The Self-Distilling Product: Limit Pricing, Experience Curves, and Endogenous Knowledge Diffusion in Frontier AI

Leon Shi*

July 2026

Abstract

Frontier AI prices look like two classical mechanisms running at once. Each lab prices just under the viability floor of the next-best model, which is limit pricing in the sense of Bain [1956], Sylos-Labini [1962] and Modigliani [1958]; and each prices below its own current cost to buy volume, which is experience-curve pricing in the sense of Wright [1936], Arrow [1962] and Spence [1981]. Stacked, these two motives are Cabral and Riordan [1994]'s increasing dominance. We add one term that no classical industry had. In AI the leader's shipped output is the channel through which rivals learn: distillation, published technique, and open-weight replication run on the frontier model's own deployment. Knowledge diffusion is therefore endogenous to the leader's own volume, $\delta = \delta(D_L)$, $\delta' > 0$. A product that trains its own competitors has no classical analogue in industrial organization, and the term changes both the pricing rule and the predicted market structure. Three results follow. First, in steady state the leader's extractable rent equals its rate of cost descent divided by the leakage rate, $G^\infty = \gamma D/\delta$: rent is proportional to velocity, not to accumulated position, and decays with half-life $\ln 2/\delta$ once descent stops. Second, the optimal price carries a positive leak premium, $p^* = c_L - \lambda_{\text{learn}} + \mu_{\text{leak}}$, which rationalises withheld frontier models, rate limits, hidden chains of thought, suppressed logprobs, and contractual bans on training against outputs as a single family of δ-suppression behaviours. Third, when the mechanism is restated as a Markov perfect equilibrium in the style of Ericson and Pakes [1995], it nests Sutton [1991] but separates from it on an observable: Sutton predicts persistent concentration with stable leaders, while endogenous diffusion predicts persistent concentration with leader churn. The empirical strategy measures the gap in capability space, where benchmark-parity lag identifies δ without cost data.

1 Introduction

The claim this paper formalises can be stated in one sentence, and it was originally stated informally: AI will always be as cheap as the next-best competitor allows it to be, and AI is simultaneously an accelerant for breakaway winners who compound efficiency and capability advantages and price so as to extract the largest gap the next competitor's threshold permits.

Read carefully, that sentence contains two distinct economic mechanisms and one tension between them. The first mechanism is a boundary condition on price: the leader's price is pinned from above by somebody else's cost. The second is a law of motion for cost: the leader's own cost

*Working draft prepared as a coauthor handoff. Comments welcome.

falls in its own cumulative volume, so the boundary condition itself moves. The tension is that both mechanisms are about the gap between the leader and the marginal producer, and they push that gap in opposite directions once we notice that in this industry the leader's volume is also what teaches the marginal producer.

What is already known. The two halves are each classical, and the combination has been studied. Limit pricing gives us the boundary condition: an incumbent with a cost advantage does not charge the monopoly price, it charges just below the entry threshold of the next-best potential producer and keeps the difference [Bain, 1956, Sylos-Labini, 1962, Modigliani, 1958]. The margin it keeps is Ricardian rent, set by the marginal producer rather than by the leader's own absolute efficiency. Experience-curve pricing gives us the law of motion: unit cost falls in cumulative output [Wright, 1936, Arrow, 1962], so price is partly an investment in future cost, and the optimal price sits below current marginal cost by the shadow value of an extra unit of experience [Spence, 1981, Fudenberg and Tirole, 1983]. Put them together and you get increasing dominance: the firm that is ahead prices more aggressively *because* it is ahead, wins more volume, and pulls further ahead [Cabral and Riordan, 1994, Dasgupta and Stiglitz, 1988].

What is new. In every classical setting, including Cabral and Riordan [1994], the follower's cost curve is exogenous to the leader's behaviour except through the volume the leader denies it. Alcoa's rivals had to learn to smelt aluminium themselves. Texas Instruments' rivals had to learn to fabricate calculators themselves. In frontier AI this is false in a specific and consequential way: the leader's deployed output is itself the training signal by which rivals close the gap. Distillation of frontier outputs, replication of published technique, and open-weight models trained on frontier-generated data all mean that

$$\text{follower's rate of catch-up} \;=\; \delta\big(D_L\big), \qquad \delta' > 0,$$

where D_L is the leader's own deployment volume. Serving the model is the leak. The leader therefore faces a novel trade-off that has no counterpart in the learning curve literature: the same volume that drives its own cost down also accelerates the erosion of the gap that makes the volume profitable.

Roadmap. Section 2 fixes notation. Section 3 derives the two classical building blocks from first principles, in full, so that a reader with no industrial organization background can follow the rest. Section 4 builds the reduced-form gap dynamics and proves the rent-on-velocity result. Section 5 introduces endogenous diffusion and derives the pricing rule with its leak premium, together with the regime structure over a product generation. Section 6 states the criticality condition that separates a treadmill from a breakaway and shows what has to be true empirically for each. Section 7 restates the whole mechanism as a Markov perfect equilibrium, which is the form a theory seminar will demand, and shows which reduced-form objects survive as derived limits. Section 8 addresses the channel through which the flywheel actually runs, which may be the capital market rather than the factory. Section 9 is the empirical strategy, and it is deliberately modest about what is identified. Section 10 lists the testable predictions and the single test that separates this model from its closest existing rival. Section 11 states what is scoped out and why. Appendix A proves that the reduced form of Section 4 is the limit of the full game under a myopic follower, so nothing in the simple version is asserted rather than derived.

2 Notation and primitives

Two firms, a leader L and a follower F, plus a competitive fringe Φ whose cost is dragged down by diffusion from the frontier. Time is continuous in Sections 4–6 for transparency and discrete in Section 7, where the equilibrium concept requires it.

