

A background image showing three people in a meeting. A man on the right is standing and writing on a glass wall with a marker. A woman in the center is sitting at a table, resting her chin on her hand and looking towards the man. A man on the left is also sitting at the table, looking towards the man writing. The setting appears to be a modern office or meeting room with large windows in the background.

Measuring the AI Economy

The case for an open, cross-provider index of AI use,
output and outcomes

BRIEFING PAPER

OCTOBER 2026

Contents

Introduction	3
1 The measurement gap	4
2 Forecasts without parameters	5
3 Which scenario are we in?	6
4 What Baseline builds	8
5 Design principles: trust before data	9
6 Working groups	11
7 Recommendations	12
Conclusion	14
About Baseline	15
Annex: Measurement definitions	16
Glossary	17
Endnotes	18

Disclaimer

This briefing paper is published by Baseline, an independent measurement initiative. It does not report Baseline estimates; all figures are drawn from the published sources listed in the endnotes. Scenario values are those of the cited authors and are illustrative, not predictions. Views expressed do not represent those of any organization named in this paper.

Version 1.0, October 2026. Comments are welcome at research@baselineindex.org and will inform future versions.

© 2026 Baseline. This paper may be shared and quoted with attribution.

Introduction

Artificial intelligence may be the largest economic shift in a generation. Yet the economy it is reshaping is being measured with assumptions rather than data.

Credible forecasts of AI's contribution to output differ by more than an order of magnitude. Estimates of its effect on employment range from negligible to severe. The labor-market evidence collected so far is mostly null or mixed, with early effects concentrated among young workers in the most exposed occupations.^{1,2} These disagreements are not primarily disagreements about economic theory. They are disagreements about a handful of quantities that nobody can currently observe across the economy: how much work AI performs, how widely it is used, how much it raises output per task, and how much of that work replaces rather than assists people.

Each major AI developer now publishes some measure of how its own systems are used. These efforts are valuable, but they use different task

taxonomies, samples and definitions, and none links usage to wages, hiring or firm output. No company can publish a comparable index jointly with its competitors, and public statistical agencies cannot see provider data.

This paper sets out the case for **Baseline**, an open, cross-provider index of the AI economy. It describes the measurement gap (Chapter 1), shows why forecasts diverge (Chapter 2), maps the parameters that would settle which economic scenario we are in (Chapter 3), and describes the measurement program, governance and working groups designed to produce them (Chapters 4 to 6). It closes with recommendations for AI providers, researchers, policymakers and firms.

KEY FINDINGS

- **The quantities that decide AI's economic effect are measurable but unmeasured.** Leading scenario models reduce the debate to a few parameters; today most are set by assumption.
- **Existing measures are single-provider.** They cannot be compared with each other or linked to outcomes such as wages, hiring and firm output.
- **A neutral, federated index can close the gap without sharing raw data.** Providers run agreed queries on their own systems; only privacy-protected aggregates leave.
- **Measurement can begin in 2027.** A common task taxonomy and output unit come first, then outcome linkage, then a quarterly scenario tracker and an open model of the AI economy.

1 The measurement gap

Measurement of AI's economic footprint has advanced quickly since 2024, but it has advanced in fragments. Developers have published analyses of their own usage data, mapping conversations to occupational tasks. Researchers have linked payroll records to measures of occupational exposure. Statistical agencies have added questions on AI use to business surveys.³⁻⁷

Each of these sources answers part of the question. None answers the whole of it, for three reasons.

Incompatible taxonomies. Each provider classifies usage with its own task categories and sampling frame, so their results cannot be added up or compared.

Usage without output. Usage counts record what people ask AI to do, not what it produces, whether the result was correct, or what it cost. Without a common output unit, usage cannot be translated into productivity.

No link to outcomes. Payroll and hiring data show what is happening to workers, and exposure measures show where AI could matter. But no dataset connects measured usage, across providers, to the outcomes of the same firms and workers.

These gaps are structural rather than technical. Commercial rivals cannot pool data without a neutral party, and public institutions do not have access to provider systems. Closing the gap requires an institution designed for that purpose.

