



트레이드 렌즈
TRADE LENS

Bringing investing into focus.

LENS INSIGHT / No.04

Knowledge production, consumption and work

Will White-Collar Jobs Disappear?

AI Ushers in the Era of Mass-Produced Knowledge

Knowledge work, consumption and jobs
through the lens of the Industrial Revolution

Lower production barriers



Wider knowledge use



Allocation of gains

Lower cost · More questions · Better conditions for work

Technology raises productivity.

Demand and distribution determine jobs and welfare.

01 / RESEARCH VIEW

A threat to jobs, or a release of supply constraints?

TradeLens view | AI matters not only because it substitutes for a fixed workload, but because it can make knowledge services available to more people.

Core conclusion

AI moves parts of knowledge work from individual manual effort into repeatable production systems. This creates a structural opportunity to broaden access. Yet greater consumption, stable employment and better working conditions are distinct outcomes; they need not arrive together.

Question	This report's answer
What is changing?	Cheaper drafting, adaptation and repeat responses can let the same workforce serve more varied customers and questions.
Why does it matter?	Analysis, education and advice that were previously too costly may become viable, expanding the boundary of the market.
Who benefits?	Consumers may gain access and affordability, firms efficiency and new demand, and workers capability and time. Actual benefits depend on how gains are allocated.
What should be tested?	Cost at a fixed quality standard, actual use, new customers, hiring, pay and working hours—not output counts alone.

The evidence is task-level; the economic question is systemic

Writing and customer-support studies show gains in particular tasks. These estimates cannot be converted into whole-job headcount reductions. Employment also depends on demand responses and the creation of new tasks. [\[4\]](#) [\[5\]](#) [\[15\]](#)

The right unit of analysis

Count **knowledge outputs that pass verification and are actually used**, not files. Faster drafting is not knowledge abundance if it merely increases review burdens or unused documents.

Analysis date: 3 October 2026. Each study retains its own population and observation period. Currency, time and elasticity examples are explicitly hypothetical.

02 / HISTORICAL LENS

Industrialization offers a mechanism, not a guarantee

The spread of factories does not establish that every worker immediately enjoyed easier or better-paid work.

Automation can reduce labor per unit while increasing output and total labor demand if lower prices elicit enough additional demand. Bessen's historical manufacturing analysis emphasizes this channel and the importance of eventual demand saturation. [1]

Dimension	The factory lens	Question for AI knowledge production
Production	Individual tasks become organized processes, equipment and specialization	Does individual drafting become a data-generation-verification-delivery workflow?
Market expansion	Lower unit costs enable wider purchasing	Do new customers use analysis, explanations and education they previously could not afford?
Remaining limits	Capacity, quality and demand constraints	Verification, attention, trust and accountability constraints
Labor outcomes	Output and labor per unit must be considered together	Measure total tasks, hours and rewards, not generation speed alone

Better work did not automatically arrive with technology

Allen emphasizes the gap between early industrial productivity gains and real wages. Crafts reassesses the size and interpretation of that gap. The timing and distribution of historical gains remain debated; neither account guarantees immediate benefits for every worker. [2] [3]

The AI analogy is about customization, not merely copying

Digital files could already be copied and transmitted in the internet era. This report focuses instead on the cost of creating outputs for different questions, languages and levels of understanding. The "knowledge factory" is an analytical analogy for **mass customization**, rather than mass duplication of identical documents.

Where the analogy stops

Employment and benefits in today's semiconductor factories are not direct evidence of AI's future effects on office workers. Industry maturity, skills, capital intensity and the allocation of gains differ.

03 / PRODUCTIVITY EVIDENCE

Real productivity gains, different boundaries

Time, throughput and quality are different metrics—even when all are reported as percentages. These studies are not a performance ranking.

Study and design	Observed result	Boundary of interpretation
Noy & Zhang, 2023 [4] 453 professionals; randomized trial	Average writing time fell 40%; rated output quality rose 18%	Bounded professional writing tasks, not whole jobs or long-run employment.
Brynjolfsson et al., 2025 [5] 5,172 support agents; staggered rollout	Issues resolved per hour rose 15% on average	One firm, with rollout data from 2020–2021. Not a 15% wage increase.
METR update, 2026 [6] Experienced developers doing real work	Early-2025 experiment: completion time rose 19%	Selection bias limits interpretation of follow-up results. Not a general estimate for 2026 tools.

What the evidence supports

The evidence does not support a uniform reduction across all work. It does support a change in the production of some knowledge tasks. The business challenge is not to select the largest headline saving, but to reproduce useful gains within the organization's own workflow. [4] [5] [6]

Separate tasks

Measure sourcing, drafting, verification and approval separately.

