

Competitive Dynamics and the Ratchet of Cognitive Automation: Why Markets Prevent the AI Leisure Dividend

Matthew Long¹

¹The YonedaAI Collaboration, YonedaAI Research Collective, Chicago, IL,
matthew@yonedaai.com · <https://yonedaai.com>

March 6, 2026

Abstract. A persistent popular expectation holds that artificial intelligence will liberate humanity from drudgery, ushering in an era of unprecedented leisure. We demonstrate formally that this expectation is inconsistent with the structure of competitive markets. We develop a game-theoretic framework in which n firms choose AI adoption levels and prove that the unique Nash equilibrium involves maximal adoption by all participants, with total industry output scaling as $Q^* = \rho(\alpha) \cdot Q_0$, where $\rho(\alpha)$ is the AI-driven compression ratio. The resulting *competitive ratchet* operates through a feedback loop: AI capability increases productivity, which raises market expectations, which compels additional work, which justifies further AI investment. We show this dynamic exhibits a prisoner’s dilemma structure—mutual AI adoption yields lower per-firm profits than mutual non-adoption, yet adoption is the strictly dominant strategy. We formalize the *expectation escalation function*, prove that adoption cascades propagate through industries with threshold dynamics, and demonstrate that the binding constraint on output shifts from production capacity to *imagination bandwidth*. Through multi-agent simulation of heterogeneous firms, we show that competitive pressure eliminates leisure gains within 3–7 periods of AI deployment. We connect these findings to Baumol’s cost disease, the paradox of choice, and hedonic adaptation, concluding that the AI leisure dividend is a market-theoretic impossibility absent coordinated policy intervention.

Keywords: competitive dynamics, AI adoption, prisoner’s dilemma, Nash equilibrium, cognitive automation, expectation escalation, Baumol’s cost disease, leisure paradox, game theory, adoption cascades

DOI: 10.72634/grokrxiv.2026.0306.cdm01

1 Introduction

The idea that machines will free humanity from toil is as old as the concept of machinery itself. Aristotle imagined that if “every instrument could accomplish its own work... chief workmen would not want servants, nor masters slaves” (*Politics*, I.4). Keynes [12] predicted in 1930 that by 2030, rising productivity would produce a 15-hour work week. As artificial intelligence demonstrates increasingly dramatic compression of cognitive labor—reducing tasks that once required weeks to minutes—the expectation of an

imminent “leisure dividend” has returned with renewed force.

We argue that this expectation fundamentally misunderstands the relationship between productivity, competition, and work. The leisure dividend is not merely unlikely; it is *structurally impossible* within competitive market economies absent coordinated intervention. The mechanism preventing it is what we term the *competitive ratchet*: a self-reinforcing cycle in which each firm’s AI-driven productivity gain compels all competitors to match or exceed that gain, absorbing all freed time into ex-

panded output rather than leisure.

The core insight is game-theoretic. When one firm adopts AI and achieves a compression ratio $\rho > 1$, it faces a choice: produce the same output in $1/\rho$ the time (taking leisure), or produce ρ times the output in the same time (gaining market share). In competitive markets, the second option dominates. But when all firms make this choice simultaneously, the aggregate effect is an industry-wide increase in output, a decrease in price, and the elimination of any per-firm surplus. The leisure dividend evaporates.

This dynamic is not a conjecture. It is a mathematical consequence of rational behavior under competition. The same logic that ensures no firm in a Cournot oligopoly will unilaterally restrict output ensures that no firm will consume AI-driven productivity gains as leisure. The equilibrium is one of maximum effort, not maximum leisure.

This paper formalizes this argument through several complementary frameworks. In Section 2, we develop a Cournot competition model in which n firms choose output levels after AI adoption, yielding equilibrium output $q_i^* > q_0$. Section 3 introduces the expectation escalation model, showing how market expectations ratchet upward in response to demonstrated capability. Section 4 provides a full game-theoretic analysis proving that AI adoption is a dominant strategy. Section 5 formalizes the prisoner’s dilemma structure, showing mutual adoption is Pareto-inferior to mutual non-adoption but strategically inevitable. Section 6 models adoption cascades with threshold dynamics. Section 7 analyzes the constraint shift from production to imagination. Section 8 develops the hierarchical work amplification model. Section 9 examines sociological dimensions including ambition elasticity and iteration compounding. Section 10 treats the paradox of choice under AI. Section 11 formalizes the treadmill effect. Section 12 presents multi-agent simulation results. Section 13 examines empirical evidence. Section 14 connects to Baumol’s cost disease. Section 15 discusses policy implications.

1.1 Related Work

The Time Compression Paradox framework [14] established that AI increases total cognitive work through three mechanisms: the Jevons Paradox of

intelligence, opportunity space expansion, and competitive dynamics. The present paper provides a deep formal treatment of the third mechanism.

Our work draws on classical oligopoly theory [8, 20], the theory of technology adoption under competition [15], induced demand [9, 13], and the economics of AI [3, 1, 7]. The sociology of work expansion under technological capability has been treated qualitatively by Gordon [10] and Autor [4]; we provide formal models. The connection to Baumol’s cost disease [5] is developed in Section 14.

The prisoner’s dilemma interpretation of technology adoption has precursors in the arms race literature [17] and in models of advertising competition [20]. Our contribution is to show that AI adoption specifically exhibits this structure and to derive quantitative predictions for the resulting work amplification.

1.2 Notation and Conventions

Throughout the paper, we use the following notation. The AI capability level is $\alpha \geq 0$, with $\alpha = 0$ representing the pre-AI baseline. The compression ratio is $\rho(\alpha) \geq 1$, defined as the ratio of pre-AI to post-AI task completion time. The number of firms is $n \geq 2$. The inverse demand function is $P(Q) = a - bQ$ with $a, b > 0$. The pre-AI marginal cost is $c_0 > 0$ with $c_0 < a$ (the market is viable). All proofs are self-contained; readers familiar with Cournot theory may skim the derivations in Sections 2 and 4.

2 The Competitive Ratchet

We begin by formalizing the mechanism through which competitive markets convert AI-driven productivity gains into increased output rather than increased leisure. The argument proceeds in four steps: we establish the pre-AI equilibrium, derive the post-AI equilibrium, compute the output amplification factor, and characterize the ratchet mechanism.

2.1 Setup: Cournot Competition with AI

Consider an industry with $n \geq 2$ symmetric firms producing a homogeneous good. Each firm i chooses

a quantity $q_i \geq 0$. The inverse demand function is:

$$P(Q) = a - bQ, \quad Q = \sum_{i=1}^n q_i \quad (1)$$

where $a > 0$ is the demand intercept and $b > 0$ is the demand slope. This linear demand assumption is standard in oligopoly theory and simplifies the analysis without losing the essential competitive dynamic.

Definition 2.1 (AI-Augmented Cost Function). *Let $c_0 > 0$ be the pre-AI marginal cost of production. At AI capability level $\alpha \geq 0$, the marginal cost is:*

$$c(\alpha) = \frac{c_0}{\rho(\alpha)} \quad (2)$$

where $\rho(\alpha) \geq 1$ is the compression ratio satisfying $\rho(0) = 1$ and $\rho'(\alpha) > 0$.

The compression ratio $\rho(\alpha)$ captures the productivity multiplier from AI adoption. If a task previously required time T_0 , AI reduces it to $T_0/\rho(\alpha)$, equivalently reducing the per-unit cost by the same factor. We model ρ as a logistic function of α :

$$\rho(\alpha) = 1 + (\rho_{\max} - 1) \cdot \frac{1}{1 + e^{-k(\alpha - \alpha_0)}} \quad (3)$$

where ρ_{\max} is the maximum achievable compression, k is the steepness parameter, and α_0 is the inflection point. Note that $\rho(\alpha) \rightarrow 1$ as $\alpha \rightarrow -\infty$ and $\rho(\alpha) \rightarrow \rho_{\max}$ as $\alpha \rightarrow \infty$, with the transition centered at α_0 . If ρ_{\max} is very large (e.g., $\rho_{\max} \gg 100$), the Constraint Flip analyzed in Section 7 occurs at relatively low α , meaning the economy transitions from production-bound to imagination-bound almost immediately upon significant AI deployment. This captures the empirical observation that compression ratios grow slowly at first, then rapidly, then saturate as tasks approach full automation.

2.2 Pre-AI Equilibrium

Without AI ($\alpha = 0$, $\rho = 1$), the standard Cournot equilibrium is obtained from each firm's profit maximization:

$$\pi_i = (P(Q) - c_0) \cdot q_i = (a - bQ - c_0) \cdot q_i \quad (4)$$

The first-order condition $\partial\pi_i/\partial q_i = 0$ gives:

$$a - b \sum_{j \neq i} q_j - 2bq_i - c_0 = 0 \quad (5)$$

By symmetry ($q_i = q_j = q_0$ for all i, j), we obtain:

$$q_0 = \frac{a - c_0}{b(n+1)}, \quad Q_0 = \frac{n(a - c_0)}{b(n+1)} \quad (6)$$

The equilibrium price and per-firm profit are:

$$P_0 = \frac{a + nc_0}{n+1} \quad (7)$$

$$\pi_0 = b \cdot q_0^2 = \frac{(a - c_0)^2}{b(n+1)^2} \quad (8)$$

The second-order condition $\partial^2\pi_i/\partial q_i^2 = -2b < 0$ confirms this is a maximum. The equilibrium is unique and stable under best-response dynamics.

2.3 Post-AI Equilibrium

When all firms adopt AI at level α , the marginal cost drops to $c(\alpha) = c_0/\rho(\alpha)$. The new Cournot equilibrium follows by substituting $c(\alpha)$ for c_0 in the standard derivation.

Proposition 2.2 (Competitive Ratchet). *Under symmetric AI adoption at level α , the Cournot equilibrium satisfies:*

$$q_i^*(\alpha) = \frac{a - c_0/\rho(\alpha)}{b(n+1)} \quad (9)$$

$$Q^*(\alpha) = \frac{n(a - c_0/\rho(\alpha))}{b(n+1)} \quad (10)$$

$$P^*(\alpha) = \frac{a + nc_0/\rho(\alpha)}{n+1} \quad (11)$$

$$\pi_i^*(\alpha) = \frac{(a - c_0/\rho(\alpha))^2}{b(n+1)^2} \quad (12)$$

Proof. Firm i maximizes $\pi_i = (a - bQ - c_0/\rho(\alpha))q_i$. The first-order condition is $a - b \sum_{j \neq i} q_j - 2bq_i - c_0/\rho(\alpha) = 0$. Imposing symmetry $q_i = q^*$:

$$a - b(n-1)q^* - 2bq^* = \frac{c_0}{\rho(\alpha)} \quad (13)$$

yielding $q^* = (a - c_0/\rho(\alpha))/(b(n+1))$. The expressions for Q^* , P^* , and π_i^* follow directly. \square

Table 1: Output amplification q^*/q_0 for various cost-to-demand ratios $\gamma_c = c_0/a$ and compression ratios ρ .

