Today

- Macro Themes
  - Productivity and Jobs
  - Inequality
  - Innovation
  - Competition/Antitrust (IO)
  - Privacy, Liability, and Bias

- AI Strategy and Value Creation
  - The Task View
  - The System View

Definition

UNIVERSITY OF TORONTO

Musk vs Zuckerberg

Fei Fei Li

Founder of ImageNet, titan in modern computer vision

“AI is a ‘real deal’, … I firmly believe AI is as ‘real’ as computing, the Internet, renewable energy, new materials, etc.” — Fei Fei Li

What is Artificial Intelligence?

Daniel Kahneman

Nobel prize winner for behavioral economics

A well-known novelist wrote me some time ago that he's planning a novel. The novel is about a love triangle between two humans and a robot and what he wanted to know is how would the robot be different from the people. I proposed three main differences.

One is obvious: the robot will be much better at statistical reasoning and less enamored with stories and narratives than people are.

The other is that the robot would have much higher emotional intelligence.

The third is that the robot would be wiser… ***I do not think that there is very much that we can do that computers will not eventually [learn] to do***."

Kahneman (2018)

UNIVERSITY OF TORONTO

Geoff Hinton

Computer scientist. “Godfather” of deep learning.

“Take any old problem where you have to predict something and you have a lot of data, and deep learning is probably going to make it work better than the existing techniques.”

https://www.youtube.com/watch?v=2HMPRXstSvQ

Defining Artificial Intelligence

- Oxford English Dictionary definition: “the theory and development of computer systems able to perform tasks normally requiring human intelligence.”

- A moving target!

Agarwal, Gans, and Goldfarb

UNIVERSITY OF TORONTO

*The current excitement around AI is driven by advances in **prediction** technology… [more specifically] a drop in the cost of prediction.*

Prediction:
Using information that you do have to generate information that you don't have

Agarwal, Gans, and Goldfarb (2018)

Machine learning has become a "General Purpose Technology"

not Generative Pre-trained Transformer as in ChatGPT

Deep learning: traditional ML with "deep" neural networks (deep as in there are multiple layers in the network).

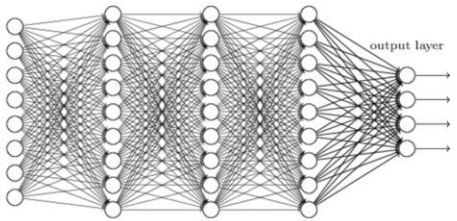

Figure 3: A five layer network, adapted from Nielsen [31].

Matt Taddy Takes it Further

UNIVERSITY OF TORONTO

AI = Domain Structure + Data Generation + General Purpose ML

Domain Structure	Data Generation	General Purpose ML
Business Expertise	Reinforcement Learning	Deep Neural Nets
Structural Econom[etr]ics	Big Data Assets	Video/Audio/Text
Relaxations and Heuristics	Sensor/Video Tracking	OOS + SGD + GPUs

Figure 1: AI systems are self-training structures of ML predictors that automate and accelerate human tasks.

ML can do fantastic things, but it is basically limited to predicting a future that looks mostly like the past… In contrast, an AI system is able to solve complex problems that have been previously reserved for humans.
[BUT] In contrast to earlier attempts at AI, the current instance of AI is *ML-driven*.

