AI, Innovation, and Economic Growth

Alexander Oettl
Georgia Institute of Technology and NBER

April 21, 2026

NBER Corporate Associates Research Symposium
New York, NY

Based on joint work with Ajay Agrawal (Toronto/NBER) and John McHale (Galway)

Why should we care about AI and Innovation?
The stakes are enormous ... and contested

Sachs

15 October 2025 | 4:48PM EDT

GLOBAL ECONOMICS ANALYST

The AI Spending Boom Is Not Too Big (Briggs)

- Announcements of further increases in an already sizable amount of AI capex spending have amplified concerns around the sustainability of AI investment. In this Global Economics Analyst, we argue that anticipated investment levels are sustainable, although the ultimate AI winners remain less clear.
- The technological backdrop remains supportive of AI capex for two reasons. First, AI applications are boosting productivity when deployed. Second, unlocking these productivity benefits requires significant computational power, especially since models are increasing in size much more quickly than computation and energy costs are falling.
- We are not concerned about the total amount of AI investment. AI investment as a share of US GDP is smaller today (<1%) than in prior large technology cycles (2-5%). Furthermore, we estimate an $8tn present-discounted value for the capital revenue unlocked by AI productivity gains in the US, with plausible estimates ranging from $5tn-$19tn. These estimates exceed current cumulative AI investment forecasts even before considering foreign profits, new profit pools, or the potential of AGI.

or the potential of AGI.

Jan Hatzius
+1(212)902-0394 | jan.hatzius@gs.com
Goldman Sachs & Co. LLC

Joseph Briggs
+1(212)902-2163 |
joseph.briggs@gs.com
Goldman Sachs & Co. LLC

Sarah Dong
+1(212)357-9741 | sarah.dong@gs.com
Goldman Sachs & Co. LLC

Megan Peters
+44(20)7051-2058 |
megan.l.peters@gs.com
Goldman Sachs International

BUSINESS • 6 MIN READ

Why this analyst says the AI bubble is 17 times bigger than the dot-com bust

OCT 18, 2025

Analysis by Allison Morrow

Is AI a transformative investment or a speculative bubble?

The answer depends on **where AI enters the production function**.

AI and Scientific Discovery: Two leading views from Economics

The Simple Macroeconomics of AI*

Daron Acemoglu
Massachusetts Institute of Technology

April 5, 2024

“…use the full framework from Acemoglu and Restrepo (2022) …upper bound on GDP effects rises to 1.6 – 1.8%... the numbers above may be overestimates …productivity gains…in tasks that are “easy-to-learn”...more modest increase in **TFP and GDP in the next 10 years that can be upper bounded by 0.55% and 0.90%, respectively.**”

American Economic Review: Insights
ISSN 2640-205X (Print) | ISSN 2640-2068 (Online)

About *AER: Insights* ⌄ Articles and Issues ⌄ Information for Authors and Reviewers ⌄

The AI Dilemma: Growth versus Existential Risk

Charles I. Jones

AMERICAN ECONOMIC REVIEW: INSIGHTS
VOL. 6, NO. 4, DECEMBER 2024
(pp. 575–90)

“Suppose that advances in AI allow computers to augment and even substitute for humans in innovation, leading to an **acceleration of economic growth to some rate g, perhaps 10 percent per year.** However, the use of this AI poses an existential risk to humanity.”

Why such different conclusions? Assumptions about the knowledge production function!

"It seems likely that AI will augment our abilities to innovate in the near term, and it is certainly within the realm of possibility that AI could match or even exceed human intelligence at many cognitive tasks and begin innovating itself. Once machines can produce ideas, the limits to growth set by the quantity and quality of researchers may no longer hold, and growth rates could speed up, potentially even leading to a so-called "singularity" with infinite consumption."

"I do not discuss how AI can have revolutionary effects by changing the process of science …because large-scale advances of this sort do not seem likely within the 10-year time frame, and many current discussions focus on automation and task complementarities."

