

Using Artificial Intelligence Well:

*Thinking **Critically** About Technology **Predictions** In Periods of **Uncertainty***

Harry Surden

Professor of Law, University of Colorado Law School
Faculty Director, Silicon Flatirons Artificial Intelligence Initiative
Associate Director, Stanford CodeX Center

Thank you



Silicon Flatirons

- **Thank you**
 - **Silicon Flatirons Staff and Faculty Directors**
 - **Silicon Flatirons Supporters and Community**
 - **CU Law School Dean, Faculty and Staff**
 - **Expert Panelists**
 - **Students**

- **All of you!**



Outline

- Overview of Artificial Intelligence Changes: Three Years In
- Uncertainty, Hope, and Anxiety about AI Predictions
- Why detailed predictions about technology and society tend to be **unreliable**
- Frameworks from Labor Economics on AI and Jobs
- How to Nudge AI in a Positive Direction

**Where We Are Today
With
Artificial Intelligence**

Taking Stock of Modern AI Changes: Three Years In (2023–2026)

•Capabilities

- Models shifted from simple text generation to reasoning, coding, and tool-using systems
- Emergence of agentic workflows that can plan and execute multi-step tasks

•Performance

- Major improvements in coding, math reasoning, and scientific problem solving
 - Many early benchmarks became quickly saturated
- New, harder benchmarks introduced to measure reasoning ability
- Rough rule of thumb: ~10% improvement per year on many knowledge tasks (compounding)

•Ecosystem

- Open-source models rapidly caught up with frontier models
- Gap between frontier and open models narrowed to roughly 3–6 months

Taking Stock of Modern AI Changes: Three Years In (2023–2026)

•Economics

- Cost for equivalent performance fell roughly 50–100× between 2023 and 2025
- AI capabilities are becoming dramatically cheaper and more accessible

•What Has Changed — and What Hasn't

- Performance progress has been steady and significant, but not uniformly explosive; jagged intelligence
- Major societal impacts remain uncertain and difficult to predict
- The pattern fits the historical norm: technology runs ahead of social adaptation

•Looking Ahead

- The trend predictions are relatively reliable: more capability, lower cost
- The labor and societal outcome predictions remain deeply uncertain

The Current Moment

Excitement

- Real capability gains
- Genuine productivity improvements
- Scientific breakthroughs — drug discovery, protein folding

Anxiety

- Job loss and displacement
- Social disruption
- Deep uncertainty about the future

But both are often driven by unreliable predictions that deserve scrutiny.

**Predictions
Reliability
And Skepticism**

Confident Predictions about AI

- Every week confident headlines and predictions about AI abundance

DeepMind CEO Demis Hassabis claims AI could cure all diseases within 10 years

Jamie Dimon just made a bold prediction about AI and your job

“Your children are going to live to 100 and not have cancer because of technology,” Dimon said, “and literally they’ll probably be working three and a half days a week.” He called it “a wonderful thing,” but he was equally direct about what comes before that future arrives.

Confident Predictions about AI

- Every week confident headlines and predictions about AI replacing jobs

Column / Behind the Curtain

Behind the Curtain: A white-collar bloodbath

Dario Amodei — CEO of Anthropic, one of the world's most powerful creators of [artificial intelligence](#) — has a blunt, scary warning for the U.S. government and all of us:

- AI could wipe out *half* of all entry-level white-collar jobs — and spike unemployment to 10-20% in the next one to five years, Amodei told us in an interview from his San Francisco

Bill Gates Says AI Will Replace Doctors, Teachers and More in Next 10 Years, Making Humans Unnecessary 'for Most Things'

Anxiety about AI and Jobs

- Every week confident headlines and predictions about AI replacing jobs

LATEST NEWS


Will AI Render Lawyers Obsolete?

SUCCESS · THE FUTURE OF WORK

100 million jobs could be wiped out from the U.S. alone thanks to AI, warns Senator Bernie Sanders

BY PRESTON FORE
SUCCESS REPORTER

October 7 2025 at 11:13 AM EDT



How Does a Researcher Approach This?