Symbol	Meaning
c_i	firm i's unit cost of serving one unit of inference
ω_i	firm i's knowledge stock (capability), entering demand as quality
G	gap, $G \equiv c_F - c_L$ in cost space or $\omega_L - \omega_F$ in capability space
p_i	firm i's price
$D_i(p_i, p_j; \omega_i, \omega_j)$	demand (deployment volume) for firm i
Q_i	cumulative volume, $\dot{Q}_i = D_i$
γ	learning intensity: cost decline per unit of own volume
ϕ	research productivity: capability gain per unit of R&D spend r_i
$\delta(\cdot)$	diffusion (leak) rate at which the follower's cost is pulled toward the leader's
κ	sunk cost of entry
b_i	firm i's cash or financing capacity
r	discount rate
λ_{learn}	shadow value of an extra unit of own volume through learning
μ_{leak}	shadow cost of an extra unit of own volume through leakage
β	elasticity of gap growth in the gap itself (returns to being ahead)

Table 1: Notation.

Assumption 1 (Learning)**.** Unit cost falls in cumulative own volume: $c_i = c_i(Q_i)$ with $c_i' = -\gamma < 0$. Equivalently $\dot{c}_i = -\gamma D_i$. This is Wright's law in differential form [Wright, 1936].

Assumption 2 (Diffusion)**.** The follower's law of motion for cost contains a term pulling it toward the leader's at a rate proportional to the gap, $-\delta\,(c_F - c_L)$, in addition to any effect of its own volume. This is the standard knowledge-diffusion specification and, mechanically, Newton's law of cooling.

Assumption 3 (Endogenous diffusion, the paper's contribution)**.** The diffusion rate is increasing in the leader's own deployment: $\delta = \delta(D_L)$ with $\delta' > 0$ and $\delta(0) = \underline{\delta} > 0$. The intercept $\underline{\delta}$ captures leakage that does not require the leader's traffic (published papers, researcher mobility, independent replication); the slope captures leakage that does (distillation of served outputs, revealed chains of thought, logprob-based imitation).

Assumption 3 is the whole paper. Everything else is either classical or a consequence.

3 The two classical building blocks, derived

This section is for the reader who has not seen either mechanism written down. Both are short derivations. Readers who know Spence [1981] and Bain [1956] can skip to Section 4.

3.1 Limit pricing: why your price is set by somebody else's cost

Take demand $D(p)$, a leader with unit cost c_L, and a potential entrant who can produce at unit cost $c_F > c_L$ but must sink κ to enter. The entrant enters if and only if it expects to recover κ, which requires an expected price above c_F for long enough. Suppose, as in the classical treatment, that the entrant believes the leader will hold its pre-entry price (the Sylos postulate). Then entry is unprofitable whenever

$$p \;\leq\; c_F + \frac{r\kappa}{D(p)} \;\equiv\; \bar{p}, \tag{1}$$

where the second term annualises the sunk cost over the volume the entrant would serve. The leader's problem is then

$$\max_p \,(p - c_L)\, D(p) \quad \text{subject to} \quad p \leq \bar{p},$$

which has two regimes. If the unconstrained monopoly price p^m satisfies $p^m \leq \bar{p}$, the constraint does not bind and the leader is a plain monopolist: its advantage is large enough that nobody could enter even at the monopoly price. If $p^m > \bar{p}$, the constraint binds and the solution is the corner

$$p^* \;=\; \bar{p} \;\approx\; c_F - \epsilon, \tag{2}$$

with per-unit margin

$$p^* - c_L \;=\; (c_F - c_L) + \frac{r\kappa}{D} \;=\; G + \frac{r\kappa}{D}. \tag{3}$$

Equation (3) is the entire economics of the boundary condition. The leader's margin is the *gap*, not its own efficiency. Halve the leader's cost with the follower's cost unchanged and the margin doubles; halve both and the margin is unchanged. This is Ricardian rent: the price is set by the marginal producer, and the inframarginal producer collects the difference. Baumol et al. [1982] pushed the logic to its limit by arguing that with costless entry and exit even a literal monopolist prices at $\bar{p}$, since the threat need never be executed to bind. We will not use contestability; Section 7 explains why and what replaces it.

Plain-language version. You are not paid for being good. You are paid for being better than whoever is second, and only up to the point where second place would find it worth trying.

3.2 Experience-curve pricing: why price is an investment

Now drop the rival and add learning. The firm's cost falls in cumulative output, $\dot{c} = -\gamma D(p)$, and it maximises the present value of profit

$$V = \int_0^\infty e^{-rt}\,\big(p(t) - c(t)\big)\, D\big(p(t)\big)\, dt \quad \text{subject to} \quad \dot{c} = -\gamma D(p),\; c(0) = c_0.$$

Form the Hamiltonian with costate λ on the state c,

$$\mathcal{H} = (p - c)D(p) \;+\; \lambda \cdot \big(-\gamma D(p)\big),$$

and note that since lower cost is better, the costate is negative; write $\lambda_{\text{learn}} \equiv -\lambda\gamma > 0$ for the shadow value of one extra unit of volume. The first-order condition $\partial\mathcal{H}/\partial p = 0$ is

$$D(p) + (p - c)D'(p) + \lambda_{\text{learn}} \frac{D'(p)}{\gamma} \cdot \gamma = 0,$$

which rearranges to the standard markup form

$$p = \underbrace{c - \lambda_{\text{learn}}}_{\text{effective marginal cost}} + \underbrace{\frac{D(p)}{-D'(p)}}_{\text{markup}}. \tag{4}$$

The content of (4) is that the relevant marginal cost is not the accounting cost c but $c - \lambda_{\text{learn}}$, because serving a unit today lowers the cost of every unit served thereafter. When the shadow value is large relative to the markup, the optimal price lies *below* accounting cost. That is not irrationality and it is not predation in the legal sense; it is investment [Spence, 1981]. The costate itself evolves as $\dot{\lambda} = r\lambda + \partial\mathcal{H}/\partial c$, and since $\partial\mathcal{H}/\partial c = -D$, the shadow value is large early (long remaining horizon, steep curve) and decays as learning is exhausted.

