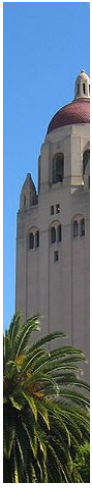

A.I. and Our Economic Future

Chad Jones

Stanford GSB Spring Reunion

May 1, 2026

A.I. and Our Economic Future

- A.I. is likely to be the most transformative technology in our lifetime
  - Latest in a line: electricity, semiconductors, the internet
- What if machines — A.I. for cognitive work and A.I. plus robots for physical work — can perform every task a human can do?
- Two scenarios:
  - A.I. accelerates economic growth
  - A.I. is “business as usual”

Scenario 1: A.I. dramatically accelerates economic growth

- Near-term productivity boosts from A.I.
  - Software: Claude Opus 4.5 performs better than any human on Anthropic's two hour coding exam
  - In the next decade: A.I. agents that can automate most coding?
  - Virtuous circle: better algorithms and A.I. agents = virtual remote workers

Scenario 1: A.I. dramatically accelerates economic growth

- Near-term productivity boosts from A.I.
  - Software: Claude Opus 4.5 performs better than any human on Anthropic's two hour coding exam
  - In the next decade: A.I. agents that can automate most coding?
  - Virtuous circle: better algorithms and A.I. agents = virtual remote workers
- Billions of virtual research assistants, running 100x faster than us — a "country of geniuses in a data center" (Dario Amodei)
  - Automate most cognitive tasks $\Rightarrow$ invent new ideas
  - E.g. better chips, robots, medical technologies, etc.
  - A.I. + robots $\Rightarrow$ automate physical tasks
- Potential to raise growth rates substantially over the next 25 years?

Scenario 2: A.I. as "Business as Usual"

- Automation has been going on for 150 years with no speed up in growth
  - Electricity, engines, semiconductors, the internet, smartphones
  - Yet growth always 2% per year
- Maybe those great ideas are what *kept* growth from slowing
  - Perhaps A.I. = latest great idea letting us maintain 2% growth for a while longer. (pessimistic view, but possible)
- Economic history $\Rightarrow$ may take longer than we expect
  - Electricity and computers changed the economy over 50 years

Average income per person in the U.S.

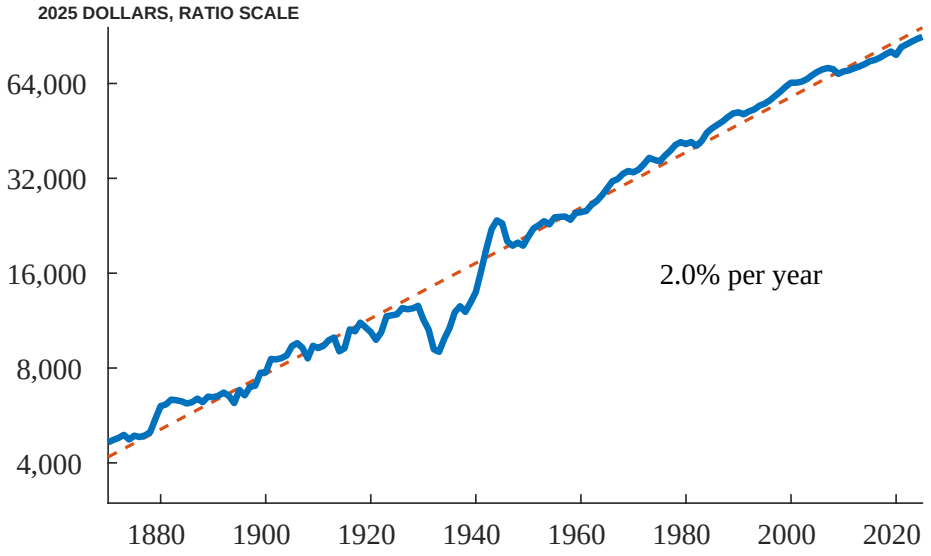

Weak Links

- Business success requires completing many tasks
  - iPhone: Design/Innovation - Manufacturing - Marketing - Retail - Customer Mgt
  - Failing at any step can reduce value considerably
  - Examples: the space shuttle Challenger's O-ring or Covid-19 supply chain issues
- "A chain is only as strong as its weakest link"
  - Making 17/20 links infinitely strong can have modest effects — bottlenecked by remaining weak links
  - 100m $\times$ more transistors in my pocket than in the early 1970s...
- Weak links are the source of scarcity $\Rightarrow$ earn high returns

What has happened to the “computer income” share of GDP?