They disagree on whether AI enters the **knowledge production function** at all.

What are the implications of AI entering the knowledge production function?

1. AI **does** change how organizations produce new knowledge.

2. It does so by *augmenting* the ability to **search vast combinatorial spaces** of questions, ideas, and designs.

3. Human judgment remains critical, but **where** judgment matters most shifts as AI reshapes each stage of the innovation process.

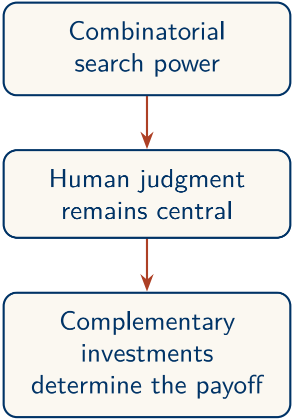

A Framework

How AI enters the knowledge production function

Innovation proceeds through four stages of combinatorial search

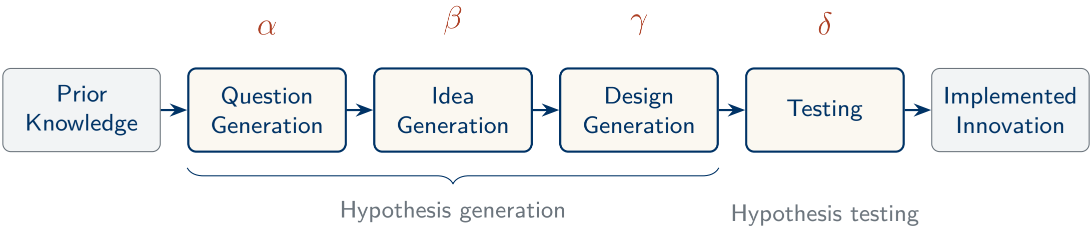

Each stage involves search over potentially vast combinatorial spaces.

In practice, substantial iteration between stages is common.

Research productivity decomposes into four stages

Knowledge production (Romer):

$$\dot{A} = \omega \cdot A \cdot S$$

$\dot{A}$ = new ideas (output)
A = stock of existing ideas
S = number of researchers
ω = research productivity

Our decomposition:

$$\omega = \alpha \cdot \beta \cdot \gamma \cdot \delta$$

α = question generation
β = idea generation
γ = design generation
δ = testing

Multiplicative structure $\Rightarrow$ AI can affect each factor differently.
A bottleneck at any stage constrains the whole pipeline.

What does this look like in practice?

Stage	**Drug Discovery** (~10^{60} molecules)	**Chip Design** (Google TPU, Synopsys)
α Question Generation	Which disease target to pursue? Which protein is malfunctioning?	What performance target? What use case?
β Idea Generation	Which therapeutic mechanism? Combine known binding pathways.	What architecture? Chiplets, tensor cores, memory hierarchy.
γ Design Generation	Which candidate molecule out of billions of possibilities?	Circuit layout and routing across billions of transistors.
δ Testing	Clinical trials (10+ years, ~$2B per drug).	Fabrication yield, reliability testing.

Each stage involves combinatorial search.

But AI is most likely to have the biggest impact when three conditions hold.

AI excels at search over combinatorial design spaces

The central challenge of frontier innovation: search over vast combinatorial spaces of questions, ideas, designs, and tests.

Design generation (γ): AI's current sweet spot

The "Hassabis Hypothesis": AI excels when:

1. A combinatorial search space too large for exhaustive search
2. A clear objective function for training
3. Sufficient data (or ability to simulate it)

Proof of concept: AlphaFold (2024 Nobel Prize), AI-guided drug discovery, materials science (2.2M new crystal structures).