Plain-language version. Early on you are not selling inference, you are buying cost reduction, and the price tag is the margin you give up.

3.3 Stacking them: increasing dominance

Now put the rival back. Cabral and Riordan [1994] show that once both firms learn, being ahead is self-reinforcing: the leader's shadow value of volume is higher precisely because it is closer to the flat part of nobody-catches-me, so it prices more aggressively, wins more volume, and widens the gap. The two motives are not additive but superadditive, and it is worth seeing why in one line. Pricing below the follower's floor pays twice: it lowers the leader's own cost through Assumption 1, *and* it starves the follower of the volume it would need to descend its own curve. Writing $V_L(c_L, c_F)$ for the leader's value function, the claim is that the cross-partial is signed,

$$\frac{\partial^2 V_L}{\partial c_L\, \partial c_F} \neq 0, \tag{5}$$

so the marginal return to own cost reduction depends on the rival's cost. The same dollar of foregone margin buys descent and stagnation at once. This is also why the intuitive alternative strategy, harvest at monopoly prices now and undercut the entrant later, is worse than it looks: during the harvest the rival accumulates both volume-learning and, in AI, time-based diffusion, so the price you must eventually meet is lower than the limit price you declined to set.

Remark 1 (Why waiting is not merely riskier but dominated here). Under classical assumptions the comparison between deterrence and harvest-then-fight is ambiguous, and Selten [1978]'s chain store paradox shows the fight is not even credible under complete information; predation is rescued only by reputation under incomplete information [Kreps and Wilson, 1982, Milgrom and Roberts, 1982], and McGee [1958] found Standard Oil mostly bought rivals rather than predating on them. Three features of frontier AI break every condition that makes waiting attractive. Entrants are a renewable resource rather than a one-shot threat, so the "restored monopoly" branch collapses. The follower's cost falls during the harvest even at zero volume, by Assumption 3, so delay raises the cost of the eventual fight monotonically. And a defeated rival's capacity does not exit: weights persist, get open-sourced on the way down, or are acqui-hired, so winning does not remove supply and may permanently lower the industry price floor. The AI industry independently rediscovered McGee's finding, absorbing rivals rather than undercutting them.

4 Reduced-form gap dynamics: rent is a velocity

Combine Assumptions 1 and 2 for the gap $G \equiv c_F - c_L$. The leader's cost falls with its own volume; the follower's cost falls both with diffusion and with its own volume, which under binding limit pricing is approximately zero because the leader holds price at the follower's floor. Then

$$\dot{c}_L = -\gamma\, D_L, \tag{6}$$

$$\dot{c}_F = -\delta\,(c_F - c_L), \tag{7}$$

and subtracting (6) from (7) gives the central law of motion

$$\boxed{\dot{G} \;=\; \underbrace{\gamma\, D_L(p)}_{\text{own learning}} \;-\; \underbrace{\delta\, G}_{\text{leak}}} \tag{8}$$

Equation (8) is a first-order linear ODE. For constant D_L its solution from $G(0) = G_0$ is

$$G(t) \;=\; \frac{\gamma D_L}{\delta} \;+\; \Big(G_0 - \frac{\gamma D_L}{\delta}\Big)e^{-\delta t}, \tag{9}$$

which converges monotonically to the unique globally stable steady state

$$\boxed{G^{\infty} \;=\; \frac{\gamma D_L}{\delta}} \tag{10}$$

Now recall from (3) that under binding limit pricing the leader's per-unit margin *is* the gap. So (10) is a statement about rent.

Proposition 1 (Rent on velocity)**.** *Under Assumptions 1–2 with binding limit pricing, the leader's steady-state per-unit rent equals its rate of cost descent divided by the leakage rate,*

$$\textit{rent}^{\infty} \;=\; \frac{\gamma D_L}{\delta} \;=\; \frac{\textit{rate of descent}}{\textit{rate of leakage}}.$$

Rent is therefore proportional to the leader's velocity *along the cost curve, not to its accumulated* position*. If descent halts at time τ ($D_L\gamma \to 0$), the gap and hence the rent decay exponentially, $G(t) = G(\tau)e^{-\delta(t-\tau)}$, with half-life*

$$t_{1/2} \;=\; \frac{\ln 2}{\delta}.$$

Proof. Immediate from (9) and (3): setting $\dot{G} = 0$ in (8) yields (10), and setting $D_L\gamma = 0$ leaves $\dot{G} = -\delta G$, whose solution is the stated exponential with the half-life obtained from $e^{-\delta t_{1/2}} = 1/2$. □

Remark 2 (Why this matters for interpreting the industry)*.* Proposition 1 says a stock of accumulated capability is worth nothing in itself. A lab with an enormous historical lead that stops improving earns the same rent as a lab that never led, after roughly $\ln 2/\delta$ of drift. The observable counterpart is that margin exists only on the current frontier tier, and every capability tier's price collapses toward cost within months of being superseded. It also implies the capital expenditure treadmill is not a strategic choice but the equilibrium requirement for holding any rent at all.

