

Artificial Intelligence and Jobs

A multidisciplinary synthesis of what research suggests about technology (incl. AI) and jobs
(excl. post-AGI/Economics of Transformative AI)

Why & how I compiled this Primer

- It is hard to open a newspaper, listen to a podcast, or scroll through Twitter/X without encountering a headline on the topic of job destruction as a result of AI.
- While research studies offer findings that can nuance such claims these subtleties rarely feature in the same headlines.
- That is why I set out to synthesize the findings on this topic into a Primer. My goal is to make the content as accessible as possible to researchers, policy makers, and anyone curious about the topic.
- The Primer seeks to bridge different disciplines (economics, history, finance, organizational behavior, technology and computer science), noting that my background mainly centers around the first two fields.
- Given the amount of writing and reading I do in my daily life, and the broad intended audience, I decided on “snappy presentations” as the medium of choice.
- To put this Primer together, I read (a task in which I have some level of expertise) several papers in full and others more selectively (typically abstracts, introductions, and results), and also sourced from meta-studies. I sought to use peer reviewed articles in as far as possible.
- The design of the slides (a task where my expertise is decidedly more limited) was a product of human-AI collaboration (Claude).

Scope & how to read this

Sidenote:

- This Primer is written for anyone curious about AI and jobs. No economics background required, though some familiarity with the debate helps. It is a literature review, not an opinion piece; the goal is to synthesize what the research says, rather than argue a position.
- The content derived from economics focuses on the *economics* of AI; the economics of transformative AI (ETAI) are not yet included. I am working through that literature and will either add it here or write a separate primer.
- The geographic focus is developed markets, particularly the US and Europe, on which most of the research studies have focused (to date). Whenever there are studies that cover developing countries, I have included those as sidenotes.
- This is a live document. Most of the material derives from my readings between Fall 2024 and Summer 2025, with selected additions since.

Navigation

00 Rationale

01 Timelines

02 Theory

03 A look back in time

04 Empirical findings

05 Outstanding questions

Historical context: Fear of technological unemployment is a recurrent theme



Just as people did two centuries ago, many fear that machines will make millions of workers redundant, causing inequality and unrest.

THE ECONOMIST (2016)

REWARD FUNCTIONS IN 1812
(UK FRAME-BREAKING ACT)



Frame-Breaking.
£.200 Reward.

WHEREAS, on Thursday Night last, about Ten o'Clock, a great Number of Men, armed with Pistols, Hammers and Clubs, entered the Dwelling-house of George Ball, Framework-Knitter, of Lenton, near Nottingham, disguised with Masks and Handkerchiefs over their Faces, and in other ways—and after striking and abusing the said George Ball, they wantonly and feloniously broke and destroyed five STOCKING FRAMES, standing in the Work-shop; four of which belonged to George Ball, and one Frame, 40 gages, belonging to Mr. Francis Headlam, hosiery, Nottingham: all of which were working at the FULL PRICE.

NOTICE IS HEREBY GIVEN,

THAT if any Person will give Information of the Offender or Offenders, or any one of them who entered such Dwelling-house and were concerned in such Felony, he or she shall receive a Reward of **£. 200,**

to be paid on Conviction, in the Proportions following:—(10s.) £50 under the King's Privy Seal, and 40s. from the Committee of the Corporation of Nottingham, and 20s. from the said Francis Headlam.



Technology is creating both new opportunities and new obligations for us. [...]

Automation is not our enemy. Our enemies are ignorance, indifference and inertia. Automation can be the ally of our prosperity if we will just look ahead. [...]

The techniques of automation are already permitting us to do many things we simply could not do otherwise.