Matt Taddy Takes it Further

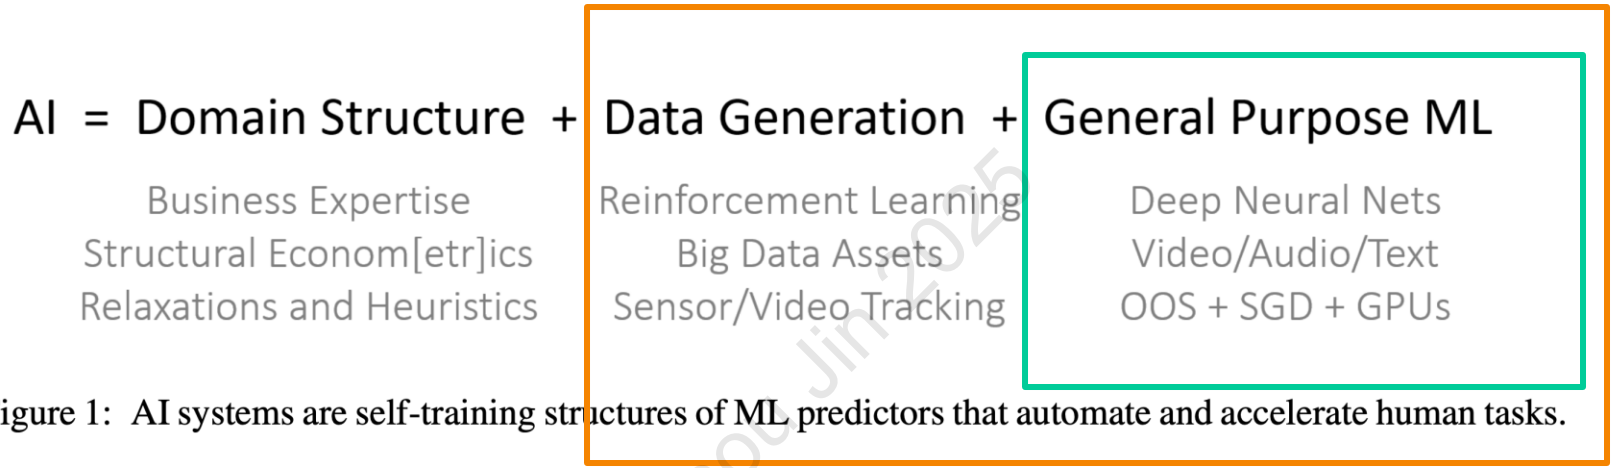

Figure 1: AI systems are self-training structures of ML predictors that automate and accelerate human tasks.

Reinforcement Learning: when ML GPT or DL algorithms actively choose the data that they consume.

Domain Structure: break a complex problem into composite tasks that can be solved with ML → defines the context of the problem

Matt Taddy Takes it Further

Statisticians care about model inference, while ML/AI cares about prediction

$$Y_i = \alpha + \beta_1 X_{1i} + \beta_2 X_{2i} + u_i$$

$$\widehat{Y} = \widehat{\alpha} + \widehat{\beta}_1 x_1 + \widehat{\beta}_2 x_2$$

model inference: beta hats

prediction: Y hat

Macro Theme 1: Productivity & Jobs

The pessimistic view

1. AI will replace jobs leaving little for humans to do.

2. AI will not be as influential as the technologies that diffused between 1870 and 1970.

"The good news is that those pessimistic predictions cannot both be right. The even better news is that they can both be wrong."—Joel Mokyr (Nobel 2025)

Assuming A.I. is productivity-enhancing…

"There are two separate questions: There's an employment question, in which the fundamental question is can we find fulfilling ways to spend our time if robots take our jobs? And there's an income question, can we find a stable and fair distribution of income."

Stevenson (2018)

Is this the end of jobs?

- No!
- Though if the change is fast, it could lead to an ugly couple of years.
- The long run can be a very long time!

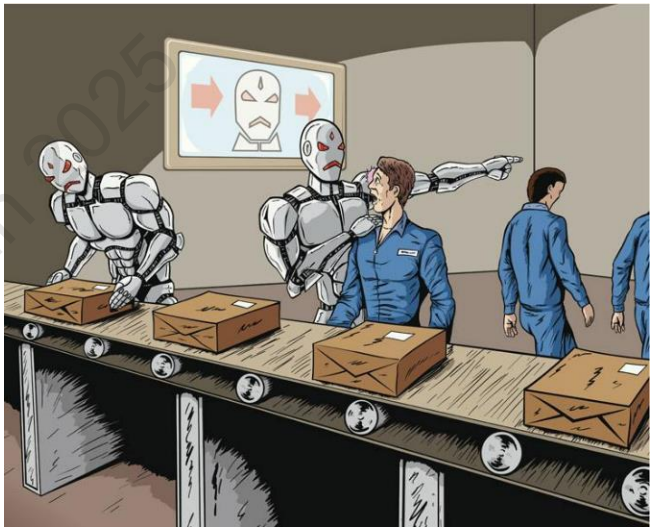

Historical perspectives

We will revisit this when talking about the "task-base" view of AI

Goolsbee (2018), Stevenson (2018), Bessen (2018), Furman (2018), Acemoglu and Restrepo (2018)

Macro Theme 2: Inequality

Will inequality get worse?

Probably.