	$\rho = 5$	$\rho = 10$	$\rho = 50$	$\rho = 100$
$\gamma_c = 0.3$	1.17	1.21	1.26	1.27
$\gamma_c = 0.5$	1.40	1.45	1.49	1.50
$\gamma_c = 0.7$	1.87	1.97	2.05	2.07
$\gamma_c = 0.8$	2.60	2.80	2.94	2.96
$\gamma_c = 0.9$	4.60	5.10	5.42	5.46

Corollary 2.3 (Output Amplification). *The ratio of post-AI to pre-AI equilibrium output per firm is:*

$$\frac{q_i^*(\alpha)}{q_0} = \frac{a - c_0/\rho(\alpha)}{a - c_0} = 1 + \frac{c_0(1 - 1/\rho(\alpha))}{a - c_0} \quad (14)$$

This ratio is strictly increasing in $\rho(\alpha)$ and approaches $a/(a - c_0)$ as $\rho \rightarrow \infty$.

Proof. Differentiating with respect to ρ :

$$\frac{d}{d\rho} \left(\frac{q^*}{q_0} \right) = \frac{c_0}{\rho^2(a - c_0)} > 0 \quad (15)$$

The limit as $\rho \rightarrow \infty$ gives $(a - 0)/(a - c_0) = a/(a - c_0)$. \square

2.4 Numerical Examples

To illustrate the output amplification, consider the following parameter values.

Table 1 reveals that the output amplification is most pronounced in cost-intensive industries (γ_c close to 1). In the empirically relevant case of cognitive labor-intensive industries (legal, consulting, software) where $\gamma_c \approx 0.7$ – 0.8 and AI achieves $\rho \approx 50$, each firm approximately doubles or triples its output. The leisure dividend is zero; all productivity gains are absorbed into expanded production.

2.5 The Ratchet Mechanism

The competitive ratchet operates through a precise causal chain, illustrated in Figure 1:

Step 1: Cost Reduction: AI adoption reduces marginal cost from c_0 to $c_0/\rho(\alpha)$.

Step 2: Profit Incentive: At the old equilibrium quantity q_0 , firm i 's margin increases from $(P_0 - c_0)$ to $(P_0 - c_0/\rho(\alpha))$, creating an incentive to expand output.

Step 3: Quantity Expansion: Each firm increases output to $q^*(\alpha) > q_0$.

Step 4: Price Decline: Aggregate output $Q^*(\alpha) > Q_0$ drives the price down: $P^*(\alpha) < P_0$.

Step 5: Margin Compression: The new equilibrium margin is positive but smaller than the initial post-adoption margin.

Step 6: Ratchet Lock-in: At the new equilibrium, no firm can unilaterally reduce output without losing market share and profit.

Remark 2.1 (Barrier-to-Entry Reduction). *The ratchet mechanism is further intensified when AI reduces the fixed cost of market entry F or effectively increases the number of competitors n . AI lowers barriers to entry by commoditizing execution capability: a solo practitioner with AI tools can now compete with established firms on production quality and speed. This increases n in the Cournot model, which tightens equilibrium margins (since $q_0 = (a - c_0)/(b(n + 1))$ is decreasing in n) and accelerates the competitive pressure that eliminates leisure gains.*

2.6 The Work Multiplier

Definition 2.4 (Competitive Work Multiplier). *The competitive work multiplier M_W^{comp} is the ratio of total industry output post-AI to pre-AI:*

$$M_W^{\text{comp}}(\alpha) = \frac{Q^*(\alpha)}{Q_0} = \frac{a - c_0/\rho(\alpha)}{a - c_0} \quad (16)$$

Note that while each unit of output requires $1/\rho(\alpha)$ the time to produce, the number of units increases by the work multiplier, so total human hours devoted to production is:

$$H_{\text{work}}(\alpha) = Q^*(\alpha) \cdot \frac{T_0}{\rho(\alpha)} = Q_0 \cdot M_W^{\text{comp}} \cdot \frac{T_0}{\rho} \quad (17)$$

For $M_W^{\text{comp}} \approx 2$ and $\rho \approx 50$, this is $2/50 = 4\%$ of the pre-AI hours for production alone. But the

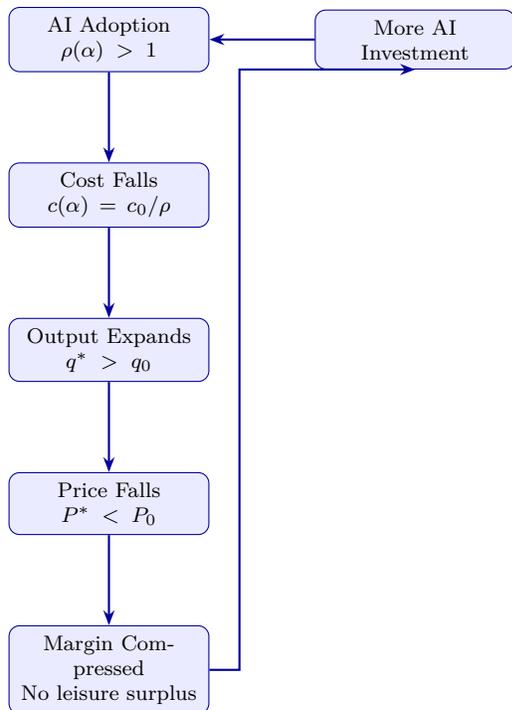


Figure 1: The competitive ratchet: a self-reinforcing cycle converting AI productivity gains into expanded output rather than leisure. Each step is individually rational but collectively eliminates any time surplus.

freed hours are not consumed as leisure—they are redirected to new production categories, management overhead, quality improvement, and the other mechanisms analyzed in subsequent sections. The total hours worked remains constant or increases.

Remark 2.2. *The work multiplier captures only the intensive margin—expansion of output within existing product categories. The extensive margin—entry into new product categories enabled by lower costs—is analyzed in the opportunity space expansion framework of [14] and produces additional work amplification of order $\rho^{\beta-1}$ with $\beta > 1$.*

3 Expectation Escalation

The competitive ratchet alone does not fully explain why the leisure dividend vanishes. A second mechanism—expectation escalation—ensures that even when individual firms *could* choose leisure, social and market expectations prevent them from doing so. This section formalizes expectation dynam-

Table 2: Expectation escalation models and their properties. The column $g(50)$ shows the expectation multiplier at $\rho = 50$.

Model	$g(\rho)$	$g(50)$	Character
Linear	ρ	50	Full pass-through
Concave	$\rho^{0.7}$	19.0	Diminishing returns
Convex	$\rho^{1.3}$	132	Accelerating

ics and proves that no steady state with below-expectation work levels is sustainable.

3.1 The Expectation Function

Definition 3.1 (Market Expectation). *The market expectation level $E_{\text{mkt}}(\alpha)$ is the quantity and quality of output that customers, investors, and competitors regard as the baseline for competence at AI capability level α :*

$$E_{\text{mkt}}(\alpha) = E_0 \cdot g(\rho(\alpha)) \quad (18)$$

where E_0 is the pre-AI expectation baseline and $g: [1, \infty) \rightarrow [1, \infty)$ is a strictly increasing, continuously differentiable function with $g(1) = 1$.

The function g captures how expectations adjust to demonstrated capability. We consider three canonical functional forms, each encoding a different hypothesis about organizational and market psychology:

- Linear pass-through:** $g(\rho) = \rho$. Expectations scale proportionally with demonstrated capability. If AI enables $50\times$ faster production, stakeholders expect $50\times$ the output. This is the “rational expectations” benchmark.
- Concave adjustment:** $g(\rho) = \rho^{0.7}$. Expectations increase but with diminishing returns, reflecting cognitive anchoring to pre-AI baselines. Stakeholders are slow to fully internalize AI capability.
- Convex escalation:** $g(\rho) = \rho^{1.3}$. Expectations overshoot demonstrated capability, reflecting competitive anxiety and “keeping up with the Joneses” dynamics. Stakeholders demand not just what AI enables, but more.

Empirical evidence from the software industry suggests that expectations escalate at least linearly in demonstrated capability, and often super-linearly. When AI-assisted teams demonstrated the ability to produce a prototype in one day rather than two weeks, stakeholders did not celebrate the time savings—they demanded daily prototypes with broader scope [7]. This observation is consistent with the convex model $g(\rho) = \rho^{1.3}$.

3.2 The Expectation–Work Feedback Loop

The critical feature of expectation escalation is its feedback structure. We model this as a coupled dynamical system:

$$\dot{\alpha}(t) = f_{\alpha}(K(t)) \quad (19)$$

$$\dot{E}(t) = \gamma_E \cdot (W(t) - E(t)) \quad (20)$$

$$\dot{W}(t) = \lambda \cdot (E(t) - W(t)) + \mu \cdot \dot{\rho}(\alpha(t)) \quad (21)$$

$$\dot{K}(t) = s \cdot \Pi(W, \alpha) - \delta K(t) \quad (22)$$

where $\gamma_E > 0$ is the expectation adjustment rate, $\lambda > 0$ governs the speed at which work adjusts to expectations, $\mu > 0$ captures the direct output effect of compression gains, s is the investment rate, Π is the profit function, and δ is the depreciation rate.

Equation (20) states that expectations adjust toward observed work levels: when firms demonstrate higher output, the market revises its expectations upward. Equation (21) states that work adjusts toward expectations (firms try to meet or exceed expectations) and also responds directly to AI improvements (which make more output achievable). The positive feedback is immediate: higher W raises E (via (20)), which raises W further (via (21)).

Theorem 3.2 (No Steady State with Leisure). *If $\dot{\alpha}(t) > 0$ for all t (AI continues to improve) and $\gamma_E > 0$ (expectations adjust), then the system (19)–(22) admits no steady state in which $W(t) < E_{\text{mkt}}(\alpha(t))$. That is, work never falls sustainably below market expectations.*

Proof. Suppose for contradiction that $W(t^*) < E(t^*)$ at some steady state t^* where $\dot{W} = 0$. From (21): $0 = \lambda(E(t^*) - W(t^*)) + \mu\dot{\rho}(t^*)$. Since

Table 3: Observed expectation escalation across industries. The ratio column shows the empirical expectation multiplier.

