



The Participation Economy

How broad access to AI agents can expand who gets to build, create, and grow

September 2, 2026

OpenAI

Executive summary

The first phase of generative AI put expertise within reach for more people, helping them solve problems, develop ideas, and work more effectively. The next phase will be about expanding both who can put that expertise to work and how much they can accomplish with it.

As the next generation of AI systems become capable of coordinating multiple agents, people and businesses will increasingly be able to take on work that used to require more capital or specialized support.

Agents can pursue tasks in parallel, work across tools, and continue while the person directing them is doing something else. People still set the goals, supply the context, judge results, and take responsibility. But they can direct more work without having to produce every component and manage every step themselves. The new economic value this unlocks extends well beyond saving time: more people can gain the capacity to turn an idea into a product, expertise into a service, or a problem into a project that a team can pursue.

This is the opening for a **Participation Economy**, in which advanced productive capabilities are available broadly enough to equip far more people to build and create. In the late industrial economy, capabilities such as specialized labor, analysis, software development, market research, design, and coordination were largely bundled inside institutions. To use them, most people had to join an organization, hire one, or raise enough capital to build one. Agents can start to unbundle those capabilities and put them directly in people's hands. More people able to build means more knowledge and ideas put to work, more problems addressed, and more solutions tested.

ChatGPT's reach shows the potential scale of that opportunity. It is now used weekly by about one in six internet-connected people worldwide and by more than 2 million businesses, twice as many as a year ago. As adoption has grown, our user base has diversified, spanning 180 countries and territories; skewing majority-women at 53%¹; and with more adult users both under and over age 35 than any other frontier lab. The challenge is to turn that broad reach into broader economic opportunity.

Access alone will not deliver that result. OpenAI has described a "capability overhang": the gap between what AI systems can do and the value most people, businesses, and countries capture from them. The next economic divide may be a **conversion gap**: differences in the ability of people, firms, and countries to translate advanced AI capabilities into useful work and economic gains. If

¹ Per our latest privacy-preserving analysis, among messages from users whose first names could be classified as typically feminine or masculine, the feminine-name share is now 53.3%.

the most capable agents are only affordable and accessible to the largest institutions, AI could deepen and compound incumbents' advantages. If some workers direct agents while others are managed by them, differences in autonomy and income could widen. And as work changes, established routes into professions may erode before new ones develop.

G20 Ministers can shape how this transition unfolds. Governments can pair connectivity and affordable access with practical, sector-specific training; help small and mid-size firms connect AI securely to their workflows; encourage experimentation and peer learning; and establish clear rules for permissions, human review, and accountability. Education and workforce policies can help people build the judgment to direct agents while creating routes into skilled work and entrepreneurship.

Ultimately, progress should be measured by what people can accomplish and who benefits. Alongside adoption, governments and companies should encourage AI that enables better work, new businesses, higher worker incomes, and greater resilience across firms and regions. The opportunity is to put the ability to build and grow within more people's reach, and give them greater agency over their work and a greater share in the value they create.

1. A new economic unit: the individual plus agents

Traditional software helps a person perform a specific action. An agent can take a goal, work out the steps, use tools, preserve context, track progress along the way, and continue toward an outcome. Multiple agents can divide the work, develop and challenge approaches, and test how the pieces fit together.


The person's role shifts from producing every component themselves to setting direction, supplying context, making judgments, and taking responsibility for the outcome. Human attention remains a scarce resource but can be applied at a higher degree of leverage. The relevant question expands from, "How quickly can I complete this task?" to, "How much work can I reliably direct, verify, and combine?"

This pattern is already emerging. As the chart below shows, between December 2025 and May 2026, the share of sampled Codex users who made at least one request estimated to correspond to work that would take one person more than 30 minutes rose to more than 80%. The share making a request that would take a person more than 60 minutes rose to 70%. The share requesting work that would take a person more than eight hours grew the fastest, from a base just above zero to over 25%. These estimates describe the work requested, not completed work or time saved².

² These thresholds are model-estimated, so they should be treated as directional rather than exact. This is also based only on individual data and based on the queries of a random sample of 0.1% of users.

Requests above human-time thresholds

30m+ 1h+ 4h+ 8h+



2. What broader participation makes possible

For this shift to create a Participation Economy, these types of advanced capabilities need to be widely available. The test is whether more people gain the means to start and carry through useful work.

Agents can broaden participation in four ways:

- **More capabilities within reach.** Coding, analysis, design, translation, market research, and coordination can become available on demand.
- **A lower cost of trying.** People can build a prototype or test a market before committing substantial capital or hiring a staff.
- **More ways to apply expertise.** Specialists can apply their knowledge across more cases and call on capabilities outside of their usual roles.
- **Geography and time can become less limiting.** Flexible access can help people outside of major hubs, people with disabilities, and those balancing caregiving or other responsibilities.