LYNDON B. JOHNSON · 36TH PRESIDENT · AUGUST 19, 1964

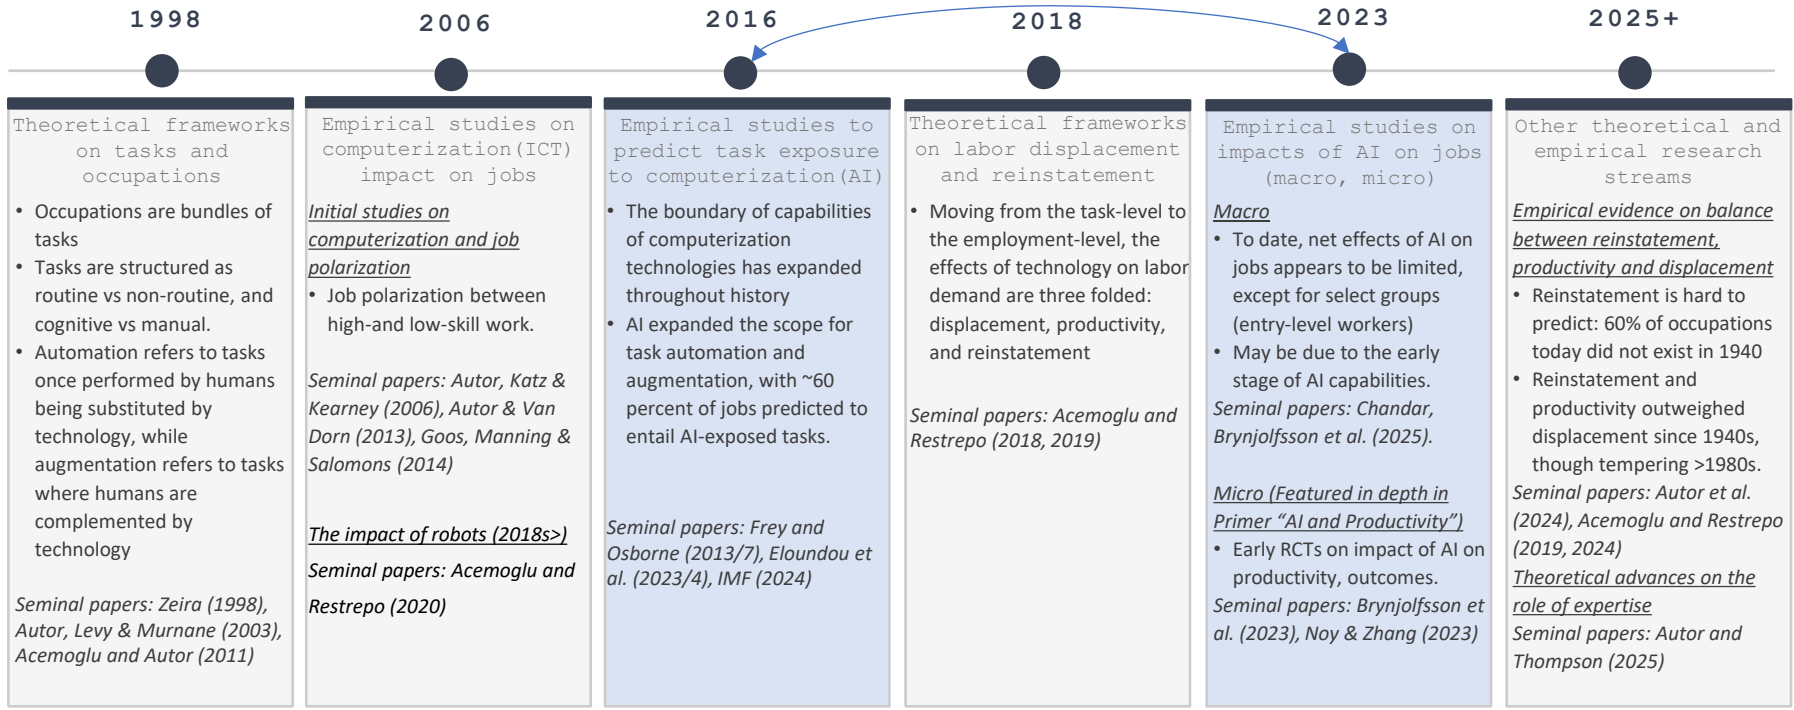
Timelines

01

Evolution of economic research on computerization & jobs

Sidenote: scope only covers research on computerization. Earlier research on impact of technology on jobs pending.

Sidenote: The former field predicts AI exposure across occupations based on tech. capabilities, while the latter (macro) measures actual impacts on jobs.



Economic theory

Theoretical frameworks from Autor et al. (2003) and Acemoglu and Autor (2011) highlight that the impact of technology (including AI) on jobs should be observed at the task level rather than the occupation/job level. Acemoglu and Restrepo distinguish labor replacement, reinstatement and productivity. (Economics)

1. Tasks are occupational bundles
2. Replacement/reinstatement/productivity

Theoretical frameworks study impacts of technology at the task level, where occupations are bundles of tasks

Definitions



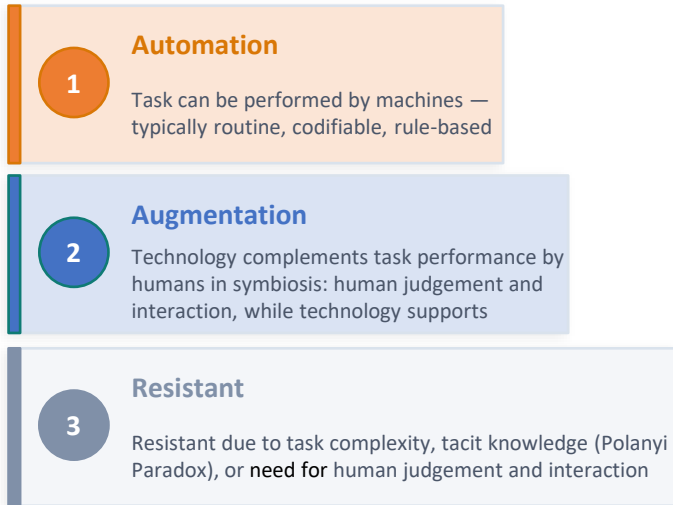
Types of tasks

	Routine tasks: codifiable, following explicit rules, and repetitive	Non-routine tasks: tacit, demand flexibility, generalized problem-solving, and communications
Cognitive (analytical, interactive)	Record-keeping, calculation, repetitive customer service	Hypothesis testing, medical diagnosis, legal writing, persuading
Manual	Picking/sorting, repetitive assembly	Janitorial services, truck driving

