Are Ideas Getting Harder to Find?

Nicholas Bloom, Charles I. Jones, John Van Reenen, and Michael Webb

Bui Duc Thinh

DEPOCEN

Octorber 24, 2025

Stable Economic Growth: Relatively Stable "Idea Output"

- Long-run growth in developed economies (e.g., US TFP growth) has been relatively stable or even declining.
- There is no sign of accelerating exponential growth.

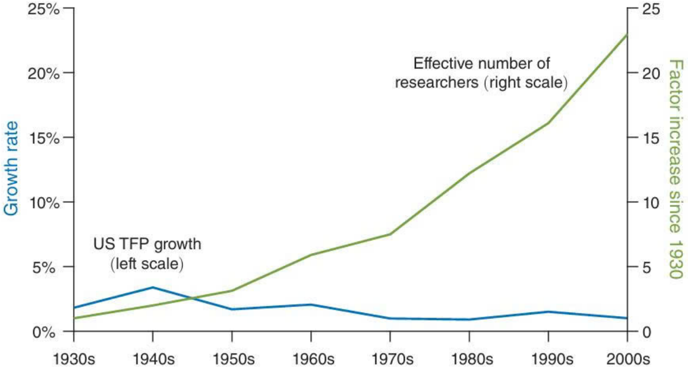

Figure 1. Aggregate Data on Growth and Research Effort

Explosive Research Effort: Explosive Growth in "Research Input"

- The number of researchers and real R&D expenditures have grown exponentially for decades.
- Why isn't explosive growth in research inputs leading to explosive economic growth?

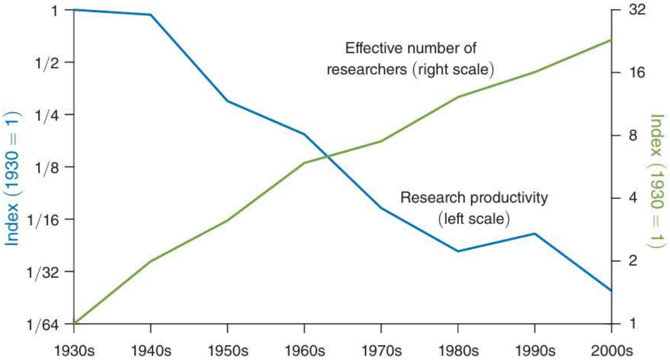

Figure 2. Aggregate Evidence on Research Productivity

Central Research Question & Main Thesis

The Question

Is it possible that the *productivity* of research itself is falling?

The Paper's Striking Thesis

Yes. Research productivity is falling sharply, everywhere we look. Ideas, and the exponential growth they imply, are getting harder and more expensive to find.

"Everywhere we look, we find that ideas are getting harder to find."
– Bloom, Jones, Van Reenen, & Webb (2020)

The Core Accounting Framework

The authors propose a simple, powerful identity to structure the analysis:

$$\underbrace{\text{Economic Growth}}_{\text{Output}} = \underbrace{\text{Research Productivity}}_{\text{Productivity}} \times \underbrace{\text{Number of Researchers}}_{\text{Input}}$$

$$g_A = \alpha_t \times S_t$$

The logic is simple:

- If Economic Growth (g_A) is stable...
- ...and the Number of Researchers (S_t) is rising exponentially...
- **...then mathematically, Research Productivity (α_t) must be falling.**

The paper's goal is to empirically measure this decline.

Presentation Outline

Relationship to the Existing Literature

1st Gen Endogenous Growth Models (Romer, Aghion-Howitt)

- Assumed a constant research productivity (α is a constant).
- The "Ideas Production Function": $\dot{A}_t = \alpha S_t A_t$
- **Key Implication:** "Permanent growth effects." Policies that increase the number of researchers (S_t) can permanently increase the long-run growth rate.
- This framework struggled to match the data (the "scale effects" puzzle).

Relationship to the Existing Literature

Semi-Endogenous Growth Models (Jones, Kortum)

- Introduced the idea that finding new ideas gets harder as the stock of knowledge (A_t) grows.
- This paper's findings provide strong empirical support for this class of models.

The "Composition Bias" Hypothesis (Howitt, Peretto)

Could the macro-level decline be misleading?

