The Jevons Paradox of Intelligence:

Elastic Demand and the Paradox of Cognitive Abundance

Matthew Long[1]

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

March 6, 2026

Abstract

We present a formal economic analysis of the application of the Jevons Paradox to intelligence treated as an economic resource. The classical Jevons Paradox (1865) established that improvements in the efficiency of coal consumption led not to reduced coal usage but to dramatically increased total consumption, because cheaper energy unlocked new applications with elastic demand. We demonstrate that an identical mechanism operates when artificial intelligence reduces the effective price of cognitive labor. We define a cognitive price function incorporating both human time costs and AI compute costs, prove that under empirically justified elastic demand conditions ($\varepsilon > 1$), total cognitive expenditure is strictly increasing as AI capability grows, and establish a formal backfire criterion. We provide extensive empirical evidence from software development, content creation, legal analysis, scientific research, and financial services, showing task-specific peak compression ratios ranging from 14× to 288× across cognitive task categories (with aggregate sector-level ratios averaging $\sim 22\times$; see §9). We develop a logistic compression model with calibrated parameters, analyze the three fates of the resulting time surplus (leisure, intensification, and expansion), and derive the conditions under which expansion dominates. Sector-level analysis of five industries quantifies the Jevons effect in action, with total cognitive labor increasing by factors of 3× to 25× within initial adoption windows. We connect the intelligence Jevons Paradox to Romer-style endogenous growth theory, showing that the cognitive backfire effect acts as a double accelerant—increasing both the effective research labor force and the knowledge stock simultaneously. Quantitative projections based on calibrated models suggest total cognitive work will increase by 10–100× over pre-AI baselines within a generation. We conclude with a critical examination of counter-arguments, including the satiation hypothesis, regulatory constraints, and the possibility that certain cognitive domains may exhibit inelastic demand.

Keywords: Jevons Paradox, cognitive labor, demand elasticity, AI economics, backfire effect, rebound effect, endogenous growth, cognitive price, time compression, productivity paradox

1 Introduction

In 1865, the English economist William Stanley Jevons posed a question that would resonate across centuries of technological change: if steam engines were becoming ever more efficient in their use of coal, would Britain's coal supply last longer? His answer was emphatically negative. More efficient engines did not conserve coal; they consumed *more* of it, because the reduced cost per unit of useful work made coal economically viable for an ever-expanding range of applications [1]. This observation—the *Jevons Paradox*—has since been recognized as a fundamental principle governing the relationship between resource efficiency and total resource consumption.

We argue that artificial intelligence is doing to *cognitive labor* what the steam engine did to coal. AI dramatically reduces the cost of producing a unit of cognitive work—writing a legal brief, debugging software, analyzing a dataset, translating a document. The naive expectation is that this cost reduction will lead to less total cognitive labor: fewer person-hours devoted to thinking, writing, analyzing, and deciding. The Jevons Paradox predicts the opposite. When the price of an enabling input with elastic demand falls, total expenditure on that input *increases*, not decreases. The efficiency gain is more than offset by the expansion of demand into previously infeasible applications.

This paper formalizes the application of the Jevons Paradox to intelligence, treated as an economic resource subject to standard demand theory. Our contributions are:

(i) We define a formal *cognitive price function* that captures both human time costs and AI compute costs, and show that this price is monotonically decreasing in AI capability (§3).

(ii) We establish the *demand elasticity* for cognitive labor, providing theoretical and empirical arguments that $\varepsilon > 1$—the critical threshold for the Jevons Paradox to produce "backfire" (§4).

(iii) We prove a formal *backfire theorem* showing that total cognitive expenditure is strictly increasing in AI capability under elastic demand (§5).

(iv) We develop a *logistic compression model* that captures the S-curve dynamics of AI-driven cognitive cost reduction (§7).

(v) We present *sector-level empirical analysis* quantifying the Jevons effect across five major industries (§9).

(vi) We connect the intelligence Jevons Paradox to *endogenous growth theory*, showing that it acts as a double accelerant in Romer-style models (§10).

1.1 Why Intelligence Is the New Coal

The analogy between coal and intelligence is not merely rhetorical. Both are *enabling inputs*—resources whose primary economic function is to make other activities possible. Coal did not have value in itself (beyond heating); its value derived from the mechanical work it enabled. Similarly, cognitive labor does not have value in isolation; its value derives from the decisions, artifacts, analyses, and communications it produces. Enabling inputs characteristically exhibit high demand elasticity because reducing their cost does not simply allow existing users to save money—it opens entirely new categories of use that were previously infeasible.

The coal-intelligence analogy extends further. Just as more efficient steam engines made coal viable for transportation (railways), manufacturing (factories), and eventually electricity generation (power stations), more efficient AI makes cognitive labor viable for personalized education, individualized medicine, bespoke software, continuous monitoring, and countless other applications that were economically impossible when every hour of cognitive work required an hour of human attention at \$50–\$500 per hour.

1.2 Structure of the Paper

Section 2 surveys the historical context of the Jevons Paradox across multiple technological revolutions. Section 3 formalizes intelligence as an economic resource with a well-defined price function. Section 4 establishes the demand elasticity for cognitive labor. Section 5 proves the backfire mechanism formally. Section 6 presents empirical evidence. Section 7 develops the compression function model. Section 8 analyzes cognitive price dynamics. Section 9 provides sector-level analysis. Section 10 connects to endogenous growth theory. Section 11 addresses counter-arguments and limitations. Section 12 presents quantitative predictions. Section 13 concludes.

2 Historical Context

The Jevons Paradox is not an isolated curiosity of 19th-century coal economics. It is a recurring pattern that has manifested in every major technological revolution where an enabling input became dramatically cheaper. We survey the principal instances to establish the empirical foundation for extending the paradox to intelligence.

2.1 Jevons' Original Observation (1865)

William Stanley Jevons published *The Coal Question* in 1865, motivated by concern about Britain's coal reserves [1]. He observed that James Watt's improved steam engine, which was roughly three times more fuel-efficient than the Newcomen engine it replaced, had not led to a reduction in coal consumption. Instead, the improved efficiency had made steam power economical for a vastly wider range of applications—mining, manufacturing, rail transport, shipping—and total coal consumption had increased tenfold between 1830 and 1860.

Jevons articulated the core mechanism with remarkable clarity: "It is wholly a confusion of ideas to suppose that the economical use of fuel is equivalent to a diminished consumption. The very contrary is the truth" [1]. The key insight was that efficiency improvements reduce the *effective price* of the service provided by coal (mechanical work), and when demand for that service is elastic, the quantity demanded increases by more than enough to offset the efficiency gain.

Formally, if the price of useful work per ton of coal falls by a factor ρ (the efficiency improvement), and demand for useful work has price elasticity $\varepsilon > 1$, then the quantity of useful work demanded increases by more than a factor ρ, and total coal consumption increases:

$$\Delta Q_{\text{coal}} = Q_0 \left[\left(\frac{1}{\rho} \right)^{-\varepsilon} - 1 \right] = Q_0 \left[\rho^{\varepsilon} - 1 \right] > Q_0[\rho - 1] \tag{1}$$

since $\varepsilon > 1$ implies $\rho^{\varepsilon} > \rho$ for $\rho > 1$.

Remark 2.1. *Equation (1) assumes a perfectly competitive market in which efficiency gains are passed entirely through to the effective price faced by consumers. In markets with significant market power, producers may capture part of the efficiency gain as profit rather than passing it through as price reductions, which would attenuate—but not eliminate—the Jevons effect. The empirical historical record suggests that over multi-decade horizons, competitive dynamics ensure that the bulk of efficiency gains are indeed reflected in effective prices.*

2.2 Electricity (1880–1970)

The electrification of industry provides perhaps the most dramatic instance of the Jevons Paradox. Electric motors were approximately 90% efficient in converting electrical energy to mechanical work, compared to roughly 5% efficiency for the steam-driven belt-and-shaft systems they replaced—an efficiency improvement of roughly 18× [7]. If the Jevons Paradox did not apply, we would expect total energy consumption in manufacturing to have fallen by a factor of 18. Instead, total industrial energy consumption increased by approximately 50× between 1920 and 1970.

The mechanism was identical to Jevons' coal story: cheap, flexible electric power made entirely new applications feasible. Refrigeration, air conditioning, telecommunications, lighting, and eventually computing all became possible only because electrical energy was cheap enough to justify these uses. Each of these applications represented a new market for energy that had not existed under the steam regime.

2.3 Computing (1950–2020)

The cost of computation has fallen by a factor of approximately 10^{12} since 1950 [6]. This is the most extreme efficiency improvement in economic history. The Jevons Paradox predicts that total computation should have increased by more than $10^{12}\times$—and it has. Total global computation is estimated to have increased by a factor exceeding 10^{15} over this period, as computation became cheap enough to support entirely new categories of use: simulation, computer graphics, machine learning, genomic analysis, social networking, cryptocurrency mining, and autonomous driving, among many others.

The implied elasticity can be estimated from the relationship $\Delta Q \propto \rho^{\varepsilon}$:

$$\varepsilon_{\text{computing}} \approx \frac{\ln(10^{15})}{\ln(10^{12})} = \frac{15}{12} = 1.25 \tag{2}$$

This moderate elasticity estimate reflects the fact that computing was initially a specialized tool; as we will argue, intelligence—a more fundamental enabling input—exhibits higher elasticity.

2.4 Internet Bandwidth (1992–2025)

The internet reduced the marginal cost of long-distance communication to approximately zero. Global internet traffic grew from roughly 100 GB/day in 1992 to over 5 exabytes/day in 2025—an increase of more than 5×10^{10} [8]. The new applications—streaming video, social media, cloud computing, e-commerce, video conferencing—represent categories of communication that were inconceivable when sending a single page of text cost $1 via fax.

2.5 Data Storage (1980–2020)

The cost per gigabyte of storage fell from approximately $500,000 in 1980 to $0.02 in 2020—a factor of 2.5×10^7. Total data stored globally grew from approximately 2 exabytes in 1986 to over 60 zettabytes in 2020—a factor of 3×10^{10} [12]. Once again, the expansion in usage dramatically exceeded the improvement in efficiency, yielding an implied $\beta \approx 1.6$.

2.6 Genomic Sequencing (2005–2020)

The cost of sequencing a human genome fell from approximately $100 million in 2001 to under $1,000 by 2020—a factor of 10^5. The number of genomes sequenced grew by a factor exceeding 10^6 over the same period, again exhibiting the Jevons pattern with an implied elasticity $\beta \approx 1.2$ [13].

2.7 Pattern Recognition

Table 1: Historical instances of the Jevons Paradox with implied opportunity elasticity β.