Hold quality fixed

Set quality standards for accuracy, evidence and relevance first.

Measure end to end

Include waiting, correction, rework and tool costs.

TradeLens interpretation

“Writing faster” is a starting point. Economic value emerges when a usable deliverable can be completed with fewer total resources.

Metrics preserve each study's original definition. Results are neither pooled nor converted into an average effect for Korean office workers.

04 / KNOWLEDGE PRODUCTION SYSTEM

Mass-produced knowledge needs quality control

The relevant unit of office automation is not a document generator. It is a workflow from the initial question to verification and actual use.




Diagram: TradeLens analytical framework, not a description of an implemented system or a company's measured performance.

Less drafting can mean a larger share of verification

NIST flags confidently incorrect content and fabricated citations, and recommends checking sources. Fast drafts still need review of original evidence, calculations, scope and responsibility. [\[14\]](#)

Stage	Human judgment to retain	Operational output
Question and evidence	What decision matters, and which evidence is valid?	Scope, original sources and dates
Generation and review	Are facts and inferences sound, and what is missing?	Evidence links, reconciled numbers and approval
Delivery and feedback	Can the user understand and apply the result?	Audience-specific explanation and feedback

Research on general-purpose technologies highlights complementary investments in training and organizational redesign. Buying a tool alone does not complete the production system. [\[13\]](#)

The bottleneck shifts

As drafts become less scarce, **good questions, reliable data, verification and accountability** become more important competitive assets. Scaling output without scaling quality control can increase the volume of unreliable outputs.

05 / KNOWLEDGE CONSUMPTION

From more information to broader understanding

Knowledge abundance should be tested through access, the range of questions served, personalization and learning—not output counts.

A 2025 study of consumer ChatGPT use found that non-work messages grew from 53% to over 70% of usage, while work use also grew. This indicates consumption beyond the workplace, not an economy-wide welfare gain. It concerns one service and includes provider-affiliated authors. [\[10\]](#)

Expansion channel · analysis	Possible change	Outcome to measure
Lower cost and time barriers	People with smaller budgets or less time can access explanations and analysis	New users, repeat use and total access cost
Customization	The same subject is explained for different languages, levels and contexts	Comprehension, independent performance and fewer misunderstandings
Small, specialized needs	Infrequent and highly specific questions become worth serving	Previously unmet problems resolved
Continuous updating	One-off documents are reinterpreted as conditions change	Current information and decision use

Education shows both the opportunity and the boundary

A study involving 194 university physics students found learning benefits from a specifically designed AI tutor. In a Turkish high-school mathematics experiment, unrestricted AI assistance was followed by 17% lower performance than the control group once the tool was removed. Different populations and designs distinguish receiving answers from learning. [\[11\]](#) [\[12\]](#)

TradeLens view

The promising form of knowledge abundance is not spending more hours reading documents. It is **obtaining relevant explanations at a lower cost of search and understanding**. More output without better comprehension or use would weaken this thesis.


06 / ILLUSTRATIVE UNIT ECONOMICS

An 80% draft saving is not an 80% total saving

Illustrative assumptions | One analysis report meeting a fixed quality standard. These are not a company's observed costs, fees or productivity forecast.

Process	Before: human hours	With AI: human hours	Assumption
Question and scope	1.0 h	1.0 h	Judgment time unchanged
Evidence and organization	2.0 h	1.2 h	Less preparation
Drafting	4.0 h	0.8 h	80% drafting saving
Verification and correction	2.0 h	2.5 h	25% more review
Editing and delivery	1.0 h	0.5 h	Less format work
Total	10.0 h	6.0 h	40% less human time

Components of direct production cost



Unit: KRW per report meeting the stated quality standard. Labor: KRW 40,000/hour; AI tools: KRW 20,000/report. Fixed setup and training, finance, taxes and error losses are excluded. Bars show direct production-cost components.

Production capacity is not sales volume

With unchanged total human hours, potential throughput becomes $10/6 = 1.667$ times the baseline, or **66.7% higher**. Sales do not rise by that amount unless demand uses the released capacity. Review skills, data rights and setup costs can also change the realized saving.

What the calculation makes visible

The key metric is total human time and cost per output at a fixed quality standard. Even here, the 35% direct-cost saving and 40% human-time saving are different measures.

07 / DEMAND AND LABOR

Consumption can rise while total labor hours fall

Illustrative assumptions | Keep human time per unit at 0.60 of baseline; the price falls 35%, to 0.65, in one market for a comparable service.


