

# Progress

---

## Class 10 - Speed and Risk

Kevin A. Bryan - Toronto Rotman - Nov 19 2025

After class notes in Red

- Why is progress in some areas slow?
  - When is progress dangerous?
- Do these necessarily go together?
  - If so, what we can we do?
  - What does this mean for AI?

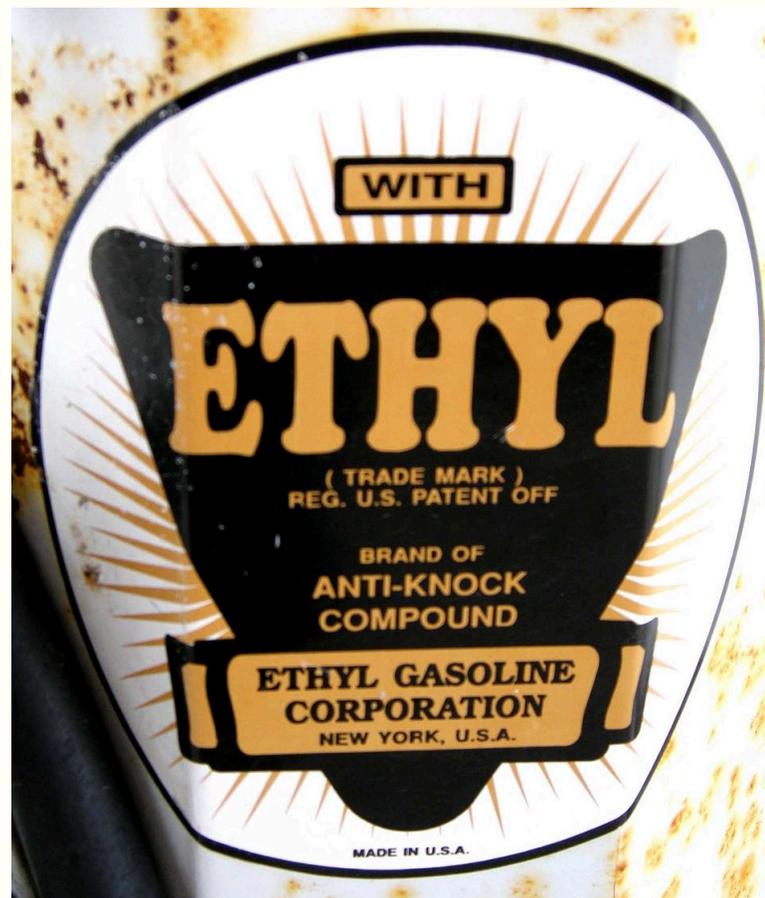
I deliberately combined speed and risk in this class. In some sectors, progress is very slow. There are many roadblocks to doing new things. Often these exist because of a bias against changing the status quo especially when there is risk of danger from the new thing.

## Two big results

1) Progress in the long-run depends on  
*bottlenecks*

2) New things *do* have risks, so does the status quo, and we can formally balance these

In *any* task, the constraint in the long-run is the bottleneck, the part of progress that is hardest to improve. Often these bottlenecks come from attempts to limit risk or danger, though not always. And these risk limitations can be rational and well-thought out, or not.



## Thomas Midgley's cautionary tale

A great early 20th century inventor. Knocking on car engines was a huge problem (and annoying) - tetraethyl lead solves it...but is poisonous. To solve a problem in safe refrigeration, he introduced the first CFC, Freon. This turns out to cause a hole in the ozone layer. He got polio and was bed-ridden in middle-age. Invented a pulley contraption to help him get out of bed...and accidentally killed himself.

*The faster the general pace of technological advance, the greater will be the increase in the overall wage level and the greater the upward pressure on costs in those industries which do not enjoy increased productivity. Faster technological progress is no blessing for the laggards, at least as far as their costs are concerned...The output per man-hour of the violinist playing a Schubert quartet in a standard concert hall is relatively fixed, and it is fairly difficult to reduce the number of actors necessary for a performance of Henry IV, Part I.*

Baumol and Bowen, On the Performing Arts

This short article on the problem with funding the arts is the origin of the Baumol Effect, sometimes called Baumol's Cost Disease. Tasks with slow productivity growth not only become relatively more expensive, but because one must attract labor away from more quickly-improving sectors, they can become *absolutely* more expensive. The famous example is a string quartet - needs four people in 1800 and today.