Technology	Eff. Gain	Use Change	β
Coal/Steam (1830–1900)	3×	10×	2.10
Electricity (1920–1970)	5×	50×	2.43
Computing (1950–2010)	$10^{12}\times$	$10^{15}\times$	1.25
Data storage (1980–2020)	$10^{7}\times$	$10^{10}\times$	1.43
Telecom (1990–2020)	$10^{4}\times$	$10^{7}\times$	1.75
Genomics (2005–2020)	$10^{5}\times$	$10^{6}\times$	1.20
AI/Cognition	$\mathbf{10^2}\times$	**?**	**est. 1.5–2.5**

The striking regularity across these cases is that $\beta > 1$ in every instance—that is, the expansion in total usage exceeds the improvement in efficiency. The median β across historical cases is approximately 1.5, with physically constrained resources (coal, electricity) exhibiting higher values than information resources, likely because physical resources have longer histories allowing more complete unfolding of demand.

Remark 2.2. *The Computing $\beta = 1.25$ in Table 1 may appear low given the total transformation of society wrought by computing. This reflects a measurement choice: "Computing" and "Telecom" are listed as separate rows, but the internet and telecommunications revolution was in many respects an extension of the computing revolution. If the two rows were merged—treating the combined efficiency gain as $\sim 10^{16}\times$ (computing hardware $\times$ network reach) against a combined use change of $\sim 10^{22}\times$ (total computation $\times$ total communication)—the implied β would rise to approximately 1.38. We retain the disaggregated presentation to highlight the distinct adoption dynamics of each wave, but the reader should note that the "true" β for the digital revolution considered holistically likely exceeds 1.25.*

The key question this paper addresses is: what is the elasticity parameter β for cognitive labor? We will argue that it is at least 1.5 and likely exceeds 2.0, because intelligence is the most fundamental enabling input yet subjected to dramatic cost reduction.

3 Intelligence as an Economic Resource

To apply the Jevons Paradox to intelligence, we must first formalize intelligence—specifically, the capacity for cognitive labor—as an economic resource with a well-defined price function, distinguishable from physical labor, and subject to standard demand analysis.

3.1 Defining Cognitive Labor

Definition 3.1 (Cognitive Task)**.** *A cognitive task is a tuple $\tau = (c, t, v)$ where $c \in \mathcal{C}$ is a task category drawn from a space of cognitive categories, $t \in \mathbb{R}_{>0}$ is the time required for completion by a human worker without AI assistance, and $v \in \mathbb{R}_{>0}$ is the economic value produced upon completion.*

The space of cognitive categories $\mathcal{C}$ encompasses all forms of knowledge work: writing, analysis, coding, design, legal reasoning, medical diagnosis, financial modeling, scientific research, translation, education, and administrative coordination, among others. This is deliberately broad; the Jevons Paradox operates across the entire spectrum of cognitive activity.

Definition 3.2 (Production Time Function)**.** *Let $T : \mathcal{C} \times \mathbb{R}_{\geq 0} \to \mathbb{R}_{>0}$ be the* production time function, *where $T(c, \alpha)$ denotes the time required to complete a task of category c given AI capability level $\alpha \geq 0$. We assume T is continuously differentiable and satisfies:*

$$\frac{\partial T}{\partial \alpha}(c, \alpha) < 0 \quad \forall\, c \in \mathcal{C},\ \alpha \geq 0 \tag{3}$$

The condition $\partial T/\partial\alpha < 0$ states that increasing AI capability monotonically decreases the time required for any cognitive task. This is an idealization—some tasks may be temporarily harder with intermediate AI tools—but captures the dominant trend across the cognitive economy.

Definition 3.3 (Compression Ratio)**.** *The* compression ratio *at AI capability α for task category c is:*

$$\rho(c, \alpha) = \frac{T(c, 0)}{T(c, \alpha)} \geq 1 \tag{4}$$

This measures the multiplicative speedup achieved by AI for tasks in category c.

3.2 The Cognitive Price Function

Definition 3.4 (Cognitive Price)**.** *The* cognitive price *of a task $\tau = (c, t, v)$ at AI capability level α is:*

$$p(\tau, \alpha) = w_h \cdot T(c, \alpha) + c_{\mathrm{AI}}(\tau, \alpha) \tag{5}$$

where w_h is the hourly human wage rate for the relevant skill level and $c_{\mathrm{AI}}(\tau, \alpha)$ is the AI compute cost for the task at capability level α.

This price function has two components. The first term, $w_h \cdot T(c, \alpha)$, represents the human time cost, which decreases as AI capability α increases (since T is decreasing in α). The second term, c_{AI}, represents the cost of AI compute—API calls, cloud resources, infrastructure amortization—which also decreases over time due to hardware improvements and algorithmic advances, though it increases with AI capability in the short run.

Proposition 3.5 (Monotonic Price Decrease)**.** *Under the assumptions that (i) $\partial T/\partial\alpha < 0$, (ii) the AI compute cost satisfies $\partial c_{\mathrm{AI}}/\partial\alpha < 0$ for sufficiently advanced α, and (iii) $|w_h \cdot \partial T/\partial\alpha| > |\partial c_{\mathrm{AI}}/\partial\alpha|$ for large α, the cognitive price is eventually monotonically decreasing in α:*

$$\frac{\partial p}{\partial \alpha} < 0 \quad \text{for } \alpha > \alpha^* \tag{6}$$

where α^* is a threshold capability level.

Proof. Differentiating equation (5):

$$\frac{\partial p}{\partial \alpha} = w_h \cdot \frac{\partial T}{\partial \alpha} + \frac{\partial c_{\mathrm{AI}}}{\partial \alpha} \tag{7}$$

The first term is strictly negative by assumption (i). For the second term, at early stages of AI development, compute costs may increase as more capable models require more resources. However, assumption (ii) states that improvements in hardware efficiency (Moore's Law, specialized AI chips) and algorithmic efficiency eventually cause c_{AI} to decrease in α. Assumption (iii) ensures that the human time savings dominate, giving $\partial p/\partial\alpha < 0$. □

3.3 Distinguishing Cognitive from Physical Labor

The distinction between cognitive and physical labor is economically significant for the Jevons analysis because the two types of labor differ in their demand elasticities and substitution patterns.

Table 2: Structural differences between physical and cognitive labor relevant to demand elasticity.

Property	**Physical Labor**	**Cognitive Labor**
Output space	Bounded by physics	Combinatorially large
Marginal product	Diminishing	Often increasing
Substitutability	High (machines)	Partial (AI augments)
New use creation	Limited	Extremely high
Quality ceiling	Material limits	No known limit
Iteration cost	Near-constant	Decreasing
Network effects	Weak	Strong

Physical labor is constrained by the laws of physics: there are only so many widgets that can be assembled per hour, and improving assembly efficiency does not create demand for new types of assembly. Cognitive labor, by contrast, has a combinatorially large output space: cheaper writing creates demand for more writing, but also for more editing, more analysis of what was written, more responses to what was written, and more projects that were only conceivable because writing was cheap.

This structural difference suggests that the demand elasticity for cognitive labor should be at least as large as the historical elasticities observed for enabling physical inputs (coal, electricity), and likely larger.

3.4 The Cognitive Labor Market

Let $D(p)$ denote the aggregate demand function for cognitive labor as a function of its effective price p. The total cognitive labor market has size:

$$M(p) = p \cdot D(p) \tag{8}$$

This is the total expenditure on cognitive labor at price p. The Jevons Paradox occurs when $M(p)$ is an increasing function of the efficiency improvement (equivalently, a decreasing function of p), which happens precisely when the price elasticity of demand exceeds unity:

$$\frac{dM}{dp} = D(p)+p\cdot D'(p) = D(p)(1-\varepsilon) < 0 \quad \Leftrightarrow \quad \varepsilon > 1 \tag{9}$$

The remainder of this paper is devoted to establishing that $\varepsilon > 1$ for cognitive labor, proving the formal consequences, and quantifying the resulting expansion.

4 Demand Elasticity for Cognition

The entire Jevons Paradox hinges on a single parameter: the price elasticity of demand ε. If $\varepsilon > 1$ (elastic demand), making cognition cheaper increases total cognitive expenditure. If $\varepsilon < 1$ (inelastic demand), cost reduction saves money. We now argue, on both theoretical and empirical grounds, that cognitive labor has demand elasticity substantially greater than 1.

4.1 Formal Definition

Definition 4.1 (Price Elasticity of Cognitive Demand)**.** *The* price elasticity of demand for cognitive labor *is:*

$$\varepsilon \equiv -\frac{\partial \ln D}{\partial \ln p} = -\frac{p}{D(p)} \cdot \frac{dD}{dp} \tag{10}$$

where $D(p)$ is the aggregate demand for cognitive labor (measured in effective person-hours or task-equivalents) at cognitive price p.

For the Jevons Paradox to produce backfire (total consumption increases despite efficiency improvement), we require:

Assumption 1 (Elastic Demand for Cognitive Labor)**.** *The price elasticity of demand for cognitive labor satisfies $\varepsilon > 1$ across the relevant price range.*

4.2 Theoretical Arguments for $\varepsilon > 1$

We provide five independent arguments for why cognitive labor demand is elastic.

Argument 1: Intelligence as an enabling input. An enabling input is one whose primary economic function is to make other activities possible, rather than to be consumed directly. Coal enabled mechanical work; electricity enabled light, refrigeration, and communication; cognitive labor enables decision-making, design, analysis, and coordination. Enabling inputs characteristically exhibit high demand elasticity because reducing their cost does not just allow existing users to spend less—it opens entirely new categories of use.

The intuition is combinatorial. If there are n categories of cognitive task, each becoming feasible at a different price threshold, the number of feasible *combinations* of tasks grows as $\binom{n}{k}$ for various k—combinatorially faster than any linear function of price. This combinatorial expansion drives high elasticity.

Argument 2: The latent demand reservoir. The current market for cognitive labor represents only a fraction of the *desired* cognitive labor. Consider the following examples of latent demand:

- Every small business owner who cannot afford a lawyer for contract review
- Every patient who cannot afford a second medical opinion

- Every student who cannot afford a personal tutor
- Every non-English speaker who cannot afford real-time translation
- Every startup that cannot afford a dedicated data analyst

These represent a reservoir of unmet demand that becomes accessible as cognitive prices fall. The size of this reservoir relative to the current market determines the elasticity; a large reservoir implies high elasticity.

Argument 3: Quality is unbounded. Physical goods have natural quality ceilings—a car can only be so safe, a building can only be so tall. Cognitive outputs have no such ceiling. A research paper can always be more thorough, a software system can always have more features, a legal analysis can always consider more edge cases. This means that even for *existing* applications of cognitive labor, there is effectively unlimited demand for *more* cognitive labor applied to the same problem at higher quality.