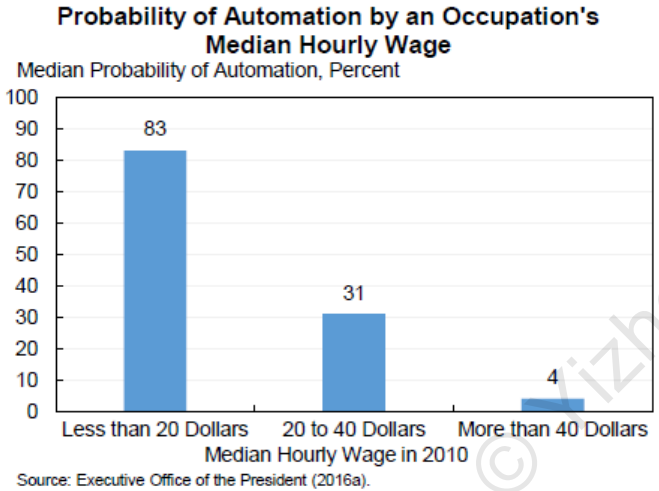

Figure 3: Probability of automation by Occupation Median Hourly Wage. Source: Furman (2018). Figure 4

- **Skill-biased technology.**
  - Automation
  - What about LLM?

We will revisit this when talking about the "task-base" view of AI

Will inequality get worse?

Probably. *Historical perspectives*

Figure 4
Share of Income Earned by Top 1 Percent, 1975-2015

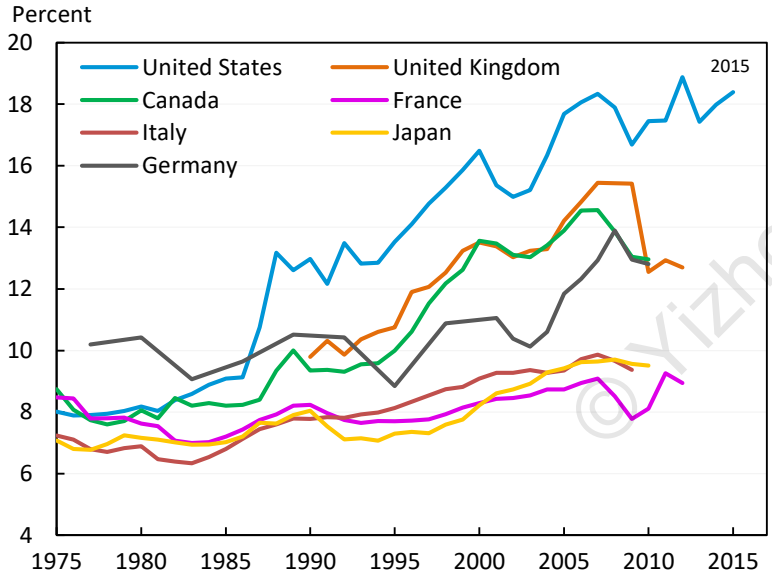

- Skill-biased technology.
- **Ownership of capital.**
  - Capital income has steadily increased vs. labor income
  - Amplified pre-existing disparities in capital / wealth

Furman (2018), Golden and Katz (2010), Deming (2017), Piketty (2013), Sachs (2018)

Policy solutions to inequality

- The social safety net

- Bill Gates called for taxation of robots.
  - BUT less investment → slower productivity growth
  - Stiglitz and Korinek (2018): A combination of finely balanced IP rights and capital taxation can limit distortions and enable distribution.

- Universal basic income
  - Big topic by itself, massive experiments ongoing

UBI

Exhibit 1

Finland set its experimental basic income at a modest level.

Average disposable income[1] by household type in Finland, 2017,
€ per month

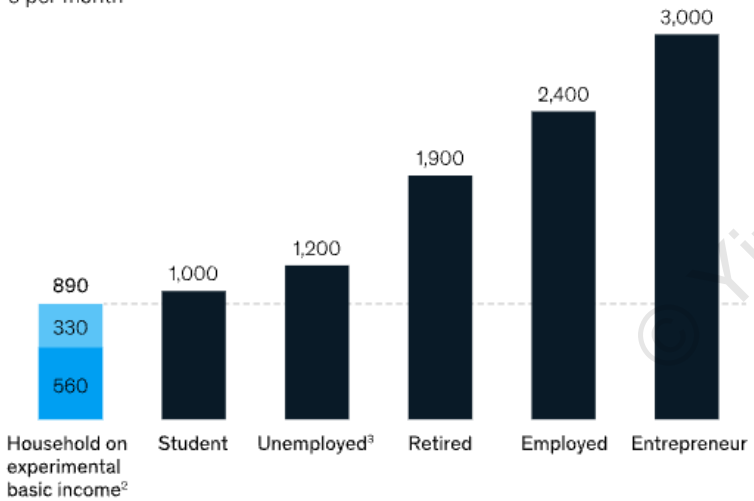

In the Finnish experiment, people on the basic income reported large and statistically significant improvements in key drivers of well-being.