Model

Quantity ratio = $0.65^{-\epsilon}$ | Human-hours ratio = $0.60 \times 0.65^{-\epsilon}$

Proportional price pass-through; fixed quality, mix and hourly labor cost. ϵ = price elasticity.

Total human-hours index

(before = 100)



Assumed elasticity ϵ	Quantity change	Total human-hours change
0.5	+24.0%	-25.6%
1.0	+53.8%	-7.7%
1.5	+90.8%	+14.5%
2.0	+136.7%	+42.0%

Maintaining labor hours requires 66.7% more demand

At 0.60 times labor per unit, quantity must rise $1/0.60 = 1.667$ -fold to hold hours constant. The price assumption implies an elasticity threshold of **1.186**. At $\epsilon = 1.0$, use grows but hours fall. [\[1\]](#)

All values are TradeLens illustrations, not estimated elasticities for Korean knowledge services. This partial-equilibrium model excludes new tasks, self-provision, income effects and shifts between markets. Total hours are not headcount.

Analytical implication

Demand growth does not automatically mean employment growth. **Actually used or purchased services of comparable quality**, not generated documents, must cross the labor-hours threshold.

08 / LABOR MARKET EVIDENCE

Stable aggregates can conceal a narrower entry route

Exposure, pay and hiring answer different questions. Evidence from different countries is not combined into a single unemployment forecast.

Evidence and scope	What it shows	What it does not show
ILO-NASK, 2025 [9] Global occupational-task exposure	About 25% of employment is in occupations with some potential exposure to generative AI	Not a prediction that 25% of workers will lose their jobs.
Danish study [7] March 2026 paper; roughly two years after launch	Linked adoption surveys and records find no clear average earnings or recorded-hours effect; effects above 2% are ruled out	Not evidence of no long-term impact or no individual losses.
US ADP sample [8] Data to June 2026; August 2026 release	Employment of ages 22-25 in highly exposed occupations is about 19% below the path implied by less-exposed peers	Not a 19% layoff rate or a causal estimate of AI's effect.

Watch hiring, not only layoffs

The US authors find no widespread economy-wide displacement, but describe the young-worker gap as operating mainly through reduced hiring. They also note sensitivity to education, earlier trends and sample representativeness. The estimate cannot be transferred to Korean employment. [8]

TradeLens interpretation | Opportunity and transition risk coexist

A firm may deliver more services with existing staff while hiring fewer newcomers. New customers and tasks may subsequently create additional labor demand. Task creation is a distinct channel alongside automation's displacement effect. [15]

What to test

Average pay among incumbents is insufficient. Track **entry-level recruitment, skill development, occupational mobility and labor demand in new services** to see whether more knowledge production also creates broader opportunity.

09 / WORKER WELFARE

Whose time is being saved?

Better work is not a software feature. It is the conversion of productivity gains into time, income and autonomy.

Weekly hours required at different output targets




Illustration: baseline 40 hours/week, 40% less human time per unit, constant staffing and quality standards. Values are hours required at different output levels, not forecasts of actual working hours.

Desired benefit	Condition that must accompany it	Measure
Less routine work and overwork	Do not absorb every saving into a higher workload target	Actual hours, after-hours contact and rework
Higher income	A route to sharing gains amid changing prices and margins	Take-home pay and hourly compensation
Greater autonomy	Retain human discretion over tools and final approval	Decision rights and authority to correct errors
Skills and entry opportunities	Train evidence judgment and verification, not just drafting	Unaided assessment, junior hiring and advancement

Work can be reorganized before earnings or hours change. The Danish evidence reinforces the distinction between changing task structure and changing economic rewards. [\[7\]](#)

TradeLens view

AI can increase the opportunity to remove tedious work, but it does not automatically improve working conditions. **Workload targets, review burdens, compensation and learning arrangements** determine what remains with workers.

10 / ECONOMIC VALUE CAPTURE

Abundance does not guarantee shared profits

TradeLens analysis | More document generation and higher supplier earnings are not the same outcome.

Value-chain layer	Source of potential value	Constraint on profit capture
Compute and models	Serving more generation and reasoning demand	Competition, inference costs and capital investment
Data and domain knowledge	Providing distinctive facts, context and usage rights	Acquisition, rights, updating and maintenance
Workflow software and distribution	Connecting tools to daily work and customer access	Adoption friction, switching and alternatives
Verification and expert services	Delivering decision-ready quality and accountability	Expert review, errors and rework
Consumers and workers	Access, capability and time savings	Price pass-through, attention and gain sharing

Assess contribution margin per service, not revenue alone

Service-unit economics | Illustrative framework

Contribution margin per service = revenue per service - variable cost per service

Unit: one service delivered after review against a consistent quality standard.