- Perhaps research productivity is constant at the *micro* (product) level.
- But as the economy grows, the number of products (N_t) increases, so researchers per product (S_t/N_t) remains constant.
- This motivates the paper's micro-level, case-study approach.

The Micro-Level Ideas Production Function

The "Composition Bias" hypothesis leads to a micro-level specification:

$$\text{Eq. (4):} \qquad \frac{\dot{A}_{it}}{A_{it}} = \alpha S_{it}$$

- $\frac{\dot{A}_{it}}{A_{it}}$: The quality growth rate of a **specific product** i.
- S_{it}: The number of scientists dedicated to improving that **specific product** i.

This leads to the paper's micro-level measure of research productivity:

$$\text{Eq. (5):} \qquad \text{Productivity} := \frac{\dot{A}_{it}/A_{it}}{S_{it}}$$

The Measurement Challenge: "Lab Equipment" Model

The Problem

In reality, inputs are total R&D expenditures (R_t), not just scientists (S_t). This is the "Lab Equipment" view.

- Naive production function: $\dot{A}_t = \alpha R_t$ (Eq. 6)
- Naive productivity measure: $(\dot{A}_t / A_t)/R_t$

Why the Naive Approach is Wrong

This measure is misleading because it ignores rising wages.

- As technology (A_t) grows, skilled wages (w_t) also grow.
- Nominal R&D (R_t) must increase just to pay for the *same number* of researchers.
- We must deflate R&D spending to find the *real quantity* of research effort.

Deriving the Correct Measure (Intuition)

The paper places the "Lab Equipment" model into a standard growth framework. After some algebra, this leads to the growth equation:

$$\text{Eq. (13):} \quad \frac{\dot{A}_t}{A_t} = \underbrace{\tilde{\alpha}_t}_{\substack{\text{Effective Research} \\ \text{Productivity}}} \times \underbrace{S_t}_{\substack{\text{Effective Research} \\ \text{Effort}}}$$

The Insight

The key is to properly translate observable R&D spending (R_t) into an economically meaningful quantity of research effort. The crucial step is the next one.

The Methodological Breakthrough: A Key Transformation

The authors perform a simple but brilliant transformation by multiplying and dividing by the nominal wage (w_t):

$$\text{Eq. (12):} \quad \frac{\dot{A}_t}{A_t} = \underbrace{\left(\frac{\alpha w_t}{A_t}\right)}_{\substack{\text{Productivity per} \\ \text{effective researcher}}} \times \underbrace{\left(\frac{R_t}{w_t}\right)}_{\substack{\textbf{Effective number} \\ \textbf{of researchers}}}$$

- $\frac{R_t}{w_t}$ deflates total R&D spending by the skilled wage.
- **Intuition:** It answers, "How many researchers could this budget hire?"
- This isolates the **real quantity** of research effort.

The Core Measurement Formulas

This theoretical derivation gives the two formulas used throughout the paper:

1. Effective Research Effort (S_t)

$$S_t = \frac{R_t}{w_t} = \frac{\text{Nominal R\&D Expenditure}}{\text{Nominal Skilled Wage}}$$

2. Research Productivity (α_t)

$$\alpha_t = \frac{\dot{A}_t / A_t}{S_t} = \frac{\text{Growth Rate of "Ideas"}}{\text{Effective Research Effort}}$$

Summary of the Methodological Approach

To test if ideas are getting harder to find, the authors will:

1. **Go Micro:** Focus on specific products/industries to avoid composition bias.
2. **Measure "Idea Output":** Find a quantifiable measure of progress for each case (e.g., transistor density, crop yields).
3. **Measure "Research Input" (S_t):**
   - Gather corresponding nominal R&D spending (R_t).
   - Deflate it by the skilled wage (w_t).
4. **Calculate Research Productivity (α_t):** Divide output growth by input (S_t) and examine its trend over time.

Case Studies

We now apply the measurement framework to four distinct domains to test the core hypothesis.

1. **Moore's Law:** The heart of the modern digital economy.
2. **Agriculture:** A traditional sector with high innovation.
3. **Health & Mortality:** An area critical to human welfare.
4. **Firm-Level Data:** To test for generality across the economy.