Importance and impact of four well-being factors

	Mental health	Physical health	Employment	Trust
Impact of basic income on well-being factor[3]	Large	Medium/large	Medium to medium/large	Large
Statistical significance of effect (confidence level), %	99.9	95	57–99[4]	99.7

McKinsey (2022)

Macro Theme 3: Innovation

A.I. as a General Purpose Technology for Innovation

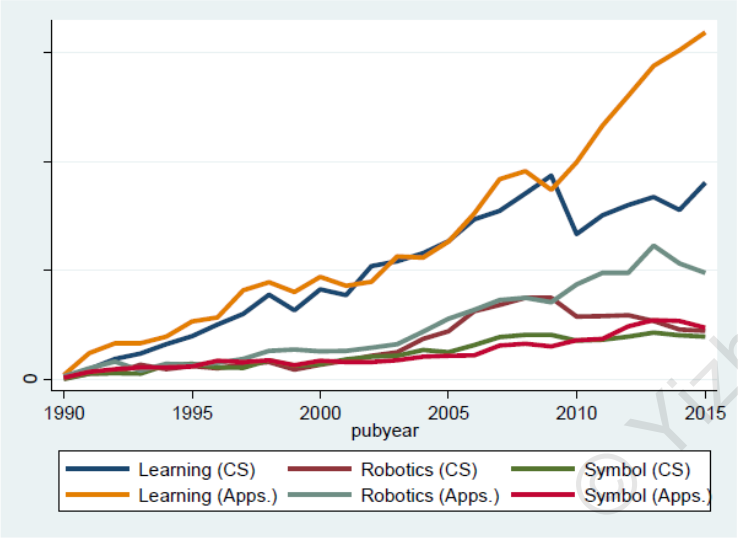

Figure 2: Publications in Computer Science versus Application Journals, by AI Field
Source: Cockburn, Henderson, and Stern (2018). Figure 4

Could lead to exponential growth in ideas and reverse the trend that scientific ideas are getting harder to find.

Aghion, Jones, and Jones (2019)
Agrawal, McHale, and Oettl (2019)
Cockburn, Henderson, and Stern (2019)

Forbes

Oct 3, 2021, 07:34pm EDT | 57,782 views

AlphaFold Is The Most Important Achievement In AI—Ever

Rob Toews Contributor ⓘ

AI

I write about the big picture of artificial intelligence.

Follow

Early Applications

Researchers at UCSF used AlphaFold to uncover previously unknown details about a key SARS-CoV-2 protein, which will advance the development of **COVID-19 therapeutics**.

Using AlphaFold, a team at the University of Colorado pinpointed a particularly tricky bacterial protein structure, a discovery that will aid their efforts to combat **antibiotic resistance**, a looming public health crisis. The team spent years unsuccessfully trying to determine this protein's structure; with AlphaFold, they learned it in 15 minutes.

Source: https://www.forbes.com/sites/robtoews/2021/10/03/alphafold-is-the-most-important-achievement-in-ai-ever/?sh=53d2c7546e0a

Early Applications

Some applications, such as the evolutionary analysis of proteins, are set to flourish because the tsunami of available genomic data might now be reliably translated into structures.

“It’s a game changer,” says Andrei Lupas, an **evolutionary biologist** at the Max Planck Institute for Developmental Biology in Tübingen, Germany.

“The model from group 427 gave us our structure in half an hour, after we had spent a decade trying everything,” Lupas says.

Source: https://www.nature.com/articles/d41586-020-03348-4

Innovation or Imitation?

"It is not obvious whether AI is a general purpose technology for innovation or a very efficient method of imitation." – Matt Mitchell

The former → less IP protection with AI since innovation becomes easier
The latter → more IP protection with AI since "true" innovation becomes less rewarding

Macro Theme 4: Competition/Antitrust (IO)

Will a few companies dominate?