TABLE 1

Selected published measures of AI's economic footprint

Measure	Data	What it shows	What it cannot show
Anthropic Economic Index	One provider's conversations	Use by task and occupation; automation versus augmentation	Use on other platforms; outcomes
How People Use ChatGPT (NBER, 2025)	One provider's messages	Mix of work and non-work use; task types	Comparability with other providers; outcomes
Microsoft, Working with AI (2025)	One provider's conversations	AI applicability by occupation	Realized output; outcomes
Google ATLAS	One provider's usage	Usage mapped to O*NET tasks	Comparability with other providers; outcomes
Stanford Digital Economy Lab, payroll studies	Payroll records	Employment changes by exposure, age and occupation	Measured AI usage
US Census Bureau, Business Trends and Outlook Survey	Firm survey	Share of firms using AI	Depth of use; output per task

Source: Baseline review of public sources; see endnotes 1 and 3 to 7. Descriptions summarize each source's scope and are not an assessment of quality.

2 Forecasts without parameters


Because the key quantities are unobserved, forecasts of AI's economic effect rest on assumptions about them, and the assumptions vary widely. Acemoglu's task-based estimate implies a total factor productivity gain of roughly 0.7% over ten years.⁸ Goldman Sachs projected that generative AI could raise global GDP by 7% over a similar horizon.⁹ A 2026 scenario framework from the Anthropic Institute spans GDP effects of 1.6% to 32% by 2030 depending on the parameters chosen.¹⁰ Other frameworks, from the OECD's cross-country estimates to compute-centric integrated assessment models, span a similarly wide range.^{11,12}

These figures are not strictly comparable: they measure different quantities over different horizons. That is part of the problem. Without a shared measurement base, estimates cannot be reconciled, updated or tested against each other.

The consequences extend beyond academic debate. Investment in AI infrastructure rests on expected productivity gains. Workforce and education policy depends on which occupations will change and how fast. Central banks must judge whether AI is raising potential output. Each of these decisions is currently made on assumed parameters.

FIGURE 1

Published estimates of AI's economic effect span more than an order of magnitude



Note: Estimates measure different quantities over different horizons and are shown together to illustrate dispersion, not to compare like with like. Sources: Acemoglu (2024); Goldman Sachs (2023); Korinek, Jones, Sacher, Cotter and McCrory (2026). See endnotes 8 to 10.

BOX 1 · WHY USAGE IS NOT OUTPUT

A count of AI conversations about writing code says little about how much code was produced, whether it worked, or what it replaced. Baseline's proposed output unit, the cost per verified task, measures what AI delivers: the cost of completing a task to a pre-registered standard, such as passing tests or human audit. It is the bridge from usage data to productivity statistics. Controlled studies already show how widely task-level gains vary by task and by worker; a common unit would make those gains comparable across settings.^{13–16} Annex A3 gives the definition and a worked example.

3 Which scenario are we in?

The most useful recent contribution to this debate reframes it. Korinek, Jones, Sacher, Cotter and McCrory build a task-based model that converts a small set of parameters into paths for GDP, wages, the labor share and unemployment through 2030.¹⁰ Their three scenarios correspond broadly to the range of published forecasts.

In the **modest** scenario, GDP in 2030 is 1.6% higher than without AI and unemployment rises by about a tenth of a point. In the **substantial** scenario, GDP is 8.3% higher. In the **extreme** scenario, GDP is 32% higher, the labor share of income falls from 60% to 45%, and nearly one in five cognitive workers is unemployed.

The authors are explicit that each parameter "is something we can potentially measure," and that the


data observed over the coming years "will tell us which scenario we are in." Today, they note, none of the three can be ruled out.¹⁰

Table 2 lists the central parameters, the values each scenario assumes, and how they are currently estimated. Most mid-2026 anchors come from a single provider's data or a firm survey; 2030 values are assumptions. A cross-provider index is the missing instrument.

Each strand of existing evidence covers one piece. Exposure studies estimate what AI could do;¹⁷ adoption surveys estimate how many people use it;¹⁸ and the economics of new-task creation explains why automation need not reduce labor demand one for one.¹⁹ Baseline would measure the pieces together, on one taxonomy.

FIGURE 3

Share of task instances performed with AI (m × d): 2026 anchor and 2030 scenario values



Note: The product of the affected share of tasks (m) and the diffusion share (d) is the share of task instances performed with AI in the source model. The scenarios differ by a factor of more than seven in 2030. Baseline calculation from Korinek, Jones, Sacher, Cotter and McCrory (2026), Table 1. See endnote 10.

