18})$$

Second, let $V_S \in \mathbb{R}^{4 \times 2}$ be any real basis matrix for E^s , and let V_{KB} consist of its first two rows. I show that

$$\det V_{KB} \neq 0. \quad (\text{S4.19})$$

This nonsingularity allows the stable manifold to select C_0, q_0 for every sufficiently close pair K_0, B_0 , rather than only for stocks on a curve. Both properties follow from the parameter restrictions, rather than being additional assumptions. I establish them by locating the stable eigenvalues from

the signs of the characteristic polynomial, then checking that their eigenvectors allow the two stocks to vary independently. Retain ν from (S4.15) and, for this linearization, define

$$\begin{aligned} a_K &= \frac{\alpha}{1 - (1 - \alpha)\omega_X}, & a_T &= \frac{(1 - \alpha)\omega_X\eta}{[1 - (1 - \alpha)\omega_X](1 - \eta)}, \\ d &= \rho - n, & P &= r^* + \delta, & h &= \rho - n + \eta g_Y^*. \end{aligned}$$

The quantities ν, a_K, a_T, d, P, h are positive, and the restriction $\eta + \omega_X < 1$ gives

$$1 - a_K - a_T = \frac{(1 - \alpha)(1 - \eta - \omega_X)}{[1 - (1 - \alpha)\omega_X](1 - \eta)} > 0. \quad (\text{S4.20})$$

For this calculation, return to the transformed state S and its shadow price p from Step 2. This gives the invertible linear change of deviations

$$\widehat{\xi} = (\xi_K, \nu\xi_B, \xi_C, \xi_q + (1 - \nu)\xi_B)'$$

The second and fourth coordinates are the log deviations of the transformed state S and its shadow price p . This changes neither the model nor the eigenvalues. In these coordinates, the reduced production equation and the research condition give

$$\log \frac{Y}{Y^*} = a_K \widehat{\xi}_K + a_T \widehat{\xi}_S, \quad \log \frac{M}{M^*} = \widehat{\xi}_S + \frac{\widehat{\xi}_p}{1 - \eta}.$$

Differentiating the capital, transformed research, Euler, and transformed costate equations at the BGP, using $U/(qB) = \rho - n + \eta g_Y^*$, gives a matrix similar to J^* :

$$\widehat{J} = \begin{pmatrix} A & B & -C & -V \\ 0 & 0 & 0 & D \\ -G & H & 0 & 0 \\ -E & F & 0 & d \end{pmatrix}, \quad (\text{S4.21})$$

where the following letters abbreviate matrix coefficients only:

$$\begin{aligned} A &= d + (1 - \alpha)\omega_X P, & B &= \frac{(1 - u^*)a_T - m^*}{k^*}, & C &= \frac{c^*}{k^*}, \\ V &= \frac{m^*}{k^*(1 - \eta)}, & D &= \frac{\eta g_Y^*}{1 - \eta}, & G &= P(1 - a_K), \\ H &= P a_T, & E &= G + h a_K, & F &= H + h(1 - a_T). \end{aligned}$$

All coefficients other than B are positive by their definitions. The BGP allocation ratios imply the

identities

$$\mathbf{L} \equiv \mathbf{B} - \frac{\mathbf{V}d}{\mathbf{D}} = \frac{(1 - \alpha)\omega_X\eta}{(1 - \eta)k^*} > 0, \quad (\text{S4.22})$$

$$\mathbf{G}\mathbf{F} - \mathbf{H}\mathbf{E} = hP(1 - a_K - a_T) > 0. \quad (\text{S4.23})$$

Set $\mathbf{N} = \mathbf{C}\mathbf{G} + \mathbf{V}\mathbf{E} > 0$ and $\mathbf{Q} = \mathbf{C}\mathbf{H} + \mathbf{V}\mathbf{F} > 0$. Expanding the determinant gives the characteristic polynomial

$$\begin{aligned} p(\ell) &= \det(\ell\mathbf{I} - \hat{\mathbf{J}}) \\ &= (\ell^2 - \mathbf{A}\ell - \mathbf{N})(\ell^2 - d\ell - \mathbf{D}\mathbf{F}) - \mathbf{D}\mathbf{E}(\mathbf{Q} - \mathbf{L}\ell). \end{aligned} \quad (\text{S4.24})$$