Let's begin with the most iconic example of exponential growth.

Case 1: Moore's Law - The Context

The "Law"

An empirical regularity stating that the number of transistors on an integrated circuit doubles approximately every two years.

- This implies a **constant exponential growth rate** of roughly 35% per year.
- It has been a key driver of economic growth for nearly half a century.
- It is not a law of nature, but the result of *intense research effort*.

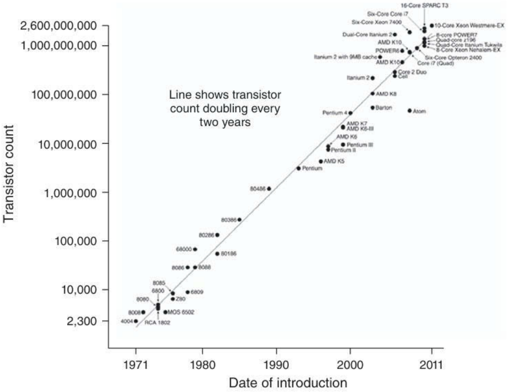

FIGURE 3. THE STEADY EXPONENTIAL GROWTH OF MOORE'S LAW

Case 1: Moore's Law - Measurement

Measuring "Idea Output" (Growth Rate of Ideas)

- The paper assumes a constant growth rate based on the doubling time.
- **Output Growth (g_A) = 35% per year.**

Measuring "Research Input" (Effective Researchers, S_t)

- **Numerator (R_t):** Sum of nominal R&D spending from all major semiconductor firms and equipment manufacturers (Intel, Samsung, TI, etc.).
- **Denominator (w_t):** Nominal wage of high-skilled workers (males with bachelor's degree or more).

Case 1: Moore's Law - The Striking Divergence

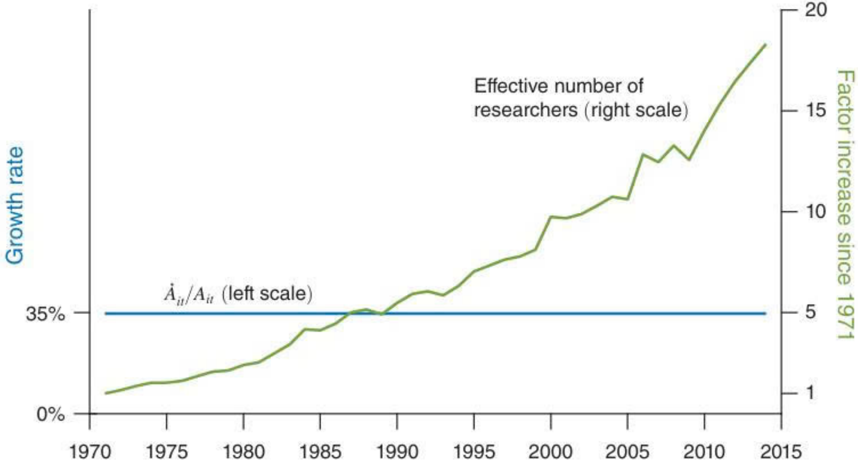

FIGURE 4. DATA ON MOORE'S LAW

Case 1: Moore's Law - Interpretation

The Numbers

- Research effort (S_t) required today is more than **18 times larger** than in the early 1970s.
- Since output growth was constant, research productivity (α_t) must have fallen by a factor of 18.
- This implies an average annual decline in research productivity of **-6.8% per year**.

The "Half-Life" of Research Productivity

At this rate, it takes approximately **10 years** for the productivity of semiconductor research to fall by half.

Table 1: Research Productivity for Moore's Law

	Factor decrease	Average growth (%)	Implied half-life (years)
Moore's Law, 1971-2014			
Baseline	18	-6.8	10.3
1. Narrow	8	-4.8	14.5
2. Narrow (downweight conglomerates)	11	-5.6	12.3
3. Broad (downweight conglomerates)	26	-7.6	9.1
4. Intel only (narrow)	347	-13.6	5.1
TFP Growth in Semiconductors, 1975-2011			
6. Narrow (no equipment R&D)	5	-3.2	21.4
7. Narrow (with equipment R&D)	7	-4.4	15.8
8. Broad (no equipment R&D)	11	-5.6	12.3
9. Broad (with equipment R&D)	13	-6.1	11.3

Key Takeaways

- Even the most conservative measure (Row 6) shows productivity falling by a factor of 5.
- The implied half-life is short, typically between 10 to 15 years, highlighting the rapid pace of decline.