Note: Autor (2003)

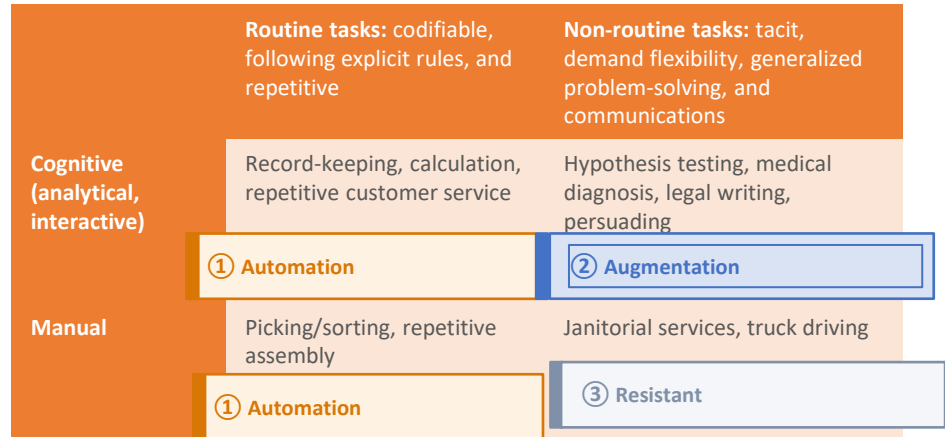
Tasks are automated (substituted), augmented (complemented), or resistant to computerization technologies

Types of impact from technology on tasks



Note: Autor (2003)

Types of tasks



From task to employment: the effects of technology on labor demand involve displacement, productivity, and reinstatement

Net effect on labor demand =

- displacement

+

+ productivity

+

+ reinstatement

DISPLACEMENT

Automation substitutes labor in tasks workers previously performed, thereby reducing labor demand.

PRODUCTIVITY

Lower costs expand output and raise real incomes, thereby increasing labor demand.

REINSTATEMENT

New tasks emerge where humans have a comparative advantage, thereby increasing labor demand.