To locate the negative eigenvalues without solving the quartic, first evaluate the polynomial at zero:

$$p(0) = \mathbf{D}\mathbf{C}hP(1 - a_K - a_T) > 0.$$

Let $\mu = (d - \sqrt{d^2 + 4\mathbf{D}\mathbf{F}})/2 < 0$, the negative root of $\ell^2 - d\ell - \mathbf{D}\mathbf{F}$. At $\ell = \mu$, the product of the two quadratic factors in $p(\ell)$ vanishes, leaving $p(\mu) = -\mathbf{D}\mathbf{E}(\mathbf{Q} - \mathbf{L}\mu) < 0$. Since $p(\ell) \rightarrow +\infty$ as $\ell \rightarrow -\infty$, there is one negative root below μ and another between μ and zero. These are the only negative roots: the cubic coefficient of p is $-(\mathbf{A} + d) < 0$, its linear coefficient is $\mathbf{A}\mathbf{D}\mathbf{F} + d\mathbf{N} + \mathbf{L}\mathbf{D}\mathbf{E} > 0$, and its constant coefficient is positive. Thus the coefficient signs of $p(-x)$, in descending powers, are $(+, +, \text{arbitrary}, -, +)$. This sequence has exactly two sign changes, regardless of the quadratic coefficient. Descartes' rule of signs therefore permits at most two negative roots of p , counting multiplicity. The two distinct roots already located are consequently the only negative roots and are simple. Write them as $\ell_1 < \mu < \ell_2 < 0$. The remaining two roots have positive product and positive sum, because $p(0) > 0$ and the sum of all four roots is $\mathbf{A} + d > 0$. If real, both are positive; if complex, their real parts are positive. This proves (S4.18).

The eigenvalue count does not by itself ensure that a convergent path can be selected for every nearby pair of initial stocks. I establish this by showing that the stock components of the stable eigenvectors are linearly independent. Write a stable eigenvector in transformed coordinates as $(x, y, z, u)'$. The component y cannot be zero: the second, fourth, and first rows of the eigenvalue equation for $\hat{\mathbf{J}}$ would successively imply $u = x = z = 0$. Normalize $y = 1$. The second and fourth rows then give

$$x(\ell) = \frac{\mathbf{D}\mathbf{F} + d\ell - \ell^2}{\mathbf{D}\mathbf{E}}.$$

The ordering of the roots implies $x(\ell_1) < 0 < x(\ell_2)$. Hence the two projected stable eigenvectors $(x(\ell_1), 1)'$ and $(x(\ell_2), 1)'$ are linearly independent. Returning to (ξ_K, ξ_B) only rescales their second component by $1/\nu > 0$, so (S4.19) follows as well.

These matrix properties now provide a selection rule for the nonlinear system. The exact normalized vector field is continuously differentiable near $\xi = 0$. The spectral property just proved and the stable-manifold theorem supply a two-dimensional local manifold whose trajectories remain

near zero and converge exponentially. The derivative of its projection onto (ξ_K, ξ_B) is V_{KB} . The nonsingularity just proved and the inverse-function theorem therefore make this manifold locally the graph of a continuously differentiable function selecting (ξ_C, ξ_q) from the two initial stock deviations. Returning from log deviations to levels is a smooth invertible change of variables, so the same holds for (C_0, q_0) as functions of (K_0, B_0) .

Step 4. Reconstruction of nearby equilibria.

I first choose a neighborhood of initial states so that the selected trajectories stay within the region where the local construction is valid. For some constants $a, H > 0$, trajectories on a sufficiently small stable manifold satisfy

$$\|\xi(t)\| \leq H e^{-at} \|\xi(0)\|, \quad t \geq 0.$$

For any sufficiently small open neighborhood $\mathcal{V}$ of zero, choose a ball contained in $\mathcal{V}$. Continuity of the graph selecting the two jump variables permits an open neighborhood $\mathcal{U}$ of the initial stock deviations for which $H\|\xi(0)\|$ is smaller than that ball's radius. Every initial stock pair in $\mathcal{U}$ therefore has a trajectory staying in $\mathcal{V}$ for all $t \geq 0$ and converging to zero. Exponentiation maps $\mathcal{U}$ to an open neighborhood of (K_0^*, B_0^*) . Shrinking these neighborhoods preserves positivity, feasibility, and existence for every $t \geq 0$. Local uniqueness follows because every trajectory remaining in this local region and converging to zero belongs to the same stable manifold, which is a graph over stocks. This does not assert convergence from arbitrary initial C_0, q_0 or exclude equilibrium trajectories that leave the region.

