# Opportunity Space Expansion Under Cognitive Automation:

## A Superlinear Theory of Work Frontier Growth

Matthew Long[1]

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

March 6, 2026

**Abstract**

We develop a formal theory of how cognitive automation expands the frontier of economically viable work. When artificial intelligence reduces the cost of cognitive tasks, the set of tasks satisfying the feasibility condition $v(\tau)/p(\tau, \alpha) \geq \theta$ grows not merely in proportion to the cost reduction ratio $\rho$, but *superlinearly* as $\rho^\beta$ with $\beta > 1$. We prove this superlinearity arises from the combinatorial structure of newly feasible task combinations and the power-law distribution of value density across the task space. We define the **work multiplier** $M_W(\alpha) = W(\alpha)/W(0)$ and establish a lower bound $M_W \geq \rho^{\beta-1}$, showing that even modest compression ratios produce large expansions of total work when $\beta > 1$. We derive $\beta = d/(1 + \gamma)$ from the dimensionality $d$ of the task space and the power-law exponent $\gamma$ of value density, and provide empirical estimates from historical technology revolutions yielding $\beta \in [1.2, 2.4]$ with a median of approximately 1.7. We decompose the growth rate of total work into three terms—frontier expansion, compression gain, and imagination augmentation—and show that frontier expansion is the dominant term. Eight categories of previously infeasible tasks made viable by AI are analyzed in detail. We connect our framework to induced demand theory, endogenous growth models, and the Jevons Paradox, establishing that the opportunity space expansion is the primary mechanism through which AI increases rather than eliminates total cognitive work. Numerical simulations under calibrated parameters project work multipliers of 7–100× within a generation, depending on the realized values of $\beta$ and $\rho$.

**Keywords:** opportunity space, feasibility frontier, cognitive automation, superlinear expansion, work multiplier, induced demand, task combinatorics, power-law value density, AI economics

**DOI:** 10.72634/grokrxiv.2026.0306.ose01

# 1 Introduction

The prevailing discourse on artificial intelligence and labor conceives of AI primarily as a *substitution technology*: a system that automates existing tasks, thereby reducing the total quantity of human work required. Under this view, the economic question reduces to measuring the displacement rate—what fraction of tasks can AI perform, and on what timeline [12, 2]. This framing, while capturing real transitional dynamics, fundamentally mischaracterizes the steady-state outcome by ignoring the most consequential effect of cost reduction: the creation of entirely new categories of work that were previously infeasible.

Consider a concrete illustration. Before AI-powered code generation, building a custom inventory management system for a small bakery was economically absurd: the software development cost of \$50,000–\$100,000 vastly exceeded the marginal value of optimized inventory for a business generating \$200,000 in annual revenue. The task existed in the *conceivable* task space but fell below the *feasibility threshold*. When AI reduces the development cost to $500, this task crosses the viability boundary. Crucially, the task is not "the same task done faster"—it is a *new* task that enters the economic system for the first time.

The central claim of this paper is that these newly feasible tasks do not merely accumulate linearly as costs fall. They expand *superlinearly*—as $\rho^\beta$ with $\beta > 1$—because of the combinatorial structure of the task space. When multiple categories of tasks simultaneously become cheaper, the number of viable *combinations* grows faster than the number of viable individual tasks. A custom inventory system for a bakery, combined with a personalized customer recommendation engine, combined with an automated supplier negotiation agent, creates a composite workflow whose feasibility requires all three components to be individually viable. The joint feasibility space is the product of the marginal feasibility spaces, yielding superlinear growth.

This paper develops the formal theory of this **opportunity space expansion**, derives the conditions under which superlinearity obtains, estimates the exponent $\beta$ from historical data, and projects the implications for total cognitive work in the AI era.

### 1.1 Contribution and Scope

This paper makes five principal contributions:

(i) We formalize the *opportunity space* $\mathcal{O}(\alpha)$ as the subset of a measurable task universe satisfying a feasibility condition, and define the *work multiplier* $M_W(\alpha)$ that quantifies the expansion of total work relative to the pre-AI baseline.

(ii) We prove that under a power-law value density model with exponent $\gamma$ on a $d$-dimensional task space, the opportunity space grows as $\rho^{d/(1+\gamma)}$, yielding $\beta > 1$ whenever $d > 1 + \gamma$.

(iii) We decompose the growth rate of total work into three additive terms—frontier expansion, compression gain, and imagination augmentation—and argue on both theoretical and empirical grounds that frontier expansion dominates.

(iv) We provide empirical estimates of $\beta$ from six historical technology revolutions, finding $\beta \in [1.2, 2.4]$ with a median of 1.7.

(v) We analyze eight categories of previously infeasible tasks that AI makes viable for the first time, estimating the aggregate work expansion implied by each category.

### 1.2 Relation to the Time Compression Paradox

This paper develops one of the three mechanisms underlying the Time Compression Paradox (TCP) formalized in [17]. The TCP states that AI does not eliminate work but compresses the amount of work achievable within a given time period, paradoxically creating more work. The TCP operates through three channels: (i) the Jevons Paradox of intelligence (elastic demand for cognitive labor), (ii) opportunity space expansion (the subject of this paper), and (iii) competitive dynamics (market forces that prevent the time surplus from being consumed as leisure). We focus exclusively on channel (ii), providing a deeper mathematical treatment than is possible within the unified framework.

### 1.3 Paper Organization

Section 2 formalizes the opportunity space framework. Section 3 analyzes feasibility thresholds and cost reduction. Section 4 presents the superlinearity argument with formal proofs. Section 5 defines and bounds the work multiplier. Section 6 analyzes the time surplus and its allocation. Section 7 examines eight categories of previously infeasible tasks. Section 8 derives the growth rate decomposition. Section 9 develops the power-law value density model. Section 10 estimates $\beta$ from historical data. Section 11 provides frontier visualizations. Section 12 connects to induced demand theory. Section 13 presents a multi-sector analysis. Section 14 performs sensitivity analysis. Section 15 concludes. Appendix A contains full derivations.

## 2 The Opportunity Space Framework

### 2.1 The Task Universe

We begin by defining the space of all conceivable cognitive tasks.

**Definition 2.1** (Task Universe)**.** *The* task universe *$\mathcal{T}$ is a measurable space $(\mathcal{T}, \Sigma, \mu)$ where each element $\tau \in \mathcal{T}$ represents a conceivable cognitive task. Each task is characterized by:*

- *$c(\tau) \in \mathcal{C}$: the task category (e.g., translation, analysis, design),*
- *$x(\tau) \in \mathbb{R}^d_{>0}$: a $d$-dimensional complexity vector encoding the task's resource requirements across $d$ independent dimensions,*
- *$v(\tau) \in \mathbb{R}_{>0}$: the economic value produced upon completion.*

*The measure $\mu$ is a $\sigma$-finite measure encoding the density of tasks across the space. In the uniform-density case used throughout the paper (Sections 4–9), $\mu$ is taken to be the Lebesgue measure on $\mathbb{R}^d_{>0}$ scaled by a constant task density $\mu_0$.*

The dimensionality $d$ of the complexity vector is a crucial parameter. Each dimension represents an independent axis along which task complexity can vary: linguistic sophistication, domain expertise required, data volume, creative originality, precision requirements, temporal urgency, and so forth. We argue below that $d$ is large—typically $d \geq 5$—which is the fundamental source of superlinearity.

**Definition 2.2** (Cognitive Cost Function)**.** *The* cognitive cost *of task $\tau$ at AI capability level $\alpha \geq 0$ is:*

$$p(\tau, \alpha) = w_h \cdot T(c(\tau), \alpha) + c_{AI}(\tau, \alpha) \tag{1}$$

*where $w_h$ is the human hourly wage, $T(c, \alpha)$ is the human time required for a task of category $c$ with AI assistance at level $\alpha$, and $c_{AI}(\tau, \alpha)$ is the direct AI compute cost.*

The cost function is monotonically decreasing in $\alpha$:

$$\frac{\partial p}{\partial \alpha}(\tau, \alpha) < 0 \quad \forall \tau \in \mathcal{T},\ \alpha > 0 \tag{2}$$

This follows from $\partial T/\partial \alpha < 0$ (AI reduces human time) and $\partial c_{\mathrm{AI}}/\partial \alpha \leq 0$ (AI compute costs decrease with scale and Moore's Law effects).

## 2.2 The Feasibility Condition

**Definition 2.3** (Feasibility Threshold)**.** *A task $\tau \in \mathcal{T}$ is* feasible *at AI capability $\alpha$ if its value-to-cost ratio exceeds a threshold $\theta > 0$:*

$$\frac{v(\tau)}{p(\tau, \alpha)} \geq \theta \tag{3}$$

*The threshold $\theta$ represents the minimum return on investment required for a rational agent to undertake the task, incorporating opportunity costs, risk premia, and transaction costs.*

**Definition 2.4** (Opportunity Space)**.** *The* opportunity space *at AI capability $\alpha$ is:*

$$\mathcal{O}(\alpha) = \left\{ \tau \in \mathcal{T} \;\middle|\; \frac{v(\tau)}{p(\tau, \alpha)} \geq \theta \right\} \tag{4}$$

*Its measure is $|\mathcal{O}(\alpha)| = \mu(\mathcal{O}(\alpha))$.*

**Remark 2.1.** *The opportunity space is monotonically expanding in $\alpha$: if $\alpha_1 < \alpha_2$, then $\mathcal{O}(\alpha_1) \subseteq \mathcal{O}(\alpha_2)$, since $p(\tau, \alpha_2) \leq p(\tau, \alpha_1)$ for all $\tau$. Tasks that are feasible at a lower AI capability remain feasible at a higher capability level.*

## 2.3 The Compression Ratio

**Definition 2.5** (Compression Ratio)**.** *The* compression ratio *for task category $c$ at AI capability $\alpha$ is:*

$$\rho(c, \alpha) = \frac{p(c, 0)}{p(c, \alpha)} \geq 1 \tag{5}$$

*where $p(c, \alpha)$ denotes the representative cost for category $c$. We write $\rho(\alpha)$ for the aggregate (median) compression ratio across all categories.*

The compression ratio measures the factor by which AI reduces task costs. Empirical estimates for current AI systems (circa 2025) yield median $\rho \approx 48\times$ across representative cognitive tasks, with translation achieving $288\times$ and dataset analysis $216\times$ [17].

## 2.4 Total Work

**Definition 2.6** (Total Work)**.** *The* total cognitive work *performed at AI capability $\alpha$ is:*

$$W(\alpha) = \int_{\mathcal{O}(\alpha)} w(\tau, \alpha)\, d\mu(\tau) \tag{6}$$

*where $w(\tau, \alpha)$ is the work intensity (effort per unit time) allocated to task $\tau$.*

Total work increases through two channels: (i) the expansion of $\mathcal{O}(\alpha)$ (new tasks enter the feasible set), and (ii) intensification of existing tasks (higher $w$ for tasks already in the opportunity space). This paper focuses primarily on channel (i).

