When code becomes free: a research compendium on labor transitions and SWE automation

The central pattern across 260 years of industrial history is remarkably consistent: when a capability layer becomes cheap, labor doesn't disappear—it migrates one abstraction layer up. Handloom weavers became factory supervisors. Human computers became programmers. Travel agents became platform engineers. Each transition followed the same arc: devastating short-term displacement, a painful intermediate period of 15–60 years, then explosive job creation in categories that didn't previously exist. David Autor's research shows that **60% of employment in 2018 consisted of job titles that didn't exist in 1940**. The question now is whether AI-driven SWE automation will follow this pattern or break it—and the evidence, while early, is pointing toward both. SWE-bench scores have leapt from ~20% to 75–80% in just 15 months, Cursor hit $1B ARR faster than any SaaS company in history, DevGraphiq and Stanford's Brynjolfsson documents a **20% employment decline for software developers aged 22–25** since late 2022. DigitalRosh Yet simultaneously, the WEF projects a net gain of **78 million jobs globally by 2030**, World Economic Forum +3 cybersecurity has **4.8 million unfilled roles**, Programs and hyperscalers are pouring **$371 billion into data center infrastructure in 2025 alone**. Deloitte Insights The bottleneck is shifting from writing code to defining problems, governing systems, and building in the physical world.

I. The Industrial Revolution: sixty years of pain before wages caught up

The handloom weavers and the Engels' Pause

The Industrial Revolution's impact on labor is best understood through wages, not just employment. Economic historian Robert Allen coined the term **"Engels' Pause"** to describe the period from roughly 1780–1840 when Wikipedia output per worker rose **46%** while real wages grew only **12%**. ResearchGate The profit rate doubled. ScienceDirect The share of profits in national income expanded from **17.2% in 1770 to 31.3% in 1860**, while landed rents collapsed from 21.8% to 8.5%—a massive redistribution from old wealth to new industrial capital.

The handloom weavers provide the era's most vivid case study of displacement. Approximately **75,000 handloom weavers** worked in Britain in 1795. Paradoxically, their numbers swelled to **over 200,000 by 1812** — Heywoodhistory Irish immigrants and desperate workers flooded into the trade even as power looms multiplied —because almost all cotton weaving was still done by hand until at least 1806. Massachusetts Institute of Tech… Wages tell the real story. During the "golden age" of the 1790s, a skilled weaver earned **30–40 shillings per week**. Average earnings stood at **240 pence/week in 1806**. By 1820, they had fallen below **100 pence/week**— Knowable Magazine a 58%+ decline. By the 1830s, wages had cratered to as little as **4–5 shillings/week**. Spartacus Educational A single power loom could produce more cotton than **10 to 20 handweavers**. Knowable Magazine An unskilled boy operating one could weave **3.5 pieces** in the time a skilled artisan wove one. Spartacus Educational

A weaver's testimony to Parliament captures the human cost: *"He had been a weaver and in prosperous times had earned from thirty or forty shillings per week; he had a wife and four children... work began to grow slack. He drew the fund he had placed in the savings-bank; he was soon exhausted... He began to sell his furniture. Before last Christmas everything had disappeared, including the Sunday clothes."* Heywoodhistory In 1807, over **130,000 weavers signed a petition** for a minimum wage. Parliament rejected it. Spartacus Educational In May 1808, **15,000 weavers** gathered at St. George's Fields in Manchester to demand minimum wages; the military dispersed them, killing one. Wikipedia

Machine-breaking as collective bargaining: the Luddites and Swing Riots

The Luddite movement erupted on **March 11, 1811**, when workers attacked a stocking factory in Arnold, near Nottingham, destroying knitting frames. EBSCO Over the next three weeks, more than **200 frames were destroyed** in Nottingham alone. Cpbml The government deployed **12,000 troops** to four northern counties by 1812—**more soldiers than Wellington had fighting Napoleon in Spain**. Cpbml The most dramatic confrontation came on **April 11, 1812**, when **George Mellor** led 100–200 men in a military-style assault on William Cartwright's Rawfolds Mill near Brighouse. Cartwright had fortified it with iron-studded doors, spiked rollers, and acid. The attack failed. Two Luddites, Samuel Hartley and John Booth, were mortally wounded. In a famous exchange, the dying Booth beckoned Reverend Hammond Roberson and asked if he could keep a secret. When Roberson said yes, Booth replied: *"So can I"*—and died without naming accomplices. Charlotte Brontë later used this attack as the basis for her novel *Shirley* (1849).

The York Special Commission opened on **January 2, 1813**. Mellor and two accomplices were hanged on January 8. On **January 16, fourteen men were hanged in a single day**—the largest mass execution at York Castle. They sang the Methodist hymn *"Behold the Saviour of Mankind"* on a scaffold deliberately set high so the crowd could witness the full convulsion.

The Swing Riots of 1830 represented the agricultural counterpart. The first threshing machine was destroyed on **August 28, 1830**, at Lower Hardres, Kent. Wikipedia Before mechanization, labor-intensive threshing employed **25% of all agricultural workers**. History of Information The movement became **the largest popular uprising in England since the Peasants' Revolt of 1381**: Coam over 3,000 acts of revolt across southern England. Repression was severe: **2,000 tried**, 252 sentenced to death (19 hanged), 644 imprisoned, and **481 transported** to Australian penal colonies.