TABLE 2

The parameters that separate the scenarios, and how Baseline would measure them

Parameter	Mid-2026 anchor	2030 value: modest / substantial / extreme	Current basis	Baseline measurement
Share of tasks affected by AI (m)	0.14	0.2 / 0.3 / 0.5	Observed exposure from one provider's usage data	Cross-provider task index on a common taxonomy (WG 1)
Diffusion: share of capable tasks where AI is used (d)	0.10	0.2 / 0.4 / 0.6	Census Business Trends and Outlook Survey	Depth-of-use measure across providers (WG 1, WG 2)
Productivity gain per task (a , log)	0.30 to 0.45	0.30 / 0.45 / 0.80	Field trials and frontier evidence	Cost per verified task (WG 2)
Automation share of AI-performed work (ψ)	n/a	0.50 / 0.75 / 0.90	Single-provider automation classification	Automation versus augmentation, classified across providers (WG 1)
Reinstatement of new tasks (ρ)	n/a	0.50 / 0.25 / 0	Historical estimates	New-task emergence in the versioned taxonomy; outcome linkage (WG 1, WG 4)

Note: The productivity anchor and the automation and reinstatement values differ by scenario in the source; automation and reinstatement are held constant within each scenario. Source: Korinek, Jones, Sacher, Cotter and McCrory (2026), Table 1, including the observed-exposure anchor attributed there to Massenkoff and McCrory (2026); Baseline mapping. See endnote 10.

FROM ASSUMPTION TO MEASUREMENT

- 1 Measure the parameters directly.** Replace assumed values with cross-provider estimates, published with uncertainty ranges.
- 2 Update every quarter.** A recurring scenario tracker reports which parameters are now measured, which remain assumed, and where the evidence points.
- 3 Build on open frameworks.** Run published models with measured inputs first; develop a fuller model as the data broaden.

4 What Baseline builds

Baseline is organized in three phases, each producing a public output and each feeding the next.

Phase 1: Measure (H1 2027)

A common task taxonomy built on O*NET and extended to multi-step and agentic work; a shared measure of depth of use; and a common output unit, the cost per verified task. Each provider runs agreed queries on its own systems. The first release is a cross-provider index of AI use by task and occupation.

Phase 2: Estimate (H2 2027)

Usage is linked to firm and worker outcomes through secure research environments, producing estimates

of output per unit of compute by sector and of the effect of AI use on wages and hiring. Identification relies on variation that firms do not choose, such as staggered rollouts, price and rate-limit changes, outages and model releases.

Phase 3: Simulate (2027 onward)

From 2027, a quarterly scenario tracker feeds measured parameters into published models and reports which scenario the evidence supports. In 2028, Baseline releases an open, calibrated model of the AI economy, combining agent-based and general-equilibrium methods, in which every parameter traces to a measured estimate.

FIGURE 2

Raw data stays with each provider; only aggregates feed the shared index



Source: Baseline.

5 Design principles: trust before data

A cross-provider index succeeds only if competitors trust it with their data and the public trusts what it publishes. Baseline's design starts from those two requirements.

Data stays with its owner

Raw data never leaves the provider. Providers run agreed queries on their own systems, using federated protocols, and only aggregates are released. Results are published only above minimum cell sizes and with differential-privacy protections, after disclosure review, following methods already used by statistical agencies.^{20,21} Outcome linkage runs inside secure research environments of the kind operated for federal statistical data.²² Annex A5 sets out the proposed rules.

Plans before data

Analysis plans are registered before any data is opened, and results are published whatever they

show. This guards against selective reporting by any party, including Baseline.

Independence

An independent steering committee of economists sets methodology and decides releases. No member organization can veto a finding.

Competition safeguards

Only aggregates pass through a neutral host, under clean-team rules. No pricing, capacity or customer-level information is exchanged, and antitrust counsel reviews the protocol before Phase 1 begins.

Open by default

Standards, code and estimates are free to use.

Membership is by invitation so that data partners can commit under agreed rules, but the outputs belong to everyone.

BOX 2 · AI IN THE MEASUREMENT PROCESS

Classifying millions of interactions into a common taxonomy is only feasible with AI classifiers running inside each provider's systems, and AI graders can check task completion at scale. Baseline uses AI to scale measurement and simulation; people set the methods, audit samples and interpret results.