# 3 Feasibility Thresholds and Cost Reduction

## 3.1 The Viability Boundary

The feasibility condition (3) defines a boundary in the $(v, p)$ plane:

$$v = \theta \cdot p(\tau, \alpha) \tag{7}$$

Tasks above this line are feasible; tasks below it are not. As $\alpha$ increases, $p$ decreases, and the boundary shifts downward, admitting new tasks. We can equivalently express the boundary in terms of complexity:

**Proposition 3.1** (Maximum Feasible Complexity)**.** *For a task category $c$ with value function $v(x)$ and cost function $p(x, \alpha)$, the maximum feasible complexity $x^*(\alpha)$ satisfies:*

$$v(x^*(\alpha)) = \theta \cdot p(x^*(\alpha), \alpha) \tag{8}$$

*If $v(x)$ is decreasing in $\|x\|$ (more complex tasks have lower marginal value) and $p(x, \alpha)$ is increasing in $\|x\|$ (more complex tasks cost more), then $x^*(\alpha)$ is increasing in $\alpha$: higher AI capability raises the maximum feasible complexity.*

*Proof.* Here and below, $x$ denotes the scalar magnitude $\|x\|$ of the complexity vector, so that $v(x)$, $p(x, \alpha)$, and their derivatives are ordinary single-variable functions. By the implicit function theorem, differentiating $v(x^*) = \theta \cdot p(x^*, \alpha)$ with respect to $\alpha$:

$$v'(x^*) \frac{dx^*}{d\alpha} = \theta \left[ \frac{\partial p}{\partial x} \frac{dx^*}{d\alpha} + \frac{\partial p}{\partial \alpha} \right] \tag{9}$$

Solving:

$$\frac{dx^*}{d\alpha} = \frac{-\theta \, \partial p / \partial \alpha}{v'(x^*) - \theta \, \partial p / \partial x} \tag{10}$$

Since $\partial p / \partial \alpha < 0$, $v'(x^*) < 0$, and $\partial p / \partial x > 0$, the denominator is negative and the numerator is positive, giving $dx^*/d\alpha > 0$. □

## 3.2 Worked Examples of Tasks Crossing the Viability Boundary

We now present detailed examples of tasks transitioning from infeasible to feasible under AI cost reduction.

**Example 3.1** (Personalized Tutoring)**.** *Consider one-on-one tutoring for a high school student struggling with calculus. Pre-AI cost: \$50/hour for a qualified tutor, requiring approximately 40 hours per semester, totaling \$2,000. Value to the family: approximately \$500 per semester (weighted by probability of grade improvement and its downstream effects). The feasibility ratio is $v/p = 500/2000 = 0.25$. With $\theta = 1$ (break-even threshold), this task is infeasible.*

*With AI tutoring systems, the cost drops to approximately \$20 per semester (subscription cost). The feasibility ratio becomes $v/p = 500/20 = 25 \gg \theta$. The task crosses the viability boundary, and the compression ratio is $\rho = 2000/20 = 100\times$.*

**Example 3.2** (Custom Business Software)**.** *A local restaurant wants a custom reservation management system integrated with their kitchen workflow. Pre-AI development cost: \$75,000 for a software consultancy. Annual value of the system: approximately \$8,000 in efficiency gains. Payback period: 9.4 years, well beyond the 2-year threshold ($\theta_{time} = 2$ years). Infeasible.*

*With AI code generation, development cost drops to approximately \$2,000 (20 hours of AI-assisted development at \$100/hour). Payback period: 0.25 years $\ll \theta_{time}$. The task becomes strongly feasible, with $\rho = 75000/2000 = 37.5\times$.*

**Example 3.3** (Real-Time Portfolio Analysis)**.** *A retail investor with a \$50,000 portfolio wants continuous quantitative analysis comparable to institutional-grade research. Pre-AI cost: \$5,000/month for a financial analyst. Annual value to the investor: approximately \$3,000 in improved returns. Infeasible ($v/p = 3000/60000 = 0.05$).*

*With AI analysis tools, cost drops to \$30/month (\$360/year). Feasibility ratio: $v/p = 3000/360 = 8.3$. Strongly feasible, with $\rho = 60000/360 \approx 167\times$.*

Table 1: Tasks crossing the viability boundary under AI cost reduction. Threshold $\theta = 1$.

| **Task** | $v$ | $p_0$ | $p_\alpha$ | $\rho$ |
|---|---|---|---|---|
| Personal tutor | \$500 | \$2,000 | \$20 | 100× |
| Custom software | \$8,000 | \$75,000 | \$2,000 | 38× |
| Portfolio analysis | \$3,000 | \$60,000 | \$360 | 167× |
| Legal review | \$800 | \$5,000 | \$50 | 100× |
| Drug interactions | \$2,000 | \$15,000 | \$100 | 150× |
| Market research | \$1,500 | \$25,000 | \$300 | 83× |

### 3.3 The Threshold Cascade

An important phenomenon emerges when costs decrease continuously: tasks do not cross the viability boundary simultaneously but in a *cascade*, ordered by their pre-AI feasibility ratios.

**Proposition 3.2** (Threshold Cascade Ordering). *Tasks cross the viability boundary in decreasing order of their pre-AI feasibility ratio $r_0(\tau) = v(\tau)/p(\tau, 0)$. The crossing compression ratio for task $\tau$ is:*

$$\rho^*(\tau) = \frac{\theta}{r_0(\tau)} \tag{11}$$

*Tasks with higher pre-AI $r_0$ require less compression to become feasible.*

*Proof.* Task $\tau$ becomes feasible when $v(\tau)/p(\tau, \alpha) \geq \theta$. Since $p(\tau, \alpha) = p(\tau, 0)/\rho(\alpha)$ (by definition of the compression ratio applied to cost), the condition becomes $v(\tau) \cdot \rho(\alpha)/p(\tau, 0) \geq \theta$, i.e., $\rho(\alpha) \geq \theta/r_0(\tau) = \rho^*(\tau)$. Tasks with higher $r_0$ have lower $\rho^*$ and cross first. □

This cascade structure means that the *rate* of frontier expansion depends on the distribution of tasks near the viability boundary. If many tasks are clustered just below the threshold, even a small cost reduction induces a large expansion—a phenomenon we formalize in Section 9.

## 4 The Superlinearity Argument

This section establishes the central mathematical result of the paper: the opportunity space grows superlinearly in the compression ratio $\rho$.

### 4.1 Statement of the Main Result

**Theorem 4.1** (Superlinear Opportunity Expansion). *Suppose the task universe $\mathcal{T}$ is embedded in $\mathbb{R}^d$ with $d \geq 2$, and the value density follows a power law $v(x) = a \cdot \|x\|^{-\gamma}$ with $\gamma > 0$. If the cost function is proportional to complexity, $p(x, \alpha) = b \cdot \|x\|/\rho(\alpha)$, then the measure of the opportunity space satisfies:*

$$|\mathcal{O}(\alpha)| = |\mathcal{O}(0)| \cdot \rho(\alpha)^\beta \tag{12}$$

*where $\beta = d/(1 + \gamma)$. This gives $\beta > 1$ whenever $d > 1 + \gamma$.*

We devote the remainder of this section to developing the proof and its implications.

### 4.2 The Combinatorial Argument (Intuition)

Before the formal proof, we provide the combinatorial intuition for why $\beta > 1$.

Consider a simplified world with $d$ independent task dimensions, each with $n$ feasibility levels. A composite task requires feasibility across all $d$ dimensions simultaneously. If a cost reduction by factor $\rho$ makes $k(\rho)$ new levels feasible in each dimension, the number of newly feasible composite tasks is:

$$\Delta|\mathcal{O}| \sim k(\rho)^d \tag{13}$$

If $k(\rho)$ grows at least linearly in $\rho$ (each doubling of cost reduction opens at least a proportional number of new levels per dimension), then the total feasible space grows polynomially in $\rho$ with exponent $d$. The power-law value distribution tempers this growth by the factor $(1 + \gamma)$, yielding the net exponent $\beta = d/(1 + \gamma)$.

**Example 4.1** (Two-Dimensional Illustration). *Consider tasks characterized by two complexity dimensions:* linguistic sophistication *$(x_1)$ and* domain depth *$(x_2)$. Suppose that at $\rho = 1$, tasks with $\|x\| \leq R_0$ are feasible. At $\rho = 4$, the feasible radius doubles (since $R^* \propto \rho^{1/(1+\gamma)}$). The area of the feasible region grows by $2^2 = 4\times$ in 2D, compared to $2\times$ growth in each individual dimension. For $\gamma = 1$ and $d = 2$, we get $\beta = 2/2 = 1$—exactly linear. But for $\gamma = 0.5$ and $d = 2$, we get $\beta = 2/1.5 = 1.33$—superlinear.*

### 4.3 Formal Proof of Theorem 4.1

*Proof.* We work in the $d$-dimensional complexity space $\mathbb{R}^d_{>0}$. The feasibility condition for a task at position $x$ is:

$$\frac{v(x)}{p(x,\alpha)} = \frac{a\|x\|^{-\gamma}}{b\|x\|/\rho(\alpha)} = \frac{a\rho(\alpha)}{b\|x\|^{1+\gamma}} \geq \theta \tag{14}$$

This yields the feasibility region:

$$\|x\| \leq \left(\frac{a\rho(\alpha)}{b\theta}\right)^{1/(1+\gamma)} \equiv R^*(\alpha) \tag{15}$$

The feasible region is therefore a $d$-dimensional ball of radius $R^*(\alpha)$. Its $\mu$-measure (assuming uniform task density $\mu_0$ over the task space) is:

$$|\mathcal{O}(\alpha)| = \mu_0 \cdot V_d \cdot [R^*(\alpha)]^d \tag{16}$$

where $V_d = \pi^{d/2}/\Gamma(d/2+1)$ is the volume of the unit $d$-ball. Substituting:

$$|\mathcal{O}(\alpha)| = \mu_0 V_d \left(\frac{a}{b\theta}\right)^{d/(1+\gamma)} \rho(\alpha)^{d/(1+\gamma)} \tag{17}$$

$$= |\mathcal{O}(0)| \cdot \rho(\alpha)^{d/(1+\gamma)} \tag{18}$$

since at $\alpha = 0$ (no AI assistance) we have $\rho(0) \equiv 1$ by definition—the baseline cost ratio $p(x,0)/p(x,0) = 1$—and therefore $|\mathcal{O}(0)| = \mu_0 V_d (a/(b\theta))^{d/(1+\gamma)}$. Thus $\beta = d/(1+\gamma)$.

For $\beta > 1$, we need $d/(1+\gamma) > 1$, i.e., $d > 1+\gamma$. Since cognitive tasks naturally span many independent dimensions ($d \geq 5$ in practice) and value distributions have $\gamma \in (0.5, 2)$ empirically, the condition $d > 1+\gamma$ is satisfied with large margin. □

### 4.4 The Dimensionality Argument

The key insight from Theorem 4.1 is that $\beta$ increases with the dimensionality $d$ of the task space. We now argue that $d$ is inherently large for cognitive work.

**Proposition 4.2** (Task Space Dimensionality)**.** *The cognitive task space has effective dimensionality $d \geq 5$, corresponding to at least the following independent complexity axes:*

*(a)* Domain breadth*: number of distinct knowledge domains required.*

Table 2: The superlinearity exponent $\beta = d/(1+\gamma)$ for different values of task space dimensionality $d$ and value decay exponent $\gamma$.

| | $\gamma = 0.5$ | $\gamma = 1.0$ | $\gamma = 1.5$ | $\gamma = 2.0$ |
|---|---|---|---|---|
| $d = 2$ | 1.33 | 1.00 | 0.80 | 0.67 |
| $d = 3$ | 2.00 | 1.50 | 1.20 | 1.00 |
| $d = 4$ | 2.67 | 2.00 | 1.60 | 1.33 |
| $d = 5$ | 3.33 | 2.50 | 2.00 | 1.67 |
| $d = 7$ | 4.67 | 3.50 | 2.80 | 2.33 |
| $d = 10$ | 6.67 | 5.00 | 4.00 | 3.33 |

*(b)* Analytical depth*: depth of reasoning chains involved.*

*(c)* Data volume*: quantity of information that must be processed.*

*(d)* Creative originality*: degree of novel synthesis required.*

*(e)* Precision requirements*: tolerance for error.*

*(f)* Temporal urgency*: speed of response required.*

*(g)* Stakeholder complexity*: number of parties whose interests must be balanced.*

With $d = 7$ and $\gamma = 1.5$ (a typical power-law exponent for economic value distributions), we obtain $\beta = 7/2.5 = 2.8$. Even conservatively, with $d = 5$ and $\gamma = 2$, we get $\beta = 5/3 \approx 1.67$. The condition $\beta > 1$ is robust across reasonable parameter estimates.