Institutions will still supply capital, infrastructure, community, and accountability. But more people may be able to enter markets, demonstrate capability, or organize meaningful work without securing an institution's backing.



OpenAI Chief Financial Officer Sarah Friar has described **abundance** not simply as more compute, but as putting more useful intelligence within reach – including capabilities that people and small businesses once had to rely on much larger organizations to provide.

Realizing that opportunity requires progress on three frontiers: **capability**, or what models can do; **diffusion**, or who can access and use them; and **conversion**, or who can turn that use into economic value. More capable models and wider access create the opening. The economic gains depend upon what people can accomplish with them.

3. Productivity through expanded participation

Much of the early evidence has focused on time saved in existing work. A randomized experiment involving 6,000 knowledge workers found that generative AI reduced email time by roughly three hours per week but did not significantly reduce meeting time – suggesting that technology changes individually controlled tasks faster than organizational processes³.

Agents could extend those gains by working across tasks and tools. Organizations will still need to adapt their permissions, decisions, and processes to make this possible.

For individuals, the productivity gains go beyond completing a familiar task faster:

- **Shorter cycles of work.** Tasks can run in parallel and continue while a person is occupied elsewhere, shrinking the time between an idea, an experiment, and the next decision.
- **Expertise available as needed.** People can draw on specialized capabilities without having to assemble a team for every function.
- **More ideas worth testing.** Lower costs can make it practical to try ideas that would otherwise be too expensive to pursue.

These gains could grow as more people participate. Each person brings different knowledge, needs, relationships, and ideas to the production process. Broad access could expand the range of problems the economy addresses, rather than simply distribute a static productivity gain among more people.

Measuring that change requires looking beyond adoption alone. In OpenAI’s privacy-preserving [analysis](#) of individual account data, 30.4% of messages were work-related. Within that subset, 45% were classified as “doing,” 32% as “asking,” and 22% as “expressing.” Because the dataset excludes enterprise and Codex use, it likely understates business and technical activity. Even so, it shows that adoption counts alone cannot reveal whether people are using AI to complete work.

³ [NBER, “Shifting Work Patterns with Generative AI”](#)

Our scorecard focuses on work accomplished and the cost of achieving it. “Useful intelligence per dollar” asks what each dependable outcome costs after review and rework⁴.

For a Participation Economy, that scorecard must also show who benefits. Are individuals and small firms gaining capabilities alongside large institutions? Are those gains reaching people across regions and income levels? Productivity matters, but so does who gains the ability to initiate and direct useful work.

4. Individuals as builders, founders, and problem-solvers

For individuals, the practical change is the ability to turn what they know into something they can use or offer to others. A nurse, mechanic, teacher, or nonprofit leader could apply their expertise to a tool or service without first assembling a technical team.

One early example comes from OpenAI’s finance team. An employee with no coding background used Codex to build a planning tool that translated monthly advertising forecasts into weekly and daily plans.⁵ We also found substantial engineering-related Codex usage across business operations, marketing, and other non-engineering functions, delivering the benefit of more efficient work cycles and faster strategic decisions. These are just some of the ways the team employs agents to expand the work they can do and the speed at which they can do it.

As a whole, the work could also help people demonstrate their capabilities beyond titles or credentials. A useful tool, an improved process, or a completed project gives others something concrete to evaluate.

But people need more than nominal access. They need capable systems, affordable usage, data, delegation and verification skills, and the freedom to use AI creatively. A worker managed by agents has not gained the same agency as one empowered to build with them.

The participation test should therefore be concrete: Does a deployment expand what people are able to initiate, decide, create, or own? Or does it primarily increase the ability of an institution to monitor, standardize, and control them?

⁴ [Sarah Friar, “A scorecard for the AI age”](#)

⁵ [Sarah Friar, “What building an AI-native finance function taught me”](#)

5. Making it easier to start and grow a firm

The same capabilities could change what it takes to start and run a business. Software development, financial planning, and market analysis often require employees or outside providers. Agents could lower some of those costs. Two effects could follow:

- **Existing firms could develop new products and enter new markets more quickly.** Agents could help them test ideas, serve new customers, and explore new geographies before committing to a larger physical footprint.
- **New firms could start with fewer resources.** Founders, independent professionals, and small businesses could take on functions that once required a larger organization.