BOX 3 · PRECEDENTS FOR SHARING WITHOUT EXPOSURE

Baseline's model has working precedents. Under the 1996 Bermuda Principles, Human Genome Project centers agreed to release sequence data publicly and rapidly, turning competitors into contributors to a shared resource.²³ In the United Kingdom, the OpenSAFELY platform runs approved analysis code inside health-record systems so that researchers receive only aggregate results, never patient records.²⁴ Baseline applies the same principle to AI usage data: the query travels to the data, and only safe aggregates travel back.

Limitations

Coverage. The index will describe participating providers. Open-weight models run on private infrastructure, and providers that do not join, will be outside it at first; releases will report coverage explicitly.

Selection. Users of AI tools differ from non-users, and enterprise and consumer use differ. Estimates will be weighted and reported separately where samples allow.

Verification. Cost per verified task is most reliable where outcomes can be checked, such as code or structured documents. For open-ended work, verification relies on sampled human review and is less precise.

Causality. Usage linked to outcomes is descriptive unless a credible source of variation is available. Baseline will label descriptive and causal estimates distinctly.

6 Working groups

Six working groups build the measurement stack, from task to world model. Each pairs provider engineers with independent economists and publishes its specifications for public comment.

TABLE 3
Baseline working groups

Phase	Working group	Scope	Delivers
Measure	1. Task ontology	A versioned crosswalk from provider logs to O*NET tasks, extended to multi-step and agentic work	The Baseline Task Taxonomy and an open classifier
Measure	2. Output and verification	Cost per verified task, depth of use, and a price deflator for AI-produced work	The output-unit specification and verification protocols
Measure	3. Privacy-preserving aggregation	Federated queries, minimum cell sizes and differential privacy	The query protocol and disclosure-review rules
Estimate	4. Causal identification	Linking use to wages, hiring and firm output through natural experiments	Pre-registered study designs and linked research panels
Estimate	5. Compute production functions	How compute converts to output by sector, with labor substitution and energy limits	Sector-level elasticity estimates with stated uncertainty
Simulate	6. The AI economy world model	Agent-based and general-equilibrium simulation across adoption regimes	A quarterly scenario tracker from 2027; the open model and scenario library in 2028

Source: Baseline. Governance and independence sit with the steering committee across all six groups.

7 Recommendations

Closing the measurement gap is a shared task. Each group of stakeholders can act now.

For AI providers

- Adopt a common task taxonomy for published usage statistics, even before pooling data.
- Pilot federated queries that return only aggregates, under neutral governance.
- Report output and verification, not only usage.

For researchers

- State forecasts in terms of measurable parameters so they can be tested.
- Pre-register designs for linking AI use to outcomes.
- Contribute to open taxonomies and verification standards.

For policymakers and statistical agencies

- Support secure research environments that can link aggregate usage to administrative labor and firm data.
- Align business-survey questions on AI with a common task taxonomy.
- Treat AI measurement as public infrastructure, as national accounts were in the twentieth century.

For firms and workers

- Track AI's effect on output and roles with comparable measures, not anecdotes.
- Share anonymized adoption experience with research programs.
- Plan for multiple scenarios rather than a single forecast.



TABLE 4
Proposed first steps, 2027

Action	Lead	Target date
Publish Baseline Task Taxonomy v0.1 for public comment	Working group 1	H1 2027
Pilot federated queries with at least two providers	Working group 3 and data partners	H1 2027
Release Scenario Tracker No. 1	Working group 6	H1 2027
Publish the first cross-provider index of AI use	Working groups 1 to 3	H1 2027
Post pre-registered outcome study designs	Working group 4	H2 2027
Open secure research environment access to approved researchers	Working groups 3 and 4	H2 2027

Source: Baseline. Timing is indicative and depends on founding data partners.

FIGURE 4
From first specification to open research access in 2027

H1 2027 · MEASURE	H2 2027 · ESTIMATE AND SIMULATE
● Task Taxonomy v0.1 Draft published for public comment ● Federated query pilot At least two providers; aggregates only ● Scenario Tracker No. 1 Which scenario the first measured parameters support ● First cross-provider index AI use by task and occupation	● Pre-registered outcome studies Designs posted before any data opens ● Secure research environments Access for approved researchers ● Quarterly scenario updates Tracker refreshed as measurement broadens ○ Toward 2028 Open AI economy world model

Source: Baseline. Timing is indicative and depends on founding data partners.