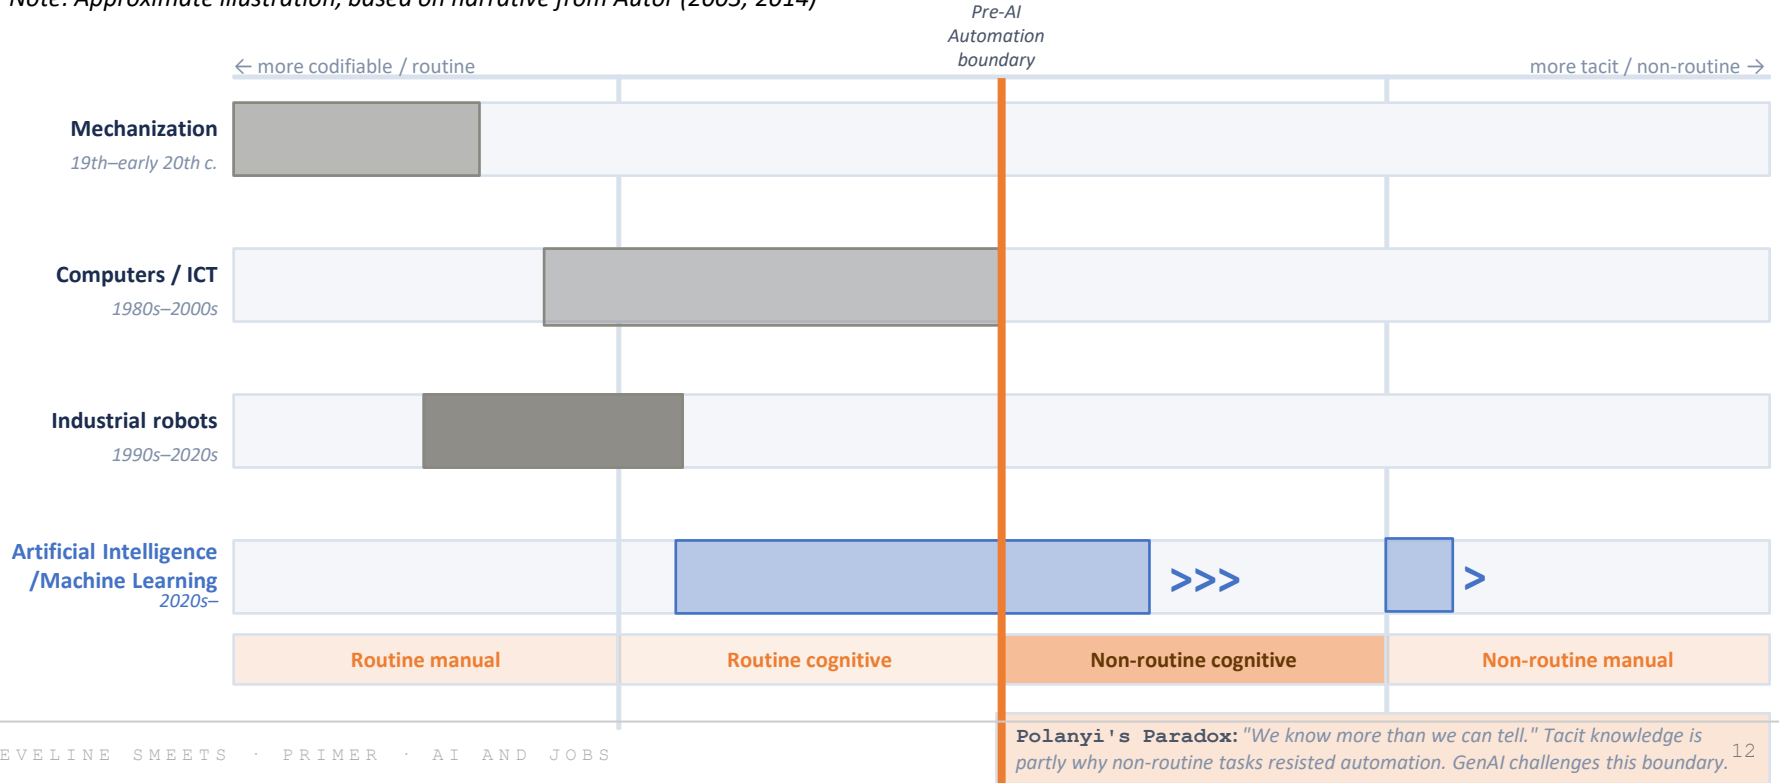
A look back in time

The full picture: what can be learnt from past technological waves
(history, economic history, economics)

03

The boundary of capabilities of computerization technologies expanded historically, and thereby the scope for automation

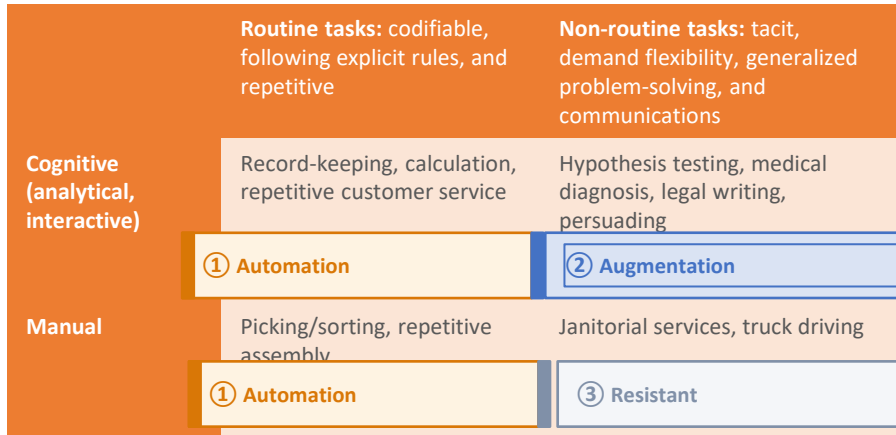
Note: Approximate illustration, based on narrative from Autor (2003, 2014)



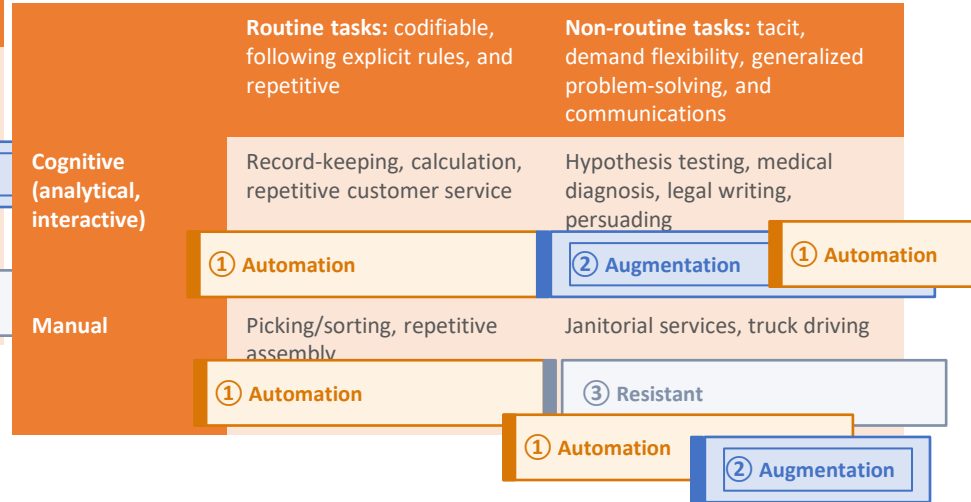
That means the types of tasks starts to look like this

Note: Approximate illustration, based on narrative from Autor (2003, 2014)

Types of tasks



Types of tasks



Since 1940s, labor reinstatement & productivity effects of technology offset displacement; however, this has been tempering since 1980s.