The transition's arc: from cottage to factory to Ford

The shift from cottage industry to factory system inverted social structures. Friedrich Engels observed that the family was *"turned upside down"* as women's factory wages undercut men's. In 1788, **two-thirds of workers in 143 water-powered cotton mills were children**. Wikipedia Agricultural employment declined from roughly World History Encyclopedia **75% of the English workforce around 1500 to 35% by 1800, 22%** Wikipedia **by 1850, and 10% by 1911**. Urban population rose from **9% in 1800 to 62% by 1900**. Between **1604 and 1914, over 5,200 enclosure acts** covered 6.8 million acres—one-fifth of England's total area.

The Jacquard loom (demonstrated 1801, patented 1804) proved pivotal not only for textiles but for computing itself. Science and Industry Museum Joseph-Marie Jacquard's use of **punched cards** to control individual warp threads— The Henry Ford a binary code system—directly inspired Charles Babbage's Analytical Engine (1837). History of Information Ada Lovelace wrote what is considered the first computer program for it. The Henry Ford Herman Hollerith adapted punch cards for the 1890 US Census; Wikipedia his Tabulating Machine Company became **IBM**. IEEE Spectrum A famous 1839 woven silk portrait of Jacquard required **24,000 punched cards**, each with over 1,000 hole positions. History of Information

Henry Ford's moving assembly line (1913) compressed the transition's logic into a single company. Assembly time for a Model T dropped from **12.5 hours to 93 minutes**, The Henry Ford eventually reaching 12 minutes by 1924. Hagerty Media But labor turnover hit **370%**—to maintain 14,000 workers, Ford hired 52,000 people in a year. Hfha His solution, announced **January 5, 1914**, was the **$5-per-day wage** for an 8-hour day—more than double the $2.34/day average. The Henry Ford On the first day, **12,000 job seekers** braved a blizzard and fire hoses at Highland Park. World Socialist Web Site The Model T's price fell from **$850 (1908) to $260 (1927)**.

The new industrial economy created new categories of work at scale. Walker Hanlon's NBER research found **439 engineers** in British biographical records born before 1850, with a "dramatic increase" between 1750 and 1850. By the 1860s, engineers produced **more patents than any other occupational group**. NBER By 1850, **a quarter-million railway navvies**—a force larger than the Army and Navy combined—had laid 3,000 miles of track. US manufacturing employment grew **fourfold from 2.5 to 10 million** between 1880 and 1920. PubMed Central The full transition from first textile innovations to widespread real wage growth took roughly **60–80 years**. Only after 1840 did wages finally catch productivity: from 1840–1900, real wages grew **123%** against 90% productivity growth. Wikipedia

---

II. The computing era: 95 times more programmers in 40 years

Disappearing occupations and the 271-occupation paradox

The early computing era (1950s–1990s) transformed office work with a striking result: of **271 occupations listed on the 1950 US census, only one—elevator operator—had been rendered obsolete by automation by 2010**. BIC Magazine This masks enormous compositional change. Clerical workers peaked at **19.3% of US employment in 1980**, Grokipedia then fell to 12% by 2019—returning to 1950s levels. In 1971, **one-third of all working women** were secretaries. Productevolution

AT&T, the largest US employer by the 1920s, employed over **350,000 telephone operators** at peak, **98% women**. By 1950, roughly **1 in every 13 working women** was a switchboard operator. BBN Times Automation was technically possible from the 1880s—the first automated system was installed in La Porte, Indiana on **November 3, 1892**— Federal Reserve Bank of Richm... but AT&T didn't install its first dial telephones until 1919, BBN Times and complete phase-out took until **1978**. The last hand-cranked system retired in **Bryant Pond, Maine, in 1983**. JSTOR Research by Feigenbaum and Gross found that when cities cut over to mechanical switching, **young female operator employment fell 50–80% immediately**— BBN Times but aggregate female employment didn't decline, as displaced workers shifted to secretarial and food service roles. NBER

The bank teller story illustrates a counter-intuitive dynamic central to the book's thesis. Bank teller employment **roughly doubled from ~300,000 in 1970 to ~600,000 by 2010**, (Wikipedia) even as **400,000+ ATMs** were installed. (BIC Magazine) James Bessen (Boston University) documented the mechanism: (CCIA) ATMs reduced tellers per branch from **20 to 13**, making branches cheaper to operate. Banks opened more branches. (IMF) Teller roles shifted from cash-handling to relationship banking. (ATM Marketplace) Only after mobile banking took hold did teller numbers finally decline to 364,100 by 2022. (Wikipedia)

Bank of America's ERMA system (operational **September 14, 1959**) demonstrates the pattern at company scale. Before automation, experienced bookkeepers processed **245 accounts per hour**; BofA's accounts were growing at **23,000 per month**, forcing banks to close by 2 PM to finish daily posting. (HandWiki) ERMA processed **550 accounts per minute** and **replaced 2,500 bookkeepers**. (Thebhc) It introduced MICR (Magnetic Ink Character Recognition), still used on checks worldwide, (Thebhc) and enabled BofA to become the world's largest bank by 1970. (HandWiki)