Conclusion

National income accounts gave the twentieth century a shared picture of its economy, and with it the ability to steer. Before them, policymakers debated the depth of the Great Depression with fragmentary data; after them, output, income and employment became common ground on which disagreements could be resolved.²⁵

AI calls for the same kind of infrastructure. The parameters that determine whether AI's economic

effect will be modest, substantial or extreme are measurable. What is missing is an institution trusted by competitors and the public to measure them consistently, publish them openly and update them as the technology changes.

Baseline is designed to be that institution: a standing, public record of what AI produces and for whom, so that its gains are visible early, measured honestly and broadly shared.



About Baseline

Baseline is an open measurement initiative building comparable, task-level data on what AI is used for, what it produces and who it affects. It brings together AI providers, data platforms and independent economists under neutral governance. Standards, code and estimates are published openly; membership is by invitation.

Baseline is forming its founding group. Founding members will be announced in 2027.

Get involved

AI labs and data platforms: help define the shared task taxonomy and contribute aggregate measures

under agreed rules. partnerships@baselineindex.org

Researchers and policy institutions: register early interest in the aggregate datasets and secure research environments opening in 2027.
research@baselineindex.org

baselineindex.org

Suggested citation

Baseline. 2026. "Measuring the AI Economy: The Case for an Open, Cross-Provider Index of AI Use, Output and Outcomes." Baseline Briefing Paper No. 1, October 2026.

Annex: Measurement definitions

These definitions are Baseline's initial proposals. The working groups will refine them and publish final specifications for public comment.

A1 · Task taxonomy

The unit of measurement is a task instance: one attempt to complete one task. Each instance is mapped to an O*NET task or detailed work activity using an open classifier run inside each provider's systems. Multi-step and agentic work is recorded as a single task instance with its component steps attached, so that delegation of whole workflows is visible. The taxonomy is versioned, and new tasks that do not map to existing categories are logged; their growth is one indicator of reinstatement (Table 2).

A2 · Depth of use

Depth of use records how much of a task AI performs, on three levels: **assist** (AI suggests, the person does the work), **co-produce** (AI and person each perform part of the work) and **delegate** (AI completes the task and a person reviews it). Releases report the distribution across levels by task and occupation. The delegate share informs the automation parameter in Table 2.

A3 · Cost per verified task

Cost per verified task (CPVT) is the full cost of producing tasks that pass a pre-registered check, divided by the number that pass:

$$\text{CPVT} = (\text{compute cost} + \text{human time} \times \text{wage} + \text{rework cost}) \div \text{tasks passing verification}$$

Verification is fixed in advance for each task family, for example passing an automated test suite, a sampled human audit or downstream acceptance.

Worked example (illustrative numbers, not estimates). A team runs 1,000 code-fix tasks with AI. Compute costs \$0.40 per task (\$400). A reviewer spends 6 minutes per task at \$60 an hour (\$6,000). Of the 1,000, 820 pass the test suite; failures are discarded rather than reworked. CPVT is $(\$400 + \$6,000) \div 820 = \$7.80$. If the same fix takes a developer 45 minutes without AI, the human-only cost is \$45 per verified task. The ratio of the two costs, measured consistently across providers and tasks, is the input that productivity estimates need.

A4 · Identification

Phase 2 estimates rely on variation that firms do not choose: staggered product rollouts (difference-in-differences with estimators robust to staggered timing), price and rate-limit changes (instrumental variables), service outages (event studies) and model releases (comparisons before and after a release date). Designs are pre-registered before data are opened.

A5 · Disclosure control

Proposed starting rules, to be finalized by working group 3: no published cell describes fewer than 10 firms or 1,000 users; each release applies differential privacy with a stated privacy-loss budget, published alongside the data; and every release passes an independent disclosure review before publication.

Glossary

Term	Meaning in this paper
Augmentation	AI performs part of a task while a person remains responsible for it.
Automation	AI performs a task that a person previously performed, with the person at most reviewing the result.
Cognitive work	Work in management, professional, sales and office occupations, the work most directly affected by current AI.
Cost per verified task	The full cost of producing tasks that pass a pre-registered check, divided by the number that pass (Annex A3).
Cross-provider index	A measure of AI use built on one taxonomy and comparable across AI providers.
Depth of use	How much of a task AI performs: assist, co-produce or delegate (Annex A2).
Differential privacy	A mathematical standard that limits how much any published statistic can reveal about a single person or firm.
Diffusion	The share of tasks AI can perform for which it is actually used.
Exposure	The share of a job's tasks that AI could, in principle, perform or speed up.
Federated query	An analysis that runs inside each data holder's systems and returns only aggregate results.
O*NET	The US Department of Labor's database of occupations and the tasks they involve.
Pre-registration	Publishing an analysis plan before seeing the data, to prevent selective reporting.
Reinstatement	The creation of new tasks for people as technology automates existing ones.
Total factor productivity	Output growth not explained by growth in labor and capital inputs.