The second effect could make entrepreneurship accessible to more people. A founder might test a market and build a prototype before hiring, earn revenue before seeking substantial financing, or serve a market too specialized for a larger firm. Already, our own research finds that at least 4 million Americans have been running, starting, or planning businesses with help from ChatGPT. Similarly, US Census figures show the number of business applications increasing since ChatGPT's launch in 2022.

That alone doesn't guarantee more competition. Advantages in proprietary data, distribution, capital, compute, or integration could allow established firms to capture most of the gains. Broad model access, interoperability, data portability, and competitive markets will influence whether agents lower barriers to entry or reinforce them – all questions for policymakers and providers to consider.

There are already large differences in how intensively firms use AI. OpenAI's enterprise analysis from June 2026 found that frontier firms (those in the top 10% of AI usage each month) now generate 8.3× as many output tokens per active user as median firms (those in the middle 10%), up from 2.6× in January. Among weekly active users, 21% at frontier firms used plugins and 19% used skills, compared with 9% and 3% at typical firms. The gap appears to come not from model access alone, but from complementary context, tools, training, and workflow redesign.

Within firms, employees need room to experiment and support to put useful results into practice. Leadership can provide shared priorities and resources while giving the people closest to a problem the freedom to solve it.

Tokenomics: what useful work costs

Who benefits from these changes will also depend on what it costs to use capable agents.

Tokens meter model activity, but they are not the economic outcome. Their value depends on capability, tool access, reliability, and conversion into completed work. Agent teams may consume more tokens because they plan, act, critique, and retry, yet still be more efficient if they finish more valuable work.

The better analogy is electricity: a metered input whose value depends on what it powers. Businesses and governments should ask how much dependable work each dollar of intelligence produces.

For the Participation Economy, three dimensions matter:

- **Abundance.** If inference is scarce or expensive, society will ration intelligence and leave valuable work undone.
- **Efficiency.** Better models, hardware, routing, context management, and product design can produce more useful work from the same compute.
- **Access.** The gains will be broadly shared only if individuals, small firms, schools, researchers, and public institutions can afford and effectively use advanced systems.

“Useful intelligence per dollar” is therefore an efficiency and inclusion metric. The question is whether a dependable outcome becomes affordable to a student, creator, local business, researcher, or first-time founder.

6. Industrial policy for the agentic era

As agents change what it takes to start and grow a firm, they will also change the work people do. Roles may shift toward judgment, relationships, accountability, and the direction of agents. Lower costs could also make new products and services viable, creating demand for work that did not previously make economic sense.

Three dynamics deserve attention.

First, productivity gains may show up as more output rather than fewer hours worked. Cheaper analysis, code, design, and experimentation could lead people and firms to take on more ambitious projects.

Second, the limits on progress may shift. As agents handle more execution, defining problems, supplying context, choosing among options, verifying results, building trust, and taking responsibility may become more valuable.

Third, established routes into certain professions could weaken. Entry-level tasks often help people develop expertise and judgment. If agents take on those tasks, employers and educators will need new ways for early-career workers to acquire that experience while learning to supervise, audit, and improve AI-assisted work.

The policy objective should be to expand people's leverage and their share of the gains. That requires training, but also mobility, credible and portable credentials, transition assistance, worker voice, and pathways to ownership and entrepreneurship. Success should be measured by whether people gain valuable skills and real choices, as well as by changes in employment.

Meeting these goals will require more than incremental updates to existing policy. But competition and entrepreneurship should remain central: people must be free to experiment, build companies, and apply AI to problems that governments or established institutions may not anticipate. But markets alone may not close the conversion gap. Early advantages in capital, infrastructure, data, and expertise could become entrenched before others have the means to participate. Industrial policy can help by building shared foundations, translating scientific breakthroughs into scaled industries, preparing people and institutions for disruption, and helping ensure that technological progress produces broad-based economic growth.

OpenAI has put forward [ideas for a new industrial policy](#) to help advance that discussion. Developed by policy and research teams, these proposals are exploratory options for governments, companies, workers, communities, and families to refine, challenge, combine, or reject through democratic debate. Experimentation, evidence, and collaboration will be essential to finding approaches that work. The objective should be clear, however: to create the conditions in which more people can convert advanced intelligence into agency, productive work, higher incomes, and meaningful participation in economic and social life.