Period	Author	Technology wave	Main finding	Displacement	Reinstatement	Real Income/ Productivity	Net job effects																									
1947–1987	Acemoglu and Restrepo (2019)	<i>Not stated, not measured (attributes labor trends not specifically to technology)</i>	Slower growth of US employment over the last three decades compared to previous decades is accounted for by: <ul style="list-style-type: none"> • Acceleration in displacement effect (automation). 	↓0.48%/yr	↑0.47%/yr	↑2.4%/yr	↑2.39%/yr																									
1987–2017	Acemoglu and Restrepo (2019)	<i>Inferred: Robots, ICT, “automation technologies” (attributes labor trends to technology and other vars)</i>	<ul style="list-style-type: none"> • Deceleration of new task creation (less reinstatement). • Slower growth of productivity. 	↓0.70%/yr	↑0.35%/yr	↑1.54%/yr	↑1.19%/yr																									
>1980s	Hotte et al. (2023), meta-review of 127 studies	ICT, Robots and definitions of “innovation” and TFP-style technology	Across 127 studies, labor displacing effect of technology appears to be more than offset by compensating mechanisms that create or reinstate labor. This holds for most technology-types, suggesting that anxieties over widespread technology-driven unemployment lack an empirical base.	<table border="1"> <caption>Stacked Bar Chart Data</caption> <thead> <tr> <th>Category</th> <th>Support</th> <th>Depends</th> <th>Weak</th> <th>No support</th> </tr> </thead> <tbody> <tr> <td>Displacement</td> <td>67%</td> <td>13%</td> <td>5%</td> <td>16%</td> </tr> <tr> <td>Reinstatement</td> <td>81%</td> <td>9%</td> <td>8%</td> <td>0%</td> </tr> <tr> <td>Real income (~Productivity)</td> <td>70%</td> <td>12%</td> <td>6%</td> <td>12%</td> </tr> <tr> <td>Net employment</td> <td>29%</td> <td>35%</td> <td>18%</td> <td>18%</td> </tr> </tbody> </table>				Category	Support	Depends	Weak	No support	Displacement	67%	13%	5%	16%	Reinstatement	81%	9%	8%	0%	Real income (~Productivity)	70%	12%	6%	12%	Net employment	29%	35%	18%	18%
Category	Support	Depends	Weak	No support																												
Displacement	67%	13%	5%	16%																												
Reinstatement	81%	9%	8%	0%																												
Real income (~Productivity)	70%	12%	6%	12%																												
Net employment	29%	35%	18%	18%																												


Sidenote: Acemoglu and Restrepo (2019) identify two reasons for the change in balance between displacement and reinstatement:

1. US may have moved along a given innovation possibilities frontier because incentives for automation have increased, while those for creating new tasks declined. For example, taxation incentives (subsidization of equipment use, taxation on labor), spirit to remove human element from production, declining gov’t support for innovation/research with longer horizons.
2. Innovation frontier linking these two types of technological change may have shifted, facilitating further innovation and making new task creation harder.

Sidenote: See section “Direction of technological change”

New occupations are created by technology (reinstatement) and demand; yet, hard to anticipate (& measure concurrently)

~60% Of employment in 2018 is in new occupational titles that did not exist in 1940 (Autor et al., 2024)