The statement also permits nearby A_0, N_0 . Equations (S4.9) and (S4.10) make the reference B_0^*, K_0^* smooth functions of these two positive initial values. After normalization by the corresponding BGP, the exact vector field depends only on the parameters and BGP ratios, not on A_0, N_0 . The set of (A_0, N_0, K_0, B_0) whose stock deviations belong to $\mathcal{U}$ is therefore open and contains the reference initial condition. From each such initial condition, the construction converges to the BGP associated with its own A_0, N_0 , not necessarily to the same levels as the original reference path.

Since $f(\xi(t)) \rightarrow 0$, the growth rates of K, B, C, q converge to their BGP values, not only normalized levels. The static equations give the same convergence for the interest rate, wage growth, and allocation ratios. The unit-elastic labor income share remains $(1 - \alpha)\omega_L$ at every date.

Exponential convergence implies that $r_t - r^*$ is integrable. Thus D_t/e^{-r^*t} tends to a finite positive constant, and both TVC expressions equal their BGP counterparts in (S4.12) and (S4.13) times factors tending to finite positive limits. Both TVCs hold. Output and all expenditure flows are bounded multiples of $e^{g_Y^*t}$, while log consumption per person is a linear function of time plus a bounded deviation. Both objective integrals therefore remain finite. The global concavity verification established above applies at these new price and allocation paths, proving agent optimality. Static optimality and market clearing then establish equilibrium, not merely a solution of the necessary conditions. $\square$

Proof of Proposition 7. I construct a sequence of research trajectories whose discounted net profit grows without bound, although each trajectory has finite expenditure and AI efficiency. Write $\theta = (1 - \alpha)/\alpha$. For a research trajectory $\{M_t\}_{t \in [0, T]}$, discounted net profit on the fixed horizon is

$$J_T(M) = \int_0^T \exp \left\{ - \int_0^t r_s ds \right\} [\pi(B_t) - M_t] dt.$$

Integrability of r bounds the discount factor above and away from zero on $[0, T]$. At the optimal service choice, equations (13) and (19) imply

$$\pi(B) = p_X X - \frac{X}{B} = e_X p_X X = (1 - \alpha) e_X s_X Y.$$

For $\sigma > 1$, marginal revenue is positive at every finite X and approaches zero as $X \rightarrow \infty$, as established in Lemma 1. The monopoly condition therefore implies $X \rightarrow \infty$ as $B \rightarrow \infty$, holding K, A, L fixed. Consequently, $s_X \rightarrow 1$ and $e_X \rightarrow \alpha$. Substituting these limits into equation (C14) gives

$$\pi(B) \sim \kappa_X K B^\theta, \quad \kappa_X = \alpha(1 - \alpha) \left[(1 - \alpha)^2 \omega_X^{\sigma/(\sigma-1)} \right]^\theta > 0.$$

The operating-profit asymptotic is uniform over dates in $[0, T]$ because K, A, L are continuous, bounded, and bounded away from zero on $[0, T]$. Thus a common lower bound proportional to B^θ applies at every date once B is sufficiently large.

The research trajectories constructed below are deviations in the developer's optimization problem, not alternative equilibrium allocations. Fix $\Delta \in (0, T)$, choose $M = \mathcal{M}$ on $[0, \Delta]$, and choose $M = 0$ afterward. Integrating the uncapped RSI technology (C1) over the research burst gives

$$B_{\mathcal{M}}(t)^{1-\eta} = B_0^{1-\eta} + (1 - \eta) \chi \mathcal{M}^\eta \min\{t, \Delta\}.$$