A participation agenda for companies

There are steps companies can take today to get started:

- **Giving people capable tools and permission to build.** Provide secure access across functions, not only to technical teams or senior leaders. Let the people closest to a problem create solutions, with clear boundaries and accountability.
- **Starting with outcomes people care about.** Identify decisions, products, and workflows where constrained capacity blocks value. Define a successful result before choosing a tool.
- **Redesigning work end-to-end.** Map the path from intent to outcome: information, actions, handoffs, approvals, and exceptions. Use agents to connect the workflow, not merely accelerate one isolated step.
- **Teaching delegation, evaluation, and orchestration.** Help employees decompose goals into parallel workstreams, assign roles, compare independent approaches, test outputs, and recognize when human judgment must take over.
- **Building a trusted-agent operating layer.** Provide governed access to company knowledge and tools, common identity and permissions, monitoring, audit trails, escalation rules, and reusable components.
- **Rewarding participation and reuse.** Recognize employees who turn domain knowledge into tools and workflows others can use. Share successful patterns across the organization rather than trapping them in isolated experiments.
- **Preserving pathways for learning and advancement.** Redesign work-based-learning and entry-level roles so workers build judgment and domain expertise while learning to supervise agentic work and leverage unique human and relational capabilities. Reinvest a share of productivity gains in training, redeployment, and mobility.
- **Measuring agency as well as efficiency.** Track cost per successful outcome, cycle time, ready-to-use rates, and review burden. Also measure who is building, how many functions and levels participate, how many new projects are attempted, and whether employees gain more authority to solve problems.

A national strategy for broad participation

Governments can extend that opportunity beyond individual workplaces. For G20 ministers, the task is to help make advanced AI affordable and widely usable, while preparing people and institutions for changes in work. A sample national strategy could include:

- **Building the physical foundation.** Expand reliable energy, data-center capacity, connectivity, and advanced-computing supply chains. Compute that cannot be powered, connected, or deployed does not become productive capacity.
- **Treating access to advanced AI as economic infrastructure.** Ensure that startups, small businesses, independent workers, researchers, schools, and public institutions – not only large incumbent firms – can obtain capable models and sufficient inference.

- **Putting agents within reach of every person.** Develop broad access models through education, libraries, workforce systems, small-business programs, and public-interest institutions. The ambition should be practical capacity, not occasional exposure.
- **Equipping people to direct AI.** Integrate delegation, verification, data literacy, tool use, and domain-specific agent orchestration into education and workforce programs. Teach people how to turn goals into well-governed systems of work.
- **Making entrepreneurship easier.** Offer compute credits, technical assistance, testbeds, shared data resources, and simpler routes to business formation and procurement. Help people translate agent-enabled prototypes into durable firms.
- **Creating new apprenticeship and credentialing models.** Support work-based learning in which people acquire domain judgment while supervising and improving AI systems. Recognize demonstrated outputs and portable skills alongside traditional degrees and titles.
- **Using government as a lead customer and participation platform.** Deploy agents in service delivery, permitting, research, procurement, and benefits administration with strong safeguards. Open appropriate tools and knowledge so citizens, researchers, and businesses can build on public capacity.
- **Promoting competition, portability, and user control.** Encourage competitive model markets and the ability for individuals and firms to move their context, data, and agent workflows between providers. Concentrated intelligence can become concentrated power.
- **Building trust into deployment.** Establish clear accountability for consequential actions, strong cybersecurity and privacy practices, incident reporting, risk-based testing, auditability, and meaningful routes for appeal. Trust is productive infrastructure.
- **Updating the national scorecard.** Measure not only inference cost and adoption by firm size and region, but also the kinds and complexity of tasks delegated, depth and repeatability of use, access to complementary tools and data, and diffusion across workers, firms, and regions. Connect those measures to realized outcomes: output, new-business formation, research throughput, occupational mobility, wages, income growth, productivity, and resilience. Track who can direct useful work—not simply who encounters AI.

Conclusion: productive agency within reach

Agents can expand what people and small firms can accomplish. More people can apply their expertise to a product, service, tool, or problem they would previously have lacked the resources to pursue.

This won't eliminate scarcity. Energy, compute, capital, data, judgment, trust, and human attention will remain constrained. As agents make more kinds of work possible, the ability to choose worthwhile goals and carry them through will matter even more.

The opportunity is to broaden both the capacity to create value, and the chance to share in it. Companies and governments can help by making capable AI accessible, providing training and support so people can use it effectively, and giving people room to put their ideas to work. The measure of a Participation Economy will be what more people can accomplish, and how that expands their opportunities.

About OpenAI

OpenAI's mission is to ensure that artificial general intelligence benefits all of humanity. We're building AI to help people solve hard problems because by helping with the hard problems, AI can benefit the most people possible – through more scientific discoveries, better healthcare and education, and improved productivity. We're off to a strong start, creating freely available intelligence used by more than 1 billion people around the world each week. We believe AI will scale human ingenuity and drive unprecedented economic growth and new freedoms that help people accomplish what we can't even imagine today.

Cover image created with ChatGPT