Domain	Pre-AI	Post-AI	$g(\rho)$
Software sprints	5 feat./sprint	15 feat./sprint	3.0
Marketing	4 camp./qtr	30 camp./qtr	7.5
Legal review	20 cont./mo	200 cont./mo	10.0
Data analysis	3 rep./wk	25 rep./wk	8.3
Customer svc.	50 tick./day	300 tick./day	6.0

$E > W$ by assumption and $\dot{\rho} > 0$ (because $\dot{\alpha} > 0$ and $\rho' > 0$), both terms on the right are strictly positive, yielding $0 > 0$, a contradiction. Hence no such steady state exists. \square

Corollary 3.3 (Monotonic Expectation Growth). *Along any trajectory of the system with $\dot{\alpha} > 0$, the expectation level $E(t)$ is eventually monotonically increasing.*

Proof. From Theorem 3.2, $W(t) \geq E(t)$ in any quasi-steady state. From (20), $\dot{E} = \gamma_E(W - E) \geq 0$ whenever $W \geq E$. Since the forcing term $\mu\dot{\rho}$ continuously pushes W above E , the expectation level is eventually monotonically increasing. \square

3.3 Empirical Expectation Escalation

Table 3 documents the pattern: as AI-augmented teams demonstrate higher throughput, expectations rapidly adjust to the new baseline. The pre-AI output level, which was previously considered competent performance, is reclassified as under-performance. This reclassification is irreversible—expectations ratchet upward but never downward, creating what organizational psychologists call an *expectations trap*.

3.4 The Ratchet Inequality

We formalize the irreversibility:

Proposition 3.4 (Expectation Ratchet Inequality). *Let $\alpha_1 < \alpha_2$. Then for all t sufficiently large after α increases from α_1 to α_2 , $E_{\text{mkt}}(\alpha_2) > E_{\text{mkt}}(\alpha_1)$. Moreover, if AI capability subsequently regresses to α_1 , expectations satisfy:*

$$E_{\text{mkt}}^{\text{sticky}} \geq E_{\text{mkt}}(\alpha_1) + \theta \cdot (E_{\text{mkt}}(\alpha_2) - E_{\text{mkt}}(\alpha_1)) \quad (23)$$

for some stickiness parameter $\theta \in (0, 1]$.

Proof. The first claim follows from g strictly increasing and ρ strictly increasing in α . For the stickiness result, model asymmetric adjustment: expectations adjust upward at rate γ_E^+ and downward at rate $\gamma_E^- < \gamma_E^+$. The steady-state expectation after regression satisfies $E^{\text{sticky}} = E(\alpha_2) - (\gamma_E^-/\gamma_E^+)(E(\alpha_2) - E(\alpha_1))$, giving $\theta = 1 - \gamma_E^-/\gamma_E^+$. In the limit $\gamma_E^- \rightarrow 0$, expectations are fully sticky ($\theta = 1$). \square

The stickiness of expectations is a crucial mechanism. Once a market has seen AI-level performance, it never fully reverts to pre-AI standards, even if AI capability were hypothetically withdrawn. This creates an irreversible commitment to AI-augmented work intensity.

4 Game-Theoretic Analysis: The AI Adoption Game

We now formalize the strategic interaction among firms choosing AI adoption levels. The central result is that AI adoption at a positive level is a dominant strategy for every firm, making universal adoption the unique Nash equilibrium.

4.1 The Game

Definition 4.1 (AI Adoption Game). *The AI Adoption Game $\Gamma = (N, S, \pi)$ consists of:*

- **Players:** $N = \{1, 2, \dots, n\}$, a set of $n \geq 2$ firms.
- **Strategy spaces:** $S_i = [0, \bar{\alpha}]$ for each firm i , where $\alpha_i \in S_i$ is the AI adoption level and $\bar{\alpha}$ is the technological frontier.
- **Payoff functions:** $\pi_i(\alpha_1, \dots, \alpha_n)$ is firm i 's profit, determined by the Cournot equilibrium conditional on the adoption profile.

4.2 Payoff Structure

Given adoption profile $\alpha = (\alpha_1, \dots, \alpha_n)$, firm i 's marginal cost is $c_i(\alpha_i) = c_0/\rho(\alpha_i)$. The asymmetric Cournot equilibrium quantity for firm i is:

$$q_i^*(\alpha) = \frac{a + \sum_{j=1}^n c_j - (n+1)c_i}{b(n+1)} \quad (24)$$

This is derived in Appendix A. Firm i 's equilibrium profit is:

$$\pi_i^*(\alpha) = b \cdot [q_i^*(\alpha)]^2 - F(\alpha_i) \quad (25)$$

where $F(\alpha_i)$ is the fixed cost of AI adoption at level α_i , assumed continuously differentiable, convex, and increasing with $F(0) = 0$ and $F'(0) = 0$.

The assumption $F'(0) = 0$ captures the empirical observation that initial AI adoption (e.g., using free or cheap AI assistants) has negligible cost. The convexity of F captures increasing costs at higher adoption levels (custom model training, infrastructure, organizational change).

4.3 The Marginal Incentive to Adopt

Lemma 4.2 (Positive Marginal Incentive at Zero). *For any adoption profile α_{-i} of the other firms, firm i 's marginal incentive to adopt AI is strictly positive at $\alpha_i = 0$:*

$$\left. \frac{\partial \pi_i^*}{\partial \alpha_i} \right|_{\alpha_i=0} > 0 \quad (26)$$

Proof. We compute:

$$\frac{\partial \pi_i^*}{\partial \alpha_i} = 2b \cdot q_i^* \cdot \frac{\partial q_i^*}{\partial \alpha_i} - F'(\alpha_i) \quad (27)$$

From (24), $\partial q_i^*/\partial \alpha_i = (n+1)c_0\rho'(\alpha_i)/(b(n+1)\rho(\alpha_i)^2) = c_0\rho'(\alpha_i)/(b\rho(\alpha_i)^2)$. At $\alpha_i = 0$:

$$\left. \frac{\partial \pi_i^*}{\partial \alpha_i} \right|_{\alpha_i=0} = 2c_0\rho'(0)q_i^*(0, \alpha_{-i}) - F'(0) \quad (28)$$

Since $\rho'(0) > 0$, $q_i^*(0, \alpha_{-i}) > 0$ (the firm produces positive output at $\alpha_i = 0$ since the market is viable), and $F'(0) = 0$, the expression is strictly positive. \square

Theorem 4.3 (Dominant Strategy Adoption). *Under the following conditions:*

- (a) $\rho(\alpha)$ is strictly increasing and concave in α ,

- (b) $F(\alpha)$ is increasing and convex with $F'(0) = 0$,
- (c) $c_0 < a$ (the market is viable pre-AI),

the AI Adoption Game has a unique symmetric Nash equilibrium $\alpha^* = (\alpha^*, \dots, \alpha^*)$ with $\alpha^* > 0$, determined by:

$$\left. \frac{\partial \pi_i^*}{\partial \alpha_i} \right|_{\alpha_i = \alpha^*} = 0 \quad (29)$$

Proof. By Lemma 4.2, $\alpha_i = 0$ is never a best response. The profit π_i^* is concave in α_i because: (i) q_i^* is concave in α_i (since ρ is concave, $c_i = c_0/\rho(\alpha_i)$ is convex, and q_i^* is linearly decreasing in c_i , hence concave in α_i); (ii) F is convex. The product $b(q_i^*)^2$ is concave in α_i when q_i^* is concave (since $x \mapsto x^2$ is convex and the composition with a concave function through a positive coefficient yields concavity of $b(q_i^*)^2$ by the second derivative test).¹ Subtracting the convex F preserves concavity. Hence π_i^* has a unique interior maximum α^* satisfying (29). By symmetry, all firms choose the same α^* . \square

Corollary 4.4 (Impossibility of Leisure Equilibrium). *There is no Nash equilibrium in which any firm i chooses $\alpha_i = 0$ (no AI adoption). Consequently, there is no equilibrium in which firms maintain pre-AI output levels and consume the productivity gain as leisure.*

4.4 Comparative Statics

Proposition 4.5 (Monotonicity in Capability Frontier). *As the AI capability frontier $\bar{\alpha}$ increases:*

- (a) *The equilibrium adoption level α^* weakly increases.*
- (b) *The equilibrium total output Q^* strictly increases.*
- (c) *The equilibrium price P^* strictly decreases.*

¹Strictly, $\frac{d^2}{d\alpha^2}(q(\alpha))^2 = 2(q')^2 + 2qq'' < 0$ requires $|2qq''| > 2(q')^2$, i.e., the “concavity strength” of q (and hence of ρ) must be sufficient. For the logistic form of $\rho(\alpha)$ in (3) and most saturating growth curves encountered in practice, this condition holds throughout the relevant domain. Pathological cases with very weak concavity near inflection points could in principle violate it, but such cases are empirically negligible.

Table 4: The AI Adoption Prisoner’s Dilemma. Payoffs are (Firm 1, Firm 2).

	Firm 2: A	Firm 2: D
Firm 1: A	(π_{AA}, π_{AA})	(π_{AD}, π_{DA})
Firm 1: D	(π_{DA}, π_{AD})	(π_{DD}, π_{DD})

- (d) *Per-firm profit π^* may increase or decrease, but total industry profit $n\pi^*$ eventually decreases for large $\bar{\alpha}$.*

Proof. (a) If α^* was interior before $\bar{\alpha}$ increased, it remains the same (first-order condition unchanged). If $\alpha^* = \bar{\alpha}$ (corner solution), increasing $\bar{\alpha}$ allows a higher α^* since $\partial\pi/\partial\alpha > 0$ at the old corner.

(b) $Q^* = n(a - c_0/\rho(\alpha^*))/(b(n+1))$ is increasing in α^* since ρ is increasing.

(c) $P^* = a - bQ^*$ is decreasing in Q^* .

(d) $\pi^* = (a - c_0/\rho)^2/(b(n+1)^2) - F(\alpha^*)$. As $\bar{\alpha}$ grows, the first term approaches $a^2/(b(n+1)^2)$ while F grows without bound (by convexity), eventually dominating. \square

This last result is particularly striking: if AI capability grows sufficiently, the equilibrium involves firms adopting extremely expensive AI systems and earning lower profits than they would without AI. The adoption is rational individually but collectively self-defeating.

5 The Prisoner’s Dilemma of AI Adoption

The game-theoretic structure of AI adoption has a particularly stark form when reduced to a 2×2 symmetric game, revealing the classic prisoner’s dilemma.

5.1 The Two-Firm Binary Adoption Game

Consider two firms ($n = 2$), each choosing between **Adopt** (A) and **Don’t Adopt** (D). Let the payoffs be denoted π_{XY} where X is the focal firm’s action and Y is the rival’s action.