Case 1: Moore's Law - Robustness Caveats

Robustness Checks (Table 1)

- The finding is robust to different definitions of "semiconductor R&D" (narrow vs. broad).
- It is also robust when using industry TFP growth as the output measure instead of Moore's Law itself.
- In all specifications, the decline is large and significant.

Potential Caveats Discussed by Authors

- **Knowledge Spillovers:** Could spillovers from other fields be declining? The authors argue this would be part of the productivity decline, not an explanation against it.
- **Measurement Error:** Did they understate early R&D? They tried to mitigate this by including older, now-defunct firms (e.g., Fairchild).

Case 2: Agriculture - The Context

Why Agriculture?

- A traditional, yet highly innovative, sector.
- Data on crop yields and agricultural R&D are relatively well-measured due to its historical importance.
- Provides a stark contrast to the high-tech world of semiconductors.

The analysis focuses on four major US crops: **corn, soybeans, cotton, and wheat.**

Case 2: Agriculture - Measurement

Measuring "Idea Output" (Growth Rate of Ideas)

- **Primary Measure:** Annualized 5-year growth rate of **yield per acre**.
- **Author's Justification:** The authors choose yield-per-acre over yield-per-farmer because it is more tightly linked to the biological and chemical R&D they can measure.

Measuring "Research Input" (Effective Researchers, S_t)

- **Numerator (R_t):** Sum of public and private R&D spending directed at each specific crop.
- Two measures are used:
  - *Narrow:* R&D for seed efficiency only (e.g., hybridization, GMO).
  - *Broad:* Includes R&D for crop protection (e.g., pesticides).

Case 2: Agriculture - Results by Crop

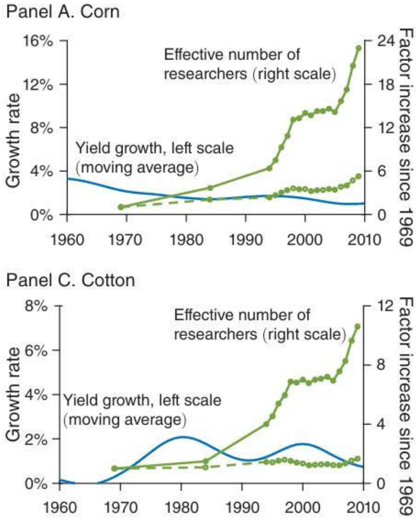

FIGURE 6. YIELD GROWTH AND RESEARCH EFFORT BY CROP

Research Productivity in Agriculture, 1969-2009

	Effective research			Research productivity	
	Factor inc.	**Avg. growth (%)**		**Factor dec.**	**Avg. growth (%)**
Research on seed efficiency only			**Research on seed efficiency only**		
Corn	23.0	7.8	Corn	52.2	-9.9
Soybeans	23.4	7.9	Soybeans	18.7	-7.3
Cotton	10.6	5.9	Cotton	3.8	-3.4
Wheat	6.1	4.5	Wheat	11.7	-6.1
Research includes crop protection			**Research includes crop protection**		
Corn	5.3	4.2	Corn	12.0	-6.2
Soybeans	7.3	5.0	Soybeans	5.8	-4.4
Cotton	1.7	1.3	Cotton	0.6	+1.3
Wheat	2.0	1.7	Wheat	3.8	-3.3

- **Massive Input Growth:** Especially for seed efficiency R&D in corn and soybeans (grew 23x).
- **Steep Productivity Decline:** With stable yield growth, productivity fell sharply, especially for corn (-9.9%/year).

Case 2: Agriculture - Robustness Caveats

Potential Caveats Discussed by Authors

- **Land Quality:** Are researchers just trying to make lower-quality land productive, thus depressing average yields?
- **Author's Rebuttal:** Total acreage for some crops has actually declined, and the required rate of decline in land quality to explain the results would have to be implausibly large and exponential.
- **Robustness Check:** The results hold when looking at state-level data.