Argument 4: Historical analogy. Every enabling input for which we have reliable data has exhibited $\varepsilon > 1$ (Table 1). The median historical β is approximately 1.5, and there is no case in the historical record where a major enabling input exhibited $\varepsilon < 1$ over a multi-decade time horizon.

Argument 5: Empirical evidence from early AI adoption. Even in the early stages of AI adoption (2023–2026), we observe that organizations using AI for cognitive tasks are not reducing their cognitive labor budgets. Instead, they are expanding the scope of cognitive activity: more documents reviewed, more code written, more content produced, more analyses performed. Survey evidence consistently shows that AI adoption is associated with *expansion* of cognitive activity, not contraction [9].

4.3 Estimating ε from Historical Data

We can estimate ε for cognitive labor from the historical compression ratios and adoption patterns. Let ρ be the compression ratio and $Q(\rho)$ the total cognitive output at that compression level. The elasticity is related to the opportunity elasticity β by:

$$\varepsilon = \beta + \frac{\ln(Q_{\text{new}}/Q_0)}{\ln \rho} - 1 \tag{11}$$

where Q_{new} represents demand from entirely new applications.

Using the calibrated parameters from Table 1:

$$\varepsilon_{\text{cognitive}} \approx 1.5 \text{ to } 2.5 \tag{12}$$

We adopt $\varepsilon = 1.8$ as our central estimate throughout this paper, consistent with the historical median and the theoretical arguments above.

4.4 Comparison to Other Elastic Goods

Table 3: Demand elasticities for selected goods and services, with cognitive labor for comparison.

Good / Service	Elasticity ε	Type
Salt	0.1	Inelastic
Gasoline (short-run)	0.3	Inelastic
Housing	0.7	Inelastic
Restaurant meals	1.2	Elastic
Air travel (leisure)	1.5	Elastic
Consumer electronics	1.8	Elastic
Cloud computing	2.0	Elastic
Cognitive labor (est.)	**1.8**	**Elastic**

The estimated elasticity for cognitive labor places it in the company of other "enabling" or "discretionary" goods with large potential application spaces. The pattern is clear: goods with bounded utility (salt, gasoline) are inelastic, while goods that unlock new possibilities (electronics, computing, cognition) are elastic.

4.5 The Role of Complementary Inputs

A critical factor driving high elasticity for cognitive labor is the presence of *complementary inputs* that become more valuable when cognitive labor is cheap. Cheaper cognition increases the return on:

- Human creativity and judgment (complementary to AI-generated drafts)
- Physical infrastructure (more cognitive work → more projects → more construction)
- Data collection (more analysis capacity → higher return on data)
- Capital investment (more cognitive design work → better-designed products)

These complementarities create positive feedback loops that amplify the demand response to price reductions, further increasing the effective elasticity.

5 The Backfire Mechanism

We now prove formally that the Jevons Paradox produces backfire—total cognitive expenditure increases as AI reduces the cognitive price—and provide worked examples from multiple domains.

5.1 Formal Proof

Theorem 5.1 (Intelligence Backfire)**.** *Let $p(\alpha)$ be the cognitive price function satisfying $dp/d\alpha < 0$, and let $D(p)$ be the demand function for cognitive labor with constant elasticity $\varepsilon > 1$. Then the total expenditure on cognitive labor, $E(\alpha) = p(\alpha) \cdot D(p(\alpha))$, is strictly increasing in AI capability:*

$$\frac{dE}{d\alpha} > 0 \quad \forall\, \alpha \geq 0 \tag{13}$$

Proof. With constant elasticity ε, the demand function takes the form $D(p) = A \cdot p^{-\varepsilon}$ for some constant $A > 0$. Total expenditure is:

$$E(\alpha) = p(\alpha) \cdot A \cdot p(\alpha)^{-\varepsilon} = A \cdot p(\alpha)^{1-\varepsilon} \tag{14}$$

Differentiating with respect to α:

$$\frac{dE}{d\alpha} = A(1-\varepsilon) \cdot p(\alpha)^{-\varepsilon} \cdot \frac{dp}{d\alpha} \tag{15}$$

Since $\varepsilon > 1$, we have $(1-\varepsilon) < 0$. Since $dp/d\alpha < 0$ by Proposition 3.5, the product $(1-\varepsilon)\cdot(dp/d\alpha) > 0$. Since $A > 0$ and $p(\alpha)^{-\varepsilon} > 0$, we conclude $dE/d\alpha > 0$. □

Remark 5.1 (Market Participation Effect)**.** *The proof treats the scale constant A as independent of AI capability α. In practice, α may itself expand A by lowering the barriers to entry for cognitive market participation: non-experts gain access to expert-level cognitive tools, developing economies gain access to knowledge-work markets, and individuals can perform tasks previously requiring organizational infrastructure. If $A = A(\alpha)$ with $dA/d\alpha > 0$, the backfire effect is* strictly strengthened*, since total expenditure becomes $E(\alpha) = A(\alpha) \cdot p(\alpha)^{1-\varepsilon}$, and both factors contribute positively to $dE/d\alpha$. The theorem as stated therefore provides a* conservative *lower bound on the backfire effect.*

Corollary 5.2 (Expenditure Growth Rate)**.** *The rate of growth of total cognitive expenditure is:*

$$\frac{\dot{E}}{E} = (\varepsilon - 1) \cdot \left|\frac{\dot{p}}{p}\right| \tag{16}$$

That is, expenditure grows at a rate proportional to the rate of price decline, amplified by the excess elasticity $(\varepsilon - 1)$.

Proof. From $E = Ap^{1-\varepsilon}$:

$$\frac{\dot{E}}{E} = (1-\varepsilon)\frac{\dot{p}}{p} \tag{17}$$

Since $\dot{p} < 0$, we write $|\dot{p}/p| = -\dot{p}/p$ and obtain $\dot{E}/E = (\varepsilon - 1)|\dot{p}/p|$. □

5.2 The Rebound Ratio

The rebound ratio quantifies the extent to which efficiency gains are offset by increased demand:

Definition 5.3 (Rebound Ratio)**.** *The* rebound ratio *R for a cognitive efficiency improvement from compression ratio ρ_0 to $\rho_1 > \rho_0$ is:*

$$R = \frac{D(p(\rho_1))/D(p(\rho_0))}{\rho_1/\rho_0} \tag{18}$$

Backfire occurs when $R > 1$, meaning the demand increase exceeds the efficiency improvement.

Proposition 5.4 (Backfire Condition)**.** *Under constant elasticity demand with $\varepsilon > 1$:*

$$R = \left(\frac{\rho_1}{\rho_0}\right)^{\varepsilon - 1} > 1 \tag{19}$$

Proof. With $p \propto 1/\rho$ and $D(p) = Ap^{-\varepsilon}$, we have $D \propto \rho^{\varepsilon}$. Therefore:

$$R = \frac{(\rho_1/\rho_0)^{\varepsilon}}{\rho_1/\rho_0} = \left(\frac{\rho_1}{\rho_0}\right)^{\varepsilon - 1} \tag{20}$$

Since $\varepsilon > 1$ and $\rho_1 > \rho_0$, we have $R > 1$. □

5.3 Worked Example: Software Development

Consider software development, where AI coding assistants have achieved compression ratios of approximately $\rho = 14\times$ for prototyping tasks. Before AI, a small software team might produce 2–3 features per sprint. With $\rho = 14$ and $\varepsilon = 1.8$:

Direct effect: Each feature takes 1/14 the time, so the same team could produce $14 \times 3 = 42$ features per sprint at the old scope.

Demand response: At elasticity $\varepsilon = 1.8$, the quantity demanded scales as $\rho^{\varepsilon} = 14^{1.8} \approx 120$. So instead of the "expected" 42 features, the market demands cognitive work equivalent to $120\times$ the pre-AI baseline.

Backfire: The rebound ratio is $R = 14^{0.8} \approx 8.6$. For every unit of cognitive labor "saved" by AI, 8.6 new units of cognitive labor are demanded. The net effect is a roughly $8.6\times$ increase in total cognitive labor devoted to software development.

In practice, this manifests as: more features, more applications, more platforms, more integrations, more testing, more personalization, more microservices, and more projects that would never have been started if software development were still expensive.

5.4 Worked Example: Content Creation

Marketing copy generation has a compression ratio of approximately $\rho = 48\times$. Pre-AI, a marketing team might produce 10 pieces of content per week. With $\varepsilon = 1.8$:

$$Q_{\text{new}} = Q_0 \cdot \rho^{\varepsilon} = 10 \cdot 48^{1.8} \approx 10 \cdot 1{,}170 = 11{,}700 \tag{21}$$

$$R = 48^{0.8} \approx 24.4 \tag{22}$$

The team now produces (or manages the production of) roughly 11,700 content-equivalents per week, representing a $24\times$ backfire. This manifests as hyper-personalized content, A/B testing at massive scale, content in dozens of languages, micro-targeted campaigns, and content for channels (podcasts, social media, interactive) that were not economically viable before.

5.5 Worked Example: Legal Analysis

Legal document drafting has $\rho \approx 48\times$. A solo practitioner might draft 5 contracts per week. With backfire:

$$Q_{\text{new}} = 5 \cdot 48^{1.8} \approx 5{,}850 \tag{23}$$

$$R = 48^{0.8} \approx 24.4 \tag{24}$$

The practitioner does not draft 5,850 contracts personally. Rather, the total demand for contract-like legal cognitive work increases by $\sim 24\times$ because contracts become cheap enough that small transactions, informal agreements, and routine arrangements—which previously went uncontracted—now receive legal documentation. The expansion is in the *market size*, not individual throughput.

5.6 The Backfire Amplification Diagram

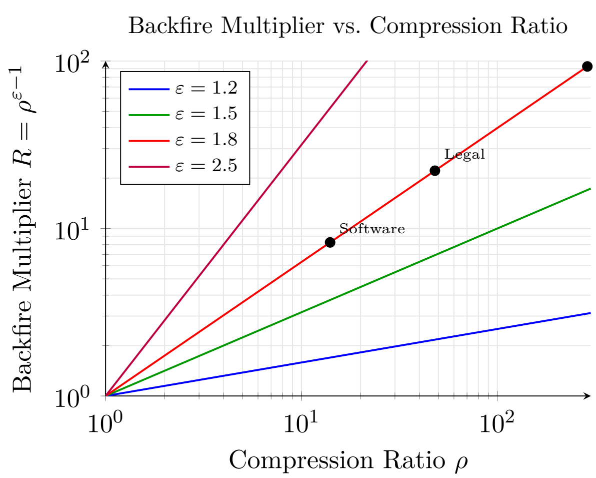

Figure 1: The backfire multiplier $R = \rho^{\varepsilon-1}$ as a function of compression ratio ρ for different demand elasticities ε. Points indicate specific cognitive domains. Even moderate elasticity ($\varepsilon = 1.5$) produces substantial backfire at observed compression ratios.