The spreadsheet revolution and job creation

The spreadsheet revolution follows the same script. VisiCalc (October 17, 1979), (Wikipedia) created by Harvard Business School student **Dan Bricklin** (LinkedIn) and programmer **Bob Frankston**, sold over **700,000 copies** (Wikipedia) and became the first "killer app"—people bought $2,000 Apple II computers to run a $100 program. Steve Jobs credited it with propelling Apple's success "more than any other single event." (The New Stack) Lotus 1-2-3 (January 26, 1983) conquered the IBM PC market, selling **175,000 copies in its first year**. (Wikipedia) The net labor impact: spreadsheets eliminated **400,000 accounting clerk positions** but created **600,000+ new accountant and auditor roles**. (The New Stack) BLS data shows 339,000 accountants in 1980 growing to **1.4 million by 2022**. (Tim Harford)

Computer specialists—the era's defining new occupation—grew **95-fold** as a share of employment between 1960 and 2000: from **12,000 workers (0.02% of employment) to 2,496,000 (1.92%)**. (bls) The term "software engineering" appeared in a June 1965 issue of *Computers and Automation* (Wikipedia) and was formalized at the **1968 NATO Software Engineering Conference** in Garmisch, Germany. (Engineering) Margaret Hamilton coined it during the Apollo program (Hack Reactor) while leading MIT's Software Engineering Division. (McNally Institute)

Grace Hopper, the ENIAC six, and NASA's human computers

Grace Hopper (1906–1992) (Wikipedia) created the first compiler (Discovertodayinhistory) (A-0 system) in 1952 while working on UNIVAC I. *"Nobody believed that,"* she said. *"I had a running compiler and nobody would touch it. They told me computers could only do arithmetic."* (Scientific Women) Her FLOW-MATIC language (1955–1959) became the predecessor to COBOL, (Wikipedia) (IEEE Computer Society) which by the 1970s was the world's most-used programming language (Yale) and still processes an estimated **95% of ATM transactions today**. The famous "bug" story dates to **September 9, 1947**, when a moth trapped in the Harvard Mark II's relays was taped into the logbook with the note: *"First actual case of bug being found."* (Discovertodayinhistory)

The six women who programmed ENIAC—Betty Holberton, Jean Bartik, Kay McNulty, Marlyn Meltzer, Frances Spence, and Ruth Teitelbaum— Eniac were selected from ~100 female "human computers." They had **no programming manuals** and no security clearance to see the hardware initially, working only from blueprints and logical diagrams. Programming involved configuring **3,000 switches** and dozens of cables. Witi They were largely forgotten for decades, IEEE Spectrum misidentified as "models" in photos, Wikipedia and not invited to ENIAC's 50th anniversary in 1996. Wikipedia

NASA's human computers exemplify adaptive response. **Dorothy Vaughan**, the first African-American supervisor at NACA (1949), U.S. National Park Service foresaw electronic computers' arrival. She taught herself FORTRAN from a library book and trained her entire West Area Computing team in programming, negotiating their transfer as a unit to the new Programming Department. Wikipedia **Katherine Johnson** verified the IBM computer's orbital equations for John Glenn's Friendship 7 mission (1962). Glenn famously asked engineers to *"get the girl"* to check the calculations: *"If she says they're good, then I'm ready to go."* NASA

III. The internet era: Jevons Paradox and the creation of the curator economy

Destruction by displacement: the Craigslist effect and retail apocalypse

The internet era (1995–2020) produced some of the most dramatic displacement metrics in economic history. Newspaper classified ad revenue peaked at **$19.6 billion in 2000**, constituting ~40% of total newspaper ad revenue. By 2012, it had fallen to **$4.6 billion—a 77% decline**. MinnPost A 2013 study by Seamans (NYU Stern) and Zhu (Harvard Business School) found Craigslist directly cost newspapers **$5 billion in lost classified revenue from 2000–2007 alone**. MediaPost Publications Wikipedia Total newspaper print ad revenue collapsed **92%**, from $73.2 billion in 2000 to $6 billion by 2023. Grokipedia Newsroom employment fell from ~71,000 in 2008 to ~31,000 by 2020, a **57% decline**. Pew Research Center Over **2,200 local newspapers closed** from 2005–2021. HopNews Wikipedia

Blockbuster's fall is the era's defining competitive parable. At its **2004 peak**, Blockbuster operated **9,094 stores** with **84,300 employees**. Wikipedia In September 2000, Reed Hastings and Marc Randolph flew to Dallas and offered to sell Netflix to CEO John Antioco for **$50 million**. ScreenGeek Executives reportedly *"struggled not to laugh."* toofab Blockbuster filed for bankruptcy on **September 23, 2010**. AARP Dish Network acquired the remnants for $320 million. Netflix was worth over **$150 billion by 2023**— Yahoo Finance 3,000× the rejected asking price.