Case 3: Health & Mortality - The Context

Why Health Research?

- Health expenditures are a huge fraction of GDP (18% in the US).
- A healthy life is one of the most valuable "goods" we consume.
- The paper focuses on two of the leading causes of death: **cancer** and **heart disease**.

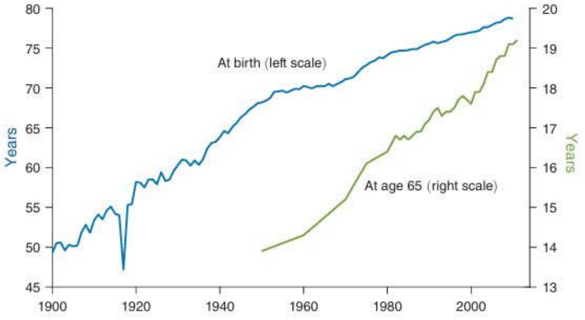

FIGURE 7. UNITED STATES LIFE EXPECTANCY

Case 3: Health & Mortality - Output Measurement

A Different Kind of Growth

- Life expectancy does not grow exponentially; it grows **linearly** (a stable number of years gained per decade).
- Therefore, using a percentage growth rate is inappropriate.

Idea Output: Years of Life Saved

- The paper measures idea output as the change in life expectancy due to reductions in disease-specific mortality.
- Derived from a formula (Eq. 16) that considers:
    1. The fraction of deaths from the disease.
    2. The remaining life expectancy.
    3. The percentage decline in the mortality rate.

Case 3: Health & Mortality - Input Measurement

The Challenge

R&D spending data is not readily available by specific disease.

A Creative Solution: PubMed Data

- The authors use the number of scientific publications as a proxy for research effort.
- Data is queried from the PubMed database using MeSH keywords (e.g., "Neoplasms" for cancer).
- Two measures are used:
  - *Broad:* All related publications.
  - *Narrow:* Publications corresponding to **clinical trials**.