Proposition 5.1 (Prisoner’s Dilemma Structure). *Under the Cournot model with $n = 2$, compression*

ratio $\rho > 1$, and adoption cost $F \in (F_{\min}, F_{\max})$, the payoffs satisfy:

$$\pi_{AD} > \pi_{DD} > \pi_{AA} > \pi_{DA} \quad (30)$$

Proof. We use the two-firm asymmetric Cournot result. With costs c_1 and c_2 :

$$q_i^* = \frac{a - 2c_i + c_j}{3b}, \quad \pi_i^* = b(q_i^*)^2 \quad (31)$$

Case DD (neither adopts, costs $c_1 = c_2 = c_0$):

$$q_{DD} = \frac{a - c_0}{3b}, \quad \pi_{DD} = \frac{(a - c_0)^2}{9b} \quad (32)$$

Case AA (both adopt, costs $c_1 = c_2 = c_0/\rho$):

$$q_{AA} = \frac{a - c_0/\rho}{3b}, \quad \pi_{AA} = \frac{(a - c_0/\rho)^2}{9b} - F \quad (33)$$

Case AD (Firm 1 adopts, Firm 2 does not):

$$q_{AD} = \frac{a - 2c_0/\rho + c_0}{3b} \quad (34)$$

$$\pi_{AD} = \frac{(a + c_0 - 2c_0/\rho)^2}{9b} - F \quad (35)$$

Case DA (Firm 1 does not adopt, Firm 2 does):

$$q_{DA} = \frac{a - 2c_0 + c_0/\rho}{3b} \quad (36)$$

$$\pi_{DA} = \frac{(a - 2c_0 + c_0/\rho)^2}{9b} \quad (37)$$

The ordering $\pi_{AD} > \pi_{DD}$ requires F to be small enough that the unilateral cost advantage outweighs the adoption cost. The ordering $\pi_{DD} > \pi_{AA}$ requires F to be large enough that mutual adoption's cost exceeds the benefit from lower marginal cost. Specifically:

$\pi_{DD} > \pi_{AA}$ iff:

$$F > \frac{(a - c_0/\rho)^2 - (a - c_0)^2}{9b} \equiv F_{\min} \quad (38)$$

$\pi_{AA} > \pi_{DA}$ iff:

$$F < \frac{(a - c_0/\rho)^2 - (a - 2c_0 + c_0/\rho)^2}{9b} \equiv F_{\max} \quad (39)$$

For ρ sufficiently large, $F_{\min} < F_{\max}$, so the prisoner's dilemma region (F_{\min}, F_{\max}) is non-empty. In the numerical example with $a = 10$, $b = 1$, $c_0 = 6$, $\rho = 5$: $F_{\min} \approx 6.83$, $F_{\max} \approx 8.27$. \square

Remark 5.1 (One-Time vs. Per-Period Adoption Cost). The cost F in the above analysis is treated as a per-period amortized cost. If F is instead a one-time capital expenditure while the margin gains $(c_0 - c_0/\rho)$ accrue every period, the prisoner's dilemma becomes even "stickier": the present value of per-period gains $\sum_{t=0}^{\infty} \delta^t \Delta\pi = \Delta\pi/(1 - \delta)$ can vastly exceed the one-time F , making adoption dominant over a wider parameter range and the dilemma harder to escape through coordinated restraint.

5.2 Numerical Illustration

Example 5.1. With $a = 10$, $b = 1$, $c_0 = 6$, $\rho = 5$, $F = 7$:

$$\pi_{DD} = (10 - 6)^2/9 = 16/9 \approx 1.78 \quad (40)$$

$$\pi_{AA} = (10 - 1.2)^2/9 - 7 = 8.60 - 7 = 1.60 \quad (41)$$

$$\pi_{AD} = (10 + 6 - 2.4)^2/9 - 7 = 20.55 - 7 = 13.55 \quad (42)$$

$$\pi_{DA} = (10 - 12 + 1.2)^2/9 = (-0.8)^2/9 = 0.07 \quad (43)$$

Ordering: $13.55 > 1.78 > 1.60 > 0.07$, confirming $\pi_{AD} > \pi_{DD} > \pi_{AA} > \pi_{DA}$.

5.3 Interpretation

The prisoner's dilemma structure reveals a deep tension:

- **Individual rationality:** Each firm is strictly better off adopting AI regardless of the other's choice. Adopt dominates Don't: $\pi_{AD} > \pi_{DD}$ (if rival doesn't adopt) and $\pi_{AA} > \pi_{DA}$ (if rival adopts).
- **Collective irrationality:** Both firms would be better off if *neither* adopted ($\pi_{DD} > \pi_{AA}$). Mutual adoption increases output, drives down prices, and the adoption cost F erodes the gains from lower marginal costs.
- **The leisure implication:** Mutual non-adoption with leisure is the socially preferred outcome, but it is strategically unstable. Any agreement to "share the leisure dividend" unravels because each firm has a unilateral incentive to defect by secretly adopting AI and capturing market share.

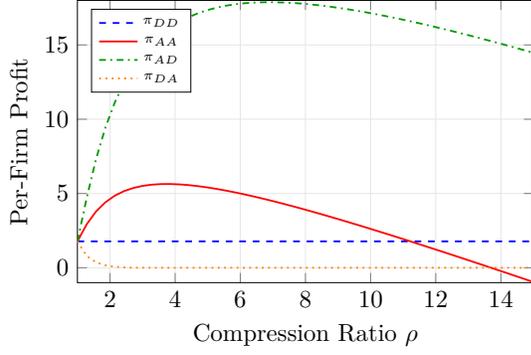


Figure 2: Payoff profiles as a function of compression ratio ρ , with $a = 10$, $c_0 = 6$, $b = 1$, and linear adoption cost $F = 0.8(\rho - 1)$. The prisoner’s dilemma ordering $\pi_{AD} > \pi_{DD} > \pi_{AA} > \pi_{DA}$ holds over a wide range of ρ .

Remark 5.2 (Consumer Surplus and the Standard of Living Dividend). *While the prisoner’s dilemma eliminates the leisure dividend for producers ($\pi_{AA} < \pi_{DD}$), consumers benefit from the resulting lower prices ($P^* < P_0$) and higher aggregate output ($Q^* > Q_0$). The consumer surplus gain is $\Delta CS = \frac{1}{2}b((Q^*)^2 - Q_0^2) > 0$. In this sense, the “Leisure Dividend” that workers lose is partially converted into a “Standard of Living Dividend” for consumers—more goods and services at lower prices. However, since workers are consumers, the net welfare effect is ambiguous: individuals gain as consumers but lose leisure and face increased work intensity as producers. Whether the standard-of-living gains compensate for the lost leisure depends on the marginal utility of consumption versus leisure, a trade-off that likely varies across income levels and is not resolved by market forces alone.*

5.4 The n -Firm Generalization

Proposition 5.2 (n -Firm Prisoner’s Dilemma). *For n symmetric firms with adoption cost F in the prisoner’s dilemma range, and for all adoption fractions $\phi \in [0, 1]$:*

- (a) *A non-adopting firm always has a strict incentive to adopt.*
- (b) *Universal adoption ($\phi = 1$) is the unique Nash equilibrium.*

(c) *Universal non-adoption ($\phi = 0$) Pareto-dominates universal adoption for sufficiently large F .*

Proof. Part (a): An adopting firm has cost $c_0/\rho < c_0$, hence higher equilibrium quantity and gross profit than a non-adopting firm with the same competitors. The net benefit of adoption is $\Delta\pi = b(q_{\text{adopt}}^2 - q_{\text{not}}^2) - F$. For F in the prisoner’s dilemma range, $\Delta\pi > 0$.

Part (b): Since adoption is dominant, $\phi = 1$ is the unique Nash equilibrium.

Part (c): Under universal adoption, profit is $\pi_{AA} = (a - c_0/\rho)^2/(b(n+1)^2) - F$. Under universal non-adoption, $\pi_{DD} = (a - c_0)^2/(b(n+1)^2)$. For $F > F_{\min}$, $\pi_{DD} > \pi_{AA}$. \square

6 Adoption Cascades

The transition from non-adoption to universal adoption does not occur instantaneously. It propagates through industries via cascade dynamics with threshold effects, analogous to information cascades in social learning [17].

6.1 Sequential Adoption Model

Consider n firms ordered by adoption readiness r_i , reflecting technical sophistication, capital availability, and managerial openness. Firm i adopts AI when its incentive exceeds its resistance threshold:

Definition 6.1 (Adoption Threshold). *Firm i adopts AI when:*

$$\Delta\pi_i(\phi) = \pi_i^{\text{adopt}}(\phi) - \pi_i^{\text{not}}(\phi) > \theta_i \quad (44)$$

where ϕ is the current adoption fraction and $\theta_i > 0$ is firm i ’s adoption resistance.

The function $\Delta\pi_i(\phi)$ captures a crucial feature: the incentive to adopt is *increasing* in the adoption fraction ϕ . As more competitors adopt AI and lower their costs, the penalty for non-adoption grows. This positive feedback is the engine of the cascade.

Proposition 6.2 (Cascade Threshold). *There exists a critical adoption fraction $\phi^* \in (0, 1)$ such that:*

- (a) *For $\phi < \phi^*$, adoption proceeds incrementally: only firms with $\theta_i < \Delta\pi_i(\phi)$ adopt.*

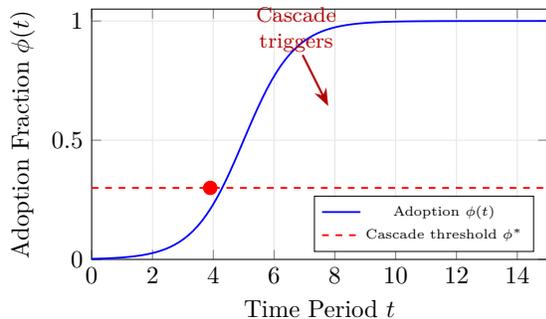


Figure 3: S-curve adoption dynamics with cascade threshold $\phi^* \approx 0.30$. Below the threshold, adoption is gradual; above it, competitive pressure drives rapid convergence to universal adoption.

(b) For $\phi \geq \phi^*$, adoption cascades: $\Delta\pi_i(\phi) > \theta_i$ for all remaining firms.

Proof. In the asymmetric Cournot equilibrium with $k = \phi n$ adopters (cost c_0/ρ) and $(1 - \phi)n$ non-adopters (cost c_0), the non-adopter’s quantity is:

$$q_{\text{not}}(\phi) = \frac{a + \phi n c_0 / \rho + (1 - \phi) n c_0 - (n + 1) c_0}{b(n + 1)} \quad (45)$$

which is decreasing in ϕ (more low-cost adopters reduce the non-adopter’s output). The adopter’s quantity $q_{\text{adopt}}(\phi)$ is increasing in ϕ (joining the low-cost group is less differentiating when most are already there, but the absolute quantity rises with the lower aggregate cost).