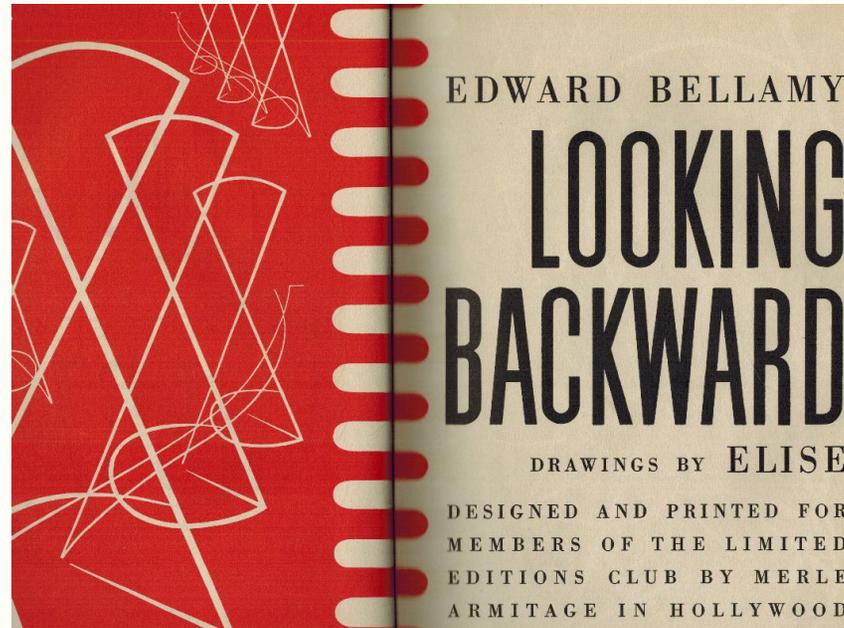
# Bottlenecks

Sector 2 doubles in productivity each period

Period	Sector 1	Sector 2	Total Hours	
$t = 0$	20	80	100	100%
$t = 1$	20	40	60	60%
$t = 2$	20	20	40	40%
$t = 3$	20	10	30	30%

Consider a task with two parts, a fast-improving one and one that's hard to improve. The fast-improving one very quickly becomes unimportant for overall productivity growth since it gets cheaper and cheaper relative to the sclerotic part. Improving farming productivity today - even free potatoes and flour! - would barely affect GDP and welfare.

# Is the Baumol Effect actually real in music production?



Consider Chapter 11 of Bellamy's famous late 1800s novel of a man who wakes up in the year 2000. Bellamy predicted many technologies, including live music over radio. His characters certainly thought there was a productivity improvement in string quartets from this! See also, electronic mixing, AI-generated music, new instruments, Spotify, etc. etc. <https://www.gutenberg.org/files/624/624-h/624-h.htm#chap11>

Perhaps: *everything is Baumol if we let it be*

Where does law, tradition, org architecture *force*  
Baumol effects that can in fact be overcome?

Sometimes a sector is Baumol - hard to reduce the amount of human labor per unit of output - because we choose to make it that way via tradition, law, etc. We *can* improve productivity in education, construction, art, but we don't (at least in part!)

Toronto

# Metrolinx doing 'late-stage' work on Eglinton Crosstown LRT but offers no new opening date

Board of directors receives update on major transit projects in GTA

Muriel Draaisma · CBC News · Posted: Dec 01, 2022 9:00 PM EST | Last Updated: December 1, 2022