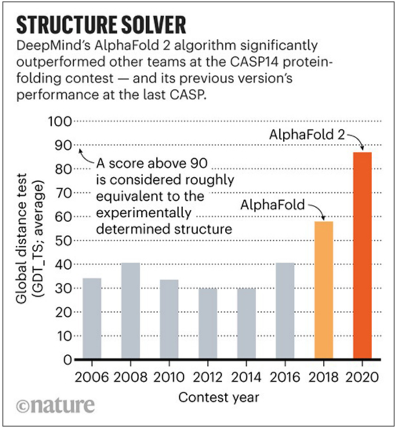

AlphaFold: 2024 Nobel Prize in Chemistry.
Years of bench work → seconds of inference.

Same impact of AI on the production function but with unique bottlenecks

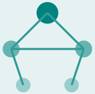

Drug Discovery

$\sim 10^{60}$ druglike molecules

- α Which disease target to pursue?
- β Which therapeutic mechanism?
- γ **Which candidate molecule?** AI screens billions
- δ Clinical trials (10+ years, ~$2B)

AI transformed γ
(AlphaFold, molecular screening)
Bottleneck moved to δ: **validation**

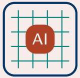

Chip Design

Google TPU, Synopsys AI EDA

- α What performance target? What use case?
- β What architecture? (chiplets, tensor cores)
- γ **Circuit layout and routing** hours instead of months
- δ Fabrication yield, reliability testing

AI transformed γ (placement and routing)
Bottleneck remains at α/β **and** δ

Same AI sweet spot (γ) but **different binding constraints.**

Human judgment remains indispensable (at least for now)

What AI does well

Search vast combinatorial spaces

Identify patterns invisible to humans

Generate candidate designs at scale

Interpolate within data-rich domains

What still requires human judgment

Abductive reasoning — the creative "guess"

Weighing competing objectives and trade-offs

Contextual and ethical sensitivity

Extrapolation in data-sparse domains

Hassabis's "Einstein Test" (2026): Train a model with a 1911 knowledge cutoff.
Can it independently derive general relativity? **Current AI cannot.**

AI is best understood as a **tool that augments** the capabilities of judgment-exercising humans, not as a full replacement for that judgment.

Why software looks different: δ is also digital

In drug discovery and chips, δ requires **physical processes**. AI cannot iterate digitally.

In software, almost **every stage** is in silico. AI cycles through $\alpha \to \beta \to \gamma \to \delta$ thousands of times without touching the physical world.

This is what happens when the testing stage (δ) has fewer bottlenecks.

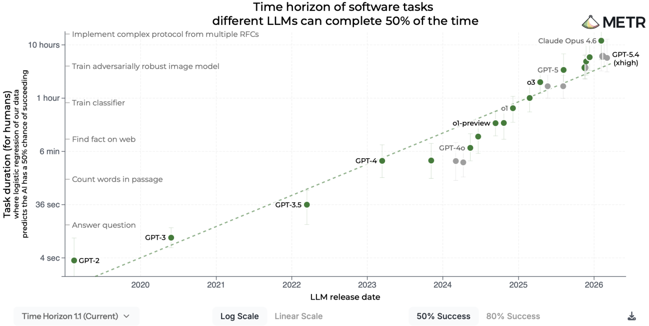

Source: METR (2026). Task capability doubling every ∼4 months.

What the Evidence Shows

Tasks or Bundles?

Early signs that AI is reshaping employment but not through mass automation

Crane & Soto (2026, Federal Reserve)

First official labor statistics indicating AI-driven employment reduction in pockets of the economy.

$\sim$500,000 fewer coders
than pre-LLM
trends predicted

Not an absolute decline but a marked slowdown in growth relative to pre-LLM trends.