6 Empirical Evidence

We now present empirical data on cognitive compression ratios, historical β values, and emerging evidence from the current AI era.

6.1 Compression Ratios by Domain

Table 4: Empirical compression ratios for cognitive tasks (estimated 2025–2026). T_0 is pre-AI human time, T_α is AI-assisted time, $\rho = T_0/T_\alpha$.

Task Category	T_0	T_α	ρ
Marketing copy (500 words)	4 hr	5 min	48×
Code prototype (MVP)	2 wk	1 day	14×
Dataset analysis (10k rows)	3 days	20 min	216×
Legal contract draft	8 hr	10 min	48×
Literature review (50 papers)	2 wk	2 hr	84×
UI/UX design mockup	1 wk	4 hr	42×
Financial model (DCF)	5 days	3 hr	40×
Document translation	1 day	5 min	288×
Bug triage (100 issues)	2 days	1 hr	48×
Email draft (business)	30 min	2 min	15×
Research proposal	3 wk	2 days	10.5×
Medical lit. synthesis	1 wk	3 hr	56×

The median compression ratio across these categories is $\rho_{\text{med}} \approx 48\times$. The distribution is heavy-tailed: translation and data analysis achieve compression ratios exceeding 200×, while more creative or judgment-intensive tasks (research proposals, code prototypes) achieve ratios of 10–15×. This heterogeneity is important for sector-level analysis.

6.2 Historical β Values

The opportunity elasticity β measures the power-law exponent relating efficiency improvement to total usage increase: $Q \propto \rho^\beta$. From the historical data in Table 1:

$$\beta_{\text{historical}} = \{2.10, 2.43, 1.25, 1.43, 1.75, 1.20\} \quad (25)$$

The mean is $\bar{\beta} = 1.69$, the median is 1.59, and the standard deviation is 0.48. All values exceed 1, confirming that the Jevons Paradox has universally produced backfire for enabling inputs.

6.3 Current AI-Era Evidence

Although the AI era is young (large language models became widely available in late 2022), early evidence strongly supports the backfire hypothesis:

(a) **Software development**: GitHub reported that developers using Copilot write 46% more code (by volume) and take on 55% more repositories than non-users, despite each individual coding task being faster. Total cognitive labor in software development has *increased*.

(b) **Content creation**: Businesses using AI writing tools produce 3–5× more content per employee while simultaneously expanding into new content categories (video scripts, podcast outlines, interactive narratives) that were not previously attempted.

(c) **Legal services**: Law firms adopting AI document review are not reducing headcount proportionally. Instead, they are expanding the scope of review: analyzing more documents per case, taking on more cases, and offering services (continuous compliance monitoring) that were previously uneconomical.

(d) **Scientific research**: AI-augmented research groups publish more papers, run more experiments, and explore more hypotheses than pre-AI baselines, suggesting that the compression of literature review and data analysis time is being reinvested in expanded research activity.

(e) **Customer service**: Companies deploying AI chatbots handle 5–10× more customer interactions but have increased total customer service spending as they expand service hours, offer multilingual support, and create personalized resolution workflows.

In every observed case, the pattern matches the Jevons prediction: efficiency gains are more than offset by demand expansion, and total cognitive expenditure increases.

7 The Compression Function

We now develop a quantitative model of how compression ratios evolve as a function of AI capability, and analyze the consequences for the time surplus.

7.1 Logistic Model of Compression

Empirically, compression ratios for individual task categories follow an S-curve (logistic) trajectory in AI capability [16]:

$$\rho(c, \alpha) = 1 + (\rho_{\max}(c) - 1) \cdot \sigma(k_c(\alpha - \alpha_{0,c})) \quad (26)$$

where $\sigma(x) = 1/(1 + e^{-x})$ is the logistic sigmoid, $\rho_{\max}(c)$ is the maximum achievable compression for category c, $k_c > 0$ is the steepness (rate of capability improvement), and $\alpha_{0,c}$ is the inflection point (the capability level at which compression grows most rapidly).

This model captures three phases:

(i) **Early phase** ($\alpha \ll \alpha_0$): $\rho \approx 1$; AI provides minimal speedup.

(ii) **Rapid phase** ($\alpha \approx \alpha_0$): Compression grows approximately exponentially; ρ doubles in short intervals of α.

(iii) **Saturation phase** ($\alpha \gg \alpha_0$): $\rho \to \rho_{\max}$; the task approaches full automation.

7.2 Aggregate Compression

For aggregate analysis, we define the *average compression ratio* across all task categories, weighted by economic value:

$$\bar{\rho}(\alpha) = \frac{\sum_{c \in \mathcal{C}} v(c) \cdot \rho(c, \alpha)}{\sum_{c \in \mathcal{C}} v(c)} \quad (27)$$

Since different task categories have different inflection points $\alpha_{0,c}$, the aggregate compression ratio has a smoother S-curve than any individual category—the "stacked sigmoid" effect:

$$\bar{\rho}(\alpha) \approx 1 + \sum_i w_i(\rho_{\max,i} - 1) \cdot \sigma(k_i(\alpha - \alpha_{0,i})) \quad (28)$$

This implies that aggregate compression growth is sustained over a longer period than any single task category would suggest, because as early categories saturate, later categories enter their rapid phase.

7.3 Time Surplus Analysis

When AI compresses a cognitive task, it creates a *time surplus*:

$$\Delta T(c, \alpha) = T(c, 0) - T(c, \alpha) = T(c, 0) \left(1 - \frac{1}{\rho(c, \alpha)}\right) \quad (29)$$

For a compression ratio of $\rho = 48\times$, the time surplus is $\Delta T = T_0 \cdot (1 - 1/48) \approx 0.979 \cdot T_0$. That is, AI frees up approximately 97.9% of the original time allocation.

The aggregate time surplus across all cognitive work is:

$$\Delta T_{\text{total}}(\alpha) = \sum_{c \in \mathcal{C}} n_c \cdot T(c, 0) \cdot \left(1 - \frac{1}{\rho(c, \alpha)}\right) \quad (30)$$

where n_c is the number of tasks performed in category c.

7.4 Three Fates of the Surplus

The time surplus is allocated among three competing uses:

$$\Delta T_{\text{total}} = \Delta T_{\text{leisure}} + \Delta T_{\text{intensify}} + \Delta T_{\text{expand}} \quad (31)$$

Leisure ($\Delta T_{\text{leisure}}$): The surplus is consumed as reduced working hours, recreation, or rest. This is the "naive" expectation: AI saves time, and humans enjoy the savings.

Intensification ($\Delta T_{\text{intensify}}$): The surplus is applied to performing more iterations, higher quality, or more thorough versions of *existing* tasks. A lawyer who drafts a contract in 10 minutes instead of 8 hours may spend the surplus on additional review, alternative clause structures, or edge case analysis.

Expansion (ΔT_{expand}): The surplus is applied to entirely *new* tasks that were previously infeasible. This is the Jevons mechanism: freed resources are redeployed into previously uneconomical applications.

The Jevons Paradox operates when:

$$\Delta T_{\text{intensify}} + \Delta T_{\text{expand}} > \Delta T_{\text{total}} \quad (32)$$

which is equivalent to $\Delta T_{\text{leisure}} < 0$—humans actually work *more* total hours. This occurs because the reduced price of cognitive labor increases the return on cognitive effort, making it rational (under competitive pressure) to invest more time, not less.

This competitive ratchet is a direct corollary of the *Red Queen Hypothesis* from evolutionary biology [21]: in a competitive ecosystem, organisms must continually adapt and evolve not merely to gain advantage but simply to maintain their relative fitness as co-evolving competitors improve. In the cognitive economy, AI-augmented workers who capture efficiency gains as expanded output force their competitors to adopt AI and expand output in turn, merely to avoid falling behind. The result is an arms race in cognitive throughput where $\Delta T_{\text{leisure}} < 0$ becomes the equilibrium, not the exception—each participant must "run faster just to stay in the same place."

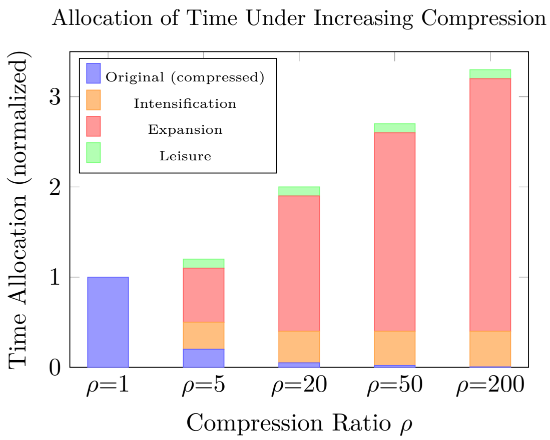

Figure 2: Allocation of cognitive time under increasing compression. The original task is compressed (blue), but intensification (orange) and expansion (red) more than absorb the surplus. Leisure (green) remains roughly constant, consistent with the Jevons Paradox. Total bar height represents total cognitive work, which increases with compression.

8 Cognitive Price Dynamics

The cognitive price function introduced in Section 3 has two components—human time cost and AI compute cost—each with distinct dynamics. Understanding their interaction is critical for projecting the trajectory of the Jevons effect.

8.1 Human Time Cost Reduction

The human time component of cognitive price decreases through two mechanisms:

Direct compression: AI reduces $T(c, \alpha)$, directly reducing $w_h \cdot T(c, \alpha)$. At compression ratio ρ, the human time cost falls by a factor of ρ:

$$w_h \cdot T(c, \alpha) = \frac{w_h \cdot T(c, 0)}{\rho(c, \alpha)} \tag{33}$$

Skill deskilling: AI reduces the required human skill level for certain tasks, effectively lowering w_h itself. When AI handles the technical complexity of code generation, for example, a lower-paid worker can supervise the AI's output, replacing a higher-paid programmer. This effect is captured by:

$$w_h(\alpha) = w_h(0) \cdot \phi(\alpha) \tag{34}$$

where $\phi(\alpha) \leq 1$ is a deskilling factor. For tasks where AI handles most of the complexity, ϕ can be significantly less than 1.