What has happened to the “computer income” share of GDP?

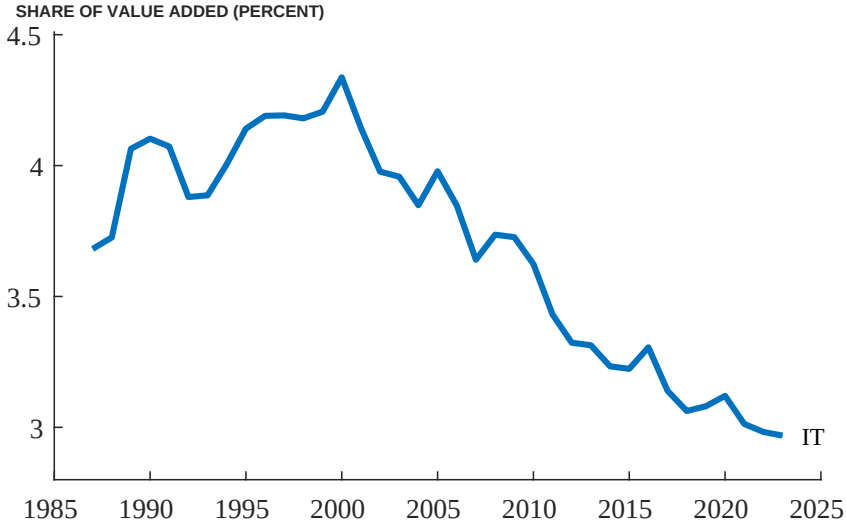

Computers are everywhere, but rapid price declines dominate — weak links

A Model to Formalize these Insights (Jones and Tonetti, 2026)

- Model
  - Ideas are the source of long-run economic growth
  - Production of goods and ideas involves weak links
  - Automation of both goods and idea production occurs endogenously over time
  - Calibrate to historical U.S. data since the 1950s
- Example: What if we had infinite software today?
  - Because of weak links, infinite amounts of a task raise GDP by that task's share of GDP.
  - Software is 2% of GDP, so we'd be just 2% richer with infinite software!

Simulating our Economic Future

- Model features both of the key ingredients from the two scenarios
  - Positive feedback: automation $\rightarrow$ new ideas $\rightarrow$ more automation
  - Weak links limit the effect of automation
- Two sets of simulations
  - A.I. as a continuation of historical patterns of automation
  - A.I. as a break with past: Moore's Law applies everywhere