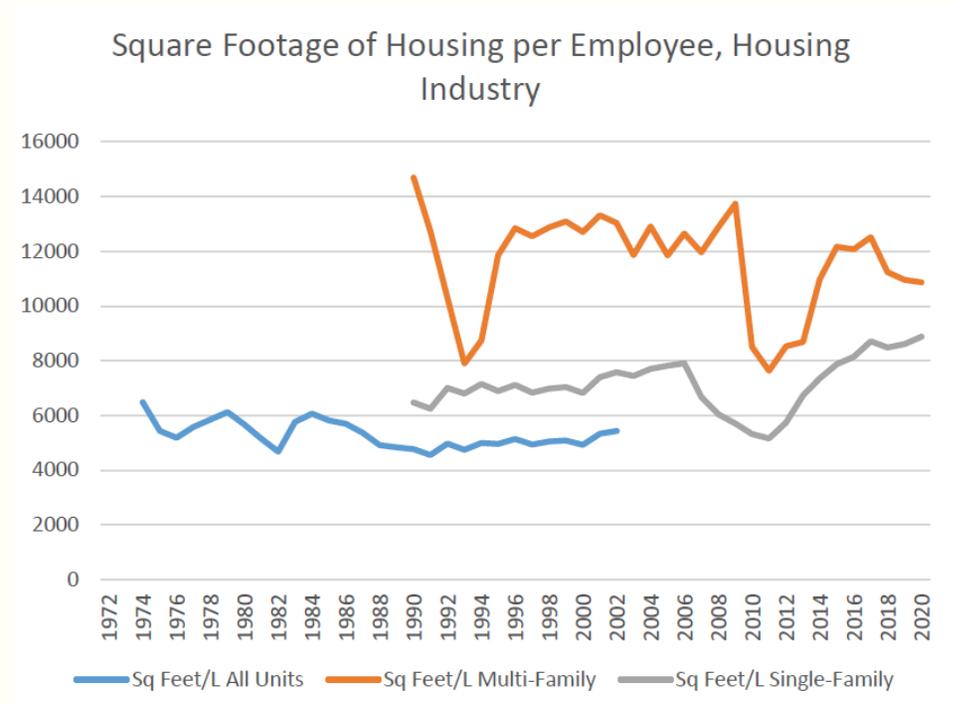
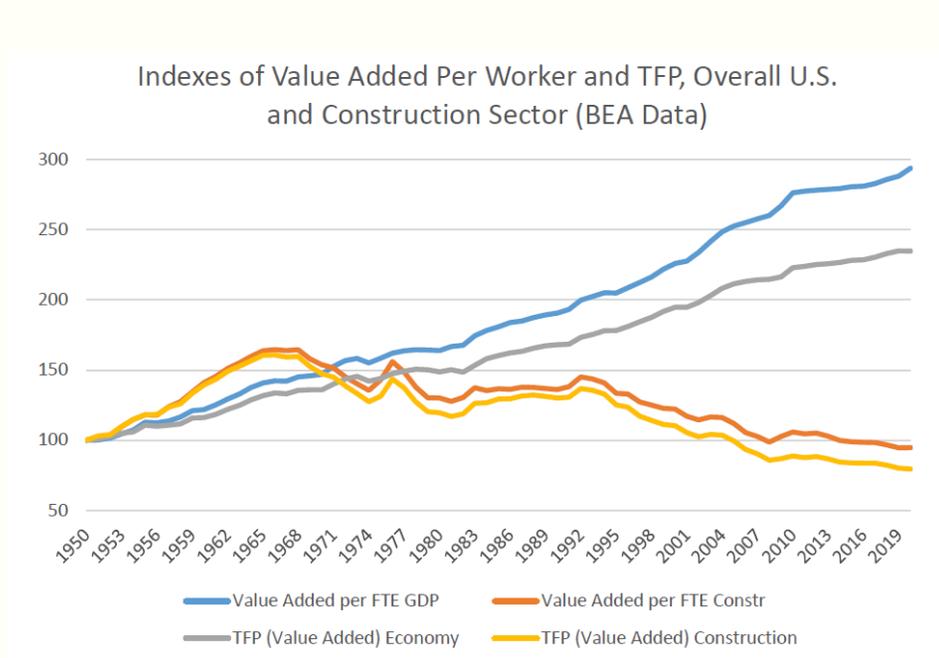
Listen to this article ⓘ

Estimated 5 minutes



Note this article is three years old, and we're still doing "final tests"! Why are construction projects getting so slow? Whole subway lines used to be built at this rate.