Corollary 1 (Comparative statics)**.** $\partial G^{\infty}/\partial\gamma > 0$*: steeper learning, larger rent.* $\partial G^{\infty}/\partial\delta < 0$*: faster diffusion, smaller rent, and the elasticity is unit-negative, so a doubling of leakage halves steady-state rent.* $\partial G^{\infty}/\partial D_L > 0$ holding δ fixed*, which is exactly the assumption Section 5 removes.*

5 Endogenous diffusion and the leak premium

Assumption 3 makes δ a function of the leader's own deployment. The steady state becomes an implicit relation rather than a formula,

$$G^\infty \;=\; \frac{\gamma\, D_L}{\delta(D_L)}, \tag{11}$$

and the sign of $\partial G^\infty/\partial D_L$ is no longer unambiguous. Writing the elasticity of leakage with respect to volume as $\eta \equiv \delta'(D_L)\, D_L/\delta(D_L)$,

$$\frac{\partial G^\infty}{\partial D_L} \;=\; \frac{\gamma}{\delta(D_L)}\,\bigl(1-\eta\bigr), \tag{12}$$

so more deployment widens the gap if and only if $\eta < 1$: leakage must be inelastic in volume. If $\eta > 1$, growth is self-defeating in the sense that the leader teaches faster than it learns, and there is an interior volume that maximises steady-state rent. This is the first substantive prediction that no classical model can generate.

5.1 The pricing rule

Solve the leader's problem with both states. The leader chooses p to maximise

$$V_L = \int_0^\infty e^{-rt}\bigl(p - c_L\bigr) D_L(p)\, dt$$

subject to $\dot c_L = -\gamma D_L(p)$ and $\dot c_F = -\delta\bigl(D_L(p)\bigr)\,(c_F - c_L)$, with the limit-pricing constraint $p \le \min\{p^m, c_F - \epsilon\}$ from Section 3. Attach costates λ to c_L and ν to c_F. Since a lower own cost is good, $\lambda < 0$; since a lower rival cost is bad, $\nu < 0$ as well, but it enters with the opposite sign because δ multiplies a term the leader wants to keep large. Define

$$\lambda_{\text{learn}} \;\equiv\; -\lambda\gamma \;>\; 0, \qquad \mu_{\text{leak}} \;\equiv\; -\nu\, \delta'\bigl(D_L\bigr)\,\bigl(c_F - c_L\bigr) \;>\; 0.$$

The Hamiltonian is

$$\mathcal{H} = (p - c_L) D_L(p) - \lambda\gamma D_L(p) - \nu\, \delta\bigl(D_L(p)\bigr)\bigl(c_F - c_L\bigr),$$

and $\partial \mathcal{H}/\partial p = 0$ gives, after collecting terms in $D_L'(p)$,

$$\boxed{p^* \;=\; c_L \;-\; \lambda_{\text{learn}} \;+\; \mu_{\text{leak}} \;+\; \frac{D_L}{-D_L'}, \qquad p^* \le \min\bigl\{p^m,\; c_F - \epsilon\bigr\}} \tag{13}$$

Proposition 2 (Leak premium)**.** *Under Assumption 3, the optimal price exceeds the pure Spence price by $\mu_{leak} > 0$. The leader optimally serves less volume than the learning motive alone would dictate, and the wedge is increasing in the elasticity η of leakage with respect to volume and in the current gap $c_F - c_L$.*

Proof. Set $\delta' = 0$ in the Hamiltonian: the third term is then independent of p, the first-order condition reduces to (4), and $\mu_{\text{leak}} = 0$. For $\delta' > 0$ the term is strictly positive because $\nu < 0$ and $c_F > c_L$ on the relevant path, and it enters (13) additively. $\square$

Corollary 2 (Non-price δ-suppression)**.** *If the leader has instruments that reduce δ at a cost lower than the margin sacrificed by raising price, it uses them first. Any action that lowers the informativeness of served output per unit of volume raises V_L at fixed D_L.*

Corollary 2 is where the model earns its keep, because it explains a cluster of frontier-lab behaviour that has no other unified account: withholding the strongest internal model from the API, rate limits on high-capability endpoints, hiding or summarising chains of thought, refusing to expose token log-probabilities, and contractual prohibitions on training competing models against outputs. In this model these are not product decisions, safety decisions, or capacity decisions. They are all the same decision: buy a lower δ without giving up the volume that drives γD_L.

5.2 Regime structure over a product generation

The two shadow values move in opposite directions over a generation, and the resulting path is close to bang-bang. Early in a generation the gap is small and the remaining learning horizon is long, so λ_{learn} is large and dominates: the leader prices below its own cost and far below the limit price. This is the Spence regime. As learning is exhausted, $\lambda_{\text{learn}} \to 0$ and the constraint in (13) binds: the leader transitions to pure limit pricing, harvesting a gap that now decays at rate δ by Proposition 1. The next frontier release resets γ and the cycle repeats. The observed industry cadence, launch cheap or free, hold price while the tier below commoditises, then launch the next frontier, is this structure at a roughly three-quarter period.

Remark 3 (The Texas Instruments failure mode)*.* When several firms each price at their own Spence optimum, the market price is set by the largest shadow value λ_{learn} in the field rather than by anyone's cost. Everyone descends the curve, nobody collects the rent, and the surplus migrates upstream to the input supplier whose own moat does not leak. This is the calculator and dynamic-memory history, and it is the natural reading of simultaneous below-cost frontier pricing today. The model's distributional prediction is that the model layer earns the semiconductor-memory outcome unless δ can be held low, which is precisely what Corollary 2 says the labs are trying to do.