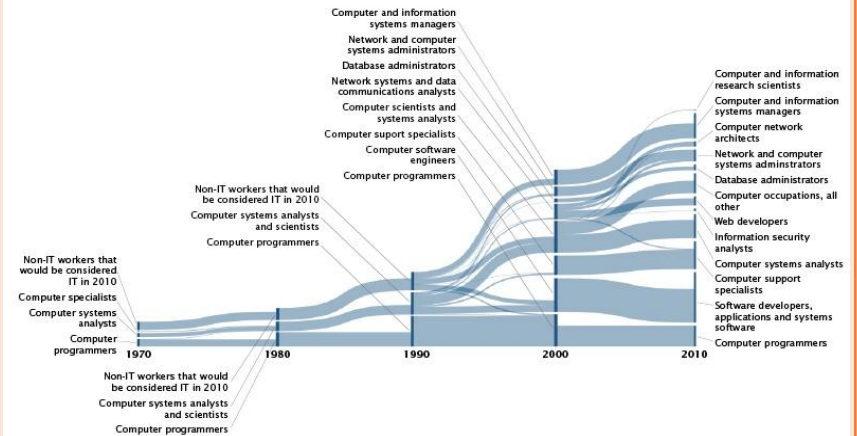


■ 2018 employment in occupations that did NOT exist in 1940

■ 2018 employment in occupations that existed in 1940

Main sources of new occupations

1. New technologies, e.g.. IT occupations since 1970s



2. Demand for new services/products derived from new tastes, aging populations and other trends, e.g. nail technicians


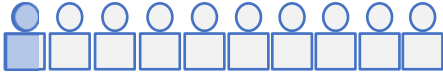
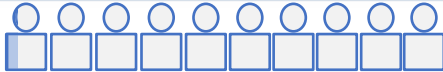

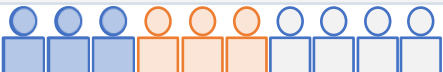
Empirics on AI's ongoing impact on jobs

AI exposure versus actual impact measurements (economics)



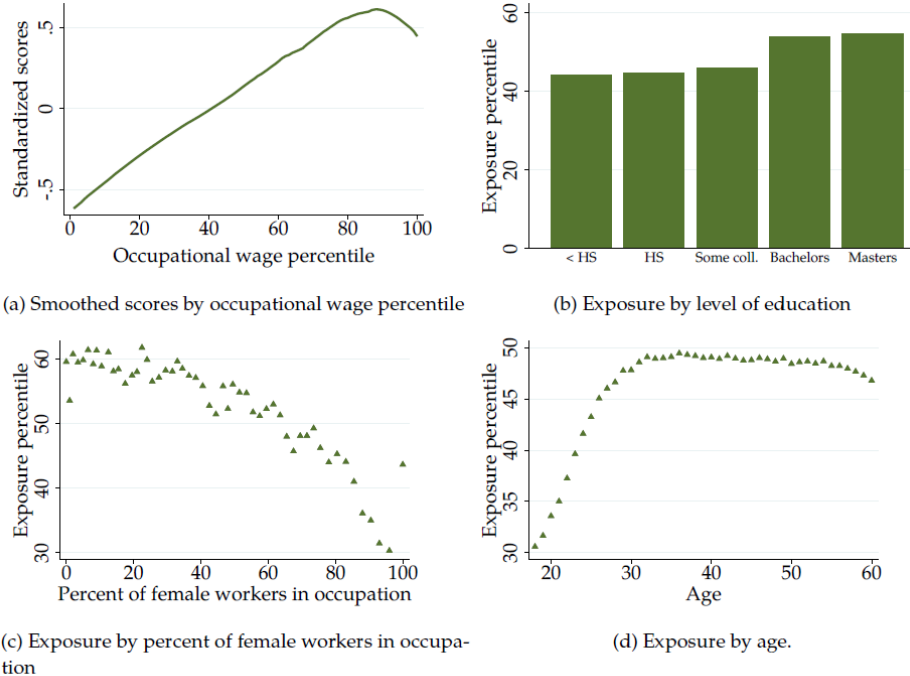
AI expanded the scope for task automation and augmentation, with ~60 percent of jobs predicted to entail AI-exposed tasks

Note: methodologies take all O*NET tasks and ask (LLM and/or expert judgement) to assess whether a machine can automate or augment that task.

AUTHORS	TECHNOLOGY	LEVEL	GEOGRAPHY	CONTRIBUTION
Frey & Osborne (2013/7)	Computerization	Occupations	USA	 47%
Arntz, Gregory & Zierahn (2016)	Computerization	Tasks	USA, OECD	 9%
Eloundou et al. (2023/4)	LLMs (Gen AI)	Employment & Tasks	USA	 1.8% of jobs have 50%+ of tasks exposed to LLMs
Eloundou et al. (2023/4)	LLMs (Gen AI) + complementary future software	Employment & Tasks	USA	 46% of jobs have 50%+ of tasks exposed to LLMs+ complementary software
IMF (2024)	Artificial Intelligence	Employment	USA, Advanced economies	 60% of employment is exposed, of which ½ by automation and ½ by augmentation

Sidenote: Other studies compare AI exposure globally, incl. ILO (2023), ILO (2025), IMF (2024), World Bank (2025). These find that AI exposure is higher in advanced economies than developing economies due to differential economic structures (sectors, formality), economic considerations (labor vs capital costs), and enabling environments (baseline digital connectivity and IT infrastructure).

Predicted AI exposure is higher for selected groups, highlighting expected distributional effects



Sources: Webb (2020)

Figure 7: Exposure to AI by demographic group

Yet, 3+ years after the launch of GPT3.5, the effect of AI on employment at the macro level appears to be fairly mixed.