# Don't blame (just) Metrolinx!



Goolsbee and Syverson 2023. And look at 1950 to 1970 vs. more recently! Barely any growth in sq ft per construction worker over time, overall construction sector productivity looks worse. During this period, look how fast manufacturing productivity grows! Question, to what extent are these slowdowns for good reason - better safety including things like air quality, worker safety, etc.

Some of the slowdown is related to a lack of economies of scale in some aspects of home production. Many building techniques would look identical to a home builder 50 years ago, so for Baumol reasons, building a brick wall might be *absolutely* more expensive, not just expensive relative to the trend of other manufactured goods. But the lack of scale economies is in part a policy choice (related to how we regulate new technology in building codes, city-specific zoning rules, etc.)

To be fair, are some things slow because they have to be, or unproductive because it's just harder?

What does "big improvement in productivity of democratic governance" mean? Of child care? Or making the elderly feel loved? How can we approve a drug that takes 5 years to show a statistical effect without running a five year trial?

It may be surprising how willing some folks are to accept big improvements here via, e.g., robots for elder care, or even AI-driven art like Xania Monet's new gospel song.

*[I]ts success was startling: development [in] 8 months of vaccines to counter a deadly coronavirus pandemic...vaccines OWS supported have been estimated to have saved some 3.2 million lives...in the USA alone by the summer of 2022*

Bonvillian, Operation Warp Speed

Let us turn to an industry that has seen slow productivity growth but also a very conscious mindset about risk: the pharma industry.

## ***Researchers Are Racing to Make a Coronavirus Vaccine. Will It Help?***

New technology and better coordination have sped up development. But a coronavirus vaccine is still months — and most likely years — away.

Share full article



28



It took researchers about 20 months to develop a vaccine against SARS in 2003 and six months when the Zika epidemic struck in 2015. Scientists want to halve that time with the coronavirus. Vido-Intervac, via Reuters

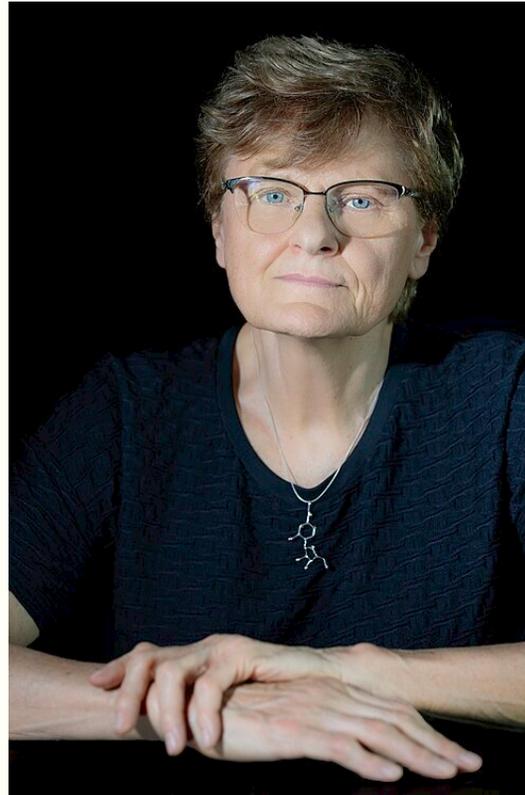
By **Knvul Sheikh** and **Katie Thomas**

Published Jan. 28, 2020 Updated June 10, 2020

*"It takes months and even years because the vaccines must undergo extensive testing in animals and humans. In the best case, it takes at least a year — and most likely longer — for any vaccine to become available to the public."*

OWS was not guaranteed from the start to be a success! That a vaccine would take years rather than 9 months was assumed in a lot of expert coverage.