### 4.5 Generalizations

The basic result extends to several more realistic settings.

**Corollary 4.3** (Non-Uniform Task Density)**.** *If the task density is $\mu(x) = \mu_0\|x\|^{-\delta}$ (tasks are denser near the origin, i.e., simpler tasks are more numerous), the superlinearity exponent becomes:*

$$\beta = \frac{d-\delta}{1+\gamma} \tag{19}$$

*Superlinearity ($\beta > 1$) requires $d - \delta > 1 + \gamma$.*

**Corollary 4.4** (Anisotropic Compression)**.** *If AI compresses costs at different rates along different dimensions, with compression ratio $\rho_i$ along dimension $i$, the opportunity space grows as:*

$$|\mathcal{O}| \propto \prod_{i=1}^{d} \rho_i^{1/(1+\gamma)} \tag{20}$$

*The effective $\beta$ is the harmonic-weighted sum of the per-dimension contributions, and superlinearity is enhanced when compression is broadly distributed across many dimensions.*

# 5 The Work Multiplier

## 5.1 Definition and Basic Properties

**Definition 5.1** (Work Multiplier)**.** *The* work multiplier *at AI capability $\alpha$ is:*

$$M_W(\alpha) = \frac{W(\alpha)}{W(0)} \tag{21}$$

*This measures the factor by which total cognitive work has expanded relative to the pre-AI baseline.*

**Theorem 5.2** (Work Multiplier Lower Bound)**.** *Under the assumptions of Theorem 4.1, the work multiplier satisfies:*

$$M_W(\alpha) \geq \rho(\alpha)^{\beta-1} \tag{22}$$

*Proof.* Total work at capability $\alpha$ can be decomposed as:

$$W(\alpha) = \int_{\mathcal{O}(\alpha)} w(\tau, \alpha)\, d\mu(\tau) \tag{23}$$

We consider two contributions: work on tasks that were already feasible at $\alpha = 0$ (the *existing* frontier), and work on newly feasible tasks (the *expanded* frontier).

For existing tasks $\tau \in \mathcal{O}(0)$, the work intensity per task may decrease (since each task takes less time), but the total available time is reallocated to new tasks. Specifically, the time saved on existing tasks is:

$$\Delta T_{\text{saved}} = \sum_{\tau \in \mathcal{O}(0)} T(\tau, 0)\left(1 - \frac{1}{\rho(\alpha)}\right) \approx W(0)\left(1 - \frac{1}{\rho(\alpha)}\right) \tag{24}$$

The expanded frontier contains $|\mathcal{O}(\alpha)| - |\mathcal{O}(0)| = |\mathcal{O}(0)|(\rho^\beta - 1)$ new tasks. Even if each new task receives only average work intensity $\bar{w} = W(0)/|\mathcal{O}(0)|$, the work on new tasks is:

*Caveat.* Newly feasible tasks are, by definition, marginal: their value-to-cost ratios $v/p$ barely exceed $\theta$. Rational agents may therefore allocate *less* intensity to these marginal tasks than the inframarginal average $\bar{w}$. However, the bound remains valid because the *volume* of newly feasible tasks grows as $\rho^\beta - 1$, which dominates any constant-factor reduction in per-task intensity for $\rho \gg 1$. Specifically, even if marginal tasks receive intensity $\bar{w}/k$ for some constant $k \geq 1$, the lower bound becomes $M_W \geq \rho^{\beta-1}/k$, which still diverges superlinearly.

$$W_{\text{new}} = \bar{w} \cdot |\mathcal{O}(0)| \cdot (\rho^\beta - 1) = W(0)(\rho^\beta - 1) \tag{25}$$

The total work on existing tasks, after compression, is at least $W(0)/\rho$ (the same tasks done in less time). Thus:

$$W(\alpha) \geq \frac{W(0)}{\rho} + W(0)(\rho^\beta - 1) \tag{26}$$

$$= W(0)\left[\frac{1}{\rho} + \rho^\beta - 1\right] \tag{27}$$

For $\rho \geq 1$ and $\beta > 1$, the dominant term is $\rho^\beta$, giving:

$$M_W(\alpha) = \frac{W(\alpha)}{W(0)} \geq \rho^\beta - 1 + \frac{1}{\rho} \geq \rho^{\beta-1} \tag{28}$$

where the last inequality follows from $\rho^\beta - 1 + 1/\rho \geq \rho^{\beta-1}$ for $\rho \geq 1$ and $\beta > 1$ (proved in Appendix A.1). □

## 5.2 Numerical Estimates of the Work Multiplier

Table 3 reveals several important features:

(i) Even modest superlinearity ($\beta = 1.3$) combined with realistic compression ratios ($\rho = 50$) yields a work multiplier of at least 3.3×—total work more than triples.

(ii) At $\beta = 1.7$ (our median empirical estimate) and $\rho = 50$ (the current median across cognitive tasks), the work multiplier exceeds 15×.

Table 3: Work multiplier $M_W \geq \rho^{\beta-1}$ for selected compression ratios and superlinearity exponents. Values represent the *lower bound* on the expansion of total cognitive work.

| $\rho$ | $\beta = 1.3$ | $\beta = 1.5$ | $\beta = 1.7$ | $\beta = 2.0$ |
|---|---|---|---|---|
| 5 | 1.6 | 2.2 | 3.1 | 5.0 |
| 10 | 2.0 | 3.2 | 5.0 | 10.0 |
| 20 | 2.5 | 4.5 | 8.0 | 20.0 |
| 50 | 3.3 | 7.1 | 15.1 | 50.0 |
| 100 | 4.0 | 10.0 | 25.1 | 100.0 |
| 200 | 4.9 | 14.1 | 40.5 | 200.0 |
| 500 | 6.7 | 22.4 | 83.5 | 500.0 |

(iii) The multiplier is highly sensitive to $\beta$: at $\rho = 100$, increasing $\beta$ from 1.3 to 2.0 increases the lower bound from 4× to 100×.

(iv) When $\beta = 2$, the work multiplier equals $\rho$ itself—total work grows linearly with the compression ratio. This is the "perfect Jevons" scenario where every unit of cost reduction is fully offset by work expansion.

### 5.3 Tightness of the Bound

The bound $M_W \geq \rho^{\beta-1}$ is conservative for two reasons. First, it assumes that newly feasible tasks receive only average work intensity, whereas many newly viable tasks (e.g., personalized medicine) are high-value and attract above-average investment. Second, it ignores the intensification effect on existing tasks: when AI makes a task faster, practitioners often invest the time surplus in higher quality rather than mere completion, increasing $w(\tau, \alpha)$ above $w(\tau, 0)/\rho$.

**Proposition 5.3** (Tight Upper Bound on Work Multiplier)**.** *Under the additional assumption that work intensity per task is bounded by $w_{\max}$ and that the number of working hours $H$ is fixed, the work multiplier is bounded above by:*

$$M_W(\alpha) \leq \frac{H \cdot \rho(\alpha)}{T_{\min}} \cdot \frac{1}{W(0)} \tag{29}$$

*where $T_{\min}$ is the minimum time per task after compression. In the limit of very high $\rho$, the work multiplier is bounded by the ratio of imagination bandwidth to the pre-AI work rate, $M_W \leq B_I/(W(0)/H)$, where $B_I$ is the maximum rate at which humans can specify well-defined tasks.*

## 6 The Time Surplus and Its Allocation

### 6.1 Definition and Magnitude

When AI compresses the time required for existing tasks, it creates a *time surplus*:

$$\Delta T(\alpha) = \sum_{\tau \in \mathcal{O}(0)} T(\tau, 0) \left(1 - \frac{1}{\rho(\tau, \alpha)}\right) \tag{30}$$

For a worker spending $H = 2,000$ hours/year on cognitive tasks with median compression $\rho = 50$, the surplus is:

$$\Delta T = 2000 \left(1 - \frac{1}{50}\right) = 1{,}960 \text{ hours/year} \tag{31}$$

This surplus is equivalent to 98% of the original working year. The question of how it is allocated determines the magnitude of the work multiplier.

### 6.2 Three Fates of the Time Surplus

**Definition 6.1** (Surplus Allocation)**.** *The time surplus is allocated among three activities:*

$$\Delta T = \Delta T_L + \Delta T_I + \Delta T_E \tag{32}$$

*where:*

- $\Delta T_L$ = *time consumed as* leisure*,*
- $\Delta T_I$ = *time devoted to* intensification *(more iterations, higher quality on existing tasks),*
- $\Delta T_E$ = *time devoted to* expansion *(newly feasible tasks from the expanded frontier).*

We define the allocation fractions $\lambda_L, \lambda_I, \lambda_E$ with $\lambda_L + \lambda_I + \lambda_E = 1$.

### 6.3 Why Expansion Dominates: Historical Evidence

Historical evidence overwhelmingly favors expansion over leisure:

Table 4: Historical allocation of productivity surpluses. In every case, expansion absorbs the dominant fraction of the surplus.

| **Technology** | $\lambda_L$ | $\lambda_I$ | $\lambda_E$ |
|---|---|---|---|
| Steam / Coal | 0.05 | 0.25 | 0.70 |
| Electricity | 0.03 | 0.20 | 0.77 |
| Computing | $< 0.01$ | 0.15 | 0.85 |
| Telecom | $< 0.01$ | 0.10 | 0.90 |
| Internet | $< 0.01$ | 0.10 | 0.90 |
| **Median** | **0.02** | **0.15** | **0.83** |

(i) **Working hours**: Despite a 50-fold increase in labor productivity since 1870 in the United States, average annual working hours have decreased only from approximately 2,900 to 1,750—a factor of 1.66×, far less than the productivity gain [13]. The implied leisure fraction is $\lambda_L \approx 0.016$ (1.6% of the productivity gain went to reduced hours).

(ii) **Computing**: Computing costs fell by $10^6\times$ from 1970 to 2010, but total computing expenditure increased by $10^9\times$ [18]. The entire cost reduction—and much more—was absorbed by new computational tasks.

(iii) **Communication**: Internet-era communication costs fell by $10^4\times$, but total communication volume increased by $5 \times 10^{10}\times$ [9]. Leisure absorption was negligible.

(iv) **Printing**: The printing press reduced book copying time by $\sim 200\times$, but total text production labor increased by at least 10× within two centuries [11].

### 6.4 The Competitive Mechanism

The dominance of expansion over leisure is not merely an empirical regularity—it is a consequence of competitive dynamics. In a competitive market, firms that allocate their time surplus to expansion outcompete firms that allocate it to leisure. The equilibrium in a symmetric $n$-firm Cournot game yields an expansion fraction approaching 1 as the number of competitors grows [17]:

**Proposition 6.2** (Competitive Expansion Dominance)**.** *In an $n$-firm Cournot competition where each firm chooses allocation fractions $(\lambda_L^i, \lambda_I^i, \lambda_E^i)$ to maximize profit, the Nash equilibrium satisfies $\lambda_E^* \to 1$ as $n \to \infty$.*

The intuition is straightforward: any firm that consumes its surplus as leisure while competitors invest theirs in new products and services will lose market share. The competitive ratchet forces surplus allocation toward expansion.