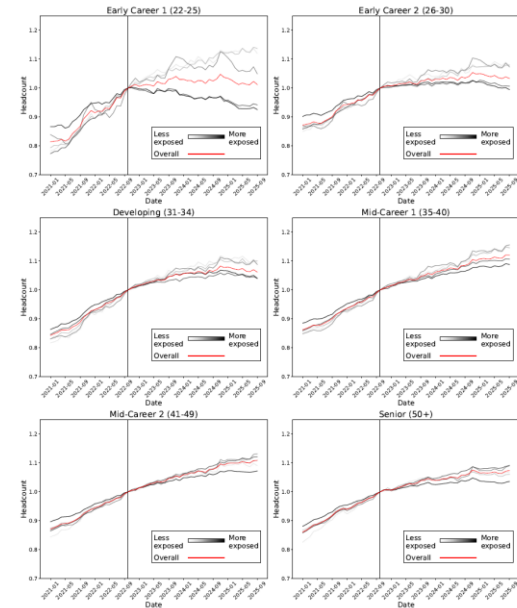
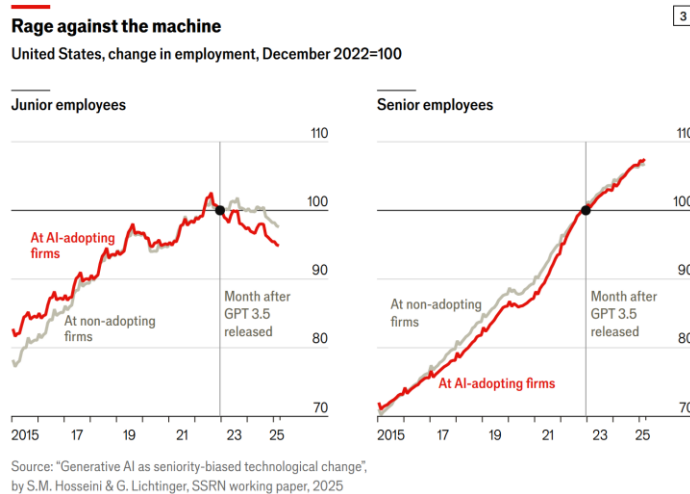
Paper	Country	Time	Findings	Finding: labor
Brynjolfsson et al. (2025)	USA	2021-2025	<ul style="list-style-type: none"> Early-career workers (ages 22-25) in AI-exposed occupations experienced 16% relative employment declines, controlling for firm-level shocks, while employment for experienced workers remained stable. Adjustments occur primarily via employment rather than compensation, with employment changes concentrated in occupations where AI automates rather than augments labor. 	(Some) Negative
Dominski & Suk Lee (2025)	USA	2022-2025	<ul style="list-style-type: none"> We find that higher exposure to AI is associated with reduced employment, higher unemployment rates, and shorter work hours. We also observe some evidence of increased secondary job holding and a decrease in full-time employment among certain demographics. These associations are more pronounced among older and younger workers (those below 30 and over 50), men, and college-educated individuals. 	(Some) Negative
Hampole et al. (2025)	USA	2010-2023	<ul style="list-style-type: none"> Two variables summarize the impact of AI on within-firm labor demand: the mean exposure of tasks within an occupation, which depresses demand, and the concentration of exposure in a few tasks, which offsets losses by enabling workers to reallocate effort. Despite strong substitution at the task level, overall employment effects are modest, as reduced demand in exposed occupations is offset by productivity-driven increases in labor demand at AI-adopting firms. 	Limited/Neutral
Humlum & Vestergaard (2025)	Denmark	2023-2024	<ul style="list-style-type: none"> Firm AI adoption is assoc. with modest productivity gain (-3% time), occupation switching and task restructure New tasks emerge (oversight, integrate) that absorb time savings. Finds precise null effects on earnings and hours worked. (based on 11 occupations) 	Limited/Neutral
Gimbel et al.	USA	2022-2025	<ul style="list-style-type: none"> In the USA, trackers (Yale Budget Lab) suggest small discernible labor market disruption since ChatGPT 	Limited/Neutral
Jiang et al. (2025)	USA	2004-2023	<ul style="list-style-type: none"> Higher AI exposure is associated with longer work hours and reduced leisure time, indicative of AI augmenting human labor rather than replacing it. This is pronounced in contexts where AI significantly enhances marginal productivity and monitoring efficiency. It is further amplified in competitive labor and product markets, where workers have limited bargaining power to retain the benefits of productivity gains, which are often captured by consumers or firms instead. 	(Some) Positive
Johnston and Makridis (2025)	USA	2017-2024	<ul style="list-style-type: none"> A 1SD increase in LLM exposure (Eloundou et al.) raises output by 7%, with effects emerging in 2021 when enterprise AI tools entered Employment effects follow the same timing but diverge by exposure type: where AI likely requires human collaboration, employment rises 4%; where AI can perform tasks independently, we find no significant employment effect. Results are robust to state-by-year and industry-by-year fixed effects and suggest AI has caused a decrease in the labor share of income. 	(Some) Positive

Sidenote: Beane (2024) argues this could disrupt the traditional skill accumulation pathway (apprentice, learn by doing)

Yet, some distributional effects may be observed; entry-level workers appear to be affected negatively.