# How so fast?



## DARPA Awards Moderna Therapeutics A Grant For Up To \$25 Million To Develop Messenger RNA Therapeutics™

Research to focus on antibody production for immune defense

NEWS PROVIDED BY  
Moderna Therapeutics →  
Oct 02, 2013, 08:00 ET



Moderna also raised a billion dollars of private money at this time. Kariko and Weissmann had spent decades working on this, as had a number of public and private companies (including Canada's Protiva). BARDA had broad authority to marshal resources for rapid vaccine production, was created in mid-2000s due to military worries about biothreats. Also, e.g., Other Transactions Authority allowed end-run around normal govt procurement. Basically, a lot of institutional architecture and basic science to make "progress" here already existed - trying to spin this up from scratch would have been very difficult.

# Pre-existing talent mattered



Health Sec Alex Azar was US head, Eli Lilly. Science lead Moncef Slaoui, head of GSK. Logistics lead Gustave Perna.

Slaoui was Immunology PhD, very concerned about vaccine threats. Perna was General in charge of US Army logistics. "Azar served as a shield, protecting it from White House, bureaucratic and legislative interference": skunkworks model? Without this background expertise in charge of OWS, could it have succeeded? Progress depends on a base of culture, knowledge, orgs, and incentives!

# Two timelines

## Covid vaccine Moderna mRNA-1273

- Genome ready 1/10/2020
- Animal study option 2/7
- First human trial 3/16
- Phase 3 7/27
- EUA approval 12/18

## Malaria vaccine BioNTech BNT165

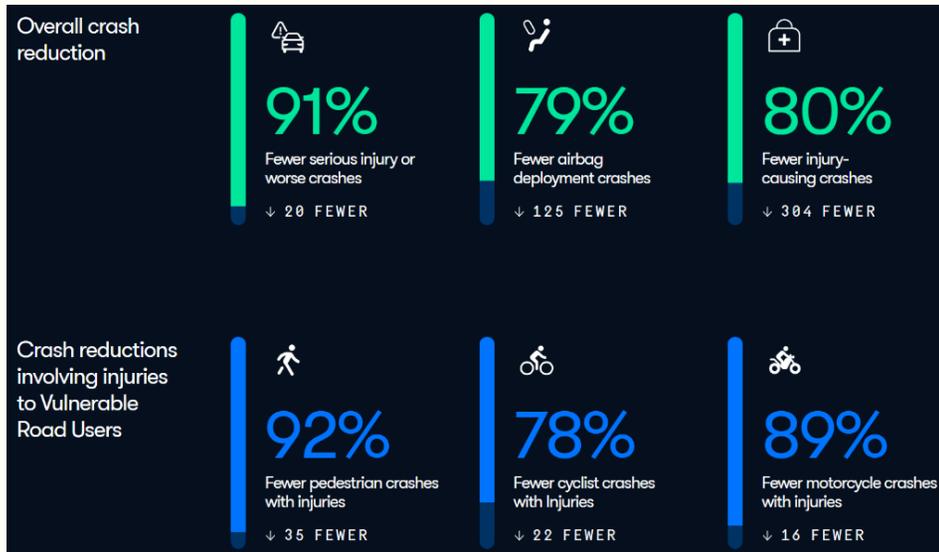
- Genome long known
- Vaccine dev from 7/26/2021
- First human trial 12/23/2022
- Phase 3 not started
- 600k will die this year, mostly kids

Note that difference in timelines. Why so different? It might just be that malaria is more difficult...but the timeline at each (successful so far!) stage suggests that we are following "normal regulation" here. That is, in an attempt to limit risks, we require substantial validation at each stage. But in the meantime, millions of kids will die. How should evaluate this speed v risk tradeoff?

*"[B]ecause of predictable features of human cognition, ordinary people deal poorly with the topic of risk."*

Sunstein, Laws of Fear

Sunstein notes that attitudes toward risk are often driven by predictable behavioral errors, especially loss aversion and availability bias.



1.2 million auto deaths/year globally, 40k in US, 2k in Canada

E.g., this particular "worry about risk" is just a crazy way of weighing costs and benefits.

Question: what is largest preventable-by-policy  
cause of death globally?

**NEWS** · NOV 19, 2018

# New Index finds air pollution reduces global life expectancy by nearly 2 years, making it the single greatest threat to human health

Loss of life expectancy is highest in Asia, exceeding 6 years in many parts of India and China; some residents of the United States still lose up to a year of life from pollution.