Case 3: Health & Mortality - Results

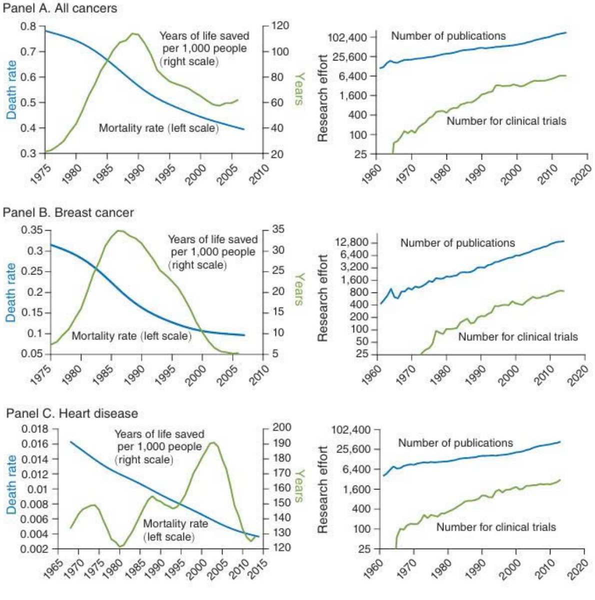

FIGURE 8. MORTALITY, YEARS OF LIFE SAVED, AND RESEARCH EFFORT

Case 3: Research Productivity for Medical Research

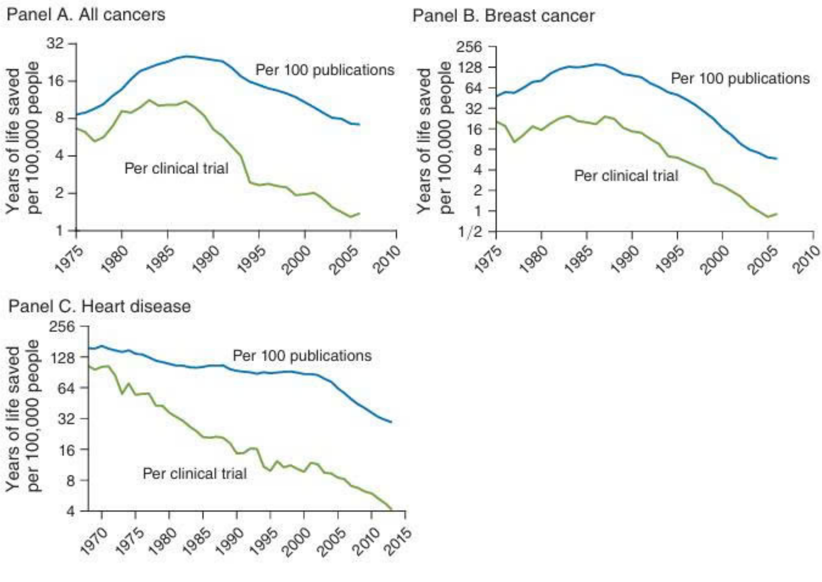

Figure 9. Research Productivity for Medical Research

Case 3: Research Productivity for Medical Research

Effective research

Disease	Factor increase	Avg. growth (%)
All publications		
Cancer, all types	3.5	4.0
Breast cancer	5.9	5.7
Heart disease	5.1	3.6
Clinical trials only		
Cancer, all types	14.1	8.5
Breast cancer	16.3	9.0
Heart disease	24.2	7.1

Research productivity

Factor decrease	Avg. growth (%)
All publications	
1.2	-0.6
8.2	-6.8
5.3	-3.7
Clinical trials only	
4.8	-5.1
22.6	-10.1
25.3	-7.2

- **Explosion in Clinical Trials:** Research effort, measured by clinical trials, grew immensely (14x to 24x).
- **Dramatic Productivity Collapse:** Productivity for breast cancer research fell by a factor of 22.6, a staggering **-10.1% per year**.

Part IV: From Case Studies to General Evidence

The Question of Generality

The case studies are powerful, but are they representative of the broader economy?

- Moore's Law is a unique high-tech example.
- Agriculture and Health have their own specific dynamics.

The Next Step

To test the generality of the findings, the authors turn to large-scale, firm-level datasets.

1. **Compustat:** Publicly traded U.S. firms across all sectors.
2. **U.S. Census of Manufacturing:** The universe of manufacturing firms.

Table 4: Research Productivity in Compustat (Selected Results)

Sample	Effective research: Factor increase	Effective research: Avg. growth (%)	Research productivity: Factor decrease	Research productivity: Avg. growth (%)
Sales revenue				
3 dec (469 firms)	3.8	6.7	9.2	-11.1
4 dec (149 firms)	13.7	8.7	40.3	-12.3
Employment				
3 dec (319 firms)	4.0	6.9	18.2	-14.5
4 dec (101 firms)	13.9	8.8	31.5	-11.5

Key Findings

- The pattern holds robustly: massive increases in R&D effort are met with sharply declining research productivity.
- The decline is **even more severe** than in the case studies, often exceeding **-10% per year**.

Figure 10: Distribution of Productivity Changes across Firms

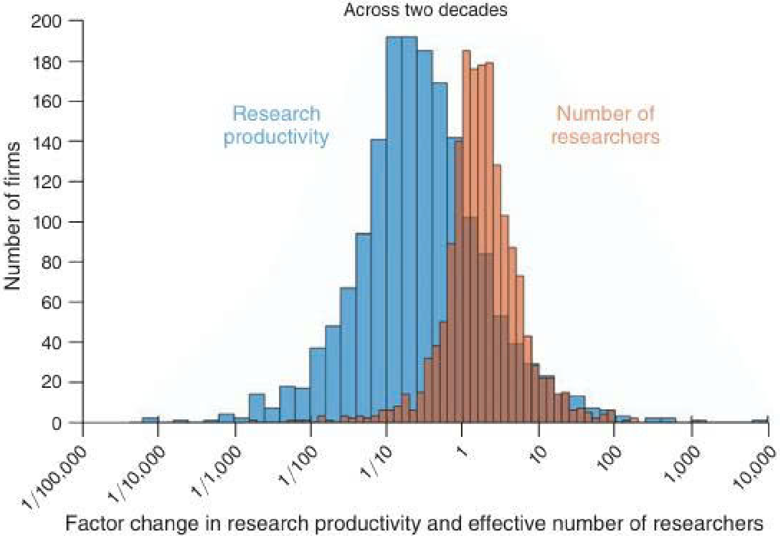

Table 6: Results from the Census of Manufacturing

Case	Factor decrease	Avg. growth (%)
1. Benchmark	2.2	-7.8
2. Winsorize g below 0.01	1.9	-6.0
3. Winsorize top/bottom	1.7	-4.9
4. Unweighted	1.9	-8.1
5. Research = scientists	**2.3**	**-6.0**