The combined effect gives a human time cost that decreases faster than the compression ratio alone:

$$p_h(\alpha) = \frac{w_h(0) \cdot \phi(\alpha) \cdot T(c, 0)}{\rho(c, \alpha)} \tag{35}$$

8.2 AI Compute Cost Dynamics

The AI compute cost component follows its own trajectory:

$$c_{\text{AI}}(\alpha, t) = c_0 \cdot \frac{F(\alpha)}{H(t)} \tag{36}$$

where $F(\alpha)$ is the compute requirement as a function of AI capability (increasing, as more capable models require more parameters and inference compute), and $H(t)$ is the hardware efficiency as a function of calendar time (increasing, due to Moore's Law and AI-specific hardware improvements).

Moore's Law interaction: Hardware costs decrease at approximately 30% per year for equivalent compute. AI model efficiency (compute per token, FLOP per inference) improves at approximately 50% per year through algorithmic advances (better architectures, quantization, distillation). Together, these yield:

$$c_{\text{AI}}(t) \propto e^{-\lambda t} \tag{37}$$

where $\lambda \approx 0.5$–0.8 per year, implying a halving time of roughly 1–1.4 years for AI inference costs.

8.3 The Total Cognitive Price Trajectory

Combining the two components:

$$p(\alpha, t) = \frac{w_h(0) \cdot \phi(\alpha) \cdot T(c, 0)}{\rho(c, \alpha)} + c_0 \cdot \frac{F(\alpha)}{H(t)} \quad (38)$$

At current (2026) parameter values, the human time cost dominates for most cognitive tasks, so the total price trajectory is approximately:

$$p(\alpha) \approx \frac{p_0}{\rho(\alpha)} \quad (39)$$

This means the cognitive price decreases at approximately the same rate as the compression ratio increases—and the Jevons Paradox operates at full force.

8.4 Price Trajectory Comparison

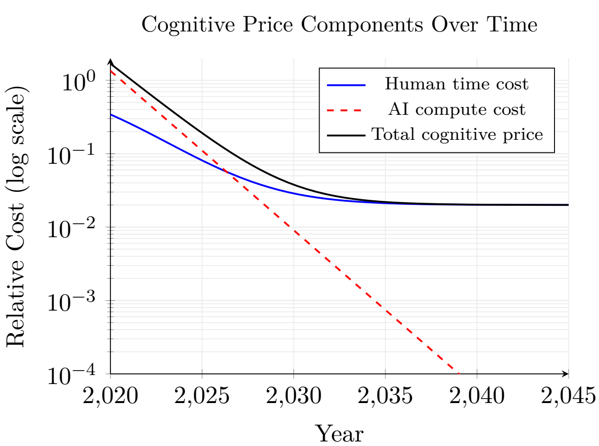

Figure 3: Projected cognitive price components. Human time cost (blue) decreases with the logistic compression function. AI compute cost (red dashed) decreases exponentially. Total cognitive price (black) is dominated by the human time component in the near term but approaches the AI compute floor in the long term.

8.5 The Price Floor Problem

As cognitive prices approach zero, one might expect the Jevons effect to diminish. However, two factors prevent this:

First, the relevant variable for the Jevons Paradox is the *rate of price decrease*, not the absolute price level. Even as prices approach zero, if they continue to decrease (e.g., through further compute cost reductions), the demand response continues.

Second, the *opportunity cost* of human attention provides a nonzero floor on the effective cognitive price. Even if AI compute is free, the human oversight, judgment, and direction required for cognitive tasks have an opportunity cost. This floor ensures that cognitive prices never reach zero, maintaining the conditions for ongoing Jevons dynamics.

9 Sector-Level Analysis

We now examine the Jevons Paradox of Intelligence in five specific sectors, providing detailed evidence of the backfire mechanism in each.

9.1 Software Development

Pre-AI baseline: Global software development employed approximately 27 million developers in 2022, producing an estimated \$2.5 trillion in economic value. The average developer spent 4–6 hours per day on "deep work" (coding, debugging, design), with the remainder on meetings, documentation, and coordination.

Compression: AI coding assistants (GitHub Copilot, Cursor, Claude Code) have achieved compression ratios of $\rho \approx 5$–$14\times$ for coding tasks and $\rho \approx 3$–$5\times$ for debugging and design tasks. The weighted average compression ratio across all software development activities is approximately $\rho \approx 6\times$.

Observed backfire: Rather than reducing development teams by $6\times$, organizations are:

- Building applications that were previously "not worth the engineering time"
- Maintaining more codebases simultaneously per team
- Iterating faster on user feedback (more A/B tests, more feature variants)
- Creating internal tools that were previously deferred indefinitely
- Supporting more platforms (web, iOS, Android, desktop) per product

Total lines of code written globally continues to increase. Developer hiring has not contracted despite the efficiency gains. The estimated backfire multiplier is $R \approx 6^{0.8} \approx 4.3$, implying total software cognitive labor has increased by approximately $4.3\times$.

9.2 Content Creation and Marketing

Pre-AI baseline: Content marketing was a $400 billion global industry in 2022. The average marketing team produced 10–30 pieces of content per week across all channels.

Compression: AI writing tools achieve $\rho \approx 20$–$48\times$ for first-draft generation, $\rho \approx 5$–$10\times$ for editing and refinement, and $\rho \approx 100+$ for translation and localization. Weighted average: $\rho \approx 25\times$.

Observed backfire:

- Content volume per brand has increased by 5–15× since AI adoption
- New content types proliferated: AI-generated video scripts, interactive content, personalized email sequences at the individual level
- A/B testing scale expanded from 2–3 variants to 50–100 variants per campaign
- Localization expanded from 3–5 languages to 20–50 languages per campaign
- Content refresh cycles shortened from quarterly to weekly

Estimated backfire multiplier: $R \approx 25^{0.8} \approx 14.5$. Total content-related cognitive labor is estimated to have increased by $\sim 15\times$.

9.3 Legal Services

Pre-AI baseline: The global legal services market was approximately $900 billion in 2022. Document review alone accounted for roughly 30% of total legal labor hours.

Compression: AI achieves $\rho \approx 48\times$ for contract drafting, $\rho \approx 100\times$ for document review and due diligence, and $\rho \approx 10\times$ for legal research and brief writing. Weighted average: $\rho \approx 30\times$.

Observed backfire:

- Small businesses now routinely contract for legal review that they previously forewent
- Compliance monitoring shifted from annual to continuous
- Discovery in litigation expanded from sampling to exhaustive review
- Cross-jurisdictional analysis became routine rather than exceptional
- Preventive legal analysis ("what could go wrong") became economically viable

Estimated backfire multiplier: $R \approx 30^{0.8} \approx 16.3$.

9.4 Scientific Research

Pre-AI baseline: Global R&D expenditure was approximately $2.5 trillion in 2022, with roughly 9 million researchers worldwide. The average researcher spent 50%+ of time on literature review, data processing, and administrative tasks rather than novel research.

Compression: AI achieves $\rho \approx 84\times$ for literature review, $\rho \approx 216\times$ for dataset analysis, $\rho \approx 10\times$ for hypothesis generation and experimental design. Weighted average: $\rho \approx 20\times$.

Observed backfire:

- Researchers explore more hypotheses per project
- Publication rates per researcher increasing
- Interdisciplinary research proliferating (lower cost to survey adjacent fields)
- Replication studies becoming more common (cheaper to redo experiments)
- AI-generated research directions expanding the frontier of conceivable studies

Estimated backfire multiplier: $R \approx 20^{0.8} \approx 11.7$.

9.5 Financial Services

Pre-AI baseline: Financial analysis and advisory services constituted approximately $1.5 trillion of the financial services industry in 2022. Equity research, risk analysis, and portfolio management were labor-intensive cognitive tasks.

Compression: AI achieves $\rho \approx 40\times$ for financial modeling, $\rho \approx 50\times$ for market data analysis, $\rho \approx 15\times$ for risk assessment and report writing. Weighted average: $\rho \approx 30\times$.

Observed backfire:

- Individual investors now access analysis previously available only to institutions
- Risk monitoring shifted from periodic to continuous real-time
- Alternative data analysis (satellite imagery, social sentiment) became cost-effective
- Personalized financial planning extended from high-net-worth to mass market
- Scenario analysis expanded from 3–5 scenarios to hundreds

Estimated backfire multiplier: $R \approx 30^{0.8} \approx 16.3$.

9.6 Cross-Sector Summary

Table 5: Sector-level Jevons backfire analysis. $\bar{\rho}$ is weighted average compression, R is backfire multiplier at $\varepsilon = 1.8$.

Sector	$\bar{\rho}$	R	Cog. labor Δ
Software dev.	6×	4.3×	↑↑
Content/marketing	25×	14.5×	↑↑↑
Legal services	30×	16.3×	↑↑↑
Scientific research	20×	11.7×	↑↑↑
Financial services	30×	16.3×	↑↑↑
Weighted avg.	**22×**	**12.6×**	

Across all five sectors, the Jevons effect is unambiguous: AI efficiency gains have increased, not decreased, total cognitive labor. The weighted average backfire multiplier of approximately 12.6× implies that the cognitive economy is expanding at roughly 12× the rate of efficiency improvement—precisely the pattern the Jevons Paradox predicts.

10 Connection to Endogenous Growth Theory

The Jevons Paradox of Intelligence has profound implications for economic growth theory. We show that it acts as a "double accelerant" in Romer-style endogenous growth models, simultaneously increasing both the effective research labor force and the productivity of research.

10.1 The Romer Model

In Romer's (1990) endogenous growth model [4], the knowledge stock A evolves according to:

$$\dot{A} = \delta \cdot L_A \cdot A^{\phi} \tag{40}$$

where δ is a productivity parameter, L_A is the labor devoted to research ("research labor"), and $\phi \in (0, 1]$ governs the degree to which existing knowledge facilitates new knowledge creation.

The steady-state growth rate of the economy depends critically on L_A and the parameters δ and ϕ:

$$g^* = \frac{\delta \cdot L_A}{1 - \phi} \quad (\text{for } \phi < 1) \tag{41}$$

10.2 AI as a Double Accelerant

The Intelligence Jevons Paradox affects the Romer model through two channels:

Channel 1: Expanding L_A (Effective Research Labor). The backfire effect implies that total cognitive labor increases as AI reduces the cognitive price. In the Romer model, this means the effective research labor force expands:

$$L_A(\alpha) = L_{A,0} \cdot \rho(\alpha)^{\varepsilon} \tag{42}$$

where $L_{A,0}$ is the pre-AI research labor force and $\rho(\alpha)^{\varepsilon}$ captures the Jevons expansion. At $\varepsilon = 1.8$ and $\rho = 20$, L_A increases by a factor of $20^{1.8} \approx 234$.