---

Specifically, people significantly overestimated highly publicized causes of death, including tornadoes, cancer, botulism, and homicide. By contrast, they underestimated the number of deaths from stroke, asthma, emphysema, and diabetes - Sunstein

# The Precautionary Principle

VS

## Cost-Benefit Analysis

*"When an activity raises threats of harm to human health or the environment, precautionary measures should be taken even if some cause and effect relationships are not fully established scientifically. In this context the proponent of an activity, rather than the public, should bear the burden of proof."*

Wingspread definition

Importantly, precautionary principle treats new things asymmetrically from old things. There may be reasons for this - recall Burke's argument from a prior class that laws and norms "summarize" past arguments whose data we have long forgotten. But it is not a cost-benefit analysis in the traditional sense.

## Arguments for Precaution:

- Irreversible changes/tipping points
- Impossible to measure cost/benefit
  - Care about the future

If species may go extinct (including us!), then the possibility of downside risk must be taken seriously. Some precautionary followers also believe that cost/benefit of new things is so uncertain that the numbers are effectively fake, and that a lot of cost/benefit analysis does not seriously account for the preferences of those not yet born.

Precaution is literally part of existing law in many places, esp. environmental law (e.g., in European Union due to Lisbon Treaty), despite serious objections from decision theorists and many philosophers.

## Arguments against precaution:

- Cost/benefit weighing is rational
  - Why privilege status quo?
- Why privilege "available" risks?
  - Practically, fewer fixes

On rationality, in many decision theoretic models cost/benefit analysis (where *all* costs and externalities are actually included!) is the *implication* of rationality, not an assumption. In particular, it avoids an asymmetry. E.g., if all cars were self-driving, precaution would not allow you to drive your car yourself. If all cars were driven by humans, precaution suggests being very slow to approve self-driving. But why does the change in the status quo affect the optimal regulation/social acceptance of a technology (of course, empirically, as Sunstein mentions, people *do* have loss aversion and availability bias).

Practically, Sunstein also argues that our ability to solve problems requires us to be able to weigh the costs and benefits of new innovations. Without it, we are stuck in a status quo morass.



## *The ALARA Principle*

The "ALARA" - as low as reasonably achievable - principle is legally, in US nuclear code, "making every reasonable effort to keep exposures as far below dose limits as practical, consistent with the purpose of the activity, and taking into account technology, economics, and societal factors, and the utilization of nuclear energy in the public interest." In practice, this often means ever-increasing stringency about radiation dosing without explicit cost/benefit analysis.

This make nuclear more expensive to build -> makes coal and oil and gas-generated electricity remain -> creates air pollution -> kills thousands of people diffusely. This is precisely the worry the "cost/benefit" fans about precaution. On the other hand, of course, you might say "we should never have split the atom - though nuclear energy is valuable, we put ourselves at serious risk, and still do, of thermonuclear war and hence human extinction." Not a crazy stance!