- Economies of scope in talent
  - Sharing of intangibles → AI conglomerates
- Economies of scale in data.
  - Entry barrier to newcomers
  - "Learning by doing"
  - But data is non-rival & there is diminishing returns
- Price discrimination (excessive markups)

Price Competition vs. Market Coordination

Figure 8: Market-Level Treatment Effects by Degree of Penetration, Submarket-Rent Quartile

(a) $Y = \log(rent)$

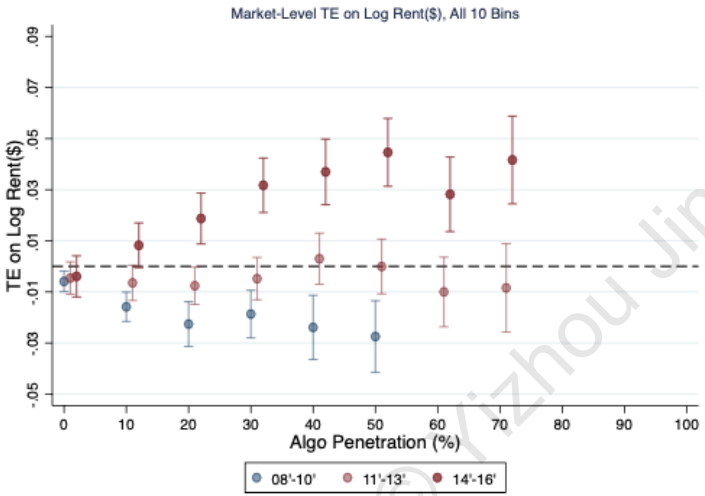

Algorithmic pricing can improve efficiency by helping firms set prices more responsive to changing market conditions. However, widespread adoption of the same algorithm could also lead to price coordination, resulting in elevated prices. In this paper, we examine the impact of algorithmic pricing on the U.S. multifamily rental housing market using hand-collected adoption decisions of property management companies merged with the data of market-rate multifamily apartments from 2005 to 2019. First, our findings suggest that algorithm adoption indeed helps building managers

Adoption of a rent-setting algo by property management companies

- In 2008-2010, increasing price competition and pushed rent down
- In 2014-2016, allowed market-wide price coordination and raised rents

Calder-Wang and Kim (2023)

Macro Theme 5: Privacy, Liability & Bias

Privacy

- In the context of A.I., privacy refers to limitations on the collection and use of data by firms and governments.
  - Consumer data as an input to A.I.
  - Data has externalities: Data created by one individual may contain data about others. (Your genetics data reveals a lot about all of your cousins)

- Regulation: too much reduces innovation and too little reduces trust/consumer sharing

Liability

- The purpose of the tort system is to deter people and companies from injuring others, and to compensate injured parties.

- Tort risk can increase or decrease innovation, depending on whether the risk is driven by new or existing products.

- Need clear liability rules. Those rules need to be strict enough for consumers to trust the technology (and for the technology to be safe!) but not so strict that companies bear too much risk.

Galasso and Luo (2018)

Discrimination and Bias

- Algorithms can generate biased results, even by accident.

- e.g. In a STEM education ad campaign on social media, ads were shown more to men than women.
  - The reason was innocuous: advertising to women cost more. So the algorithm showed it to men more.

- Should the social media network be found liable? What if human decision makers are more biased?

Lambrecht and Tucker (2018)

UNIVERSITY OF TORONTO

BREAK

How does AI Create Value? The “Task” View

TABLE 1—SUITABILITY FOR MACHINE LEARNING:

	Occupations	Tasks
Mean SML	3.47	3.47
SD of SML	0.11	0.31
Minimum SML	2.78	2.38
25th percentile SML	3.40	3.25
75th percentile SML	3.50	3.68
Max SML	3.90	4.48
Count	966	19,612

SML = “Suitability for Machine Learning”; they use O*NET (964 occupations; 18,156 tasks; 2,069 direct work activities; score each work activities (DWA) for its SML using rubric that considers 23 distinct statements evaluated on a 5-point scale varying from “strongly disagree” to “strongly agree.”

TABLE 2—LOWEST AND HIGHEST 5 SML SCORE OCCUPATIONS

Low SML occupations	SML	High SML occupations	SML
Massage therapists	2.78	Concierges	3.9
Animal scientists	3.09	Mechanical drafters	3.9
Archeologists	3.11	Morticians, undertakers, and funeral directors	3.89
Public address system and other announcers	3.13	Credit authorizers	3.78
Plasterers and stucco masons	3.14	Brokerage clerks	3.78

Source: https://www.cs.cmu.edu/~tom/pubs/AEA2018-WhatCanMachinesLearn.pdf

LLMs: GPTs are GPTs

80% of workers belong to an occupation with at least 10% of its tasks exposed to LLMs, while 19% of workers are in an occupation where over half of its tasks are labeled as exposed.