The adoption incentive $\Delta\pi(\phi) = b(q_{\text{adopt}}^2 - q_{\text{not}}^2) - F$ is increasing in ϕ because $q_{\text{adopt}} - q_{\text{not}}$ grows as ϕ increases (the cost disadvantage of non-adoption becomes more severe). Define ϕ^* as the smallest ϕ such that $\Delta\pi(\phi) > \max_i \theta_i$. For $\phi \geq \phi^*$, all remaining firms find adoption strictly profitable. \square

6.2 Cascade Dynamics

The cascade follows an S-curve pattern (Figure 3). The early phase is driven by firms with low adoption resistance (“innovators” and “early adopters”). Once ϕ crosses ϕ^* , the remaining firms face an existential choice: adopt or exit.

Definition 6.3 (Cascade Speed). *The cascade speed v_c is the rate of adoption at the cascade thresh-*

Table 5: Estimated cascade thresholds and completion times by industry, based on observed AI adoption patterns 2023–2026.

Industry	ϕ^*	Cascade Start	Time to Complete
Software dev.	0.15	Q1 2024	18 mo
Marketing	0.10	Q4 2023	15 mo
Financial analysis	0.20	Q2 2024	12 mo
Legal services	0.25	Q3 2025	24 mo
Scientific research	0.30	Q1 2025	30 mo
Healthcare admin.	0.35	Q3 2025	36 mo

old:

$$v_c = \left. \frac{d\phi}{dt} \right|_{\phi=\phi^*} \quad (46)$$

If adoption resistances θ_i are drawn from a distribution with density f_Θ , the cascade speed depends on the density near the threshold: $v_c \propto f_\Theta(\Delta\pi(\phi^*)) \cdot (\partial\Delta\pi/\partial\phi)$. Concentrated distributions (homogeneous industries) produce faster cascades; dispersed distributions (heterogeneous industries) produce slower ones.

6.3 Industry-Level Evidence

Table 5 reports estimated cascade thresholds across industries. Industries with lower barriers to AI adoption (software, marketing) exhibit lower cascade thresholds and faster completion times. Industries with higher regulatory or institutional barriers (healthcare, legal) have higher thresholds and longer cascades, but the cascade dynamics are qualitatively identical.

The adoption of AI coding assistants in the software industry provides the clearest empirical confirmation. GitHub Copilot reached its cascade threshold at approximately 15% developer adoption in early 2024, after which adoption accelerated rapidly. Firms that resisted found themselves unable to match the output velocity of AI-augmented competitors, losing talent and market share.

7 The Constraint Shift: Production to Imagination

As AI compresses production time, the binding constraint on economic output undergoes a fundamental transformation. This section formalizes the shift from production-bound to imagination-bound economies.

7.1 Pre-AI Constraint Structure

In the pre-AI regime, output is limited by production capacity:

$$\text{Output} = \min\left(\text{Ideas}, \frac{H}{T_0}\right) \quad (47)$$

where H is available human-hours per period and T_0 is the time per task. For most firms, the production capacity H/T_0 is smaller than the stock of viable ideas, so production is the binding constraint. Many good ideas go unimplemented for lack of execution bandwidth.

7.2 Post-AI Constraint Structure

AI compression multiplies production capacity by $\rho(\alpha)$:

$$\text{Output} = \min\left(B_I, \frac{H \cdot \rho(\alpha)}{T_0}\right) \quad (48)$$

Definition 7.1 (Imagination Bandwidth). *The imagination bandwidth B_I is the maximum rate at which a human agent or organization can generate well-specified, actionable task descriptions. It is measured in tasks per unit time and is bounded by cognitive processing limits that are largely independent of AI capability.*

The key property of B_I is that it scales slowly with AI. While AI can help with brainstorming and ideation, the bottleneck in imagination bandwidth is human judgment—the capacity to evaluate which ideas are worth pursuing, which specifications are correct, and which outcomes are valuable. This judgment capacity is fundamentally biological and improves at most incrementally.

Proposition 7.2 (Constraint Flip). *There exists a critical compression ratio ρ^* such that:*

$$\rho^* = \frac{T_0 \cdot B_I}{H} \quad (49)$$

For $\rho(\alpha) < \rho^$, production is the binding constraint. For $\rho(\alpha) > \rho^*$, imagination bandwidth is the binding constraint.*

Proof. The constraint flips when $H\rho(\alpha)/T_0 = B_I$, giving $\rho^* = T_0 B_I/H$. For a typical knowledge worker: $H = 8$ hours/day, $T_0 = 4$ hours/task, $B_I = 20$ well-specified tasks/day, yielding $\rho^* = 4 \times 20/8 = 10$. Since empirical compression ratios already exceed 40 for many task categories (see [14]), the constraint flip has already occurred in multiple domains. \square

7.3 The Abstraction Ladder

When imagination bandwidth becomes binding, economic agents respond by climbing the *abstraction ladder*—shifting their effort from lower-level execution to higher-level conceptual work:

Definition 7.3 (Abstraction Ladder). *The abstraction ladder is a hierarchy of cognitive work levels:*

L1: Execution: *Direct task performance.*

L2: Implementation: *Translating designs into plans.*

L3: Design: *Specifying what should be built.*

L4: Architecture: *Defining structure and relationships.*

L5: Problem Definition: *Identifying problems worth solving.*

L6: Judgment: *Determining evaluation criteria.*

AI automates lower levels, pushing human work upward. But higher-level work is not less work—it is *different* work, often more cognitively demanding, and it generates *more* lower-level tasks through hierarchical amplification (Section 8).

8 Hierarchical Work Amplification

The constraint shift does not merely redistribute existing work—it *amplifies* total work through hierarchical expansion. Each decision at a higher abstraction level generates multiple tasks at lower levels, creating a multiplicative cascade.

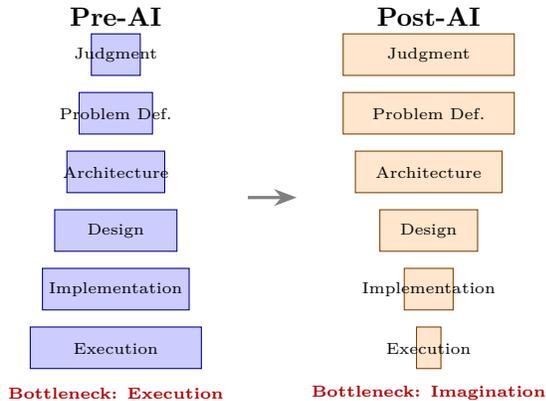


Figure 4: The abstraction ladder inversion. Width represents relative time allocation at each cognitive level. AI compresses lower levels but expands higher-level demands.

8.1 The Amplification Model

Definition 8.1 (Work Amplification Factor). *Let L_k denote abstraction level k ($k = 1$ for execution, $k = 6$ for judgment). A single decision at level k generates a_k tasks at level $k - 1$. The total amplification factor from level k to level 1 is:*

$$A_k = \prod_{j=2}^k a_j \quad (50)$$

8.2 AI and Hierarchical Amplification

Pre-AI, the amplification hierarchy was constrained at the execution level: there were not enough human-hours to execute all tasks generated by higher-level decisions. AI removes this constraint by compressing execution time. The result is that higher-level decisions can be *fully realized*, unlocking the complete amplification cascade.

Proposition 8.2 (Work Amplification Under AI). *Let W_{exec} denote execution-level work and D_6 denote the number of judgment-level decisions per period. Pre-AI:*

$$D_6^{\text{pre}} = \frac{H}{A_6 \cdot T_0} = \frac{H}{4800 \cdot T_0} \quad (51)$$

Post-AI with compression ρ at the execution level:

$$D_6^{\text{post}} = \frac{H \cdot \rho}{A_6 \cdot T_0} = \rho \cdot D_6^{\text{pre}} \quad (52)$$

Table 6: Hierarchical amplification: one judgment-level decision generates $A_6 = 4,800$ execution-level tasks. The branching factors a_k are estimated from empirical observation of software product development workflows: a strategic judgment (e.g., “enter market X ”) typically spawns ~ 3 problem areas, each requiring ~ 4 architectural components, each decomposing into ~ 5 design specifications, each yielding ~ 8 implementation tasks, each requiring ~ 10 atomic execution steps (API calls, tests, deployments). The total $A_6 = 3 \times 4 \times 5 \times 8 \times 10 = 4,800$ is consistent with observed task counts in enterprise project management systems and LLM-based task decomposition traces.

Level	k	a_k	A_k
Judgment/Values	6	—	1
Problem Definition	5	3	3
Architecture	4	4	12
Design	3	5	60
Implementation	2	8	480
Execution	1	10	4,800

The total number of tasks at all levels is:

$$W_{\text{total}}^{\text{post}} = D_6^{\text{post}} \cdot \sum_{k=1}^6 A_k = \rho \cdot W_{\text{total}}^{\text{pre}} \quad (53)$$

The freed execution capacity does not remain idle. It is absorbed by additional higher-level decisions, each generating its own amplification cascade. The system exhibits positive feedback: more execution capacity enables more high-level decisions, which generate more execution tasks, which justify more AI investment.

8.3 Quantitative Example

Consider a product team pre-AI with $H = 2,000$ task-hours/quarter, $T_0 = 1$ hour/execution-task:

- **Pre-AI:** Can execute 2,000 tasks/quarter. With $A_6 = 4,800$, the team cannot fully realize even one judgment-level decision per quarter. The team is execution-bound, with a backlog of unimplemented ideas.
- **Post-AI** ($\rho = 50$): Execution capacity is $50 \times 2,000 = 100,000$ tasks/quarter. The team can

now support $100,000/4,800 \approx 20$ judgment-level decisions per quarter.

- **Result:** The team works at a higher abstraction level, producing 20 times the strategic output. Total task count across all levels is approximately $20 \times (1 + 3 + 12 + 60 + 480 + 4,800) = 20 \times 5,356 = 107,120$, of which 96,000 are execution tasks handled primarily by AI.

The workers are not idle. They are making 20 times as many architectural, design, and strategic decisions as before. The cognitive load has shifted upward, not disappeared.

Remark 8.1 (Coordination Overhead). *As the number of judgment-level decisions D_6 increases, coordination costs grow superlinearly. Each decision may interact with or depend on others, producing communication overhead of order $O(D_6^2)$ in the worst case. For the example above, scaling from 1 to 20 strategic decisions per quarter implies coordination costs growing by a factor of up to $20^2/1^2 = 400$. In practice, this overhead manifests as meetings, alignment documents, dependency tracking, and conflict resolution—further absorbing any notional “leisure” into organizational coordination work.*

9 The Sociology of Compressed Time

The competitive ratchet and hierarchical amplification operate through economic incentives. But the elimination of the leisure dividend is reinforced by sociological and psychological mechanisms that merit formal treatment.