Endnotes

1. Brynjolfsson, E., B. Chandar and R. Chen. 2025, updated August 2026. "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence." Stanford Digital Economy Lab. digitaleconomy.stanford.edu
2. Humlum, A. and E. Vestergaard. 2025. "Large Language Models, Small Labor Market Effects." NBER Working Paper 33777. nber.org/papers/w33777
3. Anthropic. Anthropic Economic Index. anthropic.com/economic-index
4. Chatterji, A., T. Cunningham, D. Deming, Z. Hitzig, C. Ong, C. Shan and K. Wadman. 2025. "How People Use ChatGPT." NBER Working Paper 34255. nber.org/papers/w34255
5. Tomlinson, K., S. Jaffe, W. Wang, S. Counts and S. Suri. 2025. "Working with AI: Measuring the Occupational Implications of Generative AI." arXiv:2507.07935. arxiv.org/abs/2507.07935
6. Google. AI and the Economy. ai.google/economy
7. US Census Bureau. Business Trends and Outlook Survey. census.gov/hfp/btos
8. Acemoglu, D. 2024. "The Simple Macroeconomics of AI." NBER Working Paper 32487. nber.org/papers/w32487
9. Goldman Sachs. 2023. "Generative AI Could Raise Global GDP by 7%." goldmansachs.com
10. Korinek, A., C. I. Jones, S. Sacher, T. Cotter and P. McCrory. 2026. "Economic Scenarios for Transformative AI." The Anthropic Institute Working Paper No. 2026-02. anthropic.com/institute/econ-scenarios
11. Filippucci, F., P. Gal and M. Schief. 2024. "Miracle or Myth? Assessing the Macroeconomic Productivity Gains from Artificial Intelligence." OECD Artificial Intelligence Papers No. 29. OECD Publishing.
12. Erdil, E. et al. 2025. "GATE: An Integrated Assessment Model for AI Automation." arXiv:2503.04941. arxiv.org/abs/2503.04941
13. Noy, S. and W. Zhang. 2023. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *Science* 381 (6654): 187–192.
14. Brynjolfsson, E., D. Li and L. Raymond. 2025. "Generative AI at Work." *Quarterly Journal of Economics* 140 (2): 889–942.
15. Dell'Acqua, F. et al. 2023. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality." Harvard Business School Working Paper 24-013.
16. Peng, S., E. Kalliamvakou, P. Cihon and M. Demirer. 2023. "The Impact of AI on Developer Productivity: Evidence from GitHub Copilot." arXiv:2302.06590.
17. Eloundou, T., S. Manning, P. Mishkin and D. Rock. 2024. "GPTs Are GPTs: Labor Market Impact Potential of LLMs." *Science* 384 (6702): 1306–1308.
18. Bick, A., A. Blandin and D. Deming. 2024. "The Rapid Adoption of Generative AI." NBER Working Paper 32966.
19. Acemoglu, D. and P. Restrepo. 2019. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives* 33 (2): 3–30.
20. Dwork, C. and A. Roth. 2014. *The Algorithmic Foundations of Differential Privacy*. Foundations and Trends in Theoretical Computer Science 9 (3–4).
21. Abowd, J. M. 2018. "The U.S. Census Bureau Adopts Differential Privacy." *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
22. US Census Bureau. Federal Statistical Research Data Centers. census.gov/about/adrm/fsrdc
23. Human Genome Project. 1996. "Summary of Principles Agreed at the First International Strategy Meeting on Human Genome Sequencing" (the Bermuda Principles).
24. Williamson, E. J. et al. 2020. "Factors Associated with COVID-19-Related Death Using OpenSAFELY." *Nature* 584: 430–436.
25. Kuznets, S. 1934. "National Income, 1929–1932." Senate Document No. 124, 73rd Congress.



The open index of the AI economy

baselineindex.org

partnerships@baselineindex.org

research@baselineindex.org