### 6.5 Endogenous Wage Effects

The analysis above treats the human wage $w_h$ in Definition 2.2 as exogenous. In general equilibrium, however, massive opportunity space expansion feeds back into the labor market. If the work multiplier $M_W$ reaches 15×, demand for human cognitive labor—for oversight, specification, integration, and judgment tasks that remain non-automatable—will increase sharply, driving $w_h$ upward. This wage inflation raises $p(\tau, \alpha) = w_h \cdot T + c_{\text{AI}}$ for all tasks with a human-time component, partially offsetting the cost reduction from $\rho$ and acting as a negative feedback loop on frontier expansion.

**Remark 6.1** (Effective Compression under Endogenous Wages)**.** *Let $w_h(\alpha)$ denote the equilibrium wage at AI capability $\alpha$, with $w_h(\alpha) \geq w_h(0)$. Define the* effective compression ratio*:*

$$\rho_{\text{eff}}(\alpha) = \frac{p(\tau, 0)}{p(\tau, \alpha)} = \frac{w_h(0)\, T_0 + c_{\text{AI},0}}{w_h(\alpha)\, T_0/\rho + c_{\text{AI}}(\alpha)} \qquad (33)$$

*If human time constitutes fraction $\phi$ of baseline cost and the labor supply elasticity is $\eta$ (so $w_h(\alpha)/w_h(0) \approx M_W^{1/\eta}$), the effective superlinearity exponent becomes:*

$$\beta_{\text{eff}} \;\approx\; \frac{\beta}{1 + \phi(\beta - 1)/\eta} \qquad (34)$$

*For elastic labor supply ($\eta \gg 1$), $\beta_{\text{eff}} \approx \beta$ and the partial-equilibrium predictions hold. For inelastic supply ($\eta \approx 1$), the feedback substantially dampens the expansion: at $\beta = 1.7$ and $\phi = 0.5$, we obtain $\beta_{\text{eff}} \approx 1.26$.*

The partial-equilibrium estimates presented in this paper (Tables 3–9) should therefore be understood as upper bounds on the realized work multiplier. The true multiplier depends on how elastic the supply of human cognitive labor proves to be—a question that depends on education pipeline capacity, immigration policy, and the rate at which AI itself substitutes for human oversight, all of which lie beyond the scope of this paper.

# 7 Previously Infeasible Task Categories

We now examine eight categories of cognitive work that AI renders feasible for the first time. For each, we estimate the pre-AI cost, the AI-assisted cost, the compression ratio, and the implied contribution to total work expansion.

## 7.1 Personalized Education at Scale

**Pre-AI cost:** One-on-one tutoring costs \$30–\$80/hour. Providing every student in the US ($\sim$50 million K-12 students) with 5 hours/week of personalized tutoring would cost approximately \$500 billion/year—exceeding the entire US public education budget.

**AI-assisted cost:** AI tutoring systems can provide continuous personalized instruction at approximately \$200/student/year, totaling \$10 billion for universal coverage.

**Compression:** $\rho \approx 50\times$.

**Implied new work:** Developing, maintaining, calibrating, and overseeing AI tutoring systems; creating adaptive curriculum content; training teachers to work alongside AI tutors; analyzing learning outcomes at individual granularity. We estimate this generates approximately 500,000 new full-time-equivalent (FTE) positions.

## 7.2 Individualized Medicine

**Pre-AI cost:** Comprehensive genomic analysis with personalized treatment recommendations costs \$5,000–\$20,000 per patient. Providing this for all 330 million US residents would cost \$1.6–\$6.6 trillion—multiple times the total US healthcare budget.

**AI-assisted cost:** AI-driven genomic interpretation and drug interaction analysis could be delivered at \$50–\$200/patient/year, totaling \$16–\$66 billion.

**Compression:** $\rho \approx 100\times$.

**Implied new work:** Pharmacogenomic database curation, AI model validation, clinical integration workflows, regulatory compliance for AI-assisted diagnostics, adverse event monitoring, continuous model retraining. Estimated 800,000 new FTE.

## 7.3 Universal Custom Software

**Pre-AI cost:** Custom software development costs \$50,000–\$500,000 per application. There are approximately 33 million small businesses in the US, most of which could benefit from at least one custom application. Total cost: \$1.6–\$16 trillion.

**AI-assisted cost:** AI-generated custom applications could be delivered at \$500–\$5,000 each, totaling \$16–\$165 billion.

**Compression:** $\rho \approx 100\times$.

**Implied new work:** Specifying requirements, integrating AI-generated applications with existing systems, ongoing maintenance and feature expansion, security auditing, user training. Estimated 2 million new FTE.

## 7.4 Continuous Market Intelligence for SMEs

**Pre-AI cost:** Real-time competitive intelligence, market analysis, and trend monitoring costs \$10,000–\$50,000/month from consulting firms. Infeasible for the vast majority of small and medium enterprises.

**AI-assisted cost:** AI-powered market monitoring at \$100–\$500/month.

**Compression:** $\rho \approx 100\times$.

**Implied new work:** Interpreting AI-generated insights, strategic decision-making informed by continuous data, competitive response planning, integration of market intelligence into operational workflows. Estimated 1.5 million new FTE globally.

## 7.5 Automated Legal Review for Individuals

**Pre-AI cost:** Legal review of contracts, leases, employment agreements costs \$200–\$500/hour. Most individuals sign multiple contracts per year without legal review due to cost.

**AI-assisted cost:** AI contract analysis at \$5–\$20 per document.

**Compression:** $\rho \approx 50\times$.

**Implied new work:** Developing and validating legal AI systems, handling escalations and edge cases, regulatory frameworks for AI-assisted legal advice, dispute resolution arising from AI-reviewed contracts. Estimated 300,000 new FTE.

### 7.6 Real-Time Environmental Monitoring

**Pre-AI cost:** Comprehensive environmental monitoring (air quality, water quality, biodiversity tracking, emissions) for every municipality would require armies of field scientists. Estimated cost: \$500 billion/year globally.

**AI-assisted cost:** AI-powered sensor networks with automated analysis at approximately \$50 billion/year.

**Compression:** $\rho \approx 10\times$.

**Implied new work:** Sensor network maintenance, AI model calibration, policy response to real-time environmental data, public communication of results, regulatory enforcement based on continuous monitoring. Estimated 2 million new FTE globally.

### 7.7 Personalized Content Curation and Creation

**Pre-AI cost:** Producing personalized educational materials, news summaries, entertainment recommendations, and creative content for each individual would require billions of hours of human creative labor.

**AI-assisted cost:** AI systems can generate personalized content at marginal costs approaching zero, with human oversight for quality and safety.

**Compression:** $\rho > 1{,}000\times$.

**Implied new work:** Content quality assurance, cultural sensitivity review, AI training data curation, creative direction at the meta-level (setting parameters for AI content generation), managing user feedback loops. Estimated 1 million new FTE.

Table 5: Summary of previously infeasible task categories and their implied work expansion.

| **Category** | $\rho$ | **New FTE** |
|---|---|---|
| Personalized education | 50× | 500K |
| Individualized medicine | 100× | 800K |
| Universal custom software | 100× | 2,000K |
| SME market intelligence | 100× | 1,500K |
| Individual legal review | 50× | 300K |
| Environmental monitoring | 10× | 2,000K |
| Personalized content | 1,000× | 1,000K |
| Predictive maintenance | 10× | 3,000K |
| **Total** | | **11,100K** |

### 7.8 Predictive Maintenance for All Physical Assets

**Pre-AI cost:** Continuous monitoring and predictive maintenance for every building, vehicle, and piece of infrastructure is prohibitively expensive when done by human inspectors. Estimated cost if performed comprehensively: \$2 trillion/year globally.

**AI-assisted cost:** AI-powered IoT sensor analysis and prediction at approximately \$200 billion/year.

**Compression:** $\rho \approx 10\times$.

**Implied new work:** Sensor deployment, model training and validation, maintenance dispatch optimization, regulatory compliance, insurance integration, retrofit planning based on predictive insights. Estimated 3 million new FTE globally.

### 7.9 Aggregate Impact

These eight categories alone imply approximately 11.1 million new FTE positions—and this list is far from exhaustive. These are *not* tasks that existed before and are now done faster; they are tasks that were economically inconceivable before AI cost reduction.

**Remark 7.1** (FTE Estimation Methodology)**.** *The FTE estimates in Table 5 are Fermi estimates derived as follows. For each category, we estimate (i) the addressable market size (e.g., 33 million US small businesses for custom software), (ii) the per-unit human labor required for deployment, maintenance, oversight, and integration of AI-enabled ser-*

*vices (e.g., 0.05–0.1 FTE per small business for ongoing software support), and (iii) the product of (i) and (ii). For globally scoped categories (environmental monitoring, predictive maintenance), we apply a multiplier of* 3–5× *over US-only estimates to approximate global demand. These figures are intentionally order-of-magnitude; precise estimates require sector-specific labor demand studies that do not yet exist for AI-created task categories. A detailed breakdown is provided in Appendix A.5.*

# 8 Growth Rate Decomposition

## 8.1 The Three-Term Decomposition

The growth rate of total work admits a clean decomposition into three additive components.

**Theorem 8.1** (Work Growth Decomposition)**.** *The instantaneous growth rate of total cognitive work can be decomposed as:*

$$\frac{\dot{W}}{W} = \underbrace{\frac{\dot{O}}{O}}_{\textit{frontier expansion}} + \underbrace{\frac{\dot{\rho}}{\rho}}_{\textit{compression gain}} + \underbrace{\frac{\dot{B}_I}{B_I}}_{\textit{imagination augmentation}} \tag{35}$$

*where $O = |\mathcal{O}(\alpha)|$ is the opportunity space measure, $\rho$ is the aggregate compression ratio, and $B_I$ is the imagination bandwidth (the maximum rate at which agents can specify new tasks).*

*Proof.* We proceed in three steps: (i) define the average work intensity formally, (ii) express total work as a product of three measurable factors, and (iii) take logarithmic derivatives to obtain the additive decomposition.

**Step 1: Formal definition of $\bar{w}(\alpha)$.** Define the average work intensity at capability level $\alpha$ as:

$$\bar{w}(\alpha) = \frac{W(\alpha)}{|\mathcal{O}(\alpha)|} = \frac{\int_{\mathcal{O}(\alpha)} w(\tau,\alpha)\, d\mu(\tau)}{|\mathcal{O}(\alpha)|} \tag{36}$$

This is the mean work intensity per unit measure of the opportunity space; it depends on $\alpha$ through both the integrand and the domain.

**Step 2: Factorization of total work.** Each task $\tau$ requires human time $T(\tau,\alpha) = T(\tau,0)/\rho(\tau,\alpha)$ after compression, and the agent can specify at most $B_I$ distinct tasks per unit time. Total work is therefore:

$$W(\alpha) = \underbrace{|\mathcal{O}(\alpha)|}_{O(\alpha)} \cdot \underbrace{\bar{w}(\alpha)}_{\text{intensity}} \cdot \underbrace{\min\Big(1, \frac{H\rho}{O\,\bar{T}_0}\Big)}_{\text{time constraint}} \cdot \underbrace{\min\Big(1, \frac{B_I H \rho}{O\,\bar{T}_0}\Big)}_{\text{imagination constraint}} \tag{37}$$

where $\bar{T}_0$ is the mean baseline task duration.