- **Augmentation vs. Displacement**: LLMs can make many tasks more efficient (*complements*) but may replace others entirely (*substitutes*).

- **Exposure Across Income Levels**: *Higher-wage roles show greater exposure*, indicating LLMs could impact a wide range of job types.

GPTs are GPTs: An Early Look at the Labor Market Impact Potential of LLMs (Eloundou, Manning, Mishkin, and Rock 2023)

Present Key Results from an Empirical Paper

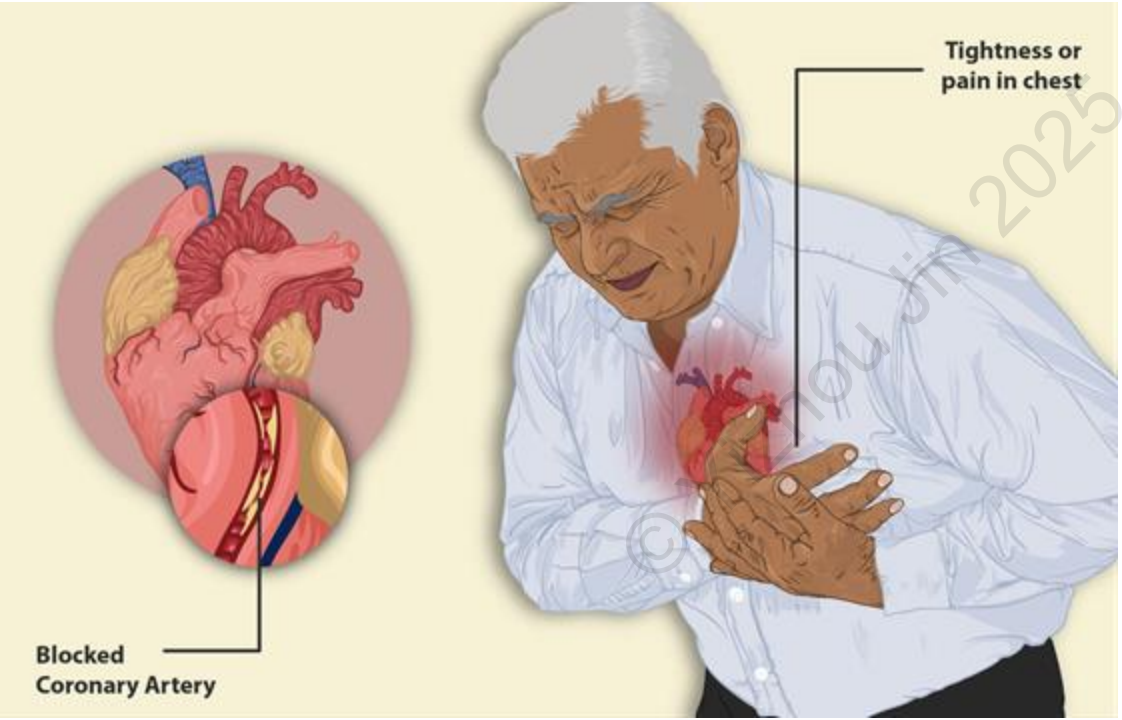

Heart attack prediction at emergency departments

Diagnosing Physician Error: A Machine Learning Approach to Low-Value Health Care

---Sendhil Mullainathan & Ziad Obermeyer (2019)

Heart Attack? Not Easy to Determine

- Patient arrives at emergency
- Some cases are obvious (e.g., severe chest pain)
- In other cases, life-threatening blockages can have subtle symptoms:
  - a subtle squeezing sensation in the chest, shortness of breath, nausea
  - can be easily attributed to other causes (e.g., acid reflux, viral infection, pinched nerve in the back)
  - key decision: whether to conduct an imaging test for heart attack

The Decision: Test or No Test

Benefit of testing

- More information on condition → reduces false negatives (should treat but didn't)

Cost of testing

- Tests are expensive and risky to patients
  - thousands USD + monitoring/hospitalization
  - single highest dose of ionizing radiation (long-term cancer risks)

Prediction and Judgment: Humans

Today, doctors receive information about a patient, make a prediction of the health risk in their heads, and then decide whether to test.