6 Criticality: treadmill or breakaway

The original question, whether a leader can run away permanently, is a question about the sign of a fixed point. Generalise the growth term in (8) to allow returns to being ahead, which is what a data or capability flywheel means: a firm that is further ahead converts a unit of effort into more gap. Write

$$\dot{G} \;=\; \alpha\, G^{\beta} \;-\; \delta\, G. \tag{14}$$

If $\beta = 1$ the equation is linear and the outcome is decided by the sign of $\alpha - \delta$ alone. The interesting case is $\beta > 1$, superlinear compounding. Then (14) has two fixed points, $G = 0$ and

$$G^* \;=\; \left(\frac{\delta}{\alpha}\right)^{\frac{1}{\beta-1}}, \tag{15}$$

and the interior one is *unstable*: differentiating the right-hand side of (14) gives $\alpha\beta G^{\beta-1} - \delta$, which at G^* equals $\delta(\beta - 1) > 0$. So G^* is a separatrix.

Proposition 3 (Phase structure)**.** *With $\beta > 1$, the industry has two basins. For $G_0 < G^*$ diffusion wins, every lead decays to zero, and the industry is a treadmill in which rents accrue only to current velocity. For $G_0 > G^*$ the gap grows without bound and the leader breaks away. The threshold rises in δ and falls in α.*

This is the same mathematical structure as a criticality condition, and it converts an argument into two measurements. Is $\beta > 1$, that is, do capability advantages compound superlinearly? And has any firm's G ever exceeded G^*? The observation that every frontier lead so far has decayed within months implies either $\beta \leq 1$ or a δ large enough to put G^* above anything achieved. Section 9 argues both are measurable in capability space.

Remark 4 (Which physics to keep)*.* Two of the physical analogies do formal work and should stay. The first is (14) itself, which is a criticality condition and nothing more. The second is genuinely a published isomorphism rather than a metaphor: Bianconi and Barabási [2001] map growing networks with heterogeneous node fitness onto a Bose gas and show two thermodynamic phases, a fit-get-rich phase where the leader holds a large but bounded share, and a condensate phase where a single node captures a finite fraction of all links permanently. The phase depends on the tail of the fitness distribution, which is the statistical-mechanics version of the same question Proposition 3 asks; and firm sizes are known to follow the required heavy-tailed regularity [Axtell, 2001]. Other analogies that came up in development, a Carnot bound on extractable rent and classical nucleation as a theory of entry, are expositionally attractive but carry no theorem this paper needs. They are omitted deliberately: in an economics venue a decorative analogy is a negative signal, and Lasry and Lions [2007] mean-field games already supply the rigorous bridge if the population version is ever wanted, since they couple forward-looking optimisation to a diffusing population distribution and therefore do not suffer the defect that firms, unlike particles, anticipate.

7 The full game: Markov perfect equilibrium

Sections 4–6 are a reduced form with a passive follower. That is the first thing a theory seminar attacks, and correctly: the follower also has a learning motive and also prices strategically. This section states the model a referee would accept. Appendix A then shows the reduced form is the limit of this game when the follower is myopic or capacity constrained, so nothing above is discarded.

7.1 Primitives

Discrete time, infinite horizon, two firms plus a fringe. The state is

$$s \;=\; \big(\omega_L,\, \omega_F,\, b_L,\, b_F\big) \in \mathcal{S}, \tag{16}$$

knowledge stocks and cash. Capability enters demand as quality and cost enters separately, which resolves the objection that a scalar gap conflates "cheaper" with "better". Demand is differentiated Bertrand with logit shares,

$$D_i(p;\omega) \;=\; M\,\frac{\exp\big(\theta\omega_i - \sigma p_i\big)}{\exp(u_0) + \sum_j \exp\big(\theta\omega_j - \sigma p_j\big)}, \tag{17}$$

where M is market size, u_0 the outside option including the open-weight fringe, and $\theta, \sigma > 0$. Each period firms choose price p_i and research spend r_i. Transitions are

$$\omega_L' = \omega_L + \gamma\, q_L + \phi\, r_L + \varepsilon_L, \tag{18}$$

$$\omega_F' = \omega_F + \gamma\, q_F + \phi\, r_F + \underbrace{\delta\big(q_L\big)\,\big(\omega_L - \omega_F\big)}_{\text{the leak}} + \varepsilon_F, \tag{19}$$

$$b_i' = b_i + \pi_i(s, p, r), \tag{20}$$

with q_i realised volume and ε_i mean-zero research shocks. This is an Ericson and Pakes [1995] industry dynamics model with one modification, the leak term in (19), which is the only place Assumption 3 enters.

7.2 Equilibrium

Definition 1 (Markov perfect equilibrium)**.** A strategy profile $\{p_i^*(s), r_i^*(s)\}$ and value functions $\{V_i(s)\}$ constitute an MPE if for each i and each $s \in \mathcal{S}$,

$$V_i(s) \;=\; \max_{p_i, r_i} \left\{ \pi_i\big(s, p_i, p_{-i}^*(s), r_i\big) \;+\; \frac{1}{1+r}\,\mathbb{E}\big[V_i(s') \mid s, p_i, p_{-i}^*(s), r_i, r_{-i}^*(s)\big] \right\},$$

with s' generated by (18)–(20).

Existence and computation follow the standard route [Pakes and McGuire, 1994], with the usual caveat that uniqueness is not guaranteed and multiplicity must be reported rather than assumed away.

7.3 Three objections that dissolve in this formulation

Credibility. Milgrom and Roberts [1982] object that under complete information a low price today commits a firm to nothing tomorrow, so limit pricing is not credible. In a state-based game there is no threat to police: a low price today *physically changes* tomorrow's state, raising ω_L through (18) and denying q_F in (19). Deterrence is then a property of the equilibrium path, and it has a clean characterisation: equilibrium price lies below the static best response if and only if the shadow values $\partial V_i / \partial \omega_i$ are positive. Nothing needs to be assumed about off-path behaviour.