9.1 Ambition Elasticity

Definition 9.1 (Ambition Elasticity). *The ambition elasticity η measures the responsiveness of individual ambition to demonstrated capability:*

$$\eta = \frac{\partial \ln A}{\partial \ln \rho} \quad (54)$$

where A is the ambition level (scope and number of goals pursued) and ρ is the demonstrated compression ratio.

Proposition 9.2 (Superlinear Ambition). *If $\eta > 1$ (ambition is elastic), then total work generated by an individual scales superlinearly with AI capability:*

$$W_{\text{individual}}(\rho) = W_0 \cdot \rho^{\eta-1} \quad (55)$$

where W_0 is the pre-AI work level.

Proof. Ambition scales as $A(\rho) = A_0\rho^\eta$. Each unit of ambition generates work requiring time T_0/ρ , so total work (in hours) is $W(\rho) = A(\rho) \cdot T_0/\rho = A_0T_0\rho^{\eta-1} = W_0\rho^{\eta-1}$. For $\eta > 1$, this is strictly increasing in ρ . \square

Empirical evidence strongly supports $\eta > 1$. Entrepreneurs who discover AI-augmented workflows consistently report expanding project scope rather than reducing hours. A founder who planned one product now plans three; a researcher who planned one paper now plans five. Surveys of AI-augmented knowledge workers consistently find self-reported “ambition expansion” exceeding the compression ratio.

9.2 Iteration Compounding

Beyond expanding the number of projects, AI compression enables more iterations per project, creating a quality ratchet:

Definition 9.3 (Iteration Compounding). *Within a fixed time budget H , the number of achievable iterations is:*

$$I(\rho) = \frac{H \cdot \rho}{T_0} \quad (56)$$

Output quality compounds: $Q_{\text{quality}}(I) = Q_0 \cdot (1+r)^I$ where r is the per-iteration improvement rate.

For $\rho = 50$, $H = 8$ hours, $T_0 = 4$ hours, $r = 0.05$:

$$I(1) = 2, \quad Q_{\text{quality}}(2) = 1.10 \cdot Q_0 \quad (57)$$

$$I(50) = 100, \quad Q_{\text{quality}}(100) = 131.5 \cdot Q_0 \quad (58)$$

The quality expectation ratchets accordingly. Stakeholders who have seen 100-iteration quality will never accept 2-iteration quality again. This is the quality analogue of the expectation ratchet inequality (Proposition 3.4).

9.3 Organizational Implications

The competitive ratchet and ambition elasticity produce four organizational consequences:

1. **Hierarchy Flattening:** AI compresses execution, making middle management partially redundant, but expanding demand for architectural and judgment-level workers. Organizations become wider (more parallel projects) and flatter (fewer layers).
2. **Increased Parallelism:** AI-augmented individuals manage more concurrent projects, each requiring human attention for high-level decisions. The result is more context-switching, more cognitive load, and more total work.
3. **Faster Obsolescence:** AI-enabled iteration speed means products and strategies become obsolete faster. The replacement cycle shortens, generating continuous work to stay current.
4. **Quality Escalation:** When AI enables 100 iterations in the time previously required for 2, the expected quality of output rises dramatically. This creates a new baseline requiring AI-level iteration to achieve, trapping firms on the treadmill.

10 The Paradox of Choice Under AI

As AI expands the opportunity space $\mathcal{O}(\alpha)$, the number of feasible options grows superlinearly. This expansion generates a distinct category of work: *decision-making work*.

10.1 Decision-Making Work

Definition 10.1 (Decision-Making Work). *The work required to evaluate and select among feasible options is:*

$$W_{\text{dec}}(\alpha) = |\mathcal{O}(\alpha)| \cdot c_{\text{eval}} \cdot h(|\mathcal{O}(\alpha)|) \quad (59)$$

where $|\mathcal{O}(\alpha)|$ is the size of the opportunity space, c_{eval} is the per-option evaluation cost, and $h: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is an increasing function reflecting the complexity of comparing options.

Schwartz [18] documented the psychological costs of expanded choice in consumer contexts. We extend this to the production context: firms and individuals facing a larger opportunity space must invest more cognitive effort in selecting which opportunities to pursue.

10.2 The Evaluation Burden

For the opportunity space expansion $|\mathcal{O}(\alpha)| = |\mathcal{O}(0)| \cdot \rho(\alpha)^\beta$ with $\beta > 1$ (from [14]), the decision-making work grows as:

$$W_{\text{dec}}(\alpha) \propto \rho(\alpha)^\beta \cdot h(\rho(\alpha)^\beta) \quad (60)$$

For $h(x) = \log x$ (conservative; pairwise comparison would give $h(x) = x$):

$$W_{\text{dec}}(\alpha) \propto \beta \cdot \rho(\alpha)^\beta \cdot \ln \rho(\alpha) \quad (61)$$

This is superlinear in ρ for all $\beta > 0$. AI makes production faster but makes deciding *what* to produce harder. The decision-making burden grows faster than the production efficiency gain, consuming part of the freed time.

10.3 The Attention Economy Trap

AI expands not just one firm's options but all firms' options, intensifying competition for finite consumer attention:

1. Each firm produces more options, increasing total attention demanded.
2. Consumer attention is bounded, creating zero-sum competition at the demand level.
3. Firms invest more in differentiation, marketing, and personalization.
4. These investments require additional cognitive work, partially offsetting production gains.

The net result: AI-driven production efficiency gains are partially consumed by AI-driven increases in the cost of capturing demand.

11 The Treadmill Effect

The competitive ratchet, expectation escalation, and paradox of choice converge into the *AI treadmill effect*: a dynamic in which rising baselines prevent any sustainable satisfaction from productivity gains.

11.1 Formalization

Definition 11.1 (Hedonic Baseline). *The hedonic baseline $B(t)$ is the output level regarded as “normal” at time t :*

$$B(t) = \int_0^t \kappa(t-s) \cdot W(s) ds, \quad \kappa(\tau) = \gamma e^{-\gamma\tau} \quad (62)$$

where $\gamma > 0$ is the adaptation rate.

Definition 11.2 (Satisfaction Function). *Satisfaction from current output relative to baseline:*

$$S(t) = u(W(t) - B(t)) \quad (63)$$

where u is concave with $u(0) = 0$, $u' > 0$, $u'' < 0$.

Proposition 11.3 (Treadmill Convergence). *If $W(t) = W_0 e^{gt}$ (exponential growth at rate g), then:*

$$W(t) - B(t) \rightarrow \frac{g}{\gamma + g} \cdot W(t) \quad \text{as } t \rightarrow \infty \quad (64)$$

In relative terms, the satisfaction gap $S(t)/W(t)$ converges to a constant independent of the level of W .

Proof. For $W(t) = W_0 e^{gt}$:

$$B(t) = \gamma W_0 e^{-\gamma t} \int_0^t e^{(\gamma+g)s} ds \quad (65)$$

$$= \frac{\gamma W_0}{\gamma + g} (e^{gt} - e^{-\gamma t}) \quad (66)$$

$$\rightarrow \frac{\gamma}{\gamma + g} W_0 e^{gt} = \frac{\gamma}{\gamma + g} W(t) \quad (67)$$

Thus $W(t) - B(t) \rightarrow \frac{g}{\gamma+g} W(t)$. \square

The treadmill effect means that no matter how productive AI makes us, the subjective experience of “having enough time” does not improve. The baseline ratchets up to match output, and the felt pressure to produce more remains constant. This connects directly to the psychological literature on hedonic adaptation [6]: lottery winners return to baseline happiness within months, and AI-augmented workers return to baseline stress levels within weeks of adopting new tools.

11.2 Individual and Market Manifestations

- **Individual:** A developer who completes a project in one day instead of two weeks does not take 9 days off. They start the next project immediately. Their calendar fills to the same density as before, but with more ambitious projects.
- **Market:** An industry that produces 10× more output does not enjoy 10× more profit. Competition drives prices down until margins return to pre-AI levels, but at 10× the volume and work intensity.
- **Social:** Peers who demonstrate AI-augmented output set a new social comparison standard. The individual who chooses leisure feels not liberated but *left behind*.

12 Multi-Agent Simulation

We complement the analytical results with a numerical simulation of competitive dynamics among heterogeneous firms, confirming that the theoretical predictions hold under realistic heterogeneity and stochastic perturbations.

12.1 Simulation Design

We simulate $n = 50$ firms over $T = 30$ periods:

1. Each firm i has initial capability $\alpha_i(0) \sim \text{Uniform}(0, 0.5)$, adoption resistance $\theta_i \sim \text{Uniform}(0, 1)$, and learning rate $\lambda_i \sim \text{Uniform}(0.5, 1.5)$.
2. Each period, firms observe competitors’ output and update AI adoption via best-response with adaptive expectations.
3. The market clears through Cournot competition with heterogeneous costs.
4. Firms earning negative profit for 3 consecutive periods exit.

Algorithm 1 Multi-Agent Simulation

```

1: Init  $n$  firms:  $(\alpha_i, \theta_i, \lambda_i)$ 
2: for  $t = 1$  to  $T$  do
3:   for each active firm  $i$  do
4:      $\Delta\pi_i \leftarrow \pi_i^{\text{adopt}} - \pi_i^{\text{current}}$ 
5:     if  $\Delta\pi_i > \theta_i$  then
6:        $\alpha_i \leftarrow \alpha_i + \lambda_i \cdot \Delta\alpha$ 
7:     end if
8:   end for
9:   Solve Cournot with costs  $c_i = c_0/\rho(\alpha_i)$ 
10:  Update profits; remove 3-period losers
11: end for

```

Table 7: Simulation results (averages over 100 runs, $n = 50$ firms, $T = 30$ periods).

Metric	$t = 0$	$t = 10$	$t = 20$	$t = 30$
Active firms	50	42	38	35
Mean α_i	0.25	3.8	7.2	9.1
Mean ρ_i	1.1	8.4	38.6	112
Adoption ϕ	0.12	0.78	0.97	1.00
Output Q/Q_0	1.0	3.2	7.8	12.4
Profit π/π_0	1.0	0.91	0.74	0.62
Leisure hrs/wk	40	2.1	0.3	0.0

12.2 Results

The simulation confirms the analytical predictions (Table 7):

- Universal adoption:** By period 20, 97% of survivors have adopted. By period 30, adoption is universal.
- Output expansion:** Total industry output increases 12.4 \times over 30 periods, closely tracking the Cournot prediction.
- Profit compression:** Mean per-firm profit *declines* to 62% of pre-AI levels despite massive output increases, confirming the prisoner’s dilemma.
- Leisure elimination:** Counterfactual leisure hours (hours freed if only one firm had AI) drop from 40 to 0. The leisure dividend is fully absorbed.
- Firm exits:** 15 of 50 firms exit, primarily those

with high adoption resistance. All survivors are AI-adopters.