- **How Does a Researcher Approach This?**
 - Should we take confident AI predictions at face value as likely outcomes?
- **Research suggests caution**
 - We should be **skeptical**
 - No matter **how confidently they are stated**
 - Or **who is making them**
 - Politicians, business leaders, technologists, Nobel prize winners
- **Why?**
 - Because detailed predictions about technology tend to be **very unreliable**
 - There are simply too many factors for reliable, confident detailed predictions about impact

Looking at historical evidence

- **As researchers**

- The best we can do is look back at past predictions
 - And compare them to what actually happened
- Assess how reliable they were and what went right or wrong
- And assess whether the level of confidence was warranted
 - Or if the confidence of the predictions did not match the uncertainty

A Practical Toolset for Evaluating Technology Predictions

- **This talk *does not offer promises of certainty***
 - Not a prediction about the future of AI, Jobs, and Society
 - That future is too complex for *anybody* to forecast in detail or with reliability
- **What this talk *does offer***
 - A practical toolset for critically evaluating technology predictions
 - This can reduce anxiety caused by extreme, but confident claims about AI
 - Can put AI and other technology predictions in perspective

Two Very Different Types of Technology Predictions

Two Different Types of Technology Predictions

- **Predictions about *Technological Trends* (1 - 3 years out)**
 - Can be reliable
- **Predictions about *Impact* of Trends on Society, Labor, Politics, etc**
 - Notoriously *unreliable*

Many confident AI predictions conflate these two

Trend vs. Detailed Impact Predictions

Predictions about *Technological Trends*

- Can be reliable when:
 - Time horizon is short (1 – 3 years)
 - Grounded in data, engineering constraints, current research, or physical phenomena
 - Examples: Moore's Law, GPU Compute Growth, AI Performance and Cost

Predictions about *Impacts of those Trends*

- Details about future effects on jobs, society, politics, institutions
- Often *unreliable* because outcomes depend upon too many factors
 - Interacting systems: labor markets, economics, social institutions, politics, psychology
 - Science and Technology itself develop in unexpected ways
- The world is simply *too complex* to make **confident, detailed** predictions
 - About future long-term technology impacts on society

Even when the Technological Premise is Right The Societal Inference is often Wrong

- **The Prediction Pattern :**
 - Asserting a likely technological *trend* will result in particular societal *outcomes*
- **How the predictive error works:**
 - Take a highly reliable technological trend (e.g., compounding AI compute)
 - Erroneously apply *overly simplistic intuitive model of the world*
 - Reach a highly confident, detailed societal conclusion that is not justified
- **The Result:** Predictions about impacts both positive and negative
 - That are notoriously unreliable
 - Ignores complexities of the way jobs, economics, society, psychology interact

**Examples
of Prediction Errors
And What We Can
Learn**

Case Study: An Optimistic Prediction

- **John Maynard Keynes**
 - Great economist who in 1930 wrote *Economic Possibilities for Our Grandchildren*
- **Correct Prediction: The Technological Trend**
 - Observed ~2% annual compounding growth in productivity (informed by data).
 - Predicted large increases in GDP and living standards by 2030
 - Predicted technology would reduce labor needed for production
- **Incorrect Prediction: The Societal Impact**
 - Predicted a 15-hour universal workweek by today
 - Predicted a society largely devoted to leisure
 - All material needs would be permanently solved
- **Why?**
 - Had an overly simplistic mental model of impact on society and labor market
 - Assumed productivity gains would translate directly into leisure, evenly distributed
 - The lesson: He correctly predicted the trend, but not a detailed societal outcome

Case Study: A Pessimistic Prediction

- **Geoffrey Hinton**
 - Father of Modern Artificial Intelligence. Won the Nobel Prize in Physics
- **In 2016, Hinton suggested that AI would replace radiologists:**
 - *"People should stop training radiologists now, because within 5 years deep learning will do better."* – Hinton (2016)
- **His intuitive (but flawed) reasoning:**
 - Radiologists examine medical images.
 - AI will be better than humans at examining medical imaging.
 - Therefore: AI will replace radiologists.
- **But in 2026, the exact opposite has happened:**
 - True: AI is better than radiologists at spotting many conditions.
 - Reality: Demand for radiologists has *increased*.
 - Pay is increasing, and there are actually widespread shortages.