As $\mathcal{M} \rightarrow \infty$, efficiency after the research burst is asymptotically proportional to $\mathcal{M}^{\eta/(1-\eta)}$. Discounted operating profit over the remaining interval therefore grows in proportion to $\mathcal{M}^{\eta\theta/(1-\eta)}$, while research expenditure grows only linearly in $\mathcal{M}$. Since optimized operating profit is nonnegative, a lower bound on discounted net profit counts only operating profits earned after research stops, while subtracting all research expenditure. For all sufficiently large $\mathcal{M}$,

$$J_T(\mathcal{M}) \geq c_\pi \mathcal{M}^{\eta(1-\alpha)/[\alpha(1-\eta)]} - c_M \mathcal{M}, \tag{S4.25}$$

where $c_\pi, c_M > 0$ are independent of $\mathcal{M}$. The exponent exceeds one because $\eta > \alpha$, so $J_T(\mathcal{M}) \rightarrow +\infty$ as $\mathcal{M} \rightarrow \infty$. Each finite $\mathcal{M}$ keeps AI efficiency and discounted net profit finite on $[0, T]$; it is the sequence of payoffs that is unbounded.

Now consider an infinite-horizon equilibrium candidate $\{\widehat{B}_t, \widehat{M}_t\}$ with finite discounted net profit $\widehat{V}_0$, and let $\widehat{J}_T$ denote its payoff through T . After T , continue the deviation using $\widehat{M}_t$. For sufficiently large $\mathcal{M}$, the deviation has $B_{\mathcal{M}}(T) > \widehat{B}_T$. Because $d(B^{1-\eta})/dt = (1 - \eta) \chi \mathcal{M}^\eta$, this efficiency advantage persists under the common subsequent expenditure. Operating profit is increasing in

efficiency by equation (20), so the continuation payoff is at least $\widehat{V}_0 - \widehat{J}_T$. The full deviation therefore yields

$$J_\infty(\mathcal{M}) \geq J_T(\mathcal{M}) + \widehat{V}_0 - \widehat{J}_T \longrightarrow +\infty.$$

No candidate with finite discounted net profit can be optimal. Moreover, every equilibrium must have finite discounted developer profit. Indeed, multiplying the household budget (4) by the discount factor and integrating, using the Euler equation (5) and transversality condition (6) gives

$$\int_0^\infty e^{-\int_0^t r_s ds} (w_t N_t + \Pi_t) dt = \frac{C_0}{\rho - n} - K_0.$$

Since labor income is nonnegative, discounted developer profit is either finite or minus infinity. A payoff of minus infinity cannot be optimal because stopping research yields nonnegative operating profit. Hence no equilibrium exists under the stated restrictions. This argument does not establish a finite-time singularity. $\square$

S5 Numerical implementation and accuracy checks

This section supplements the numerical method in Appendix B of the main paper. I describe the computational coordinates, the construction of numerical initial guesses, and the checks used to accept the simulated trajectories.

S5.1 Computational coordinates and initial guesses

For each σ , I first solve the existing-AI economy with efficiency fixed at B_0 and research unavailable. A scalar root search in $\log K_0$ solves Equation (49) using the production and monopoly equations. The resource constraint (26), with $M = 0$ and $\dot{K} = (n + \gamma)K$, gives pre-event consumption. The post-activation problem inherits the stocks but selects consumption, research, and the shadow price anew.

To represent positive variables and keep AI efficiency below its upper bound, I use the computational coordinates

$$\left(\log K, \log \frac{B}{\bar{B} - B}, \log C, \log q \right).$$

The second coordinate expresses AI efficiency relative to the remaining gap below its upper bound; it can grow without bound while AI efficiency remains bounded. I subtract the long-run growth rate of aggregate capital times t from each coordinate: $n + \gamma$ in the labor-bottleneck regime and $n + \mathcal{R}(\bar{B}) - \rho$ in the AI-dominated regime. The efficiency coordinate has the same asymptotic slope because the remaining gap $\bar{B} - B$ shrinks at that proportional rate. Removing these trends makes the numerical coordinates converge to finite limits without changing the economic variables or calendar time.

For each of the eight parameter configurations, I start near its analytical limiting point and move the initial stocks to their prescribed values through a sequence of auxiliary boundary-value problems, using each solution as the initial guess for the next. This procedure changes the initial conditions of the numerical problem, not the dates on the economic trajectory.