Prediction and Judgment: Humans

Risk Ventile	Yield (1
1	0.017
2	0.022
3	0.034
4	0.049
5	0.063
6	0.082
7	0.075
8	0.076
9	0.092
10	0.094
11	0.114
12	0.124
13	0.145
14	0.143
15	0.158
16	0.193
17	0.199
18	0.206
19	0.254
20	0.351

But there is a lot of waste

Researchers ranked patients into twenty bins with bin #1 being the lowest

They tested a random sample within each bin and Yield = probability of treatment among the tested

only 1.7% of those tested in the 1st bin needed treatment

Prediction and Judgment: Humans

Risk Ventile	Yield (SE) (1)		Cost, $ (2)
1	0.017	(0.003)	650,838
2	0.022	(0.003)	587,572
3	0.034	(0.004)	366,289
4	0.049	(0.005)	270,292
5	0.063	(0.005)	222,940
6	0.082	(0.006)	178,145
7	0.075	(0.006)	181,552
8	0.076	(0.006)	203,132
9	0.092	(0.007)	165,491
10	0.094	(0.007)	171,460
11	0.114	(0.007)	140,606
12	0.124	(0.008)	140,064
13	0.145	(0.008)	119,263
14	0.143	(0.008)	131,469
15	0.158	(0.008)	121,253
16	0.193	(0.009)	105,463
17	0.199	(0.009)	102,103
18	0.206	(0.009)	103,568
19	0.254	(0.010)	90,504
20	0.351	(0.011)	74,739

But there is a lot of waste

Researchers ranked patients into twenty bins with bin #1 being the lowest

They tested a random sample within each bin and Yield = probability of treatment among the tested

only 1.7% of those tested in the 1st bin needed treatment

cost = for each life year saved, how much was spent on testing

if you test 100 patients but only treat 1.7, then the cost per life-year saved is very high ($650K USD!) –> “low-value care”

Prediction and Judgment: Humans

Risk Ventile	Yield (SE) (1)		Cost, $ (2)	Test rate (SE) (3)	
1	0.017	(0.003)	650,838	0.015	(0.000)
2	0.022	(0.003)	587,572	0.024	(0.000)
3	0.034	(0.004)	366,289	0.030	(0.001)
4	0.049	(0.005)	270,292	0.036	(0.001)
5	0.063	(0.005)	222,940	0.042	(0.001)
6	0.082	(0.006)	178,145	0.043	(0.001)
7	0.075	(0.006)	181,552	0.048	(0.001)
8	0.076	(0.006)	203,132	0.048	(0.001)
9	0.092	(0.007)	165,491	0.052	(0.001)
10	0.094	(0.007)	171,460	0.053	(0.001)
11	0.114	(0.007)	140,606	0.056	(0.001)
12	0.124	(0.008)	140,064	0.061	(0.001)
13	0.145	(0.008)	119,263	0.064	(0.001)
14	0.143	(0.008)	131,469	0.064	(0.001)
15	0.158	(0.008)	121,253	0.070	(0.002)
16	0.193	(0.009)	105,463	0.075	(0.002)
17	0.199	(0.009)	102,103	0.079	(0.002)
18	0.206	(0.009)	103,568	0.089	(0.002)
19	0.254	(0.010)	90,504	0.103	(0.002)
20	0.351	(0.011)	74,739	0.127	(0.003)

Doctors' decisions to test are, overall, not bad: They tested 1.5% of patients in the first bin and 12.7% in the last bin. But it could have been a lot better.

Prediction and Judgment: Humans

Risk Ventile	Yield (SE) (1)		Cost, $ (2)	Test rate (SE) (3)	
1	0.017	(0.003)	650,838	0.015	(0.000)
2	0.022	(0.003)	587,572	0.024	(0.000)
3	0.034	(0.004)	366,289	0.030	(0.001)
4	0.049	(0.005)	270,292	0.036	(0.001)
5	0.063	(0.005)	222,940	0.042	(0.001)
6	0.082	(0.006)	178,145	0.043	(0.001)
7	0.075	(0.006)	181,552	0.048	(0.001)
8	0.076	(0.006)	203,132	0.048	(0.001)
9	0.092	(0.007)	165,491	0.052	(0.001)
10	0.094	(0.007)	171,460	0.053	(0.001)
11	0.114	(0.007)	140,606	0.056	(0.001)
12	0.124	(0.008)	140,064	0.061	(0.001)
13	0.145	(0.008)	119,263	0.064	(0.001)
14	0.143	(0.008)	131,469	0.064	(0.001)
15	0.158	(0.008)	121,253	0.070	(0.002)
16	0.193	(0.009)	105,463	0.075	(0.002)
17	0.199	(0.009)	102,103	0.079	(0.002)
18	0.206	(0.009)	103,568	0.089	(0.002)
19	0.254	(0.010)	90,504	0.103	(0.002)
20	0.351	(0.011)	74,739	0.127	(0.003)

Doctors' decisions to test are, overall, not bad: They tested 1.5% of patients in the first bin and 12.7% in the last bin. But it could have been a lot better.