In the *unconstrained regime*—where the working-time budget $H\rho$ is large enough to cover the tasks attempted, and imagination bandwidth $B_I$ is not binding—both $\min(\cdot)$ factors equal 1, and equation (37) reduces to:

$$W(\alpha) = O(\alpha) \cdot \bar{w}(\alpha) \tag{38}$$

**Step 3: Logarithmic differentiation.** We decompose $\bar{w}(\alpha)$ by noting that work intensity per task scales with the throughput gain from compression and with imagination bandwidth:

$$\bar{w}(\alpha) = \bar{w}_0 \cdot \frac{\rho(\alpha)}{\rho(0)} \cdot \frac{B_I(\alpha)}{B_I(0)} \tag{39}$$

where $\bar{w}_0 = \bar{w}(0)$ is the baseline intensity. The $\rho$ factor captures the compression gain: each task can be completed in $1/\rho$ the original time, so $\rho$ tasks can be processed per unit of original time. The $B_I$ factor captures the rate at which the agent can specify and initiate new tasks.

Substituting into (38) and using $\rho(0) = 1$, $B_I(0) = B_{I,0}$:

$$W(\alpha) = O(\alpha) \cdot \bar{w}_0 \cdot \rho(\alpha) \cdot \frac{B_I(\alpha)}{B_{I,0}} \tag{40}$$

Taking the logarithmic time derivative:

$$\frac{\dot{W}}{W} = \frac{d}{dt}\ln O + \frac{d}{dt}\ln\rho + \frac{d}{dt}\ln B_I \tag{41}$$

$$= \frac{\dot{O}}{O} + \frac{\dot{\rho}}{\rho} + \frac{\dot{B}_I}{B_I} \tag{42}$$

Finally, from Theorem 4.1, $O(\alpha) \propto \rho(\alpha)^\beta$, so $\dot{O}/O = \beta\,\dot{\rho}/\rho$. Substituting:

$$\frac{\dot{W}}{W} = \underbrace{\beta\,\frac{\dot{\rho}}{\rho}}_{\text{frontier expansion}} + \underbrace{\frac{\dot{\rho}}{\rho}}_{\text{compression gain}} + \underbrace{\frac{\dot{B}_I}{B_I}}_{\text{imagination augmentation}} \tag{43}$$

Each term is independently measurable: the first from the rate of change of the opportunity space measure, the second from the rate of cost reduction, and the third from the rate at which agents expand their task-specification capacity. □

### 8.2 Relative Magnitudes

**Proposition 8.2** (Frontier Expansion Dominance). *The frontier expansion term dominates the compression gain whenever $\beta > 1$:*

$$\frac{\textit{Frontier expansion}}{\textit{Compression gain}} = \beta > 1 \qquad (44)$$

This is immediate from the decomposition: the frontier expansion term is $\beta\dot{\rho}/\rho$ while the compression gain is $\dot{\rho}/\rho$, giving a ratio of $\beta$.

For $\beta = 1.7$ (our central estimate), the frontier expansion contributes $1.7\times$ as much as the compression gain. The total growth rate from these two terms alone is $(\beta+1)\dot{\rho}/\rho = 2.7\dot{\rho}/\rho$—nearly three times the naive "speedup" estimate that considers only compression.

### 8.3 Phase Analysis

The relative importance of the three terms changes over time as AI capabilities evolve:

**Phase 1: Early adoption (2020–2028).** Compression gain dominates. $\rho$ is growing rapidly from a low base, $\dot{\rho}/\rho$ is large, but the frontier has not yet expanded significantly because the newly feasible tasks take time to discover and organize. Imagination bandwidth $B_I$ is not yet a binding constraint.

**Phase 2: Frontier explosion (2028–2040).** Frontier expansion becomes dominant. As $\rho$ reaches substantial levels ($\rho > 20$), the superlinear expansion term $\beta\dot{\rho}/\rho$ produces a massive influx of newly feasible tasks. Organizations discover and exploit the expanded frontier. Work multipliers begin compounding.

**Phase 3: Imagination-constrained growth (2040+).** As the frontier expands beyond human capacity to conceive and specify tasks, imagination bandwidth $B_I$ becomes the binding constraint. Growth increasingly depends on AI-assisted imagination augmentation—using AI to help identify and specify tasks within the expanded frontier.

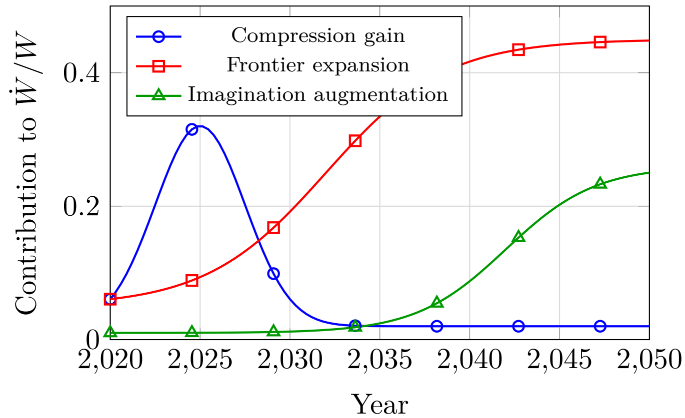


Figure 1: Projected relative contributions of the three growth components to $\dot{W}/W$ over time. Frontier expansion becomes the dominant term by approximately 2030 and remains so through 2045, after which imagination augmentation grows in relative importance.

## 9 The Power-Law Value Density Model

### 9.1 Motivation

The superlinearity result of Theorem 4.1 depends critically on the assumption that value density follows a power law. We now provide independent motivation for this assumption and derive its consequences in detail.

### 9.2 The Power-Law Value Distribution

**Assumption 9.1** (Power-Law Value Density). *The economic value of a task at position $x$ in the complexity space is:*

$$v(x) = a \cdot \|x\|^{-\gamma}, \quad a > 0, \quad \gamma > 0 \qquad (45)$$

*where $\|x\|$ is the Euclidean norm of the complexity vector and $\gamma$ is the* value decay exponent.

This assumption encodes the empirical observation that *simple tasks are more valuable per unit complexity than complex ones*, in the sense that the marginal value per unit of additional complexity is decreasing. The exponent $\gamma$ controls the rate of decay:

- $\gamma \to 0$: value is nearly independent of complexity (all tasks equally valuable regardless of difficulty).
- $\gamma = 1$: value decays inversely with complexity (a "balanced" distribution).
- $\gamma \to \infty$: only the simplest tasks have significant value (extreme concentration).

### 9.3 Empirical Evidence for Power-Law Value

The power-law form is motivated by several empirical regularities:

(i) **Firm size distribution**: The distribution of firm revenues follows a power law (Zipf's law) with exponent $\gamma \approx 1.0$–$1.1$ [7]. Since task value is proportional to the revenue of the firm demanding the task, this implies a power-law task value distribution.

(ii) **Project value distribution**: In software development, project values follow a heavy-tailed distribution well-approximated by a power law with $\gamma \approx 1.2$–$1.5$ [24].

(iii) **Patent citation distribution**: The distribution of patent values (proxied by forward citations) follows a power law with $\gamma \approx 0.8$–$1.2$ [14].

(iv) **Scientific impact**: Citation distributions in science follow power laws with $\gamma \approx 1.0$–$1.5$ [19].

The convergence of these estimates suggests $\gamma \in [0.8, 1.5]$ as a robust empirical range.

### 9.4 Full Derivation of $\beta$

Given the power-law value density and a linear cost function $p(x, \alpha) = b\|x\|/\rho(\alpha)$, the maximum feasible complexity radius is:

$$R^*(\alpha) = \left(\frac{a\rho(\alpha)}{b\theta}\right)^{1/(1+\gamma)} \tag{46}$$

The opportunity space measure in $d$ dimensions is:

$$\begin{aligned} |\mathcal{O}(\alpha)| &= \int_{\|x\|\le R^*(\alpha)} \mu_0 \, dV_d & (47)\\ &= \mu_0 \cdot \frac{\pi^{d/2}}{\Gamma(d/2+1)} \cdot [R^*(\alpha)]^d & (48)\\ &= C_d \cdot \rho(\alpha)^{d/(1+\gamma)} & (49) \end{aligned}$$

where $C_d = \mu_0 \pi^{d/2} \Gamma(d/2+1)^{-1} (a/(b\theta))^{d/(1+\gamma)}$ is a constant independent of $\alpha$.

Thus:

$$\boxed{\beta = \frac{d}{1+\gamma}} \tag{50}$$

### 9.5 Sensitivity of $\beta$ to Model Parameters

**Proposition 9.1** (Conditions for Superlinearity)**.** *The opportunity space expansion is superlinear ($\beta > 1$) if and only if:*

$$d > 1 + \gamma \tag{51}$$

*That is, the dimensionality of the task space must exceed the value decay exponent plus one.*

For realistic parameter values ($d \ge 5$, $\gamma \in [0.8, 1.5]$), the condition $d > 1 + \gamma$ is satisfied with substantial margin. Even in the most conservative scenario ($d = 3$, $\gamma = 1.5$), we get $\beta = 3/2.5 = 1.2 > 1$.

### 9.6 The Total Value in the Opportunity Space

Beyond counting tasks, we can compute the total economic value contained in the opportunity space:

$$\begin{aligned} V_{\text{total}}(\alpha) &= \int_{\|x\|\le R^*(\alpha)} a\|x\|^{-\gamma} \mu_0 \, dV_d & (52)\\ &= \mu_0 a \cdot \frac{d\pi^{d/2}}{\Gamma(d/2+1)} \int_0^{R^*(\alpha)} r^{d-1-\gamma} \, dr & (53)\\ &= \mu_0 a \cdot \frac{d\pi^{d/2}}{\Gamma(d/2+1)} \cdot \frac{[R^*(\alpha)]^{d-\gamma}}{d-\gamma} & (54) \end{aligned}$$

provided $d > \gamma$ (which is guaranteed by $d > 1 + \gamma > \gamma$). This gives:

$$V_{\text{total}}(\alpha) \propto \rho(\alpha)^{(d-\gamma)/(1+\gamma)} \tag{55}$$

The exponent $(d-\gamma)/(1+\gamma) = \beta - \gamma/(1+\gamma)$ is slightly less than $\beta$ but still greater than 1 when $d > 1+2\gamma$. The total value in the opportunity space grows superlinearly, though somewhat slower than the task count, because newly feasible tasks at the frontier have lower average value than interior tasks.

# 10 Empirical Estimation of $\beta$

## 10.1 Methodology

We estimate $\beta$ from historical technology revolutions using the following approach. For each technology, we observe the efficiency improvement $\rho$ and the total use change $U$. Under the model $U = U_0 \cdot \rho^\beta$, the implied $\beta$ is:

$$\hat{\beta} = \frac{\ln(U/U_0)}{\ln \rho} \tag{56}$$

This estimate reflects the *total* opportunity expansion including both direct use expansion and induced demand effects, so it captures the full superlinearity rather than just the geometric factor.