Borders outsourced **all online sales to Amazon** from 2001–2008, "effectively handing customers and data to its biggest competitor for six critical years." Gizoom Consulting Slate It filed bankruptcy on **February 16, 2011** Substack with $350 million in debt FourWeekMBA and closed 399 stores, laying off 10,700 employees. Slate Kodak engineer **Steve Sasson** invented the first digital camera in **1975**; Kodak shelved it. The National At peak, Kodak commanded **90% of US film sales**, The CDO TIMES employed **145,300 people**, Rochester Business Journal and had a $28–31 billion market cap. The CDO TIMES It filed bankruptcy in January 2012. Time Today it employs approximately **3,400**— The National down **97.7%**. NYC taxi medallions, which sold for over **$1 million** in 2013–2014, The Flaw crashed to $80,000–$150,000 after Uber launched, Documented an **80–92% decline**. At least 8–9 NYC taxi drivers died by suicide from 2017–2018 amid financial distress. Northeastern University Politic…

Jevons Paradox: cheap information creates more demand for curation

William Stanley Jevons observed in *The Coal Question* (1865) that efficiency improvements in coal use led to *more* coal consumption, not less. (Wikipedia) He explicitly drew the labor parallel: *"The economy of labor effected by the introduction of new machinery throws labourers out of employment for the moment. But such is the increased demand for the cheapened products, that eventually the sphere of employment is greatly widened."* (ScienceDirect)

This pattern repeated precisely in digital markets. Cheap information (Google, Wikipedia) did not reduce demand for information services—it massively increased information volume and created new demand for filtering, curation, and expertise. Google processes an estimated **8.5 billion searches per day**. The "paperless office" paradox: word processing was supposed to reduce paper use, but US paper consumption *increased* after PCs were introduced. (MAHB) **Satya Nadella explicitly cited Jevons Paradox in January 2025** in response to DeepSeek: *"As AI gets more efficient and accessible, we will see its use skyrocket, turning it into a commodity we just can't get enough of."* (NPR)

The creator economy—which didn't exist in 2000—reached **$200 billion by 2023** with projections exceeding **$600 billion by 2036**. (Passive Secrets) YouTube's ecosystem contributed **$35 billion to the US economy in 2022**, supporting **390,000 full-time equivalent jobs**. (Whop) Approximately **200–300 million content creators** operate globally. (Passive Secrets) The gig economy grew to encompass **36% of the US workforce** (58 million adults) by 2022, (TechTarget) contributing **$1.27 trillion** to the economy. (Integra Credit) Software developer employment reached **1.7 million** by 2024 with a median wage of $133,080 (U.S. Bureau of Labor Statistics) and BLS-projected growth of **15% from 2024–2034**. (U.S. Bureau of Labor Statistics)

IV. AI capabilities and the trajectory toward SWE automation

The benchmark explosion: from 20% to 80% in fifteen months

The speed of AI coding capability improvement is unprecedented. On the **SWE-bench Verified** benchmark (500 real GitHub issues from Python repositories), (Epoch AI) top model scores went from approximately **20% in August 2024 to 75–80%+ by late 2025**. Claude Opus 4.5 reached ~79.2%, Gemini 3 Pro Preview hit 77.4%, and GPT-5.2 achieved ~75.4%. (Scale AI) HumanEval, a simpler coding benchmark, is now essentially **saturated** with top models exceeding 95%.

However, the harder **SWE-bench Pro** benchmark (731 problems from copyleft and proprietary codebases, requiring edits across an average of **107.4 lines and 4.1 files**) (Scale AI) reveals a significant gap: GPT-5 scored only 23.3%, Claude Opus 4.1 scored 23.1%, (Scale AI) and even Claude 4.5 Sonnet reached only 45.8%. Models scoring 70%+ on Verified drop to 20–45% on enterprise-realistic tasks—highlighting a persistent chasm between benchmark performance and real-world complexity.

Adoption metrics and the code authorship question

AI coding tools have achieved remarkable adoption. **GitHub Copilot** reached **20 million all-time users** by July 2025, TechCrunch with **90% of Fortune 100 companies** using it. Google Copilot generates an average of **46% of code** written by active users (up from 27% at launch in 2022), reaching 61% for Java developers. Google GitHub's study with Accenture (4,800 developers) found **55% faster task completion** on JavaScript HTTP server tasks. Quantumrun PR completion time dropped from **9.6 days to 2.4 days**. Google

Cursor (Anysphere) represents the AI coding tool market's meteoric growth: revenue trajectory from $1M (2023) → $100M ARR (end 2024, fastest SaaS company ever to $100M) → **$1B+ ARR (November 2025)**. Valuation soared from $400M to **$29.3 billion** in roughly 15 months. The company had only ~40–60 employees. **Devin** (Cognition AI), the autonomous AI software engineer, grew ARR from $1M (September 2024) to **$73M (June 2025)**, swyx with 67% of its pull requests now merged. Cognition Real-world metrics include a **20× efficiency gain** on security vulnerability fixes (1.5 minutes vs. 30 minutes for humans) Cognition and **10–14× speed improvements** on code migrations. Cognition

Company-reported data on AI-generated code: **Google** reports "well over 30%" of new code is AI-generated (Sundar Pichai, April 2025 earnings call); Medium **Microsoft** reports 20–30%; Entrepreneur a **quarter of Y Combinator's** current cohort has codebases that are **95% AI-generated** (Jared Friedman, March 2025). CodeRabbit Industry-wide, an estimated **41% of all code is now AI-generated** (Q1 2025). Netcorpsoftwaredevelopment However, only ~30% of AI-suggested code is actually accepted by developers, 75% still manually review every snippet before merging, Netcorpsoftwaredevelopment and **39% of developers do not trust AI-generated code**. IT Pro

Expert predictions: optimists vs. skeptics

The prediction landscape is sharply divided. **Dario Amodei** (Anthropic CEO, March 2025) predicted AI would write 90% of code within "three to six months" IT Pro and "essentially all of the code" within 12 months. Entrepreneur **This prediction did not materialize** by its stated timeline. IT Pro **Sam Altman** predicted AI would become "the best coder in the world" by end of 2025. Fello AI **Kevin Scott** (Microsoft CTO) predicted AI would write **95% of code within five years**. CodeRabbit **Jensen Huang** advised young people to explore fields other than coding.