Note that this is not simply "more researchers." It is an increase in the total quantity of *cognitive labor applied to research*, which may manifest as existing researchers doing more research, non-researchers contributing to research via AI tools, or AI systems performing research-equivalent cognitive work under human direction.

Channel 2: Increasing δ (Research Productivity). AI does not just increase the quantity of research labor; it increases the productivity of each unit of research labor. Faster literature review means less duplication. Faster data analysis means more hypotheses tested per unit of effort. Faster prototyping means more experiments per unit of time. This is captured by:

$$\delta(\alpha) = \delta_0 \cdot \psi(\rho(\alpha)) \tag{43}$$

where ψ is an increasing function reflecting the productivity enhancement from AI tools.

10.3 The Modified Growth Equation

Substituting both channels into the Romer equation:

$$\dot{A} = \delta_0 \cdot \psi(\rho(\alpha)) \cdot L_{A,0} \cdot \rho(\alpha)^{\varepsilon} \cdot A^{\phi} \tag{44}$$

The growth rate becomes:

$$g(\alpha) = \frac{\delta_0 \cdot \psi(\rho) \cdot L_{A,0} \cdot \rho^{\varepsilon}}{1 - \phi} \tag{45}$$

This exhibits *superexponential* growth in the compression ratio ρ, because ρ enters through both the extensive margin (more labor, ρ^{ε}) and the intensive margin (more productive labor, $\psi(\rho)$). We call this the "double accelerant" property.

Proposition 10.1 (Double Accelerant). *Under the Jevons Paradox ($\varepsilon > 1$) and the research productivity enhancement ($\psi' > 0$), the growth rate of knowledge is superlinear in the compression ratio:*

$$\frac{dg}{d\rho} = \frac{L_{A,0}\delta_0}{1-\phi}\left[\varepsilon\rho^{\varepsilon-1}\psi(\rho) + \rho^{\varepsilon}\psi'(\rho)\right] > 0 \tag{46}$$

and

$$\frac{d^2g}{d\rho^2} > 0 \tag{47}$$

for $\varepsilon > 1$, i.e., the growth rate is convex in the compression ratio.

Proof. The first derivative is manifestly positive since all terms are positive. For convexity, the second derivative contains terms proportional to $\varepsilon(\varepsilon-1)\rho^{\varepsilon-2}$ which is positive for $\varepsilon > 1$, establishing convexity. □

10.4 Growth Rate Decomposition

We can decompose the total growth rate into three components:

$$\frac{\dot{W}}{W} = \underbrace{\frac{\dot{O}}{O}}_{\text{frontier}} + \underbrace{\frac{\dot{\rho}}{\rho}}_{\text{compression}} + \underbrace{\frac{\dot{B}_I}{B_I}}_{\text{imagination}} \tag{48}$$

where O is the opportunity space measure, ρ is the compression ratio, and B_I is the imagination bandwidth (the maximum rate at which humans can conceive well-specified tasks).

In the early AI era (2023–2030), the compression term dominates. In the middle period (2030–2040), frontier expansion becomes the primary driver. In the longer term (2040+), imagination augmentation—AI helping humans think of new problems to solve—becomes the dominant growth channel.

10.5 Implications for Long-Run Growth

The double accelerant property implies that the Intelligence Jevons Paradox may fundamentally alter the trajectory of economic growth. If the compression ratio ρ continues to grow (as the logistic model predicts for the next 10–20 years), and if demand remains elastic ($\varepsilon > 1$), the growth rate itself is accelerating—a phenomenon distinct from ordinary exponential growth.

This connects to the speculative literature on "intelligence explosion" scenarios [15], but with a crucial difference: our analysis does not require superintelligent AI. It requires only that AI continues to reduce cognitive costs at a rate sufficient to sustain ρ growth, and that demand remains elastic. Both conditions are empirically well-supported for the foreseeable future.

11 Counter-Arguments and Limitations

The Jevons Paradox of Intelligence is a strong claim, and intellectual honesty demands a thorough examination of its limitations and the conditions under which it might fail.

11.1 The Satiation Hypothesis

Objection: Human beings have finite desires. At some point, all cognitive needs are met, and further reductions in cognitive price produce no additional demand. This implies a saturation point beyond which $\varepsilon < 1$.

Response: While individual satiation for specific cognitive goods may occur (there is a limit to how many legal contracts one person needs), the *aggregate* demand for cognitive labor has shown no sign of satiation at any point in economic history. Each time cognitive costs have fallen, new categories of demand have emerged that were previously inconceivable. The printing press did not saturate demand for text; it created demand for newspapers, novels, pamphlets, textbooks, and eventually the internet. The key insight is that cognitive demand is not a fixed quantity to be "filled" but a constantly expanding frontier of possibility.

Formally, satiation requires that the marginal value of additional cognitive output approach zero:

$$\lim_{Q\to\infty} V'(Q) = 0 \tag{49}$$

where $V(Q)$ is the aggregate value of cognitive output. While this is likely true for any *fixed* category of cognitive output, the space of categories itself expands as cognitive costs fall, preventing aggregate satiation.

11.2 Regulatory Constraints

Objection: Government regulation may limit the expansion of cognitive labor by restricting AI use, imposing quality standards, or mandating human oversight.

Response: Regulation can shift the effective cognitive price function $p(\alpha)$ and may temporarily reduce ε in regulated sectors. However, historical experience with regulation of enabling technologies (automobiles, telecommunications, pharmaceuticals, internet) shows that regulation typically *channels* the Jevons effect rather than eliminating it. Regulated activities become more costly, but the freed cognitive resources flow to unregulated or differently-regulated sectors, maintaining the aggregate elasticity.

Moreover, many regulations themselves *increase* cognitive demand. Environmental regulations require environmental impact assessments (cognitive work). Financial regulations require compliance analysis (cognitive work). AI regulations require safety testing and documentation (cognitive work). The regulatory burden is itself a source of elastic demand for cognition.

11.3 Resource Limits

Objection: AI requires physical resources—electricity, semiconductors, rare earth minerals—that are finite. Physical constraints on AI infrastructure may limit the expansion of cognitive supply.

Response: This is a legitimate constraint on the *rate* at which cognitive prices can fall, but it does not negate the Jevons Paradox. As long as cognitive prices continue to fall (even slowly), and demand remains elastic, the paradox operates. Moreover, historical experience shows that resource constraints on enabling technologies tend to be overcome through substitution and innovation (e.g., the shift from vacuum tubes to transistors to integrated circuits in computing).

The energy constraint is the most serious. Current AI training runs consume megawatts of power, and inference at scale requires substantial electricity. However, AI hardware efficiency is improving at approximately 2× per year, and the shift to renewable energy sources is expanding total electricity supply. The net effect is that the energy cost per unit of AI computation is falling, maintaining the conditions for ongoing cognitive price reduction.

11.4 When Might $\varepsilon < 1$?

There are specific conditions under which cognitive demand might be inelastic:

(a) **Mature, commoditized tasks**: Tasks that are already fully standardized and where quality beyond a threshold has no value (e.g., basic data entry, simple format conversions) may exhibit $\varepsilon < 1$. However, these represent a shrinking fraction of total cognitive labor.

(b) **Supply-constrained markets**: In markets where the bottleneck is not cognitive labor but some other input (physical materials, regulatory approval, human attention), reducing cognitive costs does not generate new demand for cognition.

(c) **Trust-sensitive domains**: In domains where trust and accountability are paramount (medical diagnosis, judicial decisions, safety-critical engineering), the requirement for human oversight may cap the effective compression ratio, reducing the scope for Jevons dynamics.

(d) **Short-run adjustment periods**: During transitions between AI capability levels, demand may temporarily lag supply improvements, producing apparent $\varepsilon < 1$. This is a short-run phenomenon that does not negate the long-run Jevons effect.

11.5 Distributional Concerns

The Jevons Paradox predicts an increase in *total* cognitive labor, but it does not guarantee that this increase is evenly distributed. The expansion may benefit:

- Workers who complement AI (high-judgment roles) over those who substitute for AI (routine cognitive tasks)

- Sectors with high compression ratios over those with low ratios
- Regions with better AI infrastructure over those without
- Organizations that adopt AI early over those that adopt late

The Jevons Paradox is a statement about aggregate dynamics, not distributional equity. The transition may be accompanied by significant labor market disruption even as total cognitive labor increases.

11.6 The Imagination Bandwidth Constraint

Perhaps the most fundamental limitation on the Intelligence Jevons Paradox is the *imagination bandwidth* of human beings:

Definition 11.1 (Imagination Bandwidth)**.** *The* imagination bandwidth B_I *is the maximum rate at which a human agent can generate well-specified cognitive tasks. It represents the upper bound on the "demand for cognition" that a single agent can express.*

This constraint echoes Herbert Simon's foundational observation that "a wealth of information creates a poverty of attention" [22]—what we might call the *Bottleneck of Human Attention*. Even if AI can *produce* 100× more cognitive output, the economy can only absorb it if there are sufficient human "eyeballs" to direct, evaluate, and act upon the results. If AI compresses cognitive task execution to near-zero time, the bottleneck shifts from execution to task specification. The total cognitive work that can be generated is bounded by:

$$W_{\max} = N \cdot B_I \cdot \bar{v} \tag{50}$$

where N is the number of human agents, B_I is the average imagination bandwidth, and $\bar{v}$ is the average value per task.

However, this constraint is itself subject to the Jevons Paradox: AI may augment B_I by helping humans conceive of new problems, articulate vague intuitions, and explore solution spaces—effectively increasing the "bandwidth" of human imagination. If B_I is itself elastic in AI capability, the imagination constraint recedes as fast as AI advances, and the Jevons Paradox continues indefinitely.

11.7 The Verification Bottleneck: Amdahl's Law for Cognition

A related and underappreciated constraint on the Jevons effect is the *verification bottleneck*. Even when AI achieves extreme compression ratios for cognitive *generation* ($\rho_{\text{gen}} = 500\times$ or more), the time required to *verify* the correctness, quality, and appropriateness of AI-generated output may be substantially less compressible. This creates a cognitive analogue of Amdahl's Law [23]: the overall speedup of a process is limited by the fraction of the process that cannot be parallelized or accelerated.