The Future if AI = Continuing the Past: Capital Share

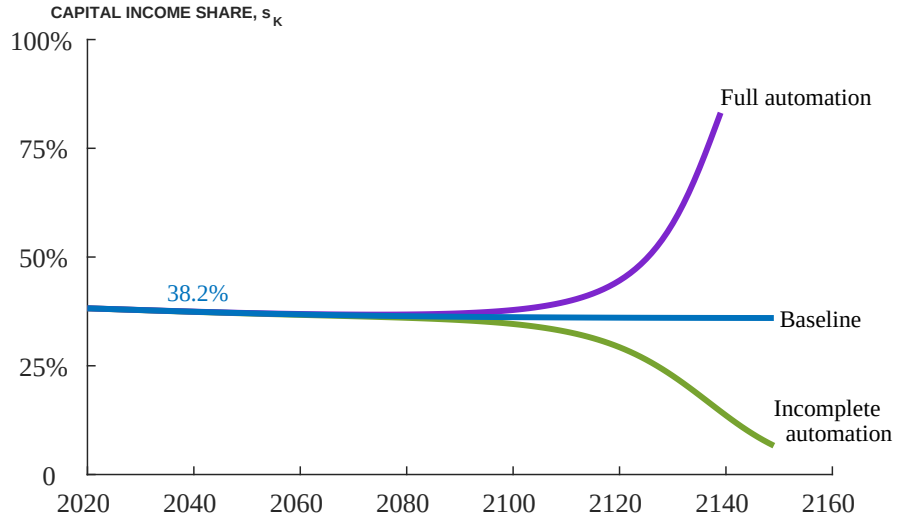

The Future if AI = Continuing the Past: Economic Growth

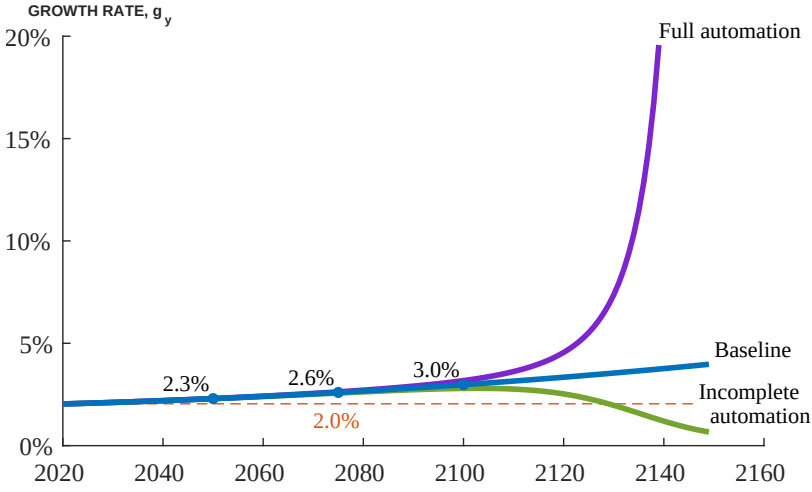

The Future if AI = Continuing the Past: GDP per Person

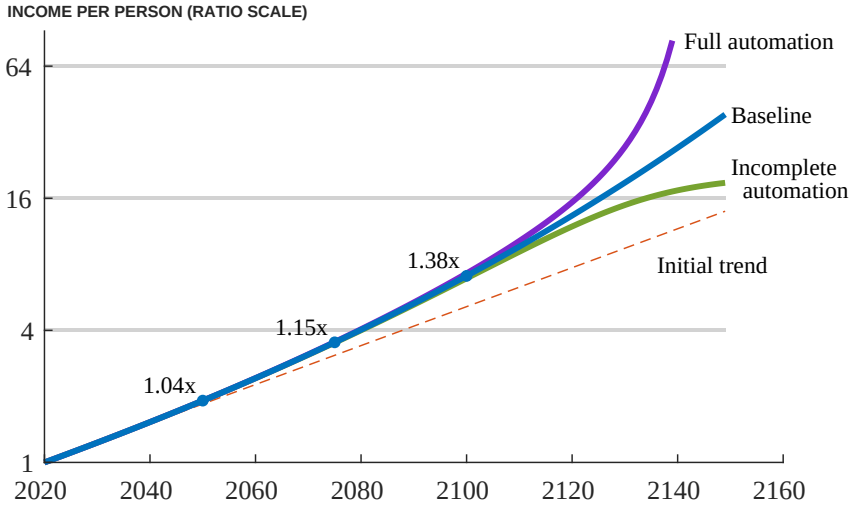

The Future if AI = Moore's Law Everywhere: Capital Share

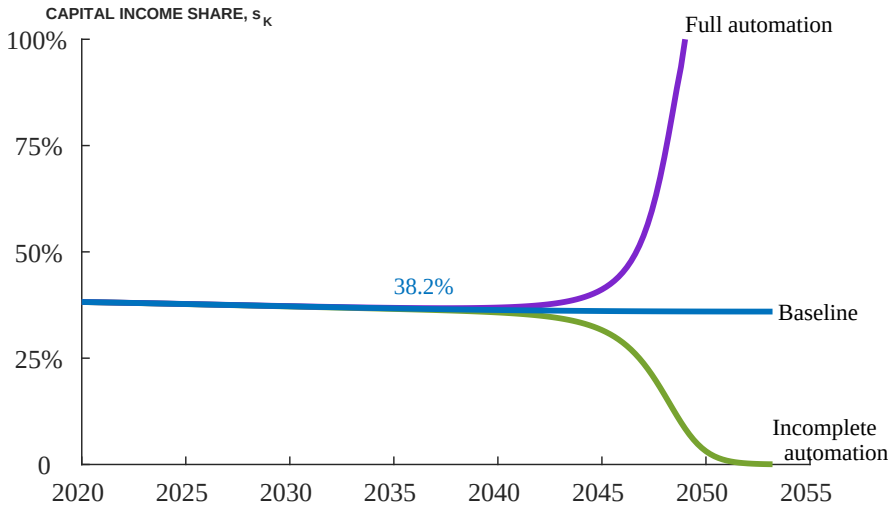

The Future if AI = Moore's Law Everywhere: Economic Growth

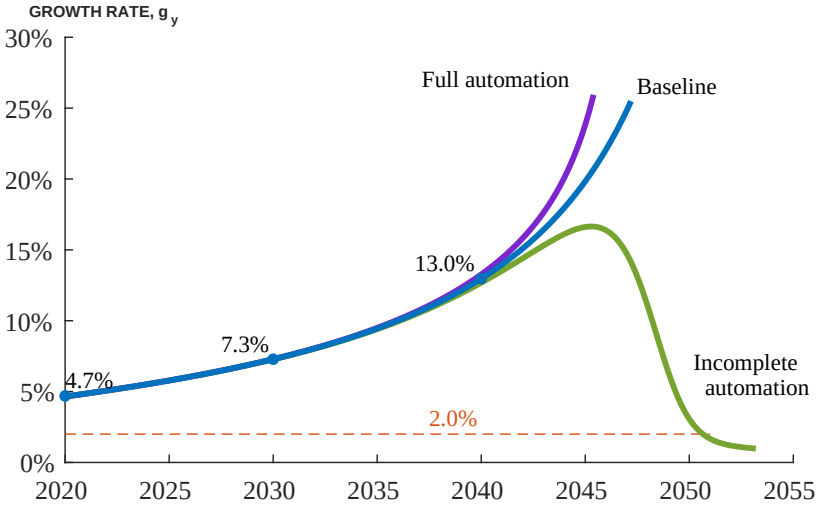

The Future if AI = Moore's Law Everywhere: GDP per Person

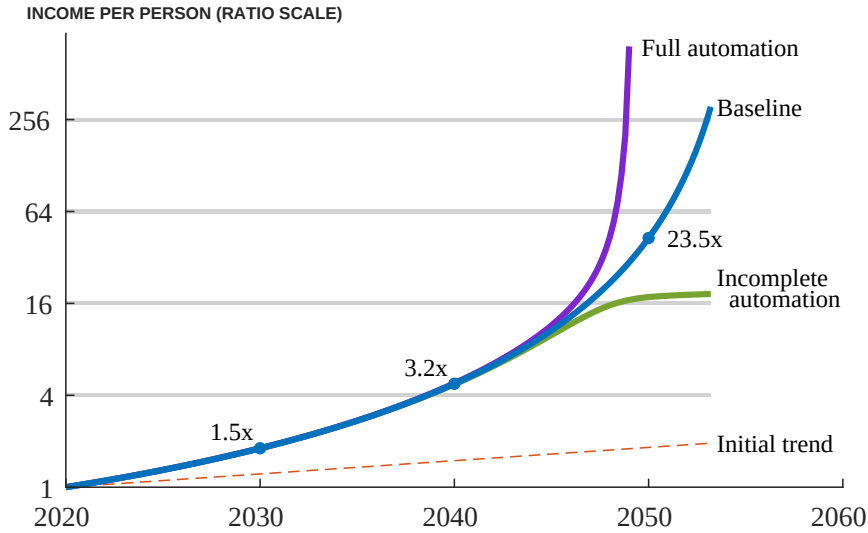

Summing Up

- Growth could indeed explode — but slowly!
  - Weak links make the transformation gradual
  - But as weak links are automated away, flywheel effects dominate
- Example: Self-driving cars — 20+ years since DARPA 2005

Jobs

- Geoff Hinton in 2016: We should stop training radiologists
  - But more radiologists today and higher pay!
- How to understand? Weak links!
  - Jobs are collections of complementary tasks — radiologists do more than just read scans
  - "Weak links" $\Rightarrow$ Automating 75% of tasks can raise wages — radiologists
  - Some jobs: A.I. automates all tasks $\Rightarrow$ wages go down — Uber drivers
- A.I. has nuanced effects on jobs and the labor market

Inequality and Meaningful Work

- Historically, labor is the main asset that many people trade to consume.
  - Could change in the future?
- The world where A.I. “changes everything” is a world where GDP is incredibly high
  - The size of the pie available for redistribution is enormous
  - Rich countries already engage in lots of redistribution, but more may be needed?
  - Transition hard?
- As we get richer, we naturally work less — this is a good thing!
- But there is also good, meaningful work
  - We may choose to value experiences involving people (arts, music, sports)
  - Retirement!