The highest dollar amount any society puts on saving a life-year is about USD 150K. This means that we should test all patients below the line and not test any patient above the line

53% of existing tests caused more harm than benefit; while many more tests should have been conducted but didn't.

Source: https://www.nber.org/system/files/working_papers/w26168/revisions/w26168.rev1.pdf?mod=article_inline *("A Machine Learning Approach to Low Value Health Care: Wasted Tests, Missed Heart Attacks, and Mis-Predictions" by Sendhil Mullainathan and Ziad Obermeyer, 2019)*

How does AI Create Value? The “System” View

The task-based view is “irrelevant”

“I focus on the highly valuable applications of [Artificial Intelligence Technologies] AITs today... My empirical conclusion about these applications is that… ***task level substitution of machine for human plays no role in these highly valuable systems****.”*

Artificial Intelligence Technologies and Aggregate Growth Prospects
(Bresnahan 2019)

Real AI Applications Function at the System Level

"I begin with AIT-based product/consumer matching engines at Amazon, Google, Facebook, Netflix... These are not demonstration projects or experiments. They are **production systems** generating revenues in the hundreds of billions of dollars."

Amazon did not use AI to automate specific tasks previously performed by humans

"Task level substitution plays no role in these applications of AIT. These very valuable early applications are ***not ones in which labor was undertaking a task and was replaced by capital.*** *Observers focus on task level substitution, not because it occurs, but because the definition of general AI includes 'tasks usually done by humans.'"*

System Thinking in AI and Heart Attack Testing

Physicians have poor incentives. They receive some private benefit from doing more tests, in the form of extra revenue or protection from malpractice risk. Also, they are not that accurate at predicting risk. Moral hazard

- Task level: AIs could help reduce low value tests and increase missed high value tests.
- System level: if AI makes a prediction that doctors, hospitals, insurers, and patients can see, rather than a system where the prediction is private in the doctor's mind. Then, the doctor must submit an explanation if they decide to administer a test when the AI predicts the risk is below a threshold.

Midterm Q

Focusing on the moral hazard issues between health insurers and the insured, which one of the following is the LEAST relevant?

A – purchasing more generous insurance coverage due to a higher likelihood of visiting the hospital

B – increased hospital visits as a result of generous insurance coverage that lowers out-of-pocket expenditures

C – more frequent skiing trips after purchasing comprehensive orthopedic coverage

D – increasing prevalence of opioid addiction due to more generous coverage of pain medication

A describes adverse selection

Taking Stock: Which Is It Then?

AI can create value on both the task and the system-level

- General AI (e.g. ChatGPT) are designed to imitate humans → better at task-level value creation
  - Wasn't widely available when Tim wrote his critique
  - We don't really know what the long-term value of human imitation will be: substitution vs complementary?
  - "The AI photo app trend has already fizzled" -*TechCrunch Feb 2023* (Lensa etc…)

- System level value creation requires economic and strategic thinking
  - Many existing systems like ranking/matching algos on Google/Tinder have already been very influential
  - Not just because they are accurate (or more accurate than humans), but because they resolved fundamental **economic frictions** and vastly **expanded** economic activities: they do things that humans couldn't or did not do before
  - e.g. auto insurance monitoring (address moral hazard)
  - e.g. ranking based on review and ad auctions (address adverse selection)

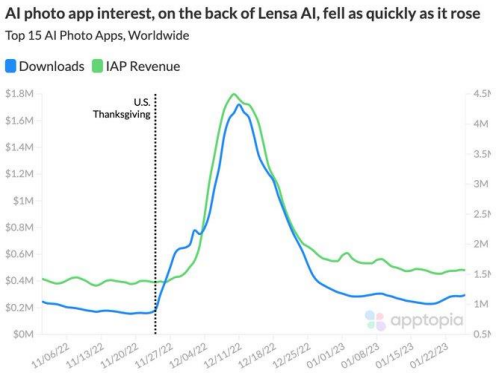

- Task-level automation without system-level thinking is dangerous
  - Zillow's algo home buying program failed due to adverse selection (lec 8)