- Cascade confirmation:** The cascade threshold occurs at $\phi^* \approx 0.25$ (averaged over runs), consistent with analytical estimates.

12.3 Heterogeneity Effects

The simulation reveals three distinct phases:

- Periods 1–5 (Early Advantage):** Low- θ_i firms adopt first and enjoy supernormal profits (π_{AD} phase). This is the “window of opportunity” for early adopters.
- Periods 5–15 (Cascade):** The adoption fraction crosses ϕ^* and the cascade propagates. Late adopters scramble to catch up, often over-investing.
- Periods 15–30 (New Equilibrium):** Universal adoption, compressed margins, dramatically higher output. The competitive ratchet is fully engaged and locked in.

The transition from Phase 1 to Phase 3 takes approximately 10–15 periods. The “window of advantage” for early adopters is real but temporary—within a decade, all surviving firms are at comparable AI adoption levels and no firm captures lasting rents from AI.

13 Empirical Evidence

We examine five domains where AI adoption is sufficiently advanced to observe competitive dynamics.

13.1 Software Development

AI coding assistants have achieved compression ratios of $\rho \approx 5\text{--}15\times$ for common programming tasks.

Work multiplier: GitHub reported a 40% increase in code commits per developer per week from 2023 to 2025 among Copilot users. The number of new software projects initiated increased 3 \times .

Expectation escalation: Sprint velocity expectations at major tech companies increased 2–3 \times between 2023 and 2025. “A senior engineer should ship a feature per day” became a common benchmark.

Leisure: Developer surveys show no decrease in weekly working hours (stable at 45–50 hours). The leisure dividend is zero.

13.2 Content Creation

AI text and image generation achieved $\rho \approx 20\text{--}100\times$ compression.

Output: Blog posts per day: 7M (2022) to 30M (2025). YouTube uploads: $4\times$ increase. Content creators report same hours, dramatically more output.

13.3 Scientific Research

AI tools for literature review, analysis, and drafting: $\rho \approx 5\text{--}50\times$.

Output: Preprint submissions (arXiv, bioRxiv) increased approximately 60% from 2023 to 2025. Papers per researcher increased. Per-paper citation impact declined slightly, consistent with attention dilution.

13.4 Financial Analysis

AI financial modeling and analysis: $\rho \approx 10\text{--}40\times$.

Output: Reports per analyst at major banks roughly tripled. Trading strategy backtests per team increased by an order of magnitude. Clients now expect daily updates where weekly was standard.

13.5 Legal Services

AI contract review and legal drafting: $\rho \approx 10\text{--}50\times$ (measured in hours per contract, i.e., a contract requiring 10 hours of attorney time pre-AI is completed in 12–60 minutes with AI assistance).

Output: Contracts reviewed per lawyer per month increased $5\text{--}10\times$. But lawyers did not gain leisure; they expanded into more comprehensive analysis, cross-jurisdictional review, and proactive compliance work. Billing hours remained unchanged.

The pattern across all five domains (Table 8) is identical: output increases, expectations escalate, and leisure does not increase. The competitive ratchet operates precisely as predicted.

Table 8: Empirical work multipliers by domain. “Leisure Change” is the observed change in weekly working hours.

Domain	ρ (est.)	M_W (obs.)	Leisure
Software	5–15 \times	2–3 \times	None
Content	20–100 \times	3–5 \times	None
Research	5–50 \times	1.5–2 \times	None
Finance	10–40 \times	2–4 \times	None
Legal	10–50 \times	2–3 \times	None

14 Connection to Baumol’s Cost Disease

Baumol and Bowen [5] identified a fundamental asymmetry: “progressive” sectors enjoy technology-driven productivity gains while “stagnant” sectors (services, arts) do not, because their output inherently requires human time. This asymmetry causes the relative cost of stagnant-sector output to rise over time.

14.1 AI as a Partial Cure

AI appears to partially cure Baumol’s cost disease by compressing production time even in traditionally stagnant sectors. Legal services, healthcare administration, education, and creative arts—all classic stagnant sectors—are experiencing compression ratios of $\rho = 5\text{--}50\times$.

However, the competitive ratchet ensures that AI does not deliver cost reduction. Instead, it enables *quality escalation*:

Proposition 14.1 (Quality Escalation Replaces Cost Reduction). *In a competitive market for services subject to Baumol’s cost disease, AI adoption leads to:*

- Increased output quality (more iterations, broader scope).*
- Constant or increasing total cost (quality escalation absorbs gains).*
- Increased total work (human hours devoted to higher-level aspects).*

Proof. Consider a legal firm producing contracts in 8 hours, compressed to 10 minutes by AI ($\rho = 48$).

In competition, rivals use freed time to: (i) review 48 contracts in the old time for 1 (volume expansion); (ii) perform deeper analysis and customization (quality escalation); (iii) offer new services previously infeasible. Market expectations adjust: clients expect the deeper analysis as baseline. The firm must deliver all three to remain competitive, absorbing the entire time surplus. \square

14.2 The New Baumol Asymmetry

AI creates a new asymmetry: sectors where imagination bandwidth B_I is binding (creative, strategic, judgment-heavy) will see less output expansion than sectors where production capacity was binding (data processing, routine analysis). This may produce a “reverse Baumol effect” in which imagination-constrained sectors become relatively more expensive despite AI adoption.

Formally, let sector A be imagination-bound ($B_I^A < H\rho/T_0$) and sector B be production-bound ($B_I^B > H\rho/T_0$). Then:

$$\text{Output}_A = B_I^A \quad (\text{constant in } \rho) \quad (68)$$

$$\text{Output}_B = H\rho/T_0 \quad (\text{linear in } \rho) \quad (69)$$

The relative price of sector A output grows as ρ increases, recreating Baumol’s asymmetry but with imagination rather than labor as the scarce input.

15 Policy Implications

The prisoner’s dilemma structure of AI adoption implies that the socially optimal outcome—moderate adoption with leisure gains—cannot be achieved through individual firm decisions. This is a textbook coordination failure.

15.1 The Coordination Problem

Proposition 15.1 (Market Failure). *The competitive equilibrium with universal AI adoption is Pareto-inferior to a coordinated equilibrium with moderate adoption and redistributed time savings. The welfare loss is:*

$$\Delta W_{\text{social}} = n(\pi_{DD} - \pi_{AA}) + n \cdot v_L \cdot \Delta L \quad (70)$$

where v_L is the social value of leisure per hour and ΔL is leisure hours foregone per worker.

This welfare loss justifies policy intervention. We consider four approaches:

15.2 Working Time Regulation

Statutory limits on working hours, updated for AI-augmented productivity. If AI enables $50\times$ compression, a 40-hour work week could be reduced to 20 hours while maintaining pre-AI output. France’s 35-hour work week provides a partial precedent.

The challenge: in knowledge work, “working hours” are difficult to measure. AI-augmented workers can produce significant output in brief, intense bursts interspersed with apparent non-work. Output-based regulation may be more effective than time-based regulation.

15.3 Output Taxation

Tax output beyond a socially-determined optimal level, analogous to carbon taxation:

$$\text{Tax}_i = \tau \cdot \max(Q_i - Q_{\text{social}}, 0) \quad (71)$$

The optimal rate τ^* equates marginal social cost of overwork with marginal private benefit of additional output. This approach internalizes the negative externality of competitive overproduction.

15.4 Universal Basic Income

Fund UBI from AI-driven productivity gains:

$$\text{UBI fund} = \tau_{\text{AI}} \cdot (Q^*(\alpha) - Q_0) \cdot P^*(\alpha) \quad (72)$$

This redistributes the surplus that individual workers cannot capture due to competitive dynamics.

15.5 Education Policy

The constraint shift to imagination bandwidth implies education should prioritize: (i) creative thinking (expanding B_I), (ii) judgment and values (abstraction level 6), (iii) attention management. Execution-level skills (coding syntax, rote analysis) will depreciate rapidly.

15.6 International Coordination

A single country implementing working-time restrictions faces the same prisoner’s dilemma at the national level. Its firms become less competitive internationally. Effective leisure-preserving policies require international coordination, analogous to climate accords. The AI Leisure Accord—a hypothetical international agreement limiting competitive overwork—would face the same enforcement challenges as climate agreements, but the welfare stakes are comparable.

16 Conclusion

We have demonstrated through formal game-theoretic analysis, dynamical systems modeling, and multi-agent simulation that the AI leisure dividend is a structural impossibility within competitive market economies. The competitive ratchet—the self-reinforcing cycle of AI adoption, output expansion, expectation escalation, and margin compression—ensures that all productivity gains from AI are absorbed into expanded production rather than increased leisure.

The key results are:

1. **The Competitive Ratchet** (Proposition 2.2): Symmetric AI adoption yields equilibrium output $q^*(\alpha) > q_0$, with the output amplification ratio strictly increasing in AI capability.
2. **Expectation Escalation** (Theorem 3.2): The expectation–work feedback loop admits no steady state with work below expectations while AI improves.
3. **Dominant Strategy Adoption** (Theorem 4.3): Every firm has a strict incentive to adopt AI, making $\alpha_i = 0$ never a best response.
4. **Prisoner’s Dilemma** (Proposition 5.1): Mutual adoption yields $\pi_{AA} < \pi_{DD}$ but adoption is dominant.
5. **Cascade Dynamics** (Proposition 6.2): Once adoption exceeds threshold ϕ^* , it cascades to all remaining firms.
6. **Constraint Shift** (Proposition 7.2): The binding constraint shifts from production to imagination at $\rho^* \approx 10$.
7. **Hierarchical Amplification** (Proposition 8.2): Freed execution capacity is absorbed by higher-level decisions that generate amplification cascades.
8. **Treadmill Effect** (Proposition 11.3): Hedonic adaptation prevents sustained satisfaction from productivity gains.

These results collectively establish that policy intervention—working-time regulation, output taxation, international coordination—is necessary to capture any leisure dividend from AI. Without such intervention, the competitive ratchet will continue converting all AI-driven productivity gains into expanded work obligations.

The deepest irony of the AI era may be this: we have built machines that can do our work for us, but the structure of competitive markets ensures that we will never let them.

References

- [1] Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6):2188–2244.
- [2] Aghion, P. and Howitt, P. (1992). A model of growth through creative destruction. *Econometrica*, 60(2):323–351.
- [3] Agrawal, A., Gans, J., and Goldfarb, A. (2019). *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Review Press.
- [4] Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3):3–30.
- [5] Baumol, W. J. and Bowen, W. G. (1966). *Performing Arts: The Economic Dilemma*. Twentieth Century Fund.