But the headline obscures the mechanism

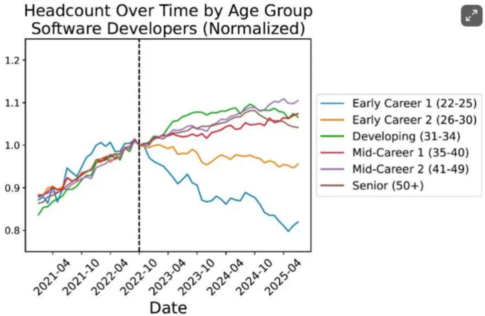

Employment for software engineers by age, normalized to 1 in October 2022.

Source: Brynjolfsson et al. (2025)

Early-career developers diverge sharply from seniors after LLMs.

Jobs are bundles of tasks and bundle structure determines AI's impact

Garicano et al. (2026, LSE): Jobs aren't lists of discrete tasks: they are **bundles** that vary in how tightly tasks are bound together.

Weak Bundles

Tasks are **discrete and separable**.

Junior coders, contractors: daily work is mainly writing code to spec.

AI conducts the coding task → the bundle **falls apart**.

Role shrinks or disappears.

Strong Bundles

Tasks are **tightly enmeshed**.

Senior devs, domain experts: coding + architecture + stakeholder judgment.

Extracting the coding part is hard → the bundle **stays intact**.

AI becomes assistant, not competitor.

The same technology **shrinks** one job while **expanding** another.

We've seen this pattern before: the spreadsheet and the bookkeeper

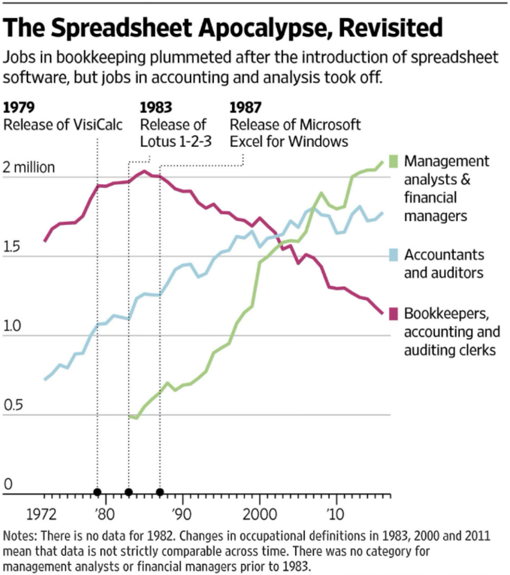

When spreadsheet software made calculation easy:

Bookkeepers (weak bundle):
Calculation *was* the job. Spreadsheets automated it.
Value **declined**. Employment fell.

Management analysts (strong bundle):
Calculation was one input into judgment-heavy work.
Spreadsheets freed time for analysis and strategy.
Value **increased**. Employment rose.

If calculation is easy, knowing what to calculate is super valuable!

The augmentation evidence: AI as tool, not replacement

What does the empirical evidence say about how AI is actually being used?

Study	Finding	Key nuance
Brynjolfsson et al. (2025)	**+15%** productivity (customer support)	Biggest gains for **lower-skilled** workers
Dell'Acqua et al. (2023)	**+40%** within AI frontier	**Worse** performance outside frontier
Peng et al. (2023)	**+55.8%** faster (GitHub Copilot)	Biggest gains for less experienced devs
Chatterji et al. (2025)	Workers use ChatGPT for **decision support**	Especially in knowledge-intensive jobs

Pattern: AI augments workers. It is a tool that enhances productivity, with the largest benefits flowing to those with **weaker initial bundles**.

The augmentation model: AI expertise is the binding constraint

Agrawal, McHale & Oettl (2026): A task-based model where AI is purely worker-augmenting. No automation.

Model setup

Two types of workers:

AI-Expert Workers
Domain knowledge + AI skills

Ordinary Workers
Domain knowledge only

AI improvement expands the share of tasks ($\bar{\tau}$) where AI expertise is valuable.

Key result

$$\frac{\partial^2 \ln \omega}{\partial \bar{\tau}\, \partial (S_{AI}/S)} > 0$$

The productivity gain from AI advancement **increases** with the share of AI-expert workers.