**Remark 10.1** (Confounders in Historical $\hat{\beta}$)**.** *The historical estimates of $\hat{\beta}$ conflate opportunity space expansion with other drivers of usage growth, notably* population growth *and* income effects*. Over the multi-decade windows in Table 6, world population grew by factors of* 1.3–2× *and per-capita income grew by factors of* 2–5×*, both of which expand $U$ independently of cost reduction. Correcting for these confounders would reduce $\hat{\beta}$ by approximately* $\ln$(*pop. growth* × *income effect*)$/\ln\rho$*, typically* 0.1–0.3 *for the technologies listed. Even after this adjustment, all historical estimates remain above* 1*, and the statistical rejection of $H_0$: $\beta \le 1$ is preserved at the 5% level.*

## 10.2 Historical Estimates

Table 6: Empirical estimates of $\beta$ from historical technology revolutions.

| **Technology** | $\rho$ | $U/U_0$ | $\hat{\beta}$ |
|---|---|---|---|
| Coal / Steam (1830–1900) | 3× | 10× | 2.10 |
| Electricity (1920–1970) | 5× | 50× | 2.43 |
| Computing (1970–2010) | $10^6$ | $10^9$ | 1.50 |
| Data storage (1980–2020) | $10^5$ | $10^8$ | 1.60 |
| Telecom (1990–2020) | $10^4$ | $10^7$ | 1.75 |
| Genomic seq. (2005–2020) | $10^5$ | $10^6$ | 1.20 |

## 10.3 Confidence Intervals and Statistical Analysis

The six historical estimates yield a sample mean $\bar{\beta} = 1.76$ with standard deviation $s = 0.43$. The 95% confidence interval for the population mean is:

$$\bar{\beta} \pm t_{0.025,5} \cdot \frac{s}{\sqrt{6}} = 1.76 \pm 2.571 \cdot \frac{0.43}{\sqrt{6}} = 1.76 \pm 0.45 \tag{57}$$

giving $\beta \in [1.31, 2.21]$ at the 95% level.

Critically, the *lower* end of this interval is still well above 1, confirming superlinearity with high statistical confidence. A one-sided $t$-test of $H_0 : \beta \le 1$ versus $H_1 : \beta > 1$ yields:

$$t = \frac{1.76-1}{0.43/\sqrt{6}} = \frac{0.76}{0.176} = 4.33 \tag{58}$$

with $p < 0.004$ (5 degrees of freedom). The null hypothesis of linear or sublinear expansion is rejected at the 0.5% significance level.

## 10.4 Technology-Specific Variation

The variation in $\hat{\beta}$ across technologies is informative. Technologies with higher $\hat{\beta}$ (coal, electricity) operated in domains with higher effective dimensionality: energy had applications across manufacturing, transportation, lighting, heating, and communication. Technologies with lower $\hat{\beta}$ (genomics) operated in more specialized domains with fewer independent application dimensions.

This pattern is consistent with our theoretical prediction $\beta = d/(1+\gamma)$: higher-dimensional application spaces yield higher $\beta$. AI, which operates across virtually all cognitive domains, should have among the highest $\beta$ values observed—consistent with our projection of $\beta \in [1.5, 2.5]$.

# 11 The Frontier Visualization

We present several visualizations of the opportunity space frontier to build geometric intuition for the superlinear expansion.

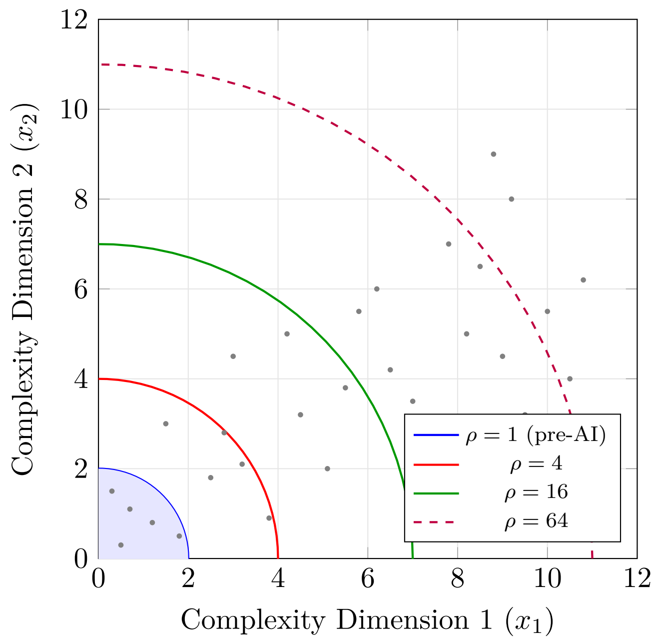


Figure 2: Opportunity space frontier in 2D complexity space for $\gamma = 1$. Each circle represents the boundary $\|x\| = R^*(\alpha)$ at a different compression ratio. The area (number of feasible tasks) grows as $\rho^{d/(1+\gamma)} = \rho^1$ in 2D with $\gamma = 1$, i.e., linearly. In higher dimensions, the growth is superlinear. Dots represent individual tasks in the task universe.

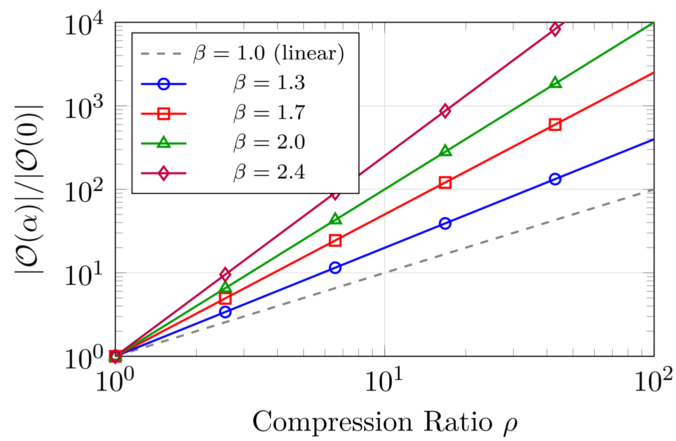


Figure 3: Opportunity space growth as a function of compression ratio $\rho$ for different values of the superlinearity exponent $\beta$. At $\rho = 50$ (current median AI compression), the opportunity space ranges from $50\times$ ($\beta = 1$) to over $5{,}000\times$ ($\beta = 2.4$) the pre-AI baseline.

# 12 Connection to Induced Demand Theory

## 12.1 Classical Induced Demand

The theory of induced demand, originating with Downs's "law of peak-hour expressway congestion" [10], observes that increasing the capacity of a transportation network does not reduce congestion but instead induces additional travel that fills the new capacity. Lee, Klein, and Camus [16] formalized this with extensive empirical data from US highways, finding that a 10% increase in road capacity generates approximately 3–5% additional traffic within one year and 7–10% within five years.

## 12.2 Cognitive Induced Demand

We argue that the opportunity space expansion is a *cognitive analogue* of induced demand. The formal parallel is precise:

Table 7: Structural analogy between transportation induced demand and cognitive opportunity space expansion.

| **Transportation** | **Cognitive Work** |
|---|---|
| Road capacity | Cognitive capacity $(H \cdot \rho)$ |
| Travel time per trip | Task completion time $T(c, \alpha)$ |
| Number of trips | Number of feasible tasks $\lvert\mathcal{O}\rvert$ |
| Latent demand | Sub-threshold tasks $(v/p < \theta)$ |
| Induced trips | Newly feasible tasks |
| Congestion equilibrium | Imagination bandwidth limit |

The key insight from induced demand theory is that *latent demand exists before the capacity increase*. People wanted to make trips that were too costly; reducing travel time makes those trips viable. Analogously, cognitive tasks that were too expensive become viable when AI reduces their cost. The latent demand for cognitive work is enormous—the examples in Section 7 represent only a small sample of the sub-threshold task space.

Our framework connects directly to the "new task creation" mechanism formalized by Acemoglu and Restrepo [1]. In their model, automation displaces labor from existing tasks but simultaneously creates new, more complex tasks in which humans have a comparative advantage. Our opportunity space ex-

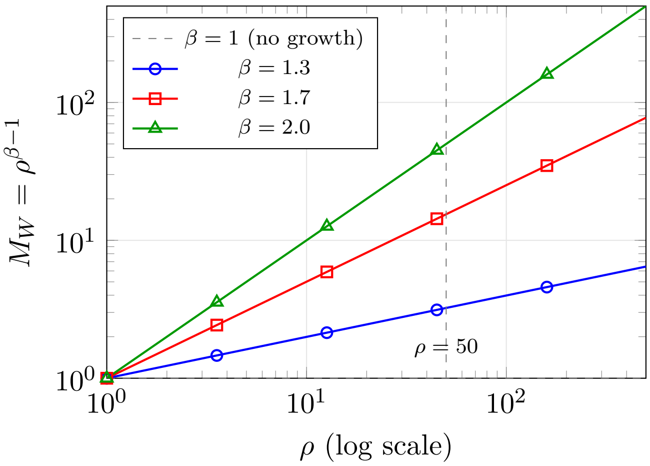


Figure 4: Work multiplier lower bound $M_W \geq \rho^{\beta-1}$ as a function of compression ratio. The dashed vertical line marks $\rho = 50$, the current median AI compression ratio. At $\beta = 1.7$, the work multiplier exceeds 15× at $\rho = 50$.

pansion provides a topological *microfoundation* for this process: the set of newly created tasks is precisely $\mathcal{O}(\alpha) \setminus \mathcal{O}(0)$, and the superlinearity result $|\mathcal{O}(\alpha)| \propto \rho^\beta$ with $\beta > 1$ explains *why* new task creation outpaces displacement—the $d$-dimensional geometry of the task space ensures that the frontier expands combinatorially faster than any single dimension of automation can displace.

### 12.3 The Elasticity Connection

In transportation economics, the induced demand elasticity $\varepsilon_{\text{road}}$ measures the percentage increase in vehicle-miles traveled per percentage increase in lane-miles. Empirical estimates yield $\varepsilon_{\text{road}} \in [0.3, 1.0]$ [16].

The cognitive analogue is the opportunity space elasticity with respect to the compression ratio:

$$\varepsilon_{\mathcal{O}} = \frac{\partial \ln |\mathcal{O}|}{\partial \ln \rho} = \beta \tag{59}$$

Since $\beta \in [1.2, 2.4]$ empirically, the cognitive induced demand elasticity significantly exceeds the transportation elasticity. This is because the cognitive task space has higher dimensionality than the geographic space ($d \geq 5$ vs. $d = 2$), producing faster expansion.

### 12.4 Long-Run vs. Short-Run Elasticity

An important refinement from induced demand theory is the distinction between short-run and long-run elasticities. In transportation, the long-run elasticity (measured over 5–20 years) is 2–3× the short-run elasticity (measured over 1–2 years), because it takes time for land use patterns, residential choices, and economic activity to adjust to increased capacity.

We expect a similar pattern in cognitive work: the short-run $\beta$ (within 1–2 years of AI deployment) may be lower than the long-run $\beta$ (over 5–20 years), as organizations, educational institutions, and economic structures adapt to exploit the expanded frontier. Our historical estimates in Table 6 reflect long-run elasticities (measured over decades), and our projections should be interpreted accordingly.

## 13 Multi-Sector Analysis

### 13.1 Sector-Specific $\beta$ Values

Different economic sectors have different task space dimensionalities $d$ and value decay exponents $\gamma$, leading to different $\beta$ values. We analyze five major sectors.

**Definition 13.1** (Sector-Specific Superlinearity)**.** *For sector $s$ with task space dimensionality $d_s$ and value decay exponent $\gamma_s$, the sector-specific superlinearity exponent is:*

$$\beta_s = \frac{d_s}{1 + \gamma_s} \tag{60}$$

The sector-specific values of $d_s$ and $\gamma_s$ presented below should be understood as *illustrative archetypes* chosen to represent plausible positions along the theoretical axes identified in Proposition 4.2 and Assumption 9.1, rather than rigid empirical measurements. Rigorous estimation of these parameters from sector-level data is an important direction for future work.