Skeptics push back forcefully. **Grady Booch** (IBM Chief Scientist for Software Engineering, UML co-creator) called Amodei's prediction *"utter BS,"* arguing he has *"a fundamental misunderstanding as to what software engineering is."* Booch views AI as another abstraction layer—like compilers replacing assembly—not a replacement for engineering judgment. **Daron Acemoglu** (MIT, 2024 Nobel laureate) projects only **0.5–0.7% total factor productivity growth** from AI over the entire decade, arguing only ~5% of economy-wide tasks can be profitably automated. mit MIT Technology Review The US **BLS projects 17% job growth** for software developers from 2023–2033. A joint **GitHub/Gartner white paper** was titled *"AI Will Not Replace Software Engineers (and May, in Fact, Require More)."*

A January 2026 Harvard Business Review study of 1,006 executives found that **companies are laying off workers "because of AI's potential—not its performance."** Forrester's Predictions 2026 report warns that half of AI-attributed layoffs will be quietly rehired, citing Klarna's experience: it replaced 700 employees with AI, quality declined, customers revolted, and the company rehired humans.

Labor market evidence: the Stanford study and company actions

The most rigorous empirical evidence to date comes from Stanford's Brynjolfsson et al. (August 2025), using ADP payroll data covering 3.5–5 million US workers. DigitalRosh **Employment for software developers aged 22–25 declined by nearly 20%** American Association for the A… from its peak in late 2022 to July 2025. Claude DigitalRosh Early-career workers in AI-exposed occupations experienced a **13% relative employment decline**, Aiworld while workers aged 30+ in the same occupations saw **6–9% growth**. DigitalRosh Adjustments are happening through decreased employment rather than lower wages. DigitalRosh

Major companies are acting on this trajectory. **Salesforce** announced zero new software engineer hires in 2025, citing 30%+ productivity from AI. All Things Talent **Microsoft** laid off 15,000 in 2025, with 40%+ targeting software engineers and product managers. **CrowdStrike** cut 500 employees (5% of workforce), directly attributing the decision to AI. **Amazon** cut 14,000 corporate roles in October 2025. Reuters reports new graduate hiring for software roles in 2023 was **50% lower than 2019**. Indeed data shows the top four roles cut in AI restructuring are: software engineers/developers, QA engineers, product managers, and project managers.

V. What becomes scarce when software becomes cheap

Problem definition emerges as the new bottleneck

When code production approaches near-zero marginal cost, the constraint shifts upstream. Rich Mironov (40-year product veteran) maps the bottleneck explicitly: *"When we speed up one step in our larger process, another step becomes the slowest one."* Mironov With 15× software development productivity, the bottleneck moves to **discovery and problem identification**. Ethan Mollick (Wharton) introduces the "jagged frontier"—AI is superhuman at some tasks but subhuman at others. One Useful Thing The bottlenecks concentrate where AI is weakest: getting organizational buy-in, understanding context, spotting anomalies, and navigating politics.

Clayton Christensen's **"Jobs to Be Done" framework** quantified the problem definition gap: **75–85% of all new products fail financially** because "they don't target a job that people are trying to get done." HBS Online When production cost approaches zero, the only differentiator becomes understanding which problem to solve. Product management roles have recovered significantly, with **over 6,000 open PM roles globally** in early 2025—up 53.6% from the 2023 bottom. Substack AI PM roles command **$130,000–$200,000**; Product School CPOs earn $186,000–$290,000. Product School

Historical analogies reinforce this pattern. When cameras became cheap (smartphones), the bottleneck shifted from image capture to curation and storytelling. When publishing distribution became near-free (internet), it shifted to content quality and audience building. When assembly lines removed the production bottleneck, logistics became the constraint. Mironov

Cybersecurity: more software means more attack surface

4.8 million cybersecurity roles remain unfilled globally (ISC2 2025 data, up from the widely cited 3.5M figure). The global workforce of ~5.5 million must increase by **87%** to meet demand. (Programs) The US alone has **514,359 active cybersecurity job postings** with 26% of roles vacant. (SQ Magazine) BLS projects **33% employment growth from 2023–2033** for information security analysts—one of the fastest growth rates across all occupations—with a median salary of **$124,910**. (California Miramar University) The global cybersecurity market is estimated at **$215 billion in 2025**, projected to reach **$697 billion by 2035**. (Business Wire)

AI introduces entirely new attack vectors. **Prompt injection** is the #1 risk on OWASP's 2025 LLM Top 10. (CrowdStrike) CrowdStrike has analyzed **300,000+ adversarial prompts** and tracks over 150 injection techniques. (CrowdStrike) Model poisoning, adversarial examples, jailbreaking, and memory poisoning represent novel threats with no pre-AI equivalents. (Wiz) AI governance roles grew **300%+ between 2022–2023** (LinkedIn data). (Tech Jacks Solutions) Technical AI governance professionals earn a median of **$221,000** (IAPP 2025). (Tech Jacks Solutions)