What went wrong? Task/Job Confusion

- **Problem 1: Confusing a *single job's task* for the *entire job***
 - Economists view jobs not as a single unit, but as a “bundle of tasks” (Acemoglu & Autor, 2011)
 - Radiologist (some tasks)
 - Patient Interaction: diagnosis, reviewing histories, imaging tests, explaining results,
 - Image Interpretation: detecting, characterizing, and comparing findings on scans, etc.
 - Communication: drafting reports, flagging urgent results, consulting with physicians or patients, etc..
- The ATM Fallacy (An Overly Simple Mental Model)
 - *Prediction*: ATMs will replace bank tellers (assuming their only job is handing out cash).
 - *Reality*: Teller employment actually *increased*; they shifted to other valuable relationship tasks.
 - *The Lesson*: Like tellers, radiologists do much more than their most obvious task.
- Similar problem with simple predictive mental model of radiologists
 - Radiologists are more than simple image readers; they diagnose and help patients through illness

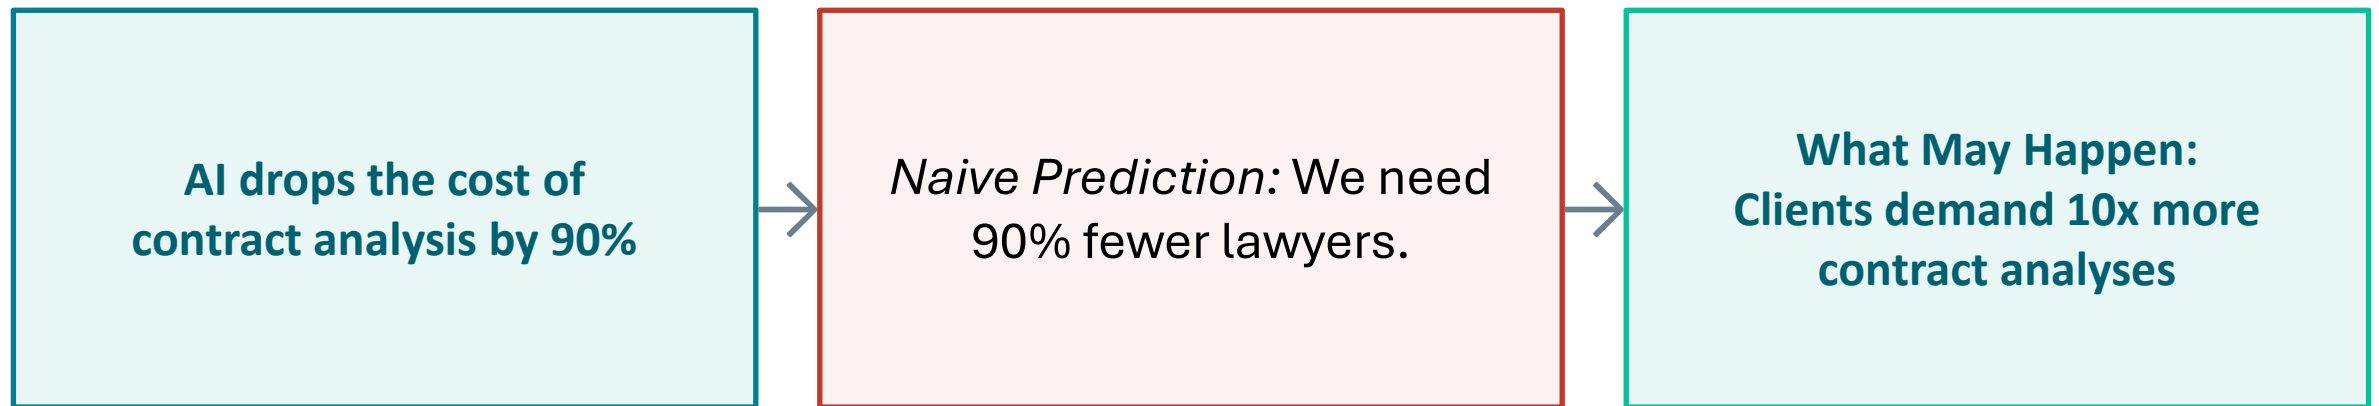
Substitutes and Complements

Problem 2: Technology Can *Complement* not Just *Substitute* a Task

- Default mental model is that technology *replaces tasks*
 - *This does happen in some cases (e.g., telephone switchboard operators, ditch diggers).*
- But in *many other skilled cases*, technology can **complement a task**
 - Making that task and workers *more valuable and more in demand*
- This may be part of what is happening in radiology
 - AI may be making radiologists better at their job, raising their value in diagnosing with patients
 - And increasing their productivity and quality and perhaps demand

Jevons Paradox: Efficiency Can Increase Demand

When technology makes something cheaper, consumers and organizations often use much more of it



- **Coal:** More efficient steam engines led to *more* coal consumption, not less.
- Cheaper legal document review → clients used more due diligence, not less
- Westlaw/Lexis → legal research got cheaper, legal work expanded
- Spreadsheet software -> more accounting and financial jobs, not fewer

The volume of work tends to expand to consume new efficiency gains.

Former Google CEO predicts AI will replace most programmers in a year

Apr 22, 2025 at 03:44 PM CDT

[Kalé Carey](#) (Anchor)



- Jevons Paradox suggests an alternative possibility
 - AI makes coding dramatically faster and cheaper
 - The cost to produce software falls significantly
 - Organizations build vastly more custom software
 - Demand for software engineers could actually *increase*

The Lump of Labor Fallacy

Lump of Labor Fallacy:

There is a **fixed amount of work** in the economy.

If AI performs a task, that's **one less job for a human**

AI and Machines Eventually Do Everything Leading to Mass Unemployment

Reality:

- Human wants and needs are highly elastic.
- When technology makes one task cheaper,
 - It tends to free up capital and generates entirely new categories of demand and work
 - new industries appear
 - entirely new kinds of jobs emerge

Historical Evidence

In 1900, 40% of Americans Worked in Agriculture

- Mechanization dramatically reduced the need for farm labor.
- Many feared mass unemployment and that jobs generally would disappear.

What Actually Happened:

Today agriculture employs only **~2% of the workforce**

- But instead of mass unemployment, new technologies created entirely new categories of jobs
 - Entirely new industries emerged
 - Total employment expanded dramatically
- Labor economists estimate that a large percentage of all jobs today did not exist in 1940.
 - Software engineer, video game design, digital animation, machine learning, etc
- Some argue that AI is different: in terms of generality, knowledge work, speed of change

Timing and Diffusion

- **Problem 3: The Timing of Technology is Hard to Predict**
 - Just because a technology like AI is starting to emerge and be used
 - Does not necessarily mean that significant labor or job impact is imminent
 - Technologies often take decades to have substantial labor market effects
 - Sometimes negative, sometimes positive, sometimes neutral
 - Electronic legal search Westlaw/Lexis was invented in the 1970s
 - Was not commonplace until around 2000, 30 years later
 - It often takes years or decades to affect overall supply/ demand or labor workflows

There are many other economic factors

- **Problem 4: Many economic factors besides technology affect jobs**
 - Technology like AI is just one component affecting jobs, tasks, and the labor force
 - Many other factors such as
 - Global supply and demand
 - Business cycles and trade
 - Changing tastes, preferences, and needs
 - Changing demographics
 - Politics and laws
 - Can influence the relative supply or demand for a job
 - The future emergence and direction of technologies like AI is hard to predict
 - Additionally, economic factors interact and are themselves hard to predict
 - The combined unpredictability makes confident predictions ***extremely unreliable***

Job Loss News Stories Can Be Misleading: Watch for the Incentives

BUSINESS - 5 MIN READ

Block lays off nearly half its staff because of AI. Its CEO said most companies will do the same

Do we know what actually caused these layoffs?