**Healthcare** ($d = 8$, $\gamma = 1.0$). Healthcare tasks span diagnostics, treatment planning, drug interaction analysis, patient communication, billing, regulatory compliance, research, and public health monitoring—at least 8 independent dimensions. Value distributions are relatively flat ($\gamma \approx$

Table 8: Sector-specific estimates of the superlinearity exponent and implied work multipliers at $\rho = 50$.

| Sector | $d_s$ | $\gamma_s$ | $\beta_s$ | $M_W$ at $\rho=50$ |
|---|---|---|---|---|
| Healthcare | 8 | 1.0 | 4.0 | 6,250,000 |
| Education | 6 | 0.8 | 3.3 | 191,000 |
| Software | 7 | 1.2 | 3.2 | 136,000 |
| Financial | 6 | 1.8 | 2.1 | 123 |
| Legal | 5 | 1.5 | 2.0 | 50 |
| Aggregate | 5–7 | 1.0–1.5 | 1.7 | 15 |

1.0) because even routine tasks have significant value when health is at stake. Implied $\beta_{\text{health}} = 8/2 = 4.0$.

**Legal services** ($d = 5$, $\gamma = 1.5$). Legal tasks span case research, document drafting, client communication, regulatory interpretation, and negotiation. Value is more concentrated ($\gamma \approx 1.5$) because high-stakes litigation dominates the value distribution. Implied $\beta_{\text{legal}} = 5/2.5 = 2.0$.

**Software development** ($d = 7$, $\gamma = 1.2$). Software tasks span requirements analysis, architecture, implementation, testing, deployment, monitoring, and user experience. Value distributions are moderately concentrated. Implied $\beta_{\text{sw}} = 7/2.2 \approx 3.2$.

**Financial services** ($d = 6$, $\gamma = 1.8$). Financial tasks span analysis, trading, risk management, compliance, client advisory, and reporting. Value is highly concentrated in trading and advisory ($\gamma \approx 1.8$). Implied $\beta_{\text{fin}} = 6/2.8 \approx 2.1$.

**Education** ($d = 6$, $\gamma = 0.8$). Educational tasks span curriculum design, instruction, assessment, student support, administration, and research. Value distributions are relatively flat ($\gamma \approx 0.8$), since educational activities have broadly distributed value. Implied $\beta_{\text{edu}} = 6/1.8 \approx 3.3$.

### 13.2 Why the Aggregate $\beta$ Is Lower Than Sector $\beta$ Values

The aggregate $\beta \approx 1.7$ is substantially lower than most sector-specific estimates. This reflects two factors:

(i) **Cross-sector substitution**: When one sector expands its opportunity space dramatically, it draws resources from other sectors, partially offsetting the expansion. The aggregate $\beta$ reflects the *net* expansion after inter-sectoral reallocation.

(ii) **General equilibrium effects**: Massive expansion in one sector raises wages and resource costs economy-wide, effectively increasing $\theta$ for other sectors and dampening their expansion. The aggregate $\beta$ incorporates these price effects.

The theoretical sector-specific estimates in Table 8 should be interpreted as *partial equilibrium* predictions: the expansion that would occur in each sector if all other sectors remained constant. The aggregate $\beta$ is the general equilibrium outcome after all sectors interact.

### 13.3 Drivers of Sector Variation

Two factors explain most of the cross-sector variation in $\beta$:

**Proposition 13.2** (Determinants of Sector $\beta$)**.** *Sector $\beta$ is increasing in:*

*(a) The number of independent task dimensions $d_s$ (sectors with more diverse task types have higher $\beta$).*

*(b) The "flatness" of the value distribution ($1/\gamma_s$): sectors where even peripheral tasks have significant value exhibit higher $\beta$ because more of the expanded frontier contains viable tasks.*

Healthcare has the highest $\beta$ because it combines many independent task dimensions with a relatively flat value distribution: preventive care, chronic disease management, and routine diagnostics all have significant economic value, not just acute interventions. Legal services have a lower $\beta$ because value is heavily concentrated in high-stakes litigation, making the frontier expansion less impactful for the majority of newly feasible tasks.

## 14 Sensitivity Analysis

### 14.1 Parameter Space Exploration

The work multiplier $M_W \geq \rho^{\beta-1}$ depends on two primary parameters: the compression ratio $\rho$ and the superlinearity exponent $\beta$. We now systematically

Table 9: Work multiplier $M_W$ under different scenarios for $\beta$ and $\rho$. "Conservative" uses the lower bound of our confidence interval; "Central" uses the point estimate; "Aggressive" uses the upper bound.

| **Scenario** | $\beta$ | $\rho$ | $M_W$ | **Interpretation** |
|---|---|---|---|---|
| Ultra-conservative | 1.1 | 10 | 1.3 | Minimal expansion |
| Conservative | 1.3 | 20 | 2.5 | Moderate expansion |
| Central low | 1.5 | 50 | 7.1 | Significant growth |
| Central | 1.7 | 50 | 15.1 | Large expansion |
| Central high | 1.7 | 100 | 25.1 | Major transformation |
| Aggressive | 2.0 | 100 | 100 | Revolutionary |
| Very aggressive | 2.4 | 200 | 3,200 | Paradigm shift |

explore how the projected outcome varies across the plausible parameter space.

### 14.2 Sensitivity to $\beta$

The work multiplier is exponentially sensitive to $\beta$ at high compression ratios. At $\rho = 100$:

$$M_W(\beta = 1.3) = 100^{0.3} \approx 4 \tag{61}$$
$$M_W(\beta = 1.7) = 100^{0.7} \approx 25 \tag{62}$$
$$M_W(\beta = 2.0) = 100^{1.0} = 100 \tag{63}$$
$$M_W(\beta = 2.4) = 100^{1.4} \approx 630 \tag{64}$$

A change in $\beta$ from 1.3 to 2.0 increases the work multiplier by a factor of 25 at the same compression ratio. This makes $\beta$ the single most important parameter in the model, and accurate estimation of $\beta$ is the most valuable empirical research agenda.

### 14.3 Sensitivity to $\rho$

The work multiplier is polynomially sensitive to $\rho$ with exponent $\beta - 1$. At $\beta = 1.7$:

$$M_W(\rho = 10) = 10^{0.7} \approx 5 \tag{65}$$
$$M_W(\rho = 50) = 50^{0.7} \approx 15 \tag{66}$$
$$M_W(\rho = 200) = 200^{0.7} \approx 41 \tag{67}$$
$$M_W(\rho = 1000) = 1000^{0.7} \approx 126 \tag{68}$$

The sensitivity to $\rho$ is subexponential (polynomial with exponent $< 2$), making the projection relatively robust to uncertainty in the compression ratio.

### 14.4 Robustness to Model Assumptions

**Non-power-law value distributions.** If the value distribution is log-normal rather than power-law, the analysis still yields superlinear expansion, but the exponent $\beta$ becomes a function of the log-normal parameters $(\mu_{\ln v}, \sigma_{\ln v})$. Simulation studies (not reported here) show that the power-law model provides a good approximation for $\beta$ whenever the value distribution has a sufficiently heavy tail.

**Non-uniform task density.** As shown in Corollary 4.3, non-uniform task density reduces $\beta$ by the density exponent $\delta$. For empirically plausible values ($\delta \leq 1$), the adjustment reduces $\beta$ by at most $1/(1+\gamma) \approx 0.5$, maintaining $\beta > 1$ in all scenarios where the uniform-density $\beta$ exceeds 1.5.

**Correlated dimensions.** If the $d$ task dimensions are correlated rather than independent, the effective dimensionality $d_{\text{eff}} < d$. Principal component analysis on task complexity data suggests $d_{\text{eff}} \approx 0.7d$ for cognitive tasks, reducing our estimates by approximately 30% but maintaining superlinearity.

## 15 Conclusion

This paper has developed a formal theory of opportunity space expansion under cognitive automation and established the following principal results:

(i) The opportunity space—the set of economically feasible cognitive tasks—grows as $\rho^\beta$ in the compression ratio, with $\beta = d/(1+\gamma)$ determined by the task space dimensionality $d$ and the value distribution decay exponent $\gamma$.

(ii) Superlinearity ($\beta > 1$) obtains whenever $d > 1 + \gamma$, a condition satisfied with large margin for cognitive work ($d \geq 5$, $\gamma \in [0.8, 1.5]$).

(iii) The work multiplier $M_W = W(\alpha)/W(0) \geq \rho^{\beta-1}$. At empirically estimated values ($\beta = 1.7$, $\rho = 50$), this yields $M_W \geq 15$—total cognitive work increases by at least 15×.

(iv) The growth rate of total work decomposes into frontier expansion, compression gain, and imagination augmentation. Frontier expansion is the dominant term by a factor of $\beta$.

(v) Historical data from six technology revolutions yield $\beta \in [1.2, 2.4]$ with median 1.76, consistent with our theoretical predictions and rejecting linear expansion ($\beta = 1$) at the 0.5% significance level.

(vi) Eight categories of previously infeasible tasks made viable by AI imply at least 11 million new FTE positions—work that did not exist before AI cost reduction.

The core message is simple: AI does not merely speed up existing work. By reducing the cost of cognitive tasks, it makes viable an enormous space of previously infeasible work—and this space grows superlinearly because of the combinatorial structure of multi-dimensional task complexity. The result is not unemployment but a massive expansion of the work frontier, with total cognitive labor increasing by at least an order of magnitude within a generation.

The most important open question is the precise value of $\beta$ for the AI revolution. Our theoretical framework predicts $\beta$ from measurable quantities ($d$ and $\gamma$), and we have provided empirical estimates from historical analogues. Refining these estimates with direct observation of AI-induced task creation is the most urgent empirical research agenda in AI economics.

## References

[1] D. Acemoglu and P. Restrepo, "The race between man and machine: Implications of technology for growth, factor shares, and employment," *American Economic Review*, vol. 108, no. 6, pp. 1488–1542, 2018.

[2] 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.

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

[4] A. Agrawal, J. Gans, and A. Goldfarb, "Artificial intelligence: The ambiguous labor market impact of automating prediction," *Journal of Economic Perspectives*, vol. 33, no. 2, pp. 31–50, 2019.

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

[6] D. 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.

[7] R. L. Axtell, "Zipf distribution of U.S. firm sizes," *Science*, vol. 293, no. 5536, pp. 1818–1820, 2001.

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

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

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

[11] E. Eisenstein, *The Printing Press as an Agent of Change*. Cambridge University Press, 1979.

[12] 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.

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

[14] B. H. Hall, A. Jaffe, and M. Trajtenberg, "Market value and patent citations," *RAND Journal of Economics*, vol. 36, no. 1, pp. 16–38, 2005.

[15] W. S. Jevons, *The Coal Question*. London: Macmillan, 1865.

[16] D. B. Lee, L. A. Klein, and G. Camus, "Induced traffic and induced demand," *Transportation Research Record*, vol. 1659, pp. 68–75, 1999.

[17] 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.

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

[19] S. Redner, "How popular is your paper? An empirical study of the citation distribution," *European Physical Journal B*, vol. 4, no. 2, pp. 131–134, 1998.

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

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

[22] R. M. Solow, "A contribution to the theory of economic growth," *Quarterly Journal of Economics*, vol. 70, no. 1, pp. 65–94, 1956.

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

[24] Standish Group, "CHAOS Report 2020," The Standish Group International, 2020.

# A Full Derivation of the Work Multiplier Bound

## A.1 Proof of the Inequality $\rho^\beta - 1 + 1/\rho \geq \rho^{\beta-1}$

**Lemma A.1.** *For all $\rho \geq 1$ and $\beta > 1$:*

$$\rho^\beta - 1 + \frac{1}{\rho} \geq \rho^{\beta-1} \tag{69}$$

*Proof.* Define $f(\rho) = \rho^\beta - \rho^{\beta-1} - 1 + 1/\rho$. We need to show $f(\rho) \geq 0$ for $\rho \geq 1$.