Physical infrastructure becomes the binding constraint

Data center construction investment has become perhaps the most dramatic physical-world consequence of the software revolution. US data center construction starts reached **$77.7 billion in 2025**—a 190% year-over-year increase. Monthly spending jumped from $500 million (mid-2021) to **$6.5 billion (December 2025)**. (Equipment World) US spending on data center construction has **officially surpassed traditional office buildings** for the first time. (Built In) The eight largest hyperscalers expect combined CapEx of **$371 billion in 2025** (+44% year-over-year). (Deloitte Insights) McKinsey estimates **$7 trillion** in global data center CapEx by 2030. AI infrastructure investment drove **over one-third of US GDP growth** in the first nine months of 2025. (Bismarckanalysis)

Power is the **#1 constraint**: 72% of executives cite grid capacity as very or extremely challenging (Deloitte 2025), (Information Week) (Deloitte Insights) with 7-year waits for some grid connections. (Deloitte Insights) US data center electricity consumption reached **183 TWh in 2024** (4%+ of total US electricity), projected to grow **133% to 426 TWh by 2030** (IEA). (Pew Research Center) The tech giants are responding with unprecedented energy deals: **Microsoft's** 20-year, $16 billion PPA with Constellation Energy to restart Three Mile Island (837 MW, (Introl) 3,400 new jobs); **Google's** order for 500 MW of small modular reactors from Kairos Power; (IEEE Spectrum) **Amazon's** $20B+ nuclear investment. (Introl) Big tech has collectively signed **10+ GW of new US nuclear capacity** in the past year. (Introl)

Construction of 2,788 announced or under-construction US data centers is projected to create roughly **4.7 million temporary construction-related jobs** (American Edge Project). (Equipment World) The CHIPS Act will create approximately **125,000 semiconductor manufacturing jobs** (NIST). TSMC's Arizona investment exceeds **$165 billion** across six fabs, (Arizona Commerce Authority) creating **6,000 direct high-tech jobs** plus tens of thousands of construction and supplier positions. (TSMC)

Baumol's cost disease and the economics of complements

Baumol's cost disease (1965) explains the economic mechanism: sectors requiring intensive human labor experience rising costs relative to sectors where technology boosts productivity. IG&H *"It still takes four musicians nine minutes to perform Beethoven's String Quartet in C minor."* UNESCO When AI makes software production cheap, everything AI *cannot* automate—physical care, trust, judgment, construction, human connection—becomes relatively more expensive and thus captures more economic value. As William Nordhaus frames it: *"Growth is determined not by what we are good at but rather by what is essential and yet hard to improve."* Medium

Complementary goods economics reinforces this logic. When the price of one good falls dramatically, demand for its complements increases. If software becomes the "cheap razor," the "blades" capturing value become domain expertise, physical implementation, security, governance, and human judgment. The WEF's Future of Jobs Report 2025 confirms this: the fastest-growing skills include AI and big data Gloat alongside **creative thinking, resilience, and flexibility**—human capabilities that complement rather than compete with AI. World Economic Forum

---

VI. Economic theories of technological unemployment

Autor's framework: tasks, polarization, and the missing middle

David Autor's **task-based framework**, developed with Levy and Murnane in their seminal 2003 paper *"The Skill Content of Recent Technological Change,"* remains the most influential model for analyzing automation's labor impact. It categorizes work into four task types: routine cognitive (bookkeeping), routine manual (assembly), non-routine abstract (problem-solving), and non-routine manual (caregiving). Computers **substitute** for routine tasks while **amplifying** the comparative advantage of workers in non-routine abstract tasks. ResearchGate

This framework predicted **job polarization**: from 1980–2010, employment grew in high-skill abstract jobs *and* low-skill service jobs, while middle-skill routine jobs declined sharply. Workers whose occupations were most exposed to routine-task automation saw *"sharp falls in real earnings from 1980 forward"*—a relationship absent before 1980, prior to large-scale commercial computerization. National Academies

Autor's 2015 paper *"Why Are There Still So Many Jobs?"* introduced two crucial concepts. **Polanyi's Paradox** (*"We know more than we can tell"*) argues many human capabilities are tacit and resist codification. The **O-Ring Principle** holds that most work requires a multiplicity of skills—automating one part doesn't make other parts less important. He explicitly addressed the ATM paradox: ATMs quadrupled from 1995–2010, yet bank teller employment rose from 500,000 to ~550,000.

His most recent work pivots to cautious optimism. In *"Applying AI to Rebuild Middle Class Jobs"* (NBER Working Paper 32140, February 2024), Autor argues AI's unique opportunity is to **"extend the relevance, reach, and value of human expertise"** and potentially restore the hollowed-out middle of the labor market. Medical care currently restricted to doctors, legal work to lawyers, software coding to engineers—AI could *"enable a larger set of workers possessing complementary knowledge to perform some of the higher-stakes decision-making tasks."* He distinguishes sharply between "automation tools" that eliminate expertise and "collaboration tools" that multiply it.