Watch for the Incentives Behind Predictions

- **Corporate Claims: AI as a Convenient Explanation**

- Companies may use "AI Impacts" as a PR shield to mask poor management decisions, over-hiring, or standard economic downturns.
- Be skeptical of corporate job-loss claims attributing everything to technology.
- We will only know once the data comes in and we do the analysis.

- **Media Claims: The Economics of Attention**

- Exaggerated promises of utopia or massive job losses due to AI get reader attention.
- Even highly trusted media like the NY Times operate on an attention economy.
- Nuance and steady transitions do not generate clicks; extreme predictions do.
- The media has an economic incentive to amplify the forecasting fallacies discussed.
- *Media rarely goes back to correct when an amplified prediction was wrong.*

Looking at historical evidence

- **What the research shows**

- Technology does change society and jobs at a broad level over time
 - While total employment often recovers, the disruption to specific workers and communities can be very real
- Detailed predictions about the impact of technology on labor force tend to be wrong
 - Predictions about: what specific occupations will decline or improve, when it will happen, etc;
- Due to the **inherent complexities and unpredictability of**
 - Emerging technology and the labor markets
- Rather, such predictions represent one of many possible outcomes
 - Some positive, some neutral, some negative
- Moreover, even seemingly intuitive and obvious predictions are often wrong
 - In some cases, the exact *opposite* of our intuitions turns out to be true

Sources: Bessen (2016), Autor (2015), Arntz, Gregory, and Zierahn (2021), Gimbel (2025)

Conclusion

- **High level, general trends no more than 3–5 years out are reasonable**
 - E.g. AI or computing will likely continue to improve year over year, and be broadly diffused
 - Some jobs will change, and some workers will likely be impacted negatively and positively
- But, **detailed** predictions about what those trends will mean for jobs or society
 - *Are highly unreliable and should be viewed with skepticism*
 - **The reason is that:**
 - The mental models used to make these predictions are often overly simplistic
 - And may ignore complex positive and negative labor market dynamics
 - More importantly,
 - The underlying reality about emerging technology and economics is just *too complex*

What Can We Do In Uncertainty?

- **Plan for multiple possible outcomes**
 - Be skeptical of confident predictions in both directions: demand evidence for extreme forecasts
 - We should plan for a wide “**menu**” of possible scenarios
 - Some positive, some neutral, some negative: Some workers will benefit, some will be disrupted
- **Measure what is actually happening**
 - Collect labor market data
 - Distinguish AI effects from normal economic changes
 - In past tech changes: ~2% of jobs disappeared, ~70% were neutral or helped, ~28% saw decreased demand
- **Adapt Policy based upon evidence**
 - Avoid premature conclusions
 - Isolate employment effects of AI from other factors
 - economic cycles, failing companies, and global competition
 - Adjust policies as real impacts emerge
 - Be realistic and transparent about limits and problems caused by new technologies like AI

Nudging AI In a Positive Direction

Build Pro-Social Institutional Norms and Governance Around AI

- **Help disrupted workers adapt**
 - Training and reskilling, transition support and safety nets
- Involve broad community input
- Build institutions that can **learn and adapt as technology and society changes**

Identify and Capture the benefits of AI

- Higher productivity
- New industries and opportunities
- Scientific Breakthroughs
- Help ensure benefits are widely and fairly distributed

Use AI to Actively Help Make the World Better

- Education, Healthcare, Scientific discovery, the Environment
- Focus on human and community flourishing and thriving
- Apply AI to improve specific problems of today