Confirmation from a Broader Dataset

- This dataset confirms the Compustat results using the universe of manufacturing firms.
- Crucially, using a **direct quantity measure** of inputs (number of scientists, Row 5) still shows a strong productivity decline of **-6.0% per year**.

Summary of Evidence: A Consistent Story Across All Domains

TABLE 7—SUMMARY OF THE EVIDENCE ON RESEARCH PRODUCTIVITY

Scope	Time period	Average annual growth rate (%)	Half-life (years)	Dynamic diminishing returns, β
Aggregate economy	1930–2015	−5.1	14	3.1
Moore's Law	1971–2014	−6.8	10	0.2
Semiconductor TFP growth	1975–2011	−5.6	12	0.4
Agriculture, US R&D	1970–2007	−3.7	19	2.2
Agriculture, global R&D	1980–2010	−5.5	13	3.3
Corn, version 1	1969–2009	−9.9	7	7.2
Corn, version 2	1969–2009	−6.2	11	4.5
Soybeans, version 1	1969–2009	−7.3	9	6.3
Soybeans, version 2	1969–2009	−4.4	16	3.8
Cotton, version 1	1969–2009	−3.4	21	2.5
Cotton, version 2	1969–2009	+1.3	−55	−0.9
Wheat, version 1	1969–2009	−6.1	11	6.8
Wheat, version 2	1969–2009	−3.3	21	3.7
New molecular entities	1970–2015	−3.5	20	...
Cancer (all), publications	1975–2006	−0.6	116	...
Cancer (all), trials	1975–2006	−5.7	12	...
Breast cancer, publications	1975–2006	−6.1	11	...
Breast cancer, trials	1975–2006	−10.1	7	...
Heart disease, publications	1968–2011	−3.7	19	...
Heart disease, trials	1968–2011	−7.2	10	...
Compustat, sales	3 decades	−11.1	6	1.1
Compustat, market cap	3 decades	−9.2	8	0.9
Compustat, employment	3 decades	−14.5	5	1.8
Compustat, sales/employment	3 decades	−4.5	15	1.1
Census of Manufacturing	1992–2012	−7.8	9	...

Implications for Growth Theory

Strong Support for Semi-Endogenous Growth

The evidence strongly favors models where finding ideas gets harder over time. The idea production function is better specified as:

$$\text{Equation (17):} \quad \frac{\dot{A}_t}{A_t} = \alpha A_t^{-\beta} S_t$$

This leads to the steady-state growth equation:

$$\text{Equation (18):} \quad g_A = \frac{g_S}{\beta}$$

- Economic growth (g_A) is proportional to the growth rate of research effort (g_S), deflated by the difficulty parameter (β).

Implications for Growth Theory: The Semiconductor Paradox

The Paradox

Why invest so heavily in semiconductors, where productivity seems to fall fastest?

The Answer Lies in β (The Difficulty Parameter)

The authors calculate $\beta = $ (productivity decline rate)/(idea growth rate)

- **Aggregate Economy:** $\beta = 5.1\%/1.5\% \approx \mathbf{3.4}$
- **Semiconductors:** $\beta = 6.8\%/35\% \approx \mathbf{0.2}$

The Insight

Productivity falls fastest in semiconductors not because it's the hardest field, but because it has the **lowest diminishing returns** (lowest β). This makes it worthwhile to grow research effort (g_S) the fastest in this sector.

Authors' Defense: Selection & Measurement Issues

The authors directly address potential critiques of their methodology in Section VII.B.

1. How high is the hurdle to overturn the main conclusion?
2. Are the inputs and outputs properly matched?
3. What about unmeasured inputs like knowledge spillovers?