Acemoglu on "the wrong kind of AI" and so-so technology

Daron Acemoglu (MIT, 2024 Nobel laureate) offers the most skeptical major economic voice. His central thesis: recent technological change has been **"biased towards automation, with insufficient focus on creating new tasks where labor can be productively employed."** The consequence: *"stagnating labor demand, declining labor share in national income, rising inequality and lower productivity growth."*

His concept of **"so-so technology"**—automation that is "just good enough to be adopted but not so much more productive than the labor it replaces"—is particularly relevant to current AI coding tools. Call-center automation exemplifies this: *"not always more productive than people; it just costs firms less."* With so-so automation, the displacement effect exists while *"powerful productivity gains contributing to labor demand are missing."*

In *"The Simple Macroeconomics of AI"* (2024), Acemoglu estimates AI will increase US GDP by **1.1–1.6% over the next 10 years**, roughly **0.05% annual productivity gain**. He bases this on findings that ~20% of US job tasks are AI-exposed, ~23% of automatable tasks could be profitably automated within 10 years, and average cost savings from AI are ~27%. *"I don't think we should belittle 0.5 percent in 10 years. That's better than zero. But it's just disappointing relative to the promises."* He identifies institutional failures: tax policies that *"subsidize capital and investment while taxing employment"* systematically bias toward automation over augmentation.

The Frey-Osborne finding and its revision

Carl Benedikt Frey and Michael Osborne's 2013 paper *"The Future of Employment"* produced the headline finding that **47% of total US employment is in the "high risk" category** for automation. Their novel methodology used Gaussian process classifiers to estimate computerization probability for **702 detailed occupations**. Three "engineering bottlenecks" limited automation: perception/manipulation, creative intelligence, and social intelligence.

The finding was substantially revised. Arntz et al. (2016, OECD) found only **9% of US jobs** at high risk using a task-based approach instead of an occupation-based approach: *"while Frey and Osborne find that 47% of US jobs are automatable, our corresponding figure is only 9%... This substantial difference is driven by the fact that even in occupations that Frey and Osborne considered to be in the high risk category, workers at least to some extent also perform tasks that are difficult to automate."* Cross-country OECD estimates ranged from 6% (Korea) to 12% (Austria).

The 2025 National Academies report notes a crucial evolution: *"unlike earlier waves of automation that mainly affected routine manual or clerical tasks, today's AI encroaches on non-routine cognitive and creative tasks once thought to be the exclusive preserve of humans."*

Keynes, Schumpeter, and the lump of labor

Keynes coined "technological unemployment" in *"Economic Possibilities for our Grandchildren"* (1930): *"unemployment due to our discovery of means of economising the use of labour outrunning the pace at which we can find new uses for labour."* He predicted living standards **4–8× higher** by 2030 (US GDP per capita has risen ~6.5×, on track for his range) and a **15-hour workweek** (which never materialized—Americans still work ~8 hours/day). His key error: underestimating insatiable human wants.

Schumpeter's **Creative Destruction** framework (*Capitalism, Socialism and Democracy*, 1942) describes capitalism as *"the process of industrial mutation that incessantly revolutionizes the economic structure from within, incessantly destroying the old one, incessantly creating a new one."* Empirical support: creative destruction accounts for **over 50% of ten-year productivity growth** in US manufacturing (1977–1987). Only five of today's hundred largest public companies were among the top hundred in 1917.

The **"lump of labor fallacy"** (Schloss, 1892) warns against assuming fixed work quantity. Historical evidence is strong: US agriculture went from 41% of workers (1900) to ~2%, yet total employment grew massively. But Daniel Susskind (*A World Without Work*) offers a potent rebuttal: *"The fallacy only suggests that there will always be more work. It doesn't suggest that humans would do the work."* Previous escapes worked because machines couldn't replicate all human capabilities. *"However, humans only have physical and cognitive abilities, and when machines become more economical in both, there is simply no role left."*

The productivity paradox: Solow's ghost haunts AI

Robert Solow wrote in 1987: *"You can see the computer age everywhere but in the productivity statistics."* Erik Brynjolfsson formalized this as the **"Productivity Paradox"** in 1993: despite massive IT investment in the 1970s–80s, labor productivity growth slowed from **over 3% in the 1960s to roughly 1% in the 1980s**. Explanations include measurement errors, adoption lags (2–5 years for IT investments to show impact), and the need for complementary organizational change. Productivity did surge in the late 1990s, attributed to IT diffusion plus organizational restructuring. As of 2025, the same paradox recurs: *"AI is everywhere except in the incoming macroeconomic data"* (Apollo Academy). MIT researchers describe a potential "Productivity J-Curve" where stagnation precedes a sharp increase.

VII. Emerging roles in a post-SWE-automation world

AI governance and the system governor concept

AI governance is emerging as one of the fastest-growing professional categories. IAPP's 2025 survey of 1,600+ professionals found that **98.5% of organizations report AI governance staffing shortages**. PwC data shows AI-related job postings have grown **257% since 2015**, with workers possessing AI skills earning **56% higher wages** than peers without AI expertise. Technical AI governance professionals earn a median of **$221,000**; combined privacy and AI governance roles pay **$169,700**. The EU AI Act, with penalties of up to **€35 million or 7% of global annual turnover**, is creating mandatory compliance infrastructure. The EU AI Office is hiring 140 employees across technology, law, economics, and administration.