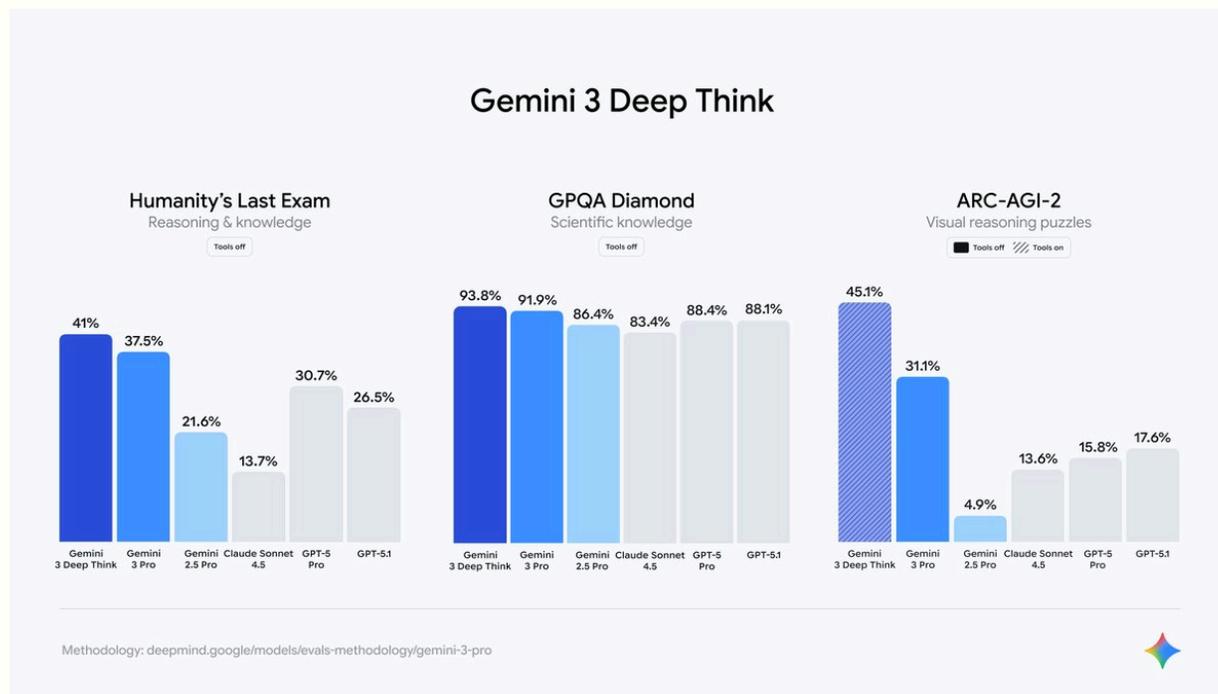
# You can see the future first in San Francisco....

*A few years ago, these people were derided as crazy—but they trusted the trendlines, which allowed them to correctly predict the AI advances of the past few years. Whether these people are also right about the next few years remains to be seen. But these are very smart people—the smartest people I have ever met—and they are the ones building this technology...Let me tell you what we see.*

Aschenbrenner, Situational Awareness, **June 2024**

This essay was written in June 2024. If anything, AI has improved *more quickly* than Leo anticipated in the past 18 months.