In AI-adopting firms, junior employment declined by 9% after the launch of ChatGPT, while senior employment did not. (Hosseini & Lichtinger, 2025)

Early-career workers (ages 22-25) in AI-exposed occupations experienced 16% relative employment declines while employment for experienced workers remained stable (Brynjolfsson et al. 2025).



Exposure \neq Impact: actual labor market impact to date falls below AI exposure predictions, for various reasons:

01

AI capabilities remain at an early stage. AI capabilities for narrow tasks (e.g., coding and writing) and performance has improved substantially vis-à-vis specific benchmarks. However, multi-step activities and tasks that involve creativity, judgement and human interaction remain more challenging, yet these represent an important component of occupations and overall economic activity.

02

Impact is contingent on the diffusion of AI among entities that employ people, which may face adoption barriers. Firms and other organizations may face challenges such as **economics**: Svanberg et al. (2024) find that, for computer vision AI, at today's costs US businesses would choose not to automate most vision tasks that have "AI Exposure," and that only 23% of worker wages being paid for vision tasks would be attractive to automate. Other adoption barriers include access to complementary inputs such as the relevant digital infrastructure and skills. There may also be **institutional barriers** such as resistance to change (Morrison, 1966), or **regulatory barriers** for industries in which professional standards require a human-in the loop (e.g., health, financial services).

03

AI exposure implies both task automation and augmentation, where the former involves job displacement and the latter rather involves reorganizing workflows and occupations. Task augmentation refers to situations where humans and AI collaborate on tasks in a more productive manner, and where AI creates new tasks. Automation and augmentation remains challenging to measure. An effort by Anthropic's Economic Index (2025) considers augmentation is 55% of task exposure, compared to 45% for automation.

04

Impacts may be delayed as firms take time to absorb AI, as per the literature on the Productivity J-curve. It typically takes time for workers to learn how to work with new technologies and for organizations to reorganize their workflows and tasks, which creates an absorption cost that could temporarily reduce impacts. (Brynjolfsson, 2021)

Sources: "AI: The Scaling Era, Chapter 2, Svanberg et al. (2024), Acemoglu (2024), Morrison (1966), Anthropic Economic Index (2025), Brynjolfsson (2021).

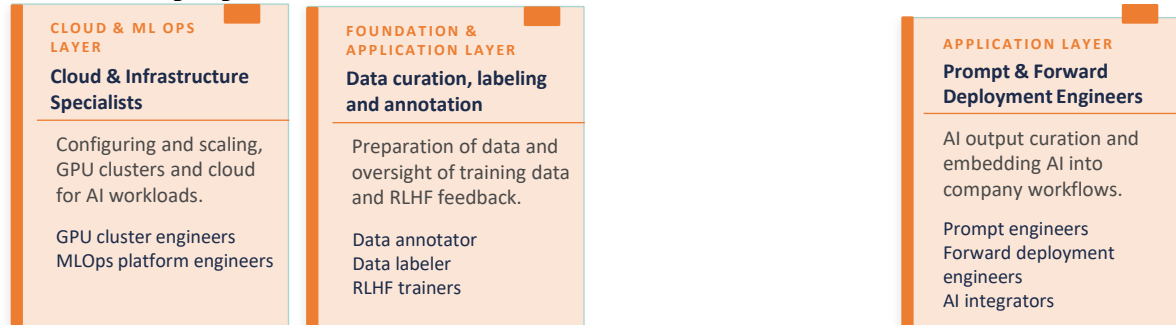
Exposure ≠ Impact: exposure is not the full story on AI's impact on jobs: labor reinstatement is hard to predict ex-ante and measure now.

Where could we situate the locus of new tasks and labor reinstatement? → From past technological waves (see slide 15), we know to focus on new industries.

Overview of the Gen AI value chain (McKinsey)

Computer Hardware	Cloud Platforms	Foundation Models	Model Hubs & MLOps	Applications	Services
Accelerator chips optimized for training & running the foundation models	Platforms to provide access to compute hardware	Core models on which generative AI applications can be built	Tools to curate, host, fine-tune, or manage foundation models (e.g., storefronts between applications and foundation models)	B2B or B2C products that use foundation models either largely as is or fine-tuned to a particular use case.	Services around specialized knowledge on how to leverage generative AI (e.g., training, feedback, reinforcement learning)

Selected new occupations emerging in the value chain (The Economist, own elaboration)



Outstanding questions

05

What makes AI different from prior waves? How does society shape the direction of technological change? How does society shape the impact of technology on jobs? (multi-disciplinary: some philosophy, computer science, etc.)

1. If technological unemployment is not observable from past technology waves, what could make AI different?

2. What influence do we have over the direction of technological change?

3. What influence do we have over the effects of technological change for jobs?

END