More AI expertise → **amplifies** productivity gains

The binding constraint is the expertise to use AI not the technology itself.

Preliminary evidence: fields with more AI expertise diverge after a foundation model shock

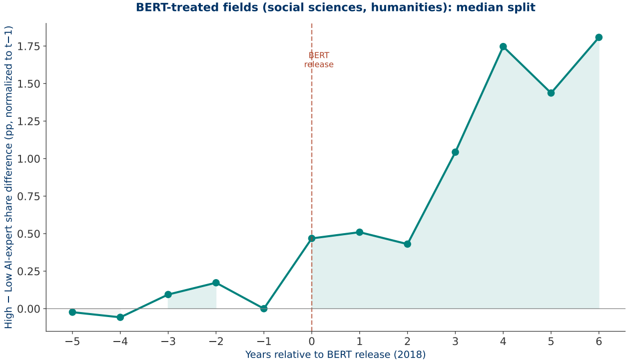

Natural experiment: BERT release (2018) as an exogenous shock to AI capabilities in social sciences and humanities.

Design: Split treated subfields by pre-treatment AI-expert share (median split). Plot the *difference* in AI-expert share growth (high − low), normalized to $t-1$.

Result: Subfields that already had more AI expertise pulled further ahead.

Consistent with the cross-partial prediction: $\frac{\partial^2 \ln \omega}{\partial \bar{\tau}\, \partial(S_{AI}/S)} > 0$

AI experts and economists disagree on how big this will be

Forecasting the Economic Effects of AI*

Ezra Karger[1,*], Otto Kuusela[2,*], Jason Abaluck[3], Kevin Bryan[4], Basil Halperin[5], Todd Jones[6], Connacher Murphy[7], Phil Trammell[7], Matt Reynolds[2], Dan Mayland[2], Ria Viswanathan[2], Ananaya Mittal[2], Rebecca Ceppas de Castro[2], Josh Rosenberg[2], and Philip E. Tetlock[8]

[1]Federal Reserve Bank of Chicago
[2]Forecasting Research Institute
[3]Yale School of Management
[4]University of Toronto
[5]University of Virginia
[6]Mississippi State University
[7]Stanford University
[8]University of Pennsylvania

March 2026

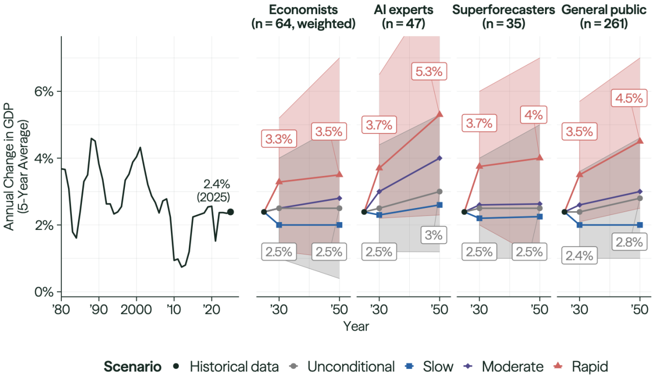

Figure 4: *Forecasts for five-year annualized change in the Gross Domestic Product (GDP).*

Why the gap? Complements are slow.

AI experts see the technology's trajectory. Economists weigh what else has to change: workforce AI expertise, workflow redesign, complementary skills, and physical bottlenecks.