Variable costs include review and rework labor, AI inference and tool usage, and other variable costs.

Business-level profit also accounts for fixed labor, customer acquisition, integration, training and fixed operating costs.

Do not count the same cost twice. This is not a company-margin estimate.

Where generation is weakly differentiated, competition may pass cost savings on through lower prices. A service combining precise problem definition, distinctive data and verification can articulate value beyond cheap production. This is a structural interpretation, not verified pricing power for a named company.

Consumers may benefit through lower prices, shorter waits and wider choice even if providers' profits do not rise much. Producer earnings, consumer value and worker rewards therefore need separate channels in any assessment of AI's economic effects.

Where scarcity may move

Distinctive data, verified outcomes and trust embedded in real workflows can matter more than generation volume itself. No layer is guaranteed higher profits simply because demand expands.

11 / BUSINESS MODEL IMPLICATIONS

The report market can extend beyond selling documents

TradeLens analysis | The opportunity is to serve previously unresolved questions, not merely sell more pages to existing customers.

Service design	Consumption-expansion logic	Quality mechanism to retain
Focused analysis for smaller firms	Support a narrow decision without requiring a large project	Defined scope, evidence and exclusions
Continuously updated research	Refresh judgments when conditions change, beyond one-off delivery	Dated evidence, change rationale and retirement of stale conclusions
Audience-specific explanation	Adapt expert material for managers, practitioners and learners	Consistent facts and conclusions; comprehension checks
Decision-support engagement	Move from reading a report to comparing concrete alternatives	Explicit assumptions, falsification conditions and responsibility

Making an existing service cheaper is not the same as expanding its market

Selling the same document more cheaply to existing clients is a cost effect. New buyers or previously unserved questions indicate market expansion. Yet self-provision can replace outsourced services: economy-wide use may rise while professional-service revenue falls.

Contract around the outcome and scope, not just the hours

A model justified mainly by drafting hours differs from one that specifies verification, accountability and regular updating. Overly broad promises can allow unlimited revisions and review work to absorb the savings. Track maintenance cost per customer alongside actual use and repeat purchase.

The business question this report proposes

Ask **who can make which new decisions**, rather than how many more reports can be produced. AI's industrial opportunity lies in expanding the users and use cases of knowledge services.

The table describes possible service designs. It does not establish observed demand, revenue, customers or profitability for these models.

12 / TESTABLE CONCLUSION

Knowledge abundance: test the outcomes

Core conclusion | AI relaxes supply constraints on knowledge services; employment and welfare depend on demand, quality and the allocation of gains.

Metric	Evidence strengthening the thesis	Evidence weakening the thesis
Cost per output at a fixed quality standard	Sustained decline including verification and rework	Review and error costs absorb drafting savings
New demand and actual use	More new customers, questions and use cases	Output grows but reading and use stagnate
Learning and decision quality	Better independent performance, understanding and choices	Weaker performance when tools are removed
Hiring and skill pathways	Newcomers gain entry through new tasks	Persistent hiring contraction and lost training routes
Time, pay and autonomy	Less routine burden and overwork, with rewards maintained	Higher targets and review burdens without worker gains

Start with comparable quality and scope

Compare like-for-like work before and after adoption, separating tool usage from total work time. Distinguish genuine new demand from duplicate generation. Track sales, actual use, errors, rework and human hours together to identify where productivity gains go.

Productivity

Was a unit completed using fewer resources?

The market

Did more people start using the resulting service?

Distribution

Whose income, time or capability improved?

The judgment to retain

Knowledge workers can move from document drafting toward designing, verifying and interpreting a knowledge-production system. That transition will not be equally accessible to everyone. **How lower production costs are connected to real consumption and better work** determines the quality of the transition.

These are tests for the thesis, not a claim that all improvements have occurred. This report does not forecast net employment change or long-term wage levels.



13 / SOURCE REGISTER

Follow the evidence to its source

Teal [numbers], source titles and “Open source” link directly to the individual source.