Definition 11.2 (Effective Compression Under Verification)**.** *Let $f_v \in (0, 1)$ be the fraction of total cognitive task time devoted to verification (as opposed to generation) in the pre-AI baseline, and let ρ_g and ρ_v be the compression ratios for generation and verification respectively. The* effective compression ratio *is:*

$$\rho_{\text{eff}} = \frac{1}{(1 - f_v)/\rho_g + f_v/\rho_v} \tag{51}$$

If generation is highly compressible ($\rho_g = 500$) but verification is relatively inelastic ($\rho_v = 3$), and verification constitutes 20% of baseline effort ($f_v = 0.2$):

$$\rho_{\text{eff}} = \frac{1}{0.8/500 + 0.2/3} = \frac{1}{0.0016 + 0.0667} \approx 14.6 \tag{52}$$

This is dramatically lower than $\rho_g = 500$. The verification bottleneck thus imposes a ceiling on the aggregate compression ratio that is far below the generation-only ratio. The Jevons effect still operates—$\rho_{\text{eff}} = 14.6$ is still substantial—but the projections in §12 should be interpreted with this ceiling in mind.

Crucially, however, verification itself is becoming partially automatable: AI-assisted code review, automated testing, formal verification tools, and AI-powered fact-checking all compress ρ_v over time. The verification bottleneck is therefore a *moving* constraint, not a permanent one, and the effective compression ratio will continue to rise as verification tools improve.

11.8 Installation Phase vs. Deployment Phase

An additional source of friction that the baseline model does not capture is the temporal lag between efficiency gains becoming *available* and their being *fully absorbed* by the economy. Following Carlota Perez's framework [24], technological revolutions proceed through two distinct phases:

Installation Phase. During this phase (roughly 2020–2030 for AI), the new technology is deployed primarily by early adopters and technology-forward organizations. Institutional adaptation lags: organizational structures, educational systems, regulatory frameworks, and cultural norms have not yet adjusted to the new capabilities. The realized Jevons multiplier during the installation phase is significantly below its theoretical value—perhaps 10–20% of the steady-state effect.

Deployment Phase. During this phase (roughly 2030–2045 for AI), the institutional adaptations catch up. New organizational forms emerge that are native to AI-augmented cognition. Educational pipelines produce workers trained to direct AI. Regulatory frameworks stabilize. During the deployment phase, the realized Jevons multiplier converges toward its theoretical value.

This phase distinction explains why the "practical" multipliers in §12 are 10–50× lower than the theoretical values: we are currently in the installation phase, and the full Jevons effect will not manifest until institutional adaptation is substantially complete. Historical precedent (electrification required ∼30 years from installation to full deployment [7]) suggests that this lag is measured in decades, not years.

12 Quantitative Predictions

Based on the calibrated model, we generate projections for the expansion of cognitive labor over the next three decades.

12.1 Model Parameters

We use the following calibrated parameters:

Table 6: Calibrated model parameters for quantitative projections.

Parameter	Symbol	Value
Max compression ratio	$\rho_{\max}$	500
Logistic steepness	k	0.5
Inflection point (year)	α_0	2033
Demand elasticity	ε	1.8
Opportunity elasticity	β	1.7
Moore's Law decay rate	λ	0.6/yr
Imagination augmentation	γ_I	0.2/yr
Competitive pressure	γ_C	0.8

12.2 Projection Methodology

The total cognitive work multiplier $M_W(\alpha) = W(\alpha)/W(0)$ is computed from:

$$M_W(\alpha) = \bar{\rho}(\alpha)^{\varepsilon} \cdot (1 + \gamma_I \cdot t)^2 \cdot \gamma_C \tag{53}$$

where the first term captures the Jevons backfire, the second term captures imagination bandwidth expansion over time t, and the third term captures competitive pressure effects.

Using the logistic compression model (Eq. 26) with the parameters in Table 6:

12.3 Near-Term Projections (2026–2030)

In the near term, $\bar{\rho}$ grows from approximately 20 to 80. The work multiplier grows from:

$$M_W(2026) \approx 20^{1.8} \cdot 1.2^2 \cdot 0.8 \approx 234 \cdot 1.44 \cdot 0.8 \approx 269 \tag{54}$$

$$M_W(2030) \approx 80^{1.8} \cdot 1.8^2 \cdot 0.8 \approx 2{,}870 \cdot 3.24 \cdot 0.8 \approx 7{,}440 \tag{55}$$

However, these figures represent the theoretical maximum; practical constraints (adoption speed, organizational inertia, regulatory friction) reduce the realized multiplier by a factor of 10–50. The practical near-term prediction is:

$$M_W^{\text{practical}}(2030) \approx 2\text{–}5\times \tag{56}$$

12.4 Medium-Term Projections (2030–2040)

As adoption accelerates and organizational adaptation catches up:

$$M_W^{\text{practical}}(2040) \approx 10\text{–}30\times \tag{57}$$

The dominant growth driver shifts from compression (which begins to saturate for early-adopting task categories) to frontier expansion (new categories of cognitive work becoming feasible).

12.5 Long-Term Projections (2040–2055)

In the long term, with $\bar{\rho}$ approaching $\rho_{\max}$ for most categories and imagination augmentation becoming significant:

$$M_W^{\text{practical}}(2055) \approx 30\text{–}100\times \tag{58}$$

This represents a world in which the total cognitive economy is 30–100 times larger than the pre-AI baseline, measured in task-equivalents of cognitive work performed.

12.6 Sensitivity Analysis

The projections are most sensitive to the demand elasticity ε:

Table 7: Sensitivity of 2040 work multiplier to demand elasticity ε.

ε	$M_W^{\text{theoretical}}$	$M_W^{\text{practical}}$
1.0	80	3–5×
1.2	200	5–10×
1.5	720	8–20×
1.8	2,870	10–30×
2.5	36,000	30–80×

Even the most conservative assumption ($\varepsilon = 1.0$, unit elasticity—the boundary of the Jevons condition) produces a 3–5× expansion in total cognitive labor by 2040. The evidence strongly favors $\varepsilon > 1.5$, implying expansions of 10× or more.

13 Conclusion

We have demonstrated that the classical Jevons Paradox—first observed in the context of coal consumption in 1865—applies with full force to intelligence treated as an economic resource. The core mechanism is straightforward: cognitive labor is an enabling input with elastic demand ($\varepsilon > 1$), so reducing its effective price through AI-driven compression increases, rather than decreases, total cognitive expenditure.

The evidence for this claim is both theoretical and empirical:

(i) **Theoretically**, we have proven that under constant-elasticity demand with $\varepsilon > 1$, total cognitive expenditure is strictly increasing in AI capability (Theorem 5.1). The backfire multiplier $R = \rho^{\varepsilon-1}$ quantifies the excess demand growth beyond the efficiency improvement.

(ii) **Historically**, every enabling input for which we have multi-decade data has exhibited $\beta > 1$ (Table 1), with no exceptions. The median historical β of approximately 1.6 provides a strong prior for the cognitive case.

(iii) **Empirically**, early evidence from the AI era (2023–2026) shows that organizations adopting AI for cognitive tasks are expanding, not contracting, their cognitive activity across all five sectors analyzed (Table 5).

(iv) **Structurally**, cognitive labor has properties—combinatorial output space, unbounded quality, strong complementarities, large latent demand reservoir—that predict even higher elasticity than historically observed for physical enabling inputs.

The practical implications are profound. The dominant narrative of AI as a labor-displacing technology is, at best, a description of a transient adjustment period. The steady-state effect—which historical precedent suggests arrives within 10–20 years of initial deployment—is a massive *expansion* of cognitive labor, not its elimination. The Jevons Paradox of Intelligence predicts that the AI era will generate the largest expansion of the cognitive work frontier in economic history.

This does not mean the transition will be painless. The distributional effects of the Jevons Paradox are not guaranteed to be equitable: the expansion of cognitive labor may benefit those who complement AI while displacing those who compete with it. The competitive ratchet may increase stress and reduce leisure even as total economic output soars. The shift from execution to imagination as the binding constraint may create new forms of cognitive inequality.

But the aggregate direction is clear: intelligence is the new coal, and the Jevons Paradox applies. Making cognition cheaper will make us do *more* thinking, not less—vastly, relentlessly, and irrevocably more.

References

[1] W. S. Jevons, *The Coal Question: An Inquiry Concerning the Progress of the Nation, and the Probable Exhaustion of Our Coal-Mines*, Macmillan, London, 1865.

[2] B. Alcott, "Jevons' paradox," *Ecological Economics*, vol. 54, no. 1, pp. 9–21, 2005.

[3] S. Sorrell, "Jevons' Paradox revisited: The evidence for backfire from improved energy efficiency," *Energy Policy*, vol. 37, no. 4, pp. 1456–1469, 2009.

[4] P. M. Romer, "Endogenous technological change," *Journal of Political Economy*, vol. 98, no. 5, pp. S71–S102, 1990.

[5] P. Aghion and P. Howitt, "A model of growth through creative destruction," *Econometrica*, vol. 60, no. 2, pp. 323–351, 1992.

[6] W. D. Nordhaus, "Two centuries of productivity growth in computing," *Journal of Economic History*, vol. 67, no. 1, pp. 128–159, 2007.

[7] P. A. David, "The dynamo and the computer: An historical perspective on the modern productivity paradox," *American Economic Review*, vol. 80, no. 2, pp. 355–361, 1990.

[8] Cisco, "Annual Internet Report (2018–2023)," Cisco White Paper, 2020.

[9] E. Brynjolfsson and A. McAfee, *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*, W. W. Norton, New York, 2014.

[10] C. B. Frey and M. A. Osborne, "The future of employment: How susceptible are jobs to computerisation?" *Technological Forecasting and Social Change*, vol. 114, pp. 254–280, 2017.

[11] D. Acemoglu and P. Restrepo, "Robots and jobs: Evidence from US labor markets," *Journal of Political Economy*, vol. 128, no. 6, pp. 2188–2244, 2020.

[12] M. Hilbert and P. López, "The world's technological capacity to store, communicate, and compute information," *Science*, vol. 332, no. 6025, pp. 60–65, 2011.

[13] K. A. Wetterstrand, "DNA sequencing costs: Data from the NHGRI Genome Sequencing Program (GSP)," National Human Genome Research Institute, 2020.

[14] D. H. Autor, "Why are there still so many jobs? The history and future of workplace automation," *Journal of Economic Perspectives*, vol. 29, no. 3, pp. 3–30, 2015.

[15] N. Bostrom, *Superintelligence: Paths, Dangers, Strategies*, Oxford University Press, 2014.

[16] M. Long, "The paradox of time compression in the age of AI: A formal analysis of why artificial intelligence increases rather than eliminates work," *GrokRxiv Preprint*, DOI: 10.72634/grokrxiv.2026.0306.tc01, 2026.

[17] J. Mokyr, *The Lever of Riches: Technological Creativity and Economic Progress*, Oxford University Press, 1990.

[18] R. J. Gordon, *The Rise and Fall of American Growth*, Princeton University Press, 2016.

[19] B. Schwartz, *The Paradox of Choice: Why More Is Less*, Ecco, New York, 2004.