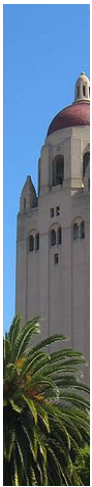

Catastrophic Risks?

Can we use economic analysis to think about the serious risks?

Two Versions of Existential Risk

- Bad actors:
  - Could use ChatGPT-8 / Opus 7 to cause harm
  - E.g. design a virus that is more lethal than Ebola and takes 3 months for symptoms
  - Nuclear weapons mangeable because so rare; if every person had them...
- Alien intelligence:
  - How would we react to a spaceship near Pluto on the way to Earth?
  - “How do we retain power over entities more powerful than us, forever?” (Stuart Russell)

How much should we spend to reduce A.I.'s catastrophic risk? (Jones 2025)

How much should we spend to reduce A.I.'s catastrophic risk? (Jones 2025)

- Covid pandemic: "spent" 4% of GDP to mitigate a mortality risk of 0.3%
  - A.I. risk is at least this large $\Rightarrow$ spend at least this much?
  - Are we massively underinvesting in mitigating this risk?

How much should we spend to reduce A.I.'s catastrophic risk? (Jones 2025)

- Covid pandemic: "spent" 4% of GDP to mitigate a mortality risk of 0.3%
  - A.I. risk is at least this large $\Rightarrow$ spend at least this much?
  - Are we massively underinvesting in mitigating this risk?
- Better intuition
  - VSL = \$10 million
  - To avoid a mortality risk of 1% $\Rightarrow$ WTP = $1\% \times \$10$ million $= \$100,000$
  - This is more than 100% of a year's per capita GDP
  - Xrisk over two decades $\Rightarrow$ **annual investment of 5% of GDP**
  - Large investments worthwhile, even with no value on future generations

How much should we spend to reduce A.I.'s catastrophic risk? (Jones 2025)

- Covid pandemic: "spent" 4% of GDP to mitigate a mortality risk of 0.3%
  - A.I. risk is at least this large $\Rightarrow$ spend at least this much?
  - Are we massively underinvesting in mitigating this risk?
- Better intuition
  - VSL = \$10 million
  - To avoid a mortality risk of 1% $\Rightarrow$ WTP = $1\% \times \$10$ million $= \$100,000$
  - This is more than 100% of a year's per capita GDP
  - Xrisk over two decades $\Rightarrow$ **annual investment of 5% of GDP**
  - Large investments worthwhile, even with no value on future generations

Incomplete: ignores the "effectiveness" of mitigation, but correct intuition; see paper.

Other Safety Considerations

- Easy to justify spending 1/3 of 1% of US GDP on safety = $100 billion!
- Weak links and safety
  - Weak links $\Rightarrow$ large upside benefits arrive slowly
  - ... but harms can come quickly!
  - Mythos and automating software engineering
- Exernalities and race dynamics: A.I. labs do not internalize the risks to all of us
- Should we tax GPUs and use the revenue to subsidize safety?

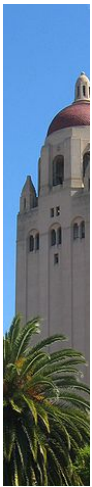

Final Thoughts

Final Thoughts

- How much did the internet change the world between 1990 and 2020?
  - How much will A.I. change things between 2015 and 2045? More or less?
  - I believe the answer is much more
  - Just because changes take 30 years instead of 5 years does not mean that the ultimate effects will not be large

- Downside risks could come sooner

We should use the intervening years to prepare for the changes

Talk based on material from several papers

- Aghion, B. Jones, and C. Jones (2019) “Artificial Intelligence and Economic Growth”
- Jones (2024 AER Insights) “The A.I. Dilemma: Growth versus Existential Risk”
- Jones (2025) “How much should we spend to reduce A.I.’s existential risk?”
- Jones and Tonetti (2026) “Past Automation and Future A.I.”
- Jones (2026) “A.I. and Our Economic Future” (for *JEP*)