01. James Bessen • NBER / University of Chicago Press | Artificial Intelligence and Jobs: The Role of Demand [Open source](#)

2019.05 • Research chapter • automation, demand elasticity and employment

02. Robert C. Allen • Explorations in Economic History | Engels’ pause: Technical change, capital accumulation, and inequality in the British industrial revolution [Open source](#)

2009.10 • Journal article abstract • industrial-era wages and productivity; full-text access terms vary

03. Nicholas Crafts • Oxford Economic Papers | Slow real wage growth during the Industrial Revolution: productivity paradox or pro-rich growth? [Open source](#)

2022.01 (online 2021.02) • Journal article • reassessing historical productivity and wage growth

04. Shakked Noy • Whitney Zhang • Science | Experimental evidence on the productivity effects of generative artificial intelligence [Open source](#)

2023.07.14 • Randomized experiment / abstract • 453 professionals; full-text access terms vary

05. Erik Brynjolfsson • Danielle Li • Lindsey Raymond • QJE | Generative AI at Work [Open source](#)

2025.05 (online 2025.02.04) • Journal article • 5,172 support agents; staggered rollout in 2020–2021

06. METR • Becker et al. | We are Changing our Developer Productivity Experiment Design [Open source](#)

2026.02.24 • Research-team update • 2025 experiment and selection bias in follow-up data

07. Anders Humlum • Emilie Vestergaard • NBER | Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI [Open source](#)

2025.05 / research version 2026.03 • Working Paper 33777 • Danish surveys and records; earnings and recorded hours

08. Brynjolfsson • Chandar • Chen • Stanford Digital Economy Lab | No Widespread Displacement, but the AI Employment Gap for Young Workers Has Widened to 19% [Open source](#)

2026.08.12 / data through 2026.06 • Authors’ research summary • age/occupation employment gaps in the US ADP sample

14 / SOURCE REGISTER

Follow the evidence to its source

Teal [numbers], source titles and “Open source” link directly to the individual source.

09. ILO • NASK One in four jobs at risk of being transformed by GenAI, new ILO–NASK Global Index shows 2025.05.20 • Official research release • potential occupational exposure, not realized job losses	Open source
10. Aaron Chatterji et al. • NBER How People Use ChatGPT 2025.09 / data 2022.11–2025.07 • Working Paper 34255 • consumer usage; provider-affiliated co-authors	Open source
11. Greg Kestin et al. • Scientific Reports AI tutoring outperforms in-class active learning: an RCT introducing a novel research-based design in an authentic educational setting 2025.06.03 / experiment 2023 • Randomized experiment • 194 university physics students; specific instructional design	Open source
12. Hamsa Bastani et al. • PNAS Generative AI without guardrails can harm learning: Evidence from high school mathematics 2025.06 / experiment 2023 • Field experiment • Turkish high-school mathematics; outcomes after tool removal	Open source
13. Brynjolfsson • Rock • Syverson • AEJ: Macroeconomics The Productivity J-Curve: How Intangibles Complement General Purpose Technologies 2021.01 • Journal article • complementary intangibles, organizational redesign and training	Open source
14. NIST Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile 2024.07 • NIST AI 600-1 • §2.2, MS-2.5-003; printed pp. 6, 31 • voluntary risk guidance	Open source
15. Daron Acemoglu • Pascual Restrepo • Journal of Economic Perspectives Automation and New Tasks: How Technology Displaces and Reinstates Labor 2019.05 • Journal article • displacement and the reinstatement of labor through new tasks	Open source

Methodology, disclaimer and permitted use

Published findings, TradeLens interpretations and illustrative assumptions are presented separately.

Item	Basis
Analysis date	3 October 2026. Publication dates and observation periods are distinguished.
Scope	Structural change in knowledge production, consumption and work under generative AI; not a valuation of a company.
Evidence	Task experiments, usage data and labor-market evidence are not pooled. Foreign results are not extrapolated to all Korean workers.
Illustrations	Pages 7–8: report labor, direct costs, prices and elasticity. Page 10: hours required at fixed staffing. Not observed fees, wages or employment forecasts.
Calculations and figures	Rounding is applied only to displayed values. Diagrams and charts are TradeLens frameworks or illustrations, not endorsements by cited institutions.

Disclaimer

This report provides general information on technology, industries and business structures, growth and profitability. It does not provide buy or sell recommendations, trading timing, fair values, target prices or personalized investment advice. Information and analysis are as of the stated date; accuracy, completeness and currency are not guaranteed. Outlooks, estimates and illustrative calculations are uncertain. Neither investment principal nor returns are guaranteed.

Investment decisions remain the user's responsibility. Except where liability is imposed by applicable law, TradeLens is not liable for investment losses arising from use of this report.

Copyright and use restrictions

The text, charts, design, editing and structure in which TradeLens holds rights are protected by applicable law. Apart from personal reading and storage, no part may be modified, edited, translated, adapted, reproduced, transmitted, reposted, redistributed or sold without prior written permission. This applies to free and paid distribution and to all formats, including files and screenshots. Removing logos, source references or rights notices, or distributing the material under a third party's name, is prohibited.

Uses permitted by law and the separate licensing terms of third-party materials remain applicable. Please share the original link to the official TradeLens publication page.