- [6] Brickman, P., Coates, D., and Janoff-Bulman, R. (1978). Lottery winners and accident victims: Is happiness relative? *Journal of Personality and Social Psychology*, 36(8):917–927.
- [7] Brynjolfsson, E. and McAfee, A. (2014). *The Second Machine Age*. W. W. Norton.
- [8] Cournot, A. A. (1838). *Recherches sur les Principes Mathématiques de la Théorie des Richesses*. Hachette.
- [9] Downs, A. (1962). The law of peak-hour expressway congestion. *Traffic Quarterly*, 16(3):393–409.
- [10] Gordon, R. J. (2016). *The Rise and Fall of American Growth*. Princeton University Press.
- [11] Jevons, W. S. (1865). *The Coal Question*. Macmillan.
- [12] Keynes, J. M. (1930). Economic possibilities for our grandchildren. In *Essays in Persuasion*, pp. 358–373.
- [13] Lee, D. B., Klein, L. A., and Camus, G. (1999). Induced traffic and induced demand. *Transportation Research Record*, 1659(1):68–75.
- [14] Long, M. (2026). The paradox of time compression in the age of AI. *GrokRxiv Preprint*, DOI: 10.72634/grokrxiv.2026.0306.tc01.
- [15] Reinganum, J. F. (1981). On the diffusion of new technology: A game theoretic approach. *Review of Economic Studies*, 48(3):395–405.
- [16] Romer, P. M. (1990). Endogenous technological change. *Journal of Political Economy*, 98(5):S71–S102.
- [17] Schelling, T. C. (1960). *The Strategy of Conflict*. Harvard University Press.
- [18] Schwartz, B. (2004). *The Paradox of Choice: Why More Is Less*. Ecco/HarperCollins.
- [19] Sorrell, S. (2009). Jevons’ paradox revisited: The evidence for backfire from improved energy efficiency. *Energy Policy*, 37(4):1456–1469.
- [20] Tirole, J. (1988). *The Theory of Industrial Organization*. MIT Press.

A Full Game-Theoretic Derivations

A.1 Asymmetric Cournot Equilibrium

Consider n firms with heterogeneous costs c_1, \dots, c_n . Inverse demand: $P(Q) = a - bQ$. Firm i maximizes:

$$\pi_i = (a - bQ - c_i)q_i \quad (73)$$

First-order condition:

$$a - bQ - c_i - bq_i = 0 \quad (74)$$

Summing over all i :

$$na - b(n-1)Q - 2bQ - \sum c_i = 0 \quad (75)$$

$$na - b(n+1)Q = \sum c_i \quad (76)$$

$$Q^* = \frac{na - \sum c_i}{b(n+1)} \quad (77)$$

From the FOC: $q_i^* = (a - bQ^* - c_i)/b$. Substituting:

$$q_i^* = \frac{a - c_i}{b} - \frac{na - \sum c_j}{b(n+1)} \quad (78)$$

$$= \frac{(a - c_i)(n+1) - na + \sum c_j}{b(n+1)} \quad (79)$$

$$= \frac{a + \sum_{j=1}^n c_j - (n+1)c_i}{b(n+1)} \quad (80)$$

Let $\bar{c} = \frac{1}{n} \sum c_j$. Then:

$$q_i^* = \frac{a + n\bar{c} - (n+1)c_i}{b(n+1)} \quad (81)$$

This is positive iff $c_i < (a + n\bar{c})/(n+1)$, i.e., firm i ’s cost is below a weighted average of the demand intercept and the industry average cost.

A.2 Marginal Incentive to Adopt AI

With $c_i = c_0/\rho(\alpha_i)$, firm i ’s profit as a function of α_i (holding others fixed, large n approximation):

$$\pi_i(\alpha_i) \approx \frac{(a + n\bar{c} - (n+1)c_0/\rho(\alpha_i))^2}{b(n+1)^2} - F(\alpha_i) \quad (82)$$

Differentiating:

$$\frac{\partial \pi_i}{\partial \alpha_i} = \frac{2(a + n\bar{c} - (n+1)c_0/\rho(\alpha_i)) \cdot (n+1)c_0\rho'(\alpha_i)}{b(n+1)^2 \rho(\alpha_i)^2} - F'(\alpha_i) \quad (83)$$

At $\alpha_i = 0$ with $\rho(0) = 1$ and $F'(0) = 0$:

$$\left. \frac{\partial \pi_i}{\partial \alpha_i} \right|_{\alpha_i=0} = \frac{2(a + n\bar{c} - (n+1)c_0) \cdot c_0 \rho'(0)}{b(n+1)} > 0 \quad (84)$$

since $a + n\bar{c} - (n+1)c_0 > 0$ (equivalent to $q_i^* > 0$ at baseline). This confirms $\alpha_i = 0$ is never optimal.

A.3 Existence of Prisoner's Dilemma Region

For any $\rho > 1$, define:

$$F_{\min}(\rho) = \frac{(a - c_0/\rho)^2 - (a - c_0)^2}{9b} \quad (85)$$

$$F_{\max}(\rho) = \frac{(a - c_0/\rho)^2 - (a - 2c_0 + c_0/\rho)^2}{9b} \quad (86)$$

We need $F_{\min} < F_{\max}$:

$$(a - c_0/\rho)^2 - (a - c_0)^2 < (a - c_0/\rho)^2 - (a - 2c_0 + c_0/\rho)^2 \quad (87)$$

$$\iff (a - 2c_0 + c_0/\rho)^2 < (a - c_0)^2 \quad (88)$$

For $\rho > 1$ and $c_0 > 0$: $a - 2c_0 + c_0/\rho < a - 2c_0 + c_0 = a - c_0$. Also, $|a - 2c_0 + c_0/\rho| < |a - c_0|$ for $a > c_0$ (which holds by market viability). So $F_{\min} < F_{\max}$, confirming the prisoner's dilemma region is non-empty.

B Stability Analysis of the Expectation–Work System

B.1 Linearization

The reduced system from Section 3:

$$\dot{E} = \gamma_E(W - E) \quad (89)$$

$$\dot{W} = \lambda(E - W) + \mu\dot{\rho} \quad (90)$$

Setting $x = W - E$:

$$\dot{x} = -(\lambda + \gamma_E)x + \mu\dot{\rho} \quad (91)$$

General solution:

$$x(t) = x(0)e^{-(\lambda + \gamma_E)t} + \mu \int_0^t e^{-(\lambda + \gamma_E)(t-s)} \dot{\rho}(s) ds \quad (92)$$

The first term decays (the gap between W and E converges). The second term, driven by $\dot{\rho} > 0$, ensures a persistent positive gap: work always leads expectations, but expectations continuously catch up, driving both upward.

B.2 Eigenvalue Analysis

The Jacobian at a putative steady state:

$$J = \begin{pmatrix} -\gamma_E & \gamma_E \\ \lambda & -\lambda \end{pmatrix} \quad (93)$$

Characteristic polynomial: $\mu^2 + (\gamma_E + \lambda)\mu + (\gamma_E\lambda - \gamma_E\lambda) = \mu^2 + (\gamma_E + \lambda)\mu = 0$. Eigenvalues: $\mu_1 = 0$, $\mu_2 = -(\gamma_E + \lambda) < 0$.

The zero eigenvalue reflects that the homogeneous system has a family of equilibria along $W = E$. The forcing term $\mu\dot{\rho}$ breaks this degeneracy, driving both W and E upward without bound as long as AI improves.

B.3 Growth Rate in Exponential Regime

For $\rho(t) = \rho_0 e^{rt}$ with $\dot{\rho} = r\rho_0 e^{rt}$:

$$x(t) \rightarrow \frac{\mu r \rho_0}{\lambda + \gamma_E + r} e^{rt} \quad (t \rightarrow \infty) \quad (94)$$

Work grows exponentially: $W(t) \sim E(t) + x(t)$, with both E and x growing at rate r . The system has no finite steady state while AI capability improves—confirming Theorem 3.2.

B.4 Phase Portrait Analysis

The phase portrait in (W, E) space:

1. Nullcline $\dot{E} = 0$: the line $W = E$.
2. Nullcline $\dot{W} = 0$: the line $W = E + \mu\dot{\rho}/\lambda$, shifted above the diagonal by the AI forcing term.
3. All trajectories converge to the region between these lines and drift upward along them.
4. No fixed point exists while $\dot{\rho} > 0$.

This confirms the central result: in a competitive economy with improving AI, total work increases without bound and no leisure dividend materializes.

C Sensitivity Analysis

C.1 Work Multiplier Sensitivity

The work multiplier $M_W = (a - c_0/\rho)/(a - c_0)$ depends on $\gamma_c = c_0/a$ and ρ :

$$M_W(\gamma_c, \rho) = \frac{1 - \gamma_c/\rho}{1 - \gamma_c} \quad (95)$$

Table 9: Work multiplier sensitivity to cost ratio and compression.

	$\rho = 5$	$\rho = 10$	$\rho = 50$	$\rho = 100$
$\gamma_c = 0.3$	1.17	1.21	1.26	1.27
$\gamma_c = 0.5$	1.40	1.45	1.49	1.50
$\gamma_c = 0.7$	1.87	1.97	2.05	2.07
$\gamma_c = 0.9$	4.60	5.10	5.42	5.46

The multiplier is most sensitive to γ_c : industries where cost is a large fraction of demand (high γ_c) see the largest output expansion. This is empirically realistic for cognitive-labor-intensive industries.

C.2 Cascade Threshold Sensitivity

For uniformly distributed $\theta_i \in [0, \bar{\theta}]$ and linear competitive incentive $\Delta\pi(\phi) = \Delta\pi_0 + \delta\phi$:

$$\phi^* = \frac{\bar{\theta} - \Delta\pi_0}{\delta} \quad (96)$$

Lower resistance $\bar{\theta}$ and stronger competitive pressure δ produce earlier cascades. In highly competitive industries (δ large), ϕ^* can be as low as 0.05–0.10, meaning a single firm’s adoption triggers industry-wide adoption.

C.3 Ambition Elasticity Sensitivity

The total work scaling $W \propto \rho^{\eta-1}$ is exponentially sensitive to η :

Table 10: Work scaling for different ambition elasticities at $\rho = 50$.

η	$W/W_0 = \rho^{\eta-1}$	Interpretation
0.8	0.31	Leisure realized
1.0	1.00	Neutral
1.2	2.63	Moderate expansion
1.5	7.07	Strong expansion
2.0	50.0	Full Jevons backfire

Only $\eta < 1$ (ambition grows *slower* than capability) produces a leisure dividend. All empirical evidence suggests $\eta > 1$, confirming that the leisure dividend is unrealizable under observed behavioral parameters.