Contestability. Baumol et al. [1982] requires costless exit, and this industry has the largest sunk capital expenditure of any in history; the critique of contestability on exactly these grounds is old and decisive [Schwartz and Reynolds, 1983]. We drop contestability entirely. Its role, an upper bound on price that does not require actual entry, is played instead by the fringe: an open-weight competitive sector whose capability is dragged up by diffusion from the frontier,

$$\omega_\Phi' \;=\; \omega_\Phi + \delta_\Phi\big(q_L\big)\big(\omega_L - \omega_\Phi\big), \tag{21}$$

priced at its own marginal cost. The fringe enters the outside option u_0 in (17) and is therefore a hard, always-present price cap that requires no entry decision at all. Limit-pricing-like behaviour now emerges from continuation values rather than being imposed, and the sunk cost κ becomes a comparative static rather than a wound.

Increasing dominance. In the reduced form it was an input; here it is a result to be checked numerically on the computed equilibrium, which is the honest status for it given that Cabral and Riordan [1994] obtain it without diffusion and the leak term works against it.

7.4 Nesting Sutton

Sutton [1991] shows that with endogenous sunk costs, escalating research expenditure produces a lower bound on concentration that does not vanish as the market grows. That result must be nested, not ignored. In this model, let market size M in (17) grow and verify that equilibrium r_i escalates so that concentration stays bounded away from zero. The distinguishing prediction then falls out of the leak term, and it is sharp: quality ladders in Sutton do not leak, so the identity of the leader is persistent; here $\delta > 0$ means the leader's own deployment continuously re-arms the follower, so concentration persists while leadership rotates.

Proposition 4 (Discriminating prediction)**.** *Both models predict persistent concentration. Sutton [1991] additionally predicts stable leader identity. This model predicts leader churn at a rate increasing in δ. Frontier leadership turnover is directly observable, so the two models are separable on existing data.*

8 Which flywheel? The financing channel

The most dangerous objection to this model does not come from economics. It comes from machine learning, and it is empirical: there may be no learning-by-doing from deployment volume at all. Capability gains in large language models come predominantly from training compute and algorithmic advance, not from cumulative served tokens. Pretraining corpora are scraped rather than user-generated, and reinforcement learning from production traffic is real but plausibly second order. If capability is a function of research capital expenditure rather than volume, then $\gamma \approx 0$, the learning channel in (18) collapses, and below-cost pricing is not a Spence investment at all.

The right response is to concede the premise and re-route the mechanism, which makes the paper stronger rather than merely safer. Separate three channels explicitly:

(i) volume $\to$ data $\to$ capability, the classical learning channel, contested and possibly small;

(ii) research spend $\to$ capability, uncontested and large, already present as ϕr_i in (18);

(iii) volume $\to$ revenue $\to$ investor beliefs $\to$ financing capacity $\to$ research spend, the channel that does not require any learning-by-doing.

Channel (iii) closes the loop through the cash state b_i in (20) and a financing constraint $r_i \leq \rho(b_i, \text{share}_i)$ with $\partial\rho/\partial\text{share} > 0$. The thesis then reads: in classical industries the experience curve runs through the factory; in frontier AI it runs through the capital market. Below-cost pricing buys the share that relaxes the financing constraint that funds the training run that produces the capability. This is deep-pockets predation in the sense of Bolton and Scharfstein [1990], and adding it also answers the referee objection that the model lacked any financing story.

The design implication for the paper is that results must be reported under both specifications, $\gamma > 0$ with ρ slack and $\gamma = 0$ with ρ binding. Any conclusion that survives both is robust to the contested empirical question; any conclusion that does not must be labelled as conditional on it. What survives untouched in either case is Assumption 3: the leak does not depend on how capability is produced, only on the fact that it is served.

9 Identification and data

An earlier version of this argument claimed every parameter was estimable. That claim was wrong and is withdrawn here, deliberately and in print, because the paper's credibility depends on it.

What is not identified. Prices are equilibrium objects, so cost curves cannot be read off the prices of firms we have just modelled as pricing strategically below cost; the simultaneity is textbook. True inference costs are private. And learning-curve identification is cursed even with good data: estimates of Wright's law conflate learning-by-doing, scale economies, and exogenous technical progress [Thompson, 2001, Benkard, 2000, Argote and Epple, 1990], and this industry is the worst case because compute prices fall exogenously and fast, so separating descent-by-volume from ambient hardware improvement requires an instrument that does not currently exist. With roughly five frontier labs over roughly five years, structural estimation of the full dynamic game in the manner of Benkard [2004] is not available. The honest ambition is calibration and stylised-fact matching, stated upfront.

What is identified, and this is the useful part. Move the gap from cost space to capability space. Benchmark parity is public. Define the parity lag for frontier release k as

$$T_k \;=\; \min\big\{\, t : \omega_\Phi(t) \geq \omega_L(t_k) \,\big\} - t_k, \tag{22}$$

the time from a frontier release until the best open-weight model matches it on a fixed evaluation suite held constant across k. Under (21) the parity lag maps directly to the leak rate,

$$\delta \;\approx\; \frac{\ln 2}{T_{1/2}}, \tag{23}$$

where $T_{1/2}$ is the half-parity lag. This requires no cost data, no demand system, and no cost-side instrument. The frontier's own improvement rate over the same evaluation suite identifies the composite $\gamma + \phi$ without decomposing it, which is sufficient for Proposition 3 because criticality depends on the ratio.

Testing the leak channel specifically. Assumption 3 says δ rises in the leader's deployment. Two designs are available. Cross-tier: a widely-served tier should exhibit a shorter parity lag than a capability-matched but access-restricted tier. Event study: treat major open-weight releases as diffusion shocks and measure the incumbent price response in a narrow window, which gives local identification of the price sensitivity in (13) without requiring cost data. Both designs use only public release dates, published prices, and fixed benchmarks.