Let's examine their arguments.

Authors' Defense: The High Hurdle Argument

Question 1: Could a different output measure change the result?

- **Observation:** The paper's measured inputs (S_t) are growing at 5-7% per year, or even faster. This means inputs are doubling every 10-15 years.
- **The Hurdle:** To keep productivity constant, an alternative output measure would also need to show growth rates that are *themselves doubling every 10-15 years*.
- **Conclusion:** This implies explosive, accelerating growth far beyond anything observed in reality. A stationary measurement error (e.g., consistently underestimating growth by 1%) would not change the conclusion at all.

The observed divergence between input and output trends is simply too massive to be explained by plausible measurement error.

Authors' Defense: Matched Measures & Missing Inputs

Question 2: Are inputs and outputs properly matched?

- The authors argue they were extremely careful.
- **Example:** They chose yield-per-acre for agriculture because it is more tightly linked to the seed R&D they can measure, unlike yield-per-farmer which is heavily influenced by capital (tractors, GPS).

Question 3: What about unmeasured inputs (spillovers, etc.)?

- The paper is measuring a "Solow residual" of idea production.
- Any unmeasured input is, by definition, captured in the productivity measure.
- The finding is that total factor productivity of research, *including spillovers*, is declining. A bias would only occur if unmeasured inputs grew systematically slower than measured ones.

Conclusion: A New "Stylized Fact" for Growth Economics

The Robust Finding

The central conclusion of the paper is unambiguous: **research productivity is falling sharply, everywhere we look.**

- Ideas are getting harder, and therefore more expensive, to find.
- The "Red Queen Effect" is real: We have to run faster and faster (double research effort every 13 years) just to keep our rate of economic growth stable.
- This finding provides overwhelming empirical support for the class of semi-endogenous growth models.

Policy Implication

If sustained growth in research effort (g_S) is the engine of economic growth, then policies that expand the supply of researchers (education, immigration) and incentivize R&D are more critical than ever.

Conclusion: A Masterclass in Empirical Research

Beyond its important findings, this paper serves as a model for modern empirical work. Its persuasive power comes not from complex econometrics, but from:

1. **A Transparent Accounting Framework:** The core logic is simple, clear, and hard to refute.
2. **Methodological Rigor in Measurement:** The theoretically-grounded definition of "effective research effort" is the paper's brilliant core.
3. **The Power of Triangulation:** The consistency of the finding across vastly different domains—from semiconductors to soybeans, from cancer research to the entire U.S. economy—makes the conclusion extraordinarily robust.
4. **Effective Data Visualization:** Simple, clear charts that make the central finding immediately apparent.

It reminds us that the highest form of empirical work is often to measure the world correctly and let the data tell its story.

Critical Perspectives: Measurement Selection

Are We Measuring the Right Things?

- **Quality of Researchers:** Is a "researcher" in 1970 equivalent to one today? Today's researchers are more specialized and rely on more capital. The wage deflator may not fully capture this.
- **The Nature of "Ideas":** Is a transistor-doubling "idea" equivalent to a cancer-curing "idea"? The paper focuses on the *rate of progress*, but the economic and social value of ideas varies enormously.

Selection Bias in Case Studies?

- The paper focuses on established technologies. What about emerging fields where productivity might be *increasing*?
- **Examples:** Artificial Intelligence (cost to train a model has fallen dramatically), CRISPR gene editing.
- Is the "hump shape" seen in health research the norm for new technologies?

Critical Perspectives: Alternative Explanations

Is declining productivity an iron law of nature, or a result of institutional factors?

"Low-Hanging Fruit" Hypothesis

- The easiest and most fundamental discoveries have been made. We are now in a phase of diminishing returns. The paper provides the best quantification of this effect to date.

Institutional Economic Factors

- **Rise of "Defensive R&D":** Are incumbent firms using R&D to create patent thickets and deter entry, rather than for true innovation?
- **Regulatory Burden:** Has the complexity and cost of getting approval (e.g., from the FDA) increased, thus lowering the measured productivity of R&D spending?
- **Shift from Basic to Applied Research:** A decline in publicly-funded basic research could be starving the innovation pipeline.

Thank you.

Questions & Discussion