The **"system governor"** concept finds its closest existing analogues in nuclear plant operators and air traffic controllers—roles where automation has advanced enormously but human oversight remains non-negotiable. US nuclear reactor operators earn a median of **$103,600** and undergo **18–24 months** of plant-specific training before licensing, with roughly **one week out of every five** spent in ongoing training. Air traffic controllers earn a median of **$144,580**. Despite sophisticated automation, the FAA maintains human controllers because **accountability chains** require human judgment for safety-critical decisions.

Site Reliability Engineers (SREs) are evolving into proto-system-governors. The emerging model defines graduated autonomy levels: Level 1 (AI co-pilot, humans control), Level 2 (AI acts within guardrails, human approval), Level 3 (autonomous with guardrails). Engineers increasingly *"define policies, specify acceptable actions, and encode business intent"* rather than perform hands-on operations. EU AI Act Article 14 mandates human oversight for high-risk AI systems, creating legal requirements for these roles. New titles appearing include: AI Operations Manager, Responsible AI Lead, AI Risk Manager, and Model Validation Specialist.

The hardware design renaissance

Custom silicon is experiencing explosive growth. **TrendForce projects custom ASIC shipments from cloud providers to grow 44.6% in 2026**, versus GPU shipments growing 16.1%. Key players: Google TPU, AWS Trainium/Inferentia, Tesla Dojo, Apple Neural Engine. The startup ecosystem is thriving: **Groq** (acquired by Nvidia for ~$20B in December 2025), **Cerebras** ($8.1B valuation, wafer-scale chip with 4 trillion transistors), and **Tenstorrent** ($2.6B valuation, led by legendary architect Jim Keller, using RISC-V open architecture). In Q1 2025 alone, 75 AI chip startups raised **$2+ billion**. US semiconductor startup funding hit record highs in 2025.

Robotics: the physical world demands human judgment

Global robotics market estimates range from **$50–108 billion in 2025**, projected to reach **$147–218 billion by 2030–2031** (17–20% CAGR). The IFR's World Robotics 2025 report recorded **542,000 industrial robots installed globally in 2024**—the fourth consecutive year above 500,000—with **4.66 million robots in operational use worldwide**. The humanoid robot segment is the most dramatic growth story: **$1.9 billion in 2025**, projected to reach **$11 billion by 2030** at 42.8% CAGR. **Figure AI** reached a **$39 billion valuation** in September 2025 (15× increase in 18 months), deploying humanoid robots at BMW's manufacturing plant. Goldman Sachs projects **1.4 million humanoid shipments by 2035**. Crucially, the IFR notes humanoid robots remain unproven for *"economically viable and scalable business cases"* in industrial applications—the physical world resists Moore's Law.

Bioengineering: promise and reckoning

Synthetic biology has experienced both extraordinary promise and sobering failures. The market is estimated at **$19–24 billion in 2025**, with growth projections ranging from 10.7–26% CAGR. **Zymergen** represents the cautionary tale: IPO'd at ~$5B valuation (April 2021), collapsed after its key product failed commercially, was acquired by Ginkgo Bioworks for $300M, then filed Chapter 11 bankruptcy (October 2023). **Ginkgo Bioworks** itself laid off one-third of its workforce in June 2024 after its stock tanked. Jay Keasling (Lawrence Berkeley National Laboratory) acknowledged *"there has been a reckoning"* in synthetic biology. Yet the AI-biology intersection remains a powerful growth vector: AlphaFold revolutionized protein structure prediction, creating demand for computational biologists. Capgemini launched a protein large language model in February 2025 reducing protein design data points by 99%.

Conclusion: the one-abstraction-layer-up thesis meets its hardest test

The historical pattern is clear and consistent. Each era's cheap capability layer—mechanical production, electronic computation, internet distribution—created devastating short-term displacement followed by explosive job creation one abstraction layer up. Handloom weavers became factory engineers. Typing pools became software developers. Travel agents became platform architects. Across all three eras, the **Jevons Paradox held**: cheaper capability meant more total demand for that capability, not less. Bank tellers doubled after ATMs. Accountants grew 4× after spreadsheets. Software development jobs grew alongside every tool designed to reduce the need for software developers.

Three structural differences make the current transition potentially unlike its predecessors. First, **speed**: SWE-bench scores tripled in 15 months, Cursor reached $1B ARR in under two years, and the Stanford study already documents measurable employment effects for junior developers. Previous transitions took decades. Second, **scope**: unlike the Jacquard loom or the spreadsheet, AI simultaneously encroaches on routine *and* non-routine cognitive tasks—what the National Academies calls tasks *"once thought to be the exclusive preserve of humans."* Third, **the substitution question**: Daniel Susskind's point that previous labor market escapes worked because *"there were still tasks in those sectors that machines couldn't perform"* has never been more relevant.

Yet the evidence also reveals powerful countervailing forces. **4.8 million cybersecurity roles sit unfilled.** Hyperscalers are spending **$371 billion** on physical infrastructure. AI governance roles have grown **300%**. The robotics industry needs humans for judgment, dexterity, and accountability that machines cannot provide. The binding constraint is shifting from *building software* to *defining problems, governing systems, and constructing physical infrastructure*—precisely the pattern predicted by Baumol's cost disease and complementary goods economics. The key variable, as Acemoglu argues, is whether institutions direct AI toward augmentation or pure automation. The historical record suggests this is not a technological question but a political and organizational one—and it is being answered, for better or worse, right now.