For the competitive-to-monopoly comparison in Section 6.4, I instead obtain K_0 from $r = \rho + \gamma$ with marginal-cost pricing $p_X = 1/B_0$; at $\sigma = 1.5$, $K_0 = 6.535$. Pre-event consumption follows from the same resource constraint. I then solve the monopoly BVP at these initial stocks, using the RSI-activation trajectories only as numerical initial guesses, not as restrictions on consumption or research. This procedure also applies to the four-elasticity comparisons in the [digital edition](#).

S5.2 Residuals and horizon checks

To check whether the computed transition depends on the terminal date or solver accuracy, I extend each computational horizon twice by 500 years and tighten the collocation tolerance from 2×10^{-6} to 10^{-8} and 10^{-9} . The final boundary tolerance is 10^{-11} . All plotted values lie within the solved interval; none is extrapolated beyond it.

I check the four dynamic equations independently of the solver's residual report. Five-point finite

differences approximate the time derivatives of the four logarithmic coordinates above; I compare these derivatives with the equations' right-hand sides expressed in the same coordinates. The dynamic residual is the largest absolute discrepancy across coordinates and checked dates, measured per year, with denser checks over the first ten years. For the monopoly and research first-order conditions, I report the absolute logarithmic deviations of $Bp_X(1 - e_X)$ and $q\chi\eta B^\eta M^{\eta-1}\psi(B)$ from one. Acceptance requires dynamic residuals below 10^{-6} and both first-order residuals below 10^{-9} .

I also require the last horizon extension to change the four detrended coordinates by less than 2×10^{-5} over years 0–500, with the same tolerance for changes in $\log C_0$ and $\log q_0$. At the terminal date, the largest coordinate deviation from the analytical limiting point must be below 10^{-4} . This last check uses the regime-specific normalized coordinates of Appendix A.2, with logarithms for coordinates having positive limits and with $\tau = (AN)^{-1}$ scaled by $\bar{d}/\bar{B}$. Here $\bar{d}$ is the limiting value of the normalized efficiency gap d in the corresponding regime. Unlike the horizon comparison, this check measures proximity to the limiting equilibrium, not agreement between two numerical solutions. I also check the pre-event equilibrium conditions and the continuity of production quantities and prices at RSI activation.

Small equation residuals do not establish that the developer prefers the computed research path to every feasible alternative. I therefore check feasibility and the sufficient conditions for global optimality in Lemma 2. The parameter condition (A2) is verified in Appendix B.2 of the main paper for all eight configurations.

To assess the behavior of the paths beyond the computational horizon, I check proximity to the analytical limiting equilibrium at the terminal date and use the asymptotic decay of both transversality expressions at rate $n - \rho < 0$. These checks support the numerical approximation over an infinite horizon; they do not constitute a proof of equilibrium existence from arbitrary initial stocks. The competitive-to-monopoly paths satisfy the same residual, horizon-refinement, transversality, and developer-optimality criteria.

Table S3: Numerical accuracy of the illustrative RSI comparisons

χ	σ	Final horizon	Dynamic residual	Horizon change
7.5	0.90	5,428.0	1.18×10^{-9}	3.06×10^{-8}
7.5	1.00	4,944.5	1.38×10^{-9}	3.46×10^{-8}
7.5	1.10	4,828.8	1.66×10^{-9}	3.73×10^{-8}
7.5	1.50	3,013.0	1.69×10^{-9}	1.56×10^{-8}
1.5	0.90	6,047.0	1.02×10^{-9}	4.30×10^{-8}
1.5	1.00	5,563.5	1.11×10^{-9}	3.80×10^{-8}
1.5	1.10	5,447.8	1.19×10^{-9}	4.08×10^{-8}
1.5	1.50	3,632.0	1.23×10^{-9}	1.80×10^{-8}

Horizons are measured in years. Dynamic residuals are the largest absolute errors in the derivatives of the four logarithmic coordinates, across the full-horizon and dense first-decade checks. Horizon change compares the last two solutions in detrended logarithmic coordinates over years 0–500. Reported error maxima are rounded upward. All eight paths pass the equation, event-continuity, developer-optimality and long-run-continuation checks. All eight configurations satisfy the sufficient parameter condition in Lemma 2.