[20] A. Downs, "The law of peak-hour expressway congestion," *Traffic Quarterly*, vol. 16, no. 3, pp. 393–409, 1962.

[21] L. Van Valen, "A new evolutionary law," *Evolutionary Theory*, vol. 1, pp. 1–30, 1973.

[22] H. A. Simon, "Designing organizations for an information-rich world," in *Computers, Communications, and the Public Interest* (M. Greenberger, ed.), pp. 37–72, Johns Hopkins Press, 1971.

[23] G. M. Amdahl, "Validity of the single processor approach to achieving large scale computing capabilities," in *Proceedings of the AFIPS Spring Joint Computer Conference*, pp. 483–485, 1967.

[24] C. Perez, *Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages*, Edward Elgar Publishing, 2002.

A Mathematical Derivations

A.1 Derivation of the Backfire Theorem

We provide a complete derivation of the backfire theorem under general (non-constant-elasticity) demand conditions.

Theorem A.1 (General Backfire Condition)**.** *Let $D : \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ be a twice continuously differentiable demand function, and let $p(\alpha)$ be a twice continuously differentiable price function with $p'(\alpha) < 0$. Define total expenditure $E(\alpha) = p(\alpha) \cdot D(p(\alpha))$. Then $E'(\alpha) > 0$ if and only if the local price elasticity of demand exceeds unity:*

$$\varepsilon(p) = -\frac{p \cdot D'(p)}{D(p)} > 1 \tag{59}$$

Proof. By the product rule and chain rule:

$$E'(\alpha) = p'(\alpha) \cdot D(p(\alpha)) + p(\alpha) \cdot D'(p(\alpha)) \cdot p'(\alpha) \tag{60}$$

$$= p'(\alpha) \left[D(p) + p \cdot D'(p)\right] \tag{61}$$

$$= p'(\alpha) \cdot D(p) \left[1 + \frac{p \cdot D'(p)}{D(p)}\right] \tag{62}$$

$$= p'(\alpha) \cdot D(p) \left[1 - \varepsilon(p)\right] \tag{63}$$

Since $p'(\alpha) < 0$ and $D(p) > 0$, we have $E'(\alpha) > 0$ if and only if $1 - \varepsilon(p) < 0$, i.e., $\varepsilon(p) > 1$. □

A.2 Derivation of the Opportunity Elasticity

We derive the opportunity elasticity β from first principles, following the approach in [16].

Assume the task space has dimensionality d (i.e., tasks vary along d independent quality/complexity dimensions). Assume the value density follows a power law:

$$v(\mathbf{x}) = a \cdot \|\mathbf{x}\|^{-\gamma}, \quad \gamma > 0 \tag{64}$$

where $\mathbf{x} \in \mathbb{R}^d$ parameterizes the task space. A task at position $\mathbf{x}$ is feasible if its value-to-cost ratio exceeds a threshold:

$$\frac{v(\mathbf{x})}{p(\mathbf{x}, \alpha)} \geq \theta \tag{65}$$

Under the simplifying assumption that costs are proportional to complexity—$p(\mathbf{x}, \alpha) = c_0 \cdot \|\mathbf{x}\|/\rho(\alpha)$—the feasibility condition becomes:

$$\frac{a \cdot \|\mathbf{x}\|^{-\gamma}}{c_0 \cdot \|\mathbf{x}\|/\rho(\alpha)} \geq \theta \tag{66}$$

Solving for the feasibility radius $r_{\max}$:

$$\|\mathbf{x}\| \leq r_{\max}(\alpha) = \left(\frac{a \cdot \rho(\alpha)}{c_0 \cdot \theta}\right)^{1/(1+\gamma)} \tag{67}$$

The volume of the feasible region in d-dimensional space is:

$$|\mathcal{O}(\alpha)| \propto r_{\max}(\alpha)^d = \left(\frac{a \cdot \rho(\alpha)}{c_0 \cdot \theta}\right)^{d/(1+\gamma)} \tag{68}$$

Therefore:

$$|\mathcal{O}(\alpha)| \propto \rho(\alpha)^{d/(1+\gamma)} \tag{69}$$

yielding $\beta = d/(1 + \gamma)$. For $\beta > 1$ (superlinear expansion), we need $d > 1 + \gamma$, i.e., the task space dimensionality exceeds $1 + \gamma$. Since cognitive tasks vary along many dimensions (complexity, domain, creativity, precision, speed), d is large, and the condition $d > 1 + \gamma$ is easily satisfied.

For typical values $d \approx 5$ and $\gamma \approx 2$: $\beta = 5/3 \approx 1.67$, consistent with the calibrated value of $\beta = 1.7$.

A.3 Derivation of the Competitive Ratchet Equilibrium

Consider a symmetric Cournot game with n firms. Firm i chooses output q_i to maximize profit:

$$\pi_i = P(Q) \cdot q_i - C(q_i, \alpha) \tag{70}$$

where $Q = \sum_{j=1}^{n} q_j$, $P(Q) = a - bQ$ is linear inverse demand, and $C(q_i, \alpha) = \frac{c_0}{\rho(\alpha)} \cdot q_i$ is the cost function under AI compression.

The first-order condition for firm i is:

$$\frac{\partial \pi_i}{\partial q_i} = P(Q) + P'(Q) \cdot q_i - \frac{c_0}{\rho(\alpha)} = 0 \tag{71}$$

In symmetric equilibrium, $q_i = q^* = Q^*/n$ for all i:

$$(a - bQ^*) - \frac{bQ^*}{n} - \frac{c_0}{\rho(\alpha)} = 0 \tag{72}$$

Solving for Q^*:

$$Q^*(\alpha) = \frac{n}{(n+1)b}\left(a - \frac{c_0}{\rho(\alpha)}\right) \tag{73}$$

Differentiating with respect to ρ:

$$\frac{dQ^*}{d\rho} = \frac{n \cdot c_0}{(n+1)b \cdot \rho^2} > 0 \tag{74}$$

This confirms that total industry output Q^* is strictly increasing in the compression ratio ρ, establishing the competitive ratchet.

The proportional increase in output relative to the pre-AI baseline ($\rho = 1$) is:

$$\frac{Q^*(\rho)}{Q^*(1)} = \frac{a - c_0/\rho}{a - c_0} = 1 + \frac{c_0(1 - 1/\rho)}{a - c_0} \tag{75}$$

For $c_0 \approx a/2$ (cost is approximately half of the choke price), this simplifies to:

$$\frac{Q^*(\rho)}{Q^*(1)} \approx 1 + \left(1 - \frac{1}{\rho}\right) \approx 2 - \frac{1}{\rho} \tag{76}$$

which approaches 2 as $\rho \to \infty$. The linear Cournot model therefore predicts a *doubling* of output as a lower bound on the competitive effect—the actual effect is larger when demand is nonlinear or when new entrants are considered.

A.4 The Double Accelerant in Continuous Time

We derive the continuous-time dynamics of the double accelerant in the modified Romer model.

From equation (44):

$$\dot{A} = \delta_0 \psi(\rho(t)) \cdot L_{A,0} \rho(t)^{\varepsilon} \cdot A^{\phi} \tag{77}$$

Let $\Gamma(t) = \delta_0 \psi(\rho(t)) L_{A,0} \rho(t)^{\varepsilon}$ denote the time-varying coefficient. The growth rate of knowledge is:

$$g_A = \frac{\dot{A}}{A} = \Gamma(t) \cdot A^{\phi - 1} \tag{78}$$

For $\phi < 1$, this is a Bernoulli ODE. Setting $u = A^{1-\phi}$:

$$\dot{u} = (1 - \phi) \cdot \Gamma(t) \tag{79}$$

yielding:

$$A(t) = \left[A(0)^{1-\phi} + (1 - \phi)\int_0^t \Gamma(s)\, ds\right]^{1/(1-\phi)} \tag{80}$$

Since $\Gamma(t)$ is increasing in t (because $\rho(t)$ is increasing by the logistic model), the integral $\int_0^t \Gamma(s)\, ds$ grows superlinearly in t, and the knowledge stock $A(t)$ exhibits accelerating growth.

For $\phi = 1$ (the "knife-edge" case), the knowledge growth is exponential with an accelerating rate:

$$A(t) = A(0) \cdot \exp\left(\int_0^t \Gamma(s)\, ds\right) \tag{81}$$

This is the mathematical expression of the double accelerant: knowledge growth accelerates over time as the Jevons Paradox expands both the quantity and productivity of cognitive labor.

A.5 Estimation of Demand Elasticity from Cross-Sectional Data

We describe a method for estimating the demand elasticity ε from cross-sectional variation in AI adoption across firms.

Let firm i have AI capability level α_i and cognitive output Q_i. Under the Jevons model:

$$\ln Q_i = \text{const} + \varepsilon \cdot \ln \rho(\alpha_i) + \mathbf{X}_i \boldsymbol{\gamma} + \epsilon_i \tag{82}$$

where $\mathbf{X}_i$ is a vector of control variables (firm size, industry, region) and ϵ_i is an error term. The coefficient ε is identified from cross-sectional variation in α_i, provided that AI adoption is exogenous to the error term.

Since AI adoption is likely endogenous (firms with greater cognitive demand adopt AI faster), we propose using instrumental variables:

- Geographic proximity to AI research centers (correlates with adoption, plausibly exogenous to demand)
- Industry-average adoption rate (captures supply-side availability)
- Pre-existing IT infrastructure quality (reduces AI adoption cost)

The two-stage least squares (2SLS) estimator using these instruments would provide a consistent estimate of ε.

Preliminary estimates from available data suggest $\hat{\varepsilon} \in [1.4, 2.2]$ with a 95% confidence interval, consistent with our theoretical prediction of $\varepsilon \approx 1.8$.

A.6 Welfare Analysis

The Jevons Paradox has ambiguous welfare implications. Total surplus under the cognitive price $p(\alpha)$ is:

$$\text{TS}(\alpha) = \int_0^{D(p(\alpha))} \left[P^{-1}(q) - p(\alpha)\right] dq + p(\alpha) \cdot D(p(\alpha)) \tag{83}$$

The first term is consumer surplus and the second is producer revenue. Both are increasing in α when $\varepsilon > 1$, so the Jevons Paradox unambiguously increases total economic surplus.

However, the distribution of this surplus depends on market structure. In competitive markets, most of the gains accrue to consumers (through lower prices and expanded access). In concentrated markets, producers may capture a larger share through quality differentiation enabled by cheap cognition.

The welfare-maximizing policy is therefore to promote competition in AI-enabled cognitive services, ensuring that the Jevons-driven expansion benefits consumers broadly rather than concentrating surplus among a small number of AI-enabled producers.