The Dynamo Lesson

Why complements take time

“You can see the computer age everywhere but in the productivity statistics”

Robert Solow, *New York Times Book Review*, July 12, 1987

Paul David (1990): The Dynamo and the Computer

The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox

By PAUL A. DAVID*

Many observers of recent trends in the industrialized economies of the West have been perplexed by the conjecture of rapid technological innovation with disappointingly slow gains in measured productivity. A generation of economists who were brought up to identify increases in total factor productivity indexes with "technical progress" has found it quite paradoxical for the growth accountants' residual measure of "the advance of knowledge" to have vanished at the very same time that a wave of major innovations was appearing—in microelectronics, in communications technologies based on lasers and fiber optics, in composite materials, and in biotechnology. Disappointments with "the computer revolution" and the newly dawned "information age" in this regard have been keenly felt. Indeed, the notion that there is something anomalous about the prevailing state of affairs has drawn much of its appeal from the apparent failure of the wave of innovations based on the microprocessor and the memory chip to elicit a surge of growth in productivity from the sectors of the U.S. economy that recently have been investing so heavily in electronic data processing equipment (see, for example, Stephen Roach, 1987, 1988; Martin Baily and Robert Gordon, 1988). This latter aspect of the so-called "productivity paradox" attained popular currency in the succinct formulation attributed to Robert Solow: "We see the computers everywhere but in the productivity statistics."

If, however, we are prepared to approach the matter from the perspective afforded by the economic history of the large technical systems characteristic of network industries, and to keep in mind a time-scale appropriate for thinking about transitions from established technological regimes to their respective successor regimes, many features of the so-called productivity paradox will be found to be neither so unprecedented nor so puzzling as they might otherwise appear.

I

My aim here simply is to convince modern economic analysts (whether perplexed by the productivity slowdown, or not) of the immediate relevance of historical studies that trace the evolution of techno-economic regimes formed around general purpose engines.[1] The latter, typically, are key functional components embodied in hardware that can be applied as elements or modular units of the engineering designs developed for a wide variety of specific operations or processes. Accordingly, they are found ubiquitously distributed throughout such systems when the latter have attained their mature, fully elaborated state. James Watt's (separate condenser) steam engine design springs to mind readily as an example of an innovation that fulfilled this technological role in the first industrial revolution. My particular line of argument will be better served, however, by directing notice to the parallel between the modern computer and another general purpose engine, one that figured prominently in what sometimes is called the "second Industrial Revolution"—namely, the electric dynamo. (But, see also Herbert Simon, 1986.)

Although the analogy between information technology and electrical technology

*Department of Economics, Encina Hall, Stanford University, Stanford, CA 94305. Discussions with Paul Rhode were particularly helpful early in the research. I am grateful for comments from Steve Broadberry, Jonathan Cave, Nick Crafts, among the participants in the Economic History Summer Workshop held at Warwick University, July 10–28, 1989; from Timothy Taylor; and from Shane Greenstein, Avner Greif, Edward Steinmueller, and other participants in the Technology and Productivity Workshop at Stanford, October 1989.

[1]This paper draws upon material developed in a longer work—my 1989 paper.

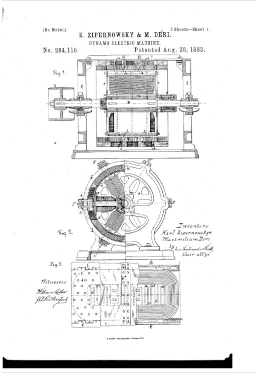

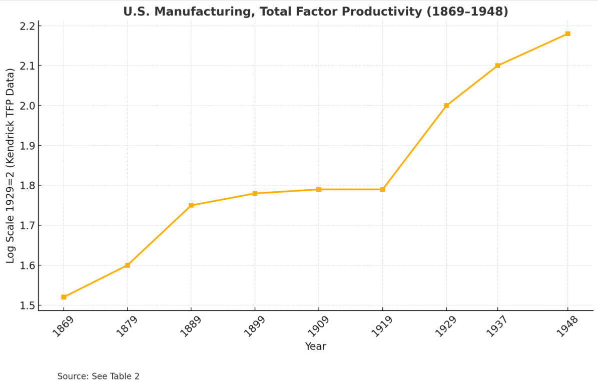

Productivity gains didn't arrive until **30–40 years** after the technology.