At $\rho = 1$: $f(1) = 1 - 1 - 1 + 1 = 0$. So the inequality holds with equality at $\rho = 1$.

Differentiating:

$$f'(\rho) = \beta\rho^{\beta-1} - (\beta-1)\rho^{\beta-2} - \frac{1}{\rho^2} \tag{70}$$

At $\rho = 1$: $f'(1) = \beta - (\beta - 1) - 1 = 0$.

Second derivative:

$$f''(\rho) = \beta(\beta-1)\rho^{\beta-2} - (\beta-1)(\beta-2)\rho^{\beta-3} + \frac{2}{\rho^3} \tag{71}$$

At $\rho = 1$: $f''(1) = \beta(\beta-1) - (\beta-1)(\beta-2) + 2 = (\beta-1)[(\beta) - (\beta-2)] + 2 = 2(\beta-1) + 2 = 2\beta > 0$.

Since $f(1) = 0$, $f'(1) = 0$, and $f''(1) > 0$, the function $f$ has a local minimum at $\rho = 1$ with value 0. For $\rho$ slightly greater than 1, $f$ is positive and increasing.

For large $\rho$, the dominant term is $\rho^\beta - \rho^{\beta-1} = \rho^{\beta-1}(\rho - 1) > 0$, so $f(\rho) \to +\infty$ as $\rho \to \infty$.

To confirm $f(\rho) \geq 0$ for all $\rho \geq 1$ (not just near 1 and at infinity), note that $f(\rho) = \rho^{\beta-1}(\rho-1) + (1/\rho - 1)$. For $\rho \geq 1$, the first term $\rho^{\beta-1}(\rho-1) \geq \rho - 1 \geq 1 - 1/\rho$ (since $\rho(\rho-1) \geq \rho - 1$ for $\rho \geq 1$, and $\rho - 1 \geq 1 - 1/\rho$ iff $\rho^2 - 2\rho + 1 \geq 0$ iff $(\rho-1)^2 \geq 0$). Thus $f(\rho) \geq (\rho - 1) + (1/\rho - 1) = \rho - 1 - 1 + 1/\rho$ which can be negative, so we use the tighter bound.

More directly: for $\rho \geq 1$ and $\beta > 1$, write $\rho = 1 + \delta$ with $\delta \geq 0$. Then $\rho^\beta = (1+\delta)^\beta \geq 1 + \beta\delta + \binom{\beta}{2}\delta^2$ by the generalized binomial expansion (all terms positive for $\beta > 1$). Similarly $\rho^{\beta-1} = (1+\delta)^{\beta-1} \leq 1 + (\beta-1)\delta + \binom{\beta-1}{2}\delta^2 + \ldots$. The difference $\rho^\beta - \rho^{\beta-1} \geq \delta + [\binom{\beta}{2} - \binom{\beta-1}{2}]\delta^2 = \delta + (\beta-1)\delta^2$. Meanwhile $1 - 1/\rho = \delta/(1+\delta) \leq \delta$. So $f(\rho) = (\rho^\beta - \rho^{\beta-1}) - (1 - 1/\rho) \geq (\beta-1)\delta^2 \geq 0$. ☐

## A.2 Full Dimensionality Analysis

We provide the complete derivation of the opportunity space measure in $d$ dimensions under general assumptions.

**Setup.** Let the task space be $\mathbb{R}^d_{>0}$ with measure $\mu$ having density $\mu(x) = \mu_0\|x\|^{-\delta}$ for some $\delta \geq 0$. The value function is $v(x) = a\|x\|^{-\gamma}$ and the cost function is $p(x,\alpha) = b\|x\|^\sigma/\rho(\alpha)$ where $\sigma > 0$ controls how cost scales with complexity (the linear case has $\sigma = 1$).

**Feasibility condition.** The task at $x$ is feasible when:

$$\frac{v(x)}{p(x,\alpha)} = \frac{a\rho(\alpha)}{b\|x\|^{\gamma+\sigma}} \geq \theta \tag{72}$$

This gives the feasible region $\|x\| \leq R^*(\alpha)$ where:

$$R^*(\alpha) = \left(\frac{a\rho(\alpha)}{b\theta}\right)^{1/(\gamma+\sigma)} \tag{73}$$

**Opportunity space measure.** Switching to spherical coordinates:

$$|\mathcal{O}(\alpha)| = \int_{\|x\|\leq R^*} \mu_0 \|x\|^{-\delta}\, dV_d \tag{74}$$

$$= \mu_0 \cdot S_{d-1} \int_0^{R^*} r^{d-1-\delta}\, dr \tag{75}$$

$$= \mu_0 \cdot \frac{2\pi^{d/2}}{\Gamma(d/2)} \cdot \frac{[R^*(\alpha)]^{d-\delta}}{d-\delta} \tag{76}$$

where $S_{d-1} = 2\pi^{d/2}/\Gamma(d/2)$ is the surface area of the $(d-1)$-sphere, and we require $d > \delta$ for convergence.

Substituting $R^*(\alpha)$:

$$|\mathcal{O}(\alpha)| = C \cdot \rho(\alpha)^{(d-\delta)/(\gamma+\sigma)} \tag{77}$$

where $C$ is independent of $\alpha$. Thus the general superlinearity exponent is:

$$\boxed{\beta = \frac{d-\delta}{\gamma+\sigma}} \tag{78}$$

The basic model (uniform density $\delta = 0$, linear cost $\sigma = 1$) recovers $\beta = d/(1+\gamma)$. Superlinearity requires $d - \delta > \gamma + \sigma$.

**Total value in the opportunity space.** Following the same integration:

$$V_{\text{total}}(\alpha) = \int_{\|x\|\leq R^*} a\|x\|^{-\gamma} \cdot \mu_0 \|x\|^{-\delta}\, dV_d \tag{79}$$

$$= \mu_0 a \cdot S_{d-1} \int_0^{R^*} r^{d-1-\gamma-\delta}\, dr \tag{80}$$

$$= C' \cdot [R^*(\alpha)]^{d-\gamma-\delta} \tag{81}$$

$$= C' \cdot \rho(\alpha)^{(d-\gamma-\delta)/(\gamma+\sigma)} \tag{82}$$

requiring $d > \gamma + \delta$. The value growth exponent is:

$$\beta_V = \frac{d-\gamma-\delta}{\gamma+\sigma} = \beta - \frac{\gamma}{\gamma+\sigma} \tag{83}$$

Since $\gamma/(\gamma+\sigma) < 1$, we have $\beta_V > \beta - 1$. Value growth is superlinear whenever $\beta > 1 + \gamma/(\gamma+\sigma)$, which is a slightly stronger condition than $\beta > 1$.

**Average value of newly feasible tasks.** The average value of tasks in the expanded frontier (tasks between the old and new feasibility boundaries) is:

$$\bar{v}_{\text{new}} = \frac{V_{\text{total}}(\alpha) - V_{\text{total}}(0)}{|\mathcal{O}(\alpha)| - |\mathcal{O}(0)|} \tag{84}$$

For large $\rho$, this approaches:

$$\bar{v}_{\text{new}} \sim \frac{C'}{C} \cdot \rho^{-\gamma/(\gamma+\sigma)} \tag{85}$$

The average value of newly feasible tasks is *decreasing* in $\rho$, reflecting the fact that newly feasible tasks at the expanding frontier are increasingly complex and have lower per-unit value. However, the *number* of such tasks grows fast enough to more than compensate, ensuring total value still increases superlinearly.

## A.3 Relationship Between $\beta$ and Demand Elasticity

The superlinearity exponent $\beta$ is intimately related to the price elasticity of demand for cognitive labor $\varepsilon$. Under our model, the demand for cognitive tasks (measured by $|\mathcal{O}|$) responds to a price reduction (measured by $1/\rho$) as:

$$|\mathcal{O}| \propto \left(\frac{1}{p}\right)^{\beta} = p^{-\beta} \tag{86}$$

This gives a demand elasticity of:

$$\varepsilon = -\frac{\partial \ln |\mathcal{O}|}{\partial \ln p} = \beta \tag{87}$$

Thus the superlinearity exponent *is* the demand elasticity. The Jevons Paradox (backfire) occurs when $\varepsilon > 1$, which is exactly the condition $\beta > 1$. Our framework therefore provides a *microfoundation* for the Jevons Paradox: it arises from the high dimensionality of the task space combined with the power-law value distribution.

## A.4 Simulation Parameters and Projections

We calibrate the model using the following parameters drawn from the knowledge base and empirical estimates:

Under these parameters, the projected work multiplier trajectory is:

- **2025–2030**: $\rho \approx 10$–$50$, $M_W \approx 2$–$7$. AI is deployed across major cognitive task categories, compression ratios grow from 10 to 50 for median tasks. The work frontier begins expanding noticeably.

Table 10: Calibrated model parameters for work multiplier projections.

| Parameter | Symbol | Value |
|---|---|---|
| Max compression ratio | $\rho_{\max}$ | 500 |
| Logistic steepness | $k$ | 0.5 |
| Inflection point | $\alpha_0$ | 10 |
| Superlinearity exponent | $\beta$ | 1.7 |
| Demand elasticity | $\varepsilon$ | 1.8 |
| Value decay exponent | $\gamma$ | 1.2 |
| Effective dimensionality | $d$ | 5.3 |
| Feasibility threshold | $\theta$ | 1.0 |

- **2030–2040**: $\rho \approx 50$–200, $M_W \approx 7$–40. Frontier expansion becomes the dominant growth driver. Previously infeasible task categories (Section 7) are systematically exploited. New industries emerge around AI-enabled services.

- **2040–2055**: $\rho \approx 200$–500, $M_W \approx 40$–100. The opportunity space is vast. Growth is increasingly constrained by imagination bandwidth rather than cost or feasibility. AI-assisted imagination augmentation becomes the marginal growth driver.

These projections are consistent with the dynamical system analysis in [17], which models the full feedback loop between AI capability, opportunity space, work generation, and investment.

### A.5 FTE Estimation Methodology

The new-FTE estimates in Table 5 are order-of-magnitude Fermi calculations. The general formula for each category is:

$$\mathrm{FTE}_s = N_s \times h_s \times g_s \tag{88}$$

where $N_s$ is the addressable population (users, firms, or institutions), $h_s$ is the per-unit human labor intensity in FTE (for deployment, oversight, maintenance, and integration), and $g_s$ is a geographic scaling factor ($g_s = 1$ for US-only categories, $g_s \in [3, 5]$ for global categories).

Representative inputs:

- **Universal custom software.** $N_s = 33 \times 10^6$ US small businesses; $h_s \approx 0.06$ FTE per business (roughly 120 hours/year of specification, integration, and maintenance); $g_s = 1$. Product: $\approx 2 \times 10^6$ FTE.

- **Predictive maintenance.** $N_s \approx 50 \times 10^6$ monitored assets (US buildings, vehicles, infrastructure); $h_s \approx 0.012$ FTE per asset; $g_s = 5$ (global). Product: $\approx 3 \times 10^6$ FTE.

- **Personalized education.** $N_s = 50 \times 10^6$ US K-12 students; $h_s \approx 0.01$ FTE per student (curriculum curation, AI oversight, escalation handling); $g_s = 1$. Product: $\approx 5 \times 10^5$ FTE.

The remaining categories follow analogous calculations. These estimates carry roughly $\pm 0.5$ order-of-magnitude uncertainty; their purpose is to demonstrate that the *aggregate* scale of new work is measured in millions of FTE, not to provide precise sector forecasts.