Data appendix (to assemble). Frontier and open-weight release dates by lab; published per-token prices by tier and date; a fixed evaluation suite with scores for every model at release; access-restriction attributes per tier (rate limits, chain-of-thought visibility, logprob availability, terms-of-service clauses on training against outputs); and announced capital expenditure or funding events for the financing channel of Section 8.

10 Testable predictions

Prediction	Source	Where it is refuted
Margin exists only on the current frontier tier and decays with half-life $\ln 2/\delta$ after supersession	Prop. 1	Persistent margin on superseded tiers
Access restrictions cluster on the highest-capability tiers and tighten as the gap widens	Cor. 2	Restrictions uncorrelated with tier capability
Widely-served tiers show shorter benchmark-parity lags than capability-matched restricted tiers	Assn. 3	Equal lags, i.e. $\delta' = 0$
Price path within a generation is below own cost early, then flat at the fringe cap	§5	Monotone price paths
Concentration persists while leadership rotates	Prop. 4	Stable leader identity, favouring Sutton [1991]
If leakage is elastic ($\eta > 1$), leaders throttle volume below the revenue optimum	Eq. (12)	Unconstrained volume expansion
Rents migrate to the layer whose advantage does not leak	Rem. 3	Durable model-layer margin

Table 2: Predictions and their refutations.

The literature this speaks to is activating now: Azoulay et al. [2024] on whether classical moats apply to model-based firms, and Lehr and Restrepo [2025] on diffusion of general-purpose technology. Neither contains a mechanism in which the product itself is the diffusion channel.

11 Scope and known gaps

Scoped out of this paper, with reasons. The vertical layer is excluded. If upstream compute suppliers price strategically, they tax the labs' rent through input prices, and the claim that surplus migrates upstream becomes an assumption rather than a prediction. Modelling it properly requires the labs' compute contracts, which are not public, so the treatment here is a fixed input price and an explicit acknowledgement that Remark 3 is conjecture rather than result. Multi-firm extension beyond duopoly plus fringe is left to computation. Antitrust welfare analysis is omitted entirely: with $\delta > 0$ the standard predation tests behave strangely, since below-cost pricing accelerates diffusion to rivals, and that deserves its own paper.

Live weaknesses. The learning channel in (18) is contested and Section 8 is the hedge, not a resolution. Uniqueness of the MPE is not established. The mapping from capability-space gaps to price in (13) requires a calibrated quality-to-willingness-to-pay coefficient θ, and that calibration is the weakest empirical link in the chain. The parity-lag measure in (22) is sensitive to benchmark saturation, so the evaluation suite must be fixed ex ante and reported with alternatives.

12 What to do next

In priority order. First, assemble the release-price-benchmark panel of Section 9 and estimate parity lags T_k; this is a self-contained empirical note that is publishable on its own and settles whether δ is large enough for Proposition 3 to bite. Second, run the cross-tier test of Assumption 3, since the entire contribution rests on $\delta' > 0$ and it is the cheapest thing to falsify. Third, compute the MPE of Section 7 on a coarse state grid and check whether increasing dominance survives the leak. Fourth, write the financing variant of Section 8 and confirm the predictions in Section 10 hold in both specifications. Venue realism: this is a *RAND Journal*, *Journal of Industrial Economics*, or *Management Science* paper, with an economics-of-AI meeting or ACM EC as the fastest route to the right readers. Top-five placement would require either a general theorem or identification nobody currently has, and claiming otherwise wastes referee goodwill.

The market-structure question, treadmill or breakaway, will be written by somebody within roughly eighteen months regardless. The part of it that is defensibly novel here is narrow and should be claimed narrowly: a product that trains its own competitors, formalised as $\delta(D_L)$, with the deployment-restraint predictions that follow.

A The reduced form as a limit of the game

The claim is that (8) is not an ad hoc simplification but the limit of Definition 1 under two restrictions on the follower.

Lemma 1 (Myopic-follower limit)**.** *Suppose (i) the follower is myopic, in the sense that it attaches no shadow value to its own knowledge stock ($\partial V_F/\partial \omega_F = 0$) and therefore prices at its static best response, and (ii) the follower is capacity constrained at $\bar{q}_F \to 0$, so its learning term γq_F vanishes. Then in the continuous-time limit of* (18)–(19)*, the capability gap $G \equiv \omega_L - \omega_F$ obeys*

$$\dot{G} \;=\; \gamma\, q_L + \phi\, r_L \;-\; \delta(q_L)\, G,$$

which is (8) *with the growth term γD_L reinterpreted as the composite $\gamma q_L + \phi r_L$.*

Proof. Under (i) the follower's price solves the static first-order condition, so its volume q_F is a function of the current state only and carries no shadow value. Under (ii) $\gamma q_F \to 0$, so (19) reduces to $\omega_F' = \omega_F + \phi r_F + \delta(q_L)(\omega_L - \omega_F)$. Subtracting from (18), taking the period length to zero, and setting the follower's research spend to zero for the pure limit-pricing case gives the stated equation. Retaining $\phi r_F > 0$ adds a constant $-\phi r_F$ to the drift and shifts G^∞ down without changing its structure. □

Two consequences. Proposition 1 holds in the full game whenever the follower's shadow values are small relative to the leader's, which is exactly the increasing-dominance region. And the growth term is a composite, so the empirical strategy of Section 9 that identifies $\gamma + \phi$ jointly is sufficient for every result that depends on the ratio to δ, which is all of them.