The factory was designed around the steam shaft

Old Factory: Central Power Shaft

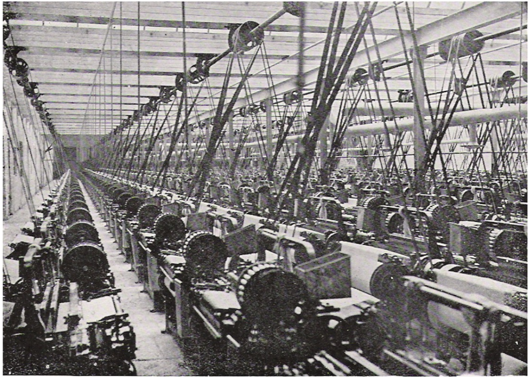

Belt-driven, multi-story, rigid layout.

Bolting on electric motors → **minimal gains**.

New Factory: Distributed Electric Power

Single-story, flexible, assembly line.

Redesigned from scratch → **transformational gains**.

The technology alone wasn't enough. **The old factory had to become obsolete.**

Complementary investments in organizational redesign were the binding constraint.

AI as General Purpose Meta-Technology: the complements that matter

AI is not just a GPT. It is a **meta**-technology: a technology for producing new technologies.

Upstream
Scaling, algorithms, compute
The "science of AI"

→

AI as GPMT
Cockburn, Henderson & Stern (2019)

→

Downstream
Skills training, workflow redesign, org. change

As AI removes a bottleneck, value shifts to the complement:

- AI accelerates design → value shifts to **testing & validation**
- AI handles routine coding → value shifts to **judgment & strategy**
- AI lowers skill barriers → value shifts to **higher-order skills**

David's lesson for AI:

- Short-term expectations too high. **Long-term too low.**
- Most adoption is Stage 1: bolting AI onto existing workflows
- Staff around the **new** bottleneck, not the old one
- Mollick: "These are the **worst** AI tools you will ever use."

Implications

For growth and for firms

The AI skill paradox: individual gains, collective costs?

Hao et al. (2026), analyzing 41 million papers in the natural sciences:

Individual productivity:

- AI adopters publish **3× more papers**
- Receive **4.8× more citations**
- Become project leaders **1.4 years earlier**

Collective exploration:

- **4.6% contraction** in breadth of topics studied
- **22% reduction** in follow-on scientific engagement

The paradox:

Individual productivity ↑

Collective exploration ↓

Are we getting better at answering known questions while asking fewer new ones?

AI may be concentrating attention on established problems rather than fostering exploration.

Realizing AI's full potential means investing in the complements:broad AI expertise, organizational redesign, and the judgment to direct AI toward new questions.

Three takeaways

1.

Focus on the knowledge production function.

Task automation so far has yielded modest gains. But if AI accelerates *innovation*, even slightly, the effects compound. Ask where the Hassabis conditions hold.

Your R&D pipeline is where AI has the most leverage over long-run value creation.

2.

Invest in AI expertise, not just AI tools.

The productivity gain from AI advancement increases with the share of workers who can actually use it. Tools alone deliver linear gains at best.

Training, hiring, and promotion criteria should make AI expertise a rewarded capability.

3.

Don't bolt AI onto the steam shaft.

Electricity took 30+ years to diffuse. The gains came from redesigning the factory, not from swapping the motor.

Most enterprise AI adoption today is Stage 1 substitution. Stage 2 reorganization, while more difficult, is ultimately more valuable.

AI transforms *which* constraints bind.

The returns will go to the organizations
that build the new factory around it.

"Compress the 21st century of biology into 5–10 years." (Amodei)

"Anyone in 2035 should have the intellectual capacity of everyone in 2025." (Altman)

Even a fraction of these gains would be transformational.

Capturing them requires **complementary investments** in skills, organization, and judgment.

Thank you.

Alex Oettl · OETTL@gatech.edu