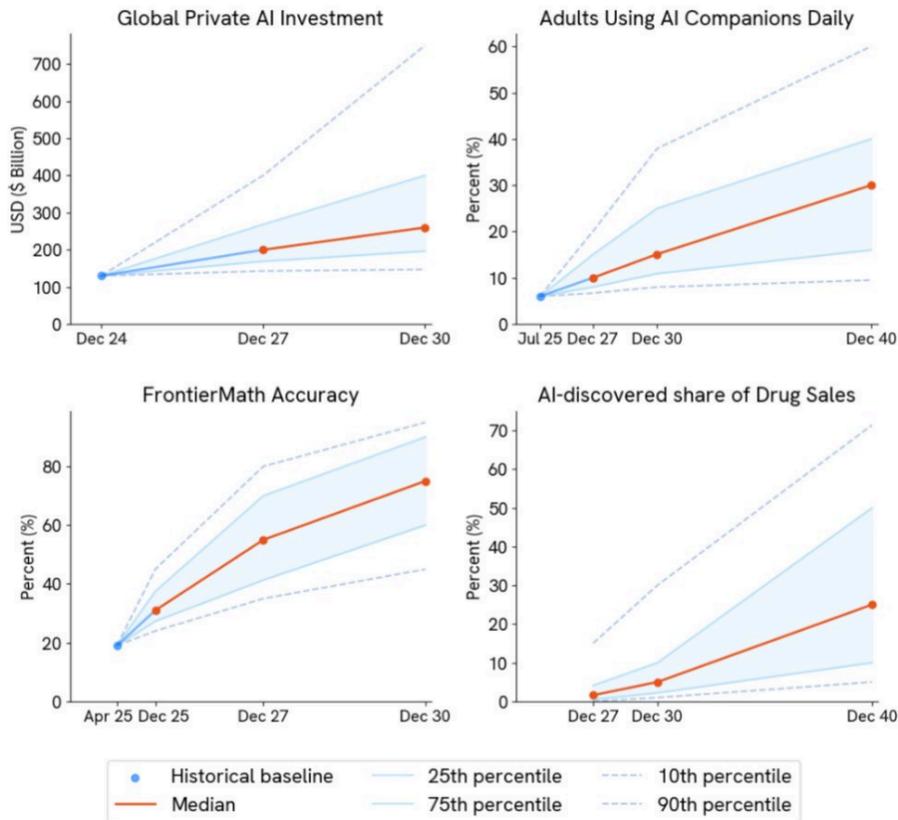


Aschenbrenner: "Claude 3 Opus currently gets ~60%, compared to in-domain PhDs who get ~80%—and I expect this benchmark to fall as well, in the next generation or two."

And now we're well into the 90s. Again, 18 months later. This is true of almost any benchmark you want.

# What are 'scaling laws' and what do they imply?

As compute power during training, training data, and compute power at inference time improve, performance of LLMs on standard benchmarks improves in predictable ways. This of course won't hold forever - but it has held for five years now. And that's without needing major conceptual breakthroughs - which we might get, particularly via self-improving AI.



### Expert Sampling

We target four expert communities. First, we include computer scientists researching topics in AI by including top-cited authors, stratified by age, and the authors of the top-rated papers at leading AI and ML conferences. Second, we identify leading economists, both across fields and within the subfield of economics focused on the economic effects of AI and new technology. We include top-cited authors of papers on AI and technology, members of the U.S. Economic Experts Panel (Clark Center 2025), and attendees of economics conferences on AI. Third, we include industry professionals, identified via their contributions to frontier models or employment at AI-related companies with extensive fundraising. Fourth, we identify institutions leading the discussion on AI development, policy, and impacts and invite research staff.

We sample from two other sources and sort them into one of the four communities above. First, we invited the honorees from TIME's 100 Most Influential People in AI in 2023 and 2024 (Barker Bonomo and Javed 2024). Second, we allowed invited respondents to recommend other qualified candidates for the survey, yielding 172 additional invitees. In order to filter this group, we require that an individual:

- meets the requirements of another sampling category;
- has over 1,000 academic citations; or
- has over 300 academic citations if a PhD student or postgraduate researcher.

These requirements excluded only 7 of the recommended candidates that ultimately enrolled. After exclusion, the referred group has 11.6 years of experience on average and 75% have a postgraduate degree. Like other expert sample expansions, referred contacts are not included in frame targets, but their responses do receive positive weights through the reweighting process.

The LEAP AI panel asks computer scientists, social scientists working on AI policy, and researchers in top private sector AI labs to make predictions about AI capabilities (I'm part of the panel). Note that the bottom left - performance on Frontier Math. This survey is a month old and we are almost certainly underpredicting performance on that measure already.

**NEW YORK TIMES BESTSELLER**

**IF ANYONE  
BUILDS IT,  
EVERYONE  
DIES**

**WHY  
SUPERHUMAN AI  
WOULD  
KILL US ALL**

**ELIEZER  
YUDKOWSKY &  
NATE SOARES**

So if AI improves such that it can largely invent the next frontier model, and then invent improvements to sensors and actuators and batteries for robots, and do so in a way that is above human intelligence: is this too dangerous? Should we allow it? How do balance "what if the AI turns us all into goo" against "what if the AI cures cancer"?

Bottlenecks to progress due to tradition, avoidance of risk, etc - maybe necessary in some cases, but we should be very careful that we are not just subject to status quo bias and other factors Sunstein discusses in "Laws of Fear".

So, a survey. You are asked by a friend (or policymaker!) "how important is fixing problems in slow-growing sectors (perhaps due to attempts to limit risk) to progress?"

Let's vote: 1-5, where 1 is "not at all" and 5 is "I imagine the most important factor we'll discuss this term". We'll keep track of these votes as we go!

3.48/5. Student responses: 2/5 "As long as the other factors we discussed work, someone will find a way around these barriers." 5/5 "But finding a way around these problems is still important".

# Next Week

## *Big Projects*

When do top-down attempts to create progress work? How did Asia rise?  
What exactly does DARPA do? What can we learn from all this?