B Referee objections and their disposition

The model above is the product of a deliberate adversarial pass. Recording the objections and where each is answered is part of the handoff, since a coauthor will raise them again.

Objection	Source	Disposition
Reduced form with a passive follower is not a model	IO theory	Restated as MPE, §7; reduced form derived as a limit, App. A
Low price commits to nothing, so limit pricing is not credible	Milgrom and Roberts [1982]	Dissolves in a state-based game: price changes the state, §7
Contestability requires costless exit; this industry has the largest sunk capex in history	Schwartz and Reynolds [1983]	Contestability dropped; diffusion-fed fringe supplies the price cap, Eq. (21)
This is just endogenous sunk costs	Sutton [1991]	Nested, then separated on leader churn, Prop. 4
No vertical layer, so upstream rent capture is assumed not derived	IO theory	Scoped out explicitly with the conjecture flagged as such, §11
No financing story, yet below-cost pricing is investor-funded	Bolton and Scharfstein [1990]	Cash state and financing constraint, §8
"Every parameter is estimable" is false: prices are equilibrium objects	Econometrics	Claim withdrawn; capability-space identification only, §9
Wright's-law estimates conflate learning, scale, and exogenous progress	Thompson [2001], Benkard [2000]	Composite $\gamma + \phi$ never decomposed; only its ratio to δ is used
Capability may not come from deployment volume at all	Machine learning	Conceded; flywheel re-routed through the capital market, §8
Scalar gap conflates cheaper with better	Machine learning	Capability enters as logit quality, cost as a separate state, Eq. (17)
Decorative physics analogies are a negative signal	All referees	Criticality and Bianconi and Barabási [2001] retained; Carnot and nucleation cut, Rem. 4

Table 3: Adversarial pass and disposition.

References

Linda Argote and Dennis Epple. Learning curves in manufacturing. *Science*, 247(4945):920–924, 1990.

Kenneth J. Arrow. The economic implications of learning by doing. *Review of Economic Studies*, 29(3):155–173, 1962.

Robert L. Axtell. Zipf distribution of U.S. firm sizes. *Science*, 293(5536):1818–1820, 2001.

Pierre Azoulay, Joshua L. Krieger, and Abhishek Nagaraj. Old moats for new models: Openness, control, and competition in generative AI. Working Paper 32474, National Bureau of Economic Research, 2024.

Joe S. Bain. *Barriers to New Competition.* Harvard University Press, Cambridge, MA, 1956.

William J. Baumol, John C. Panzar, and Robert D. Willig. *Contestable Markets and the Theory of Industry Structure*. Harcourt Brace Jovanovich, New York, 1982.

C. Lanier Benkard. Learning and forgetting: The dynamics of aircraft production. *American Economic Review*, 90(4):1034–1054, 2000.

C. Lanier Benkard. A dynamic analysis of the market for wide-bodied commercial aircraft. *Review of Economic Studies*, 71(3):581–611, 2004.

Ginestra Bianconi and Albert-László Barabási. Bose-einstein condensation in complex networks. *Physical Review Letters*, 86(24):5632–5635, 2001.

Patrick Bolton and David S. Scharfstein. A theory of predation based on agency problems in financial contracting. *American Economic Review*, 80(1):93–106, 1990.

Luís M. B. Cabral and Michael H. Riordan. The learning curve, market dominance, and predatory pricing. *Econometrica*, 62(5):1115–1140, 1994.

Partha Dasgupta and Joseph Stiglitz. Learning-by-doing, market structure and industrial and trade policies. *Oxford Economic Papers*, 40(2):246–268, 1988.

Richard Ericson and Ariel Pakes. Markov-perfect industry dynamics: A framework for empirical work. *Review of Economic Studies*, 62(1):53–82, 1995.

Drew Fudenberg and Jean Tirole. Learning-by-doing and market performance. *Bell Journal of Economics*, 14(2):522–530, 1983.

David M. Kreps and Robert Wilson. Reputation and imperfect information. *Journal of Economic Theory*, 27(2):253–279, 1982.

Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260, 2007.

Nils H. Lehr and Pascual Restrepo. The price of intelligence: How should socially-minded firms price and deploy AI? Working paper, 2025.

John S. McGee. Predatory price cutting: The standard oil (N.J.) case. *Journal of Law and Economics*, 1:137–169, 1958.

Paul Milgrom and John Roberts. Limit pricing and entry under incomplete information: An equilibrium analysis. *Econometrica*, 50(2):443–459, 1982.

Franco Modigliani. New developments on the oligopoly front. *Journal of Political Economy*, 66 (3):215–232, 1958.

Ariel Pakes and Paul McGuire. Computing markov-perfect nash equilibria: Numerical implications of a dynamic differentiated product model. *RAND Journal of Economics*, 25(4):555–589, 1994.

Marius Schwartz and Robert J. Reynolds. Contestable markets: An uprising in the theory of industry structure: Comment. *American Economic Review*, 73(3):488–490, 1983.

Reinhard Selten. The chain store paradox. *Theory and Decision*, 9(2):127–159, 1978.

A. Michael Spence. The learning curve and competition. *Bell Journal of Economics*, 12(1): 49–70, 1981.

John Sutton. *Sunk Costs and Market Structure: Price Competition, Advertising, and the Evolution of Concentration*. MIT Press, Cambridge, MA, 1991.

Paolo Sylos-Labini. *Oligopoly and Technical Progress*. Harvard University Press, Cambridge, MA, 1962.

Peter Thompson. How much did the liberty shipbuilders learn? new evidence for an old case study. *Journal of Political Economy*, 109(1):103–137, 2001.

T. P. Wright. Factors affecting the cost of airplanes. *Journal of the Aeronautical Sciences*, 3(4): 122–128, 1936.