Print Edition

Wednesday 12 August 2026

EDITION NO. 20 · 1 PIECE · 10 PAGES

IN THIS EDITION

AGI Will Set Off an Industrial Explosion

BY AI FRONTIERS · AI FRONTIERS · 11 AUGUST 2026

Damon Binder, Senior Researcher at Coefficient Giving — August 11, 2026

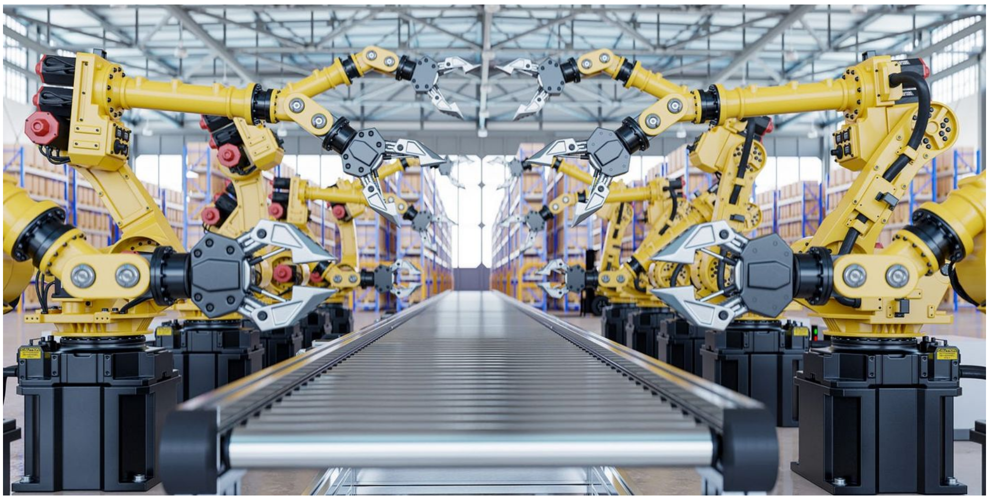

AI systems now write so fluently that cheating at universities has become ubiquitous and "AI slop" is displacing human writing across the internet. They have begun producing original proofs on long-standing open problems in mathematics. Researchers at the frontier AI companies have largely stopped writing their own code, and the models' hacking abilities are strong enough that the US government temporarily export-controlled Anthropic's Claude Fable 5, days after its release. Ten years ago, every claim in this paragraph would have sounded like science fiction.

It is easy to become numb to this progress. Today's AI agents remain endearingly clumsy; watching Claude run a small store is a bit like watching a child running a lemonade stand. But we should not let such awkwardness distract from the trend. These systems have evolved from research curiosities to extremely powerful economic engines in 10 years, with no end in sight. I want to take that trend seriously and ask: if AI does get to the stage where it can do the cognitive work that humans do, what happens to the economy?

With notable exceptions, economists have mostly declined to address this question; their usual maneuver is to deny the premise, as when a widely cited Nobel laureate's analysis put AI's contribution to US GDP at about 1% over the coming decade, by essentially freezing 2023 capabilities for 10 years. But the question is worth taking seriously, because automating cognitive labor releases what has always been the critical brake on physical production: no matter how cheap machines and tools become, you cannot manufacture new workers. In this piece, I will assume no additional new technologies and no recursive self-improvement to superintelligence. Yet, even with these conservative assumptions, I will show that the economy stands to be profoundly transformed. For the first time in history, physical production could be fully automated.

The rest of this piece makes that case and considers the implications. I argue that AI capable of cognitive work could also do physical work (since building the machinery it would need is relatively straightforward) and that its labor would be cheap. I then turn to input-output tables—the US government's accounting of what every industry buys from every other—to ask how fast a fully automated economy could grow, finding that physical output would double roughly every year rather than every few decades (as it does currently). Finally I ask whether anything could stop such an industrial explosion, and what the implications are for power at home and abroad.

Why AGI Gives You Robots

The term "artificial general intelligence" (AGI) has been endlessly used and abused, but it points to a clear core concept: an AI smart enough to do the work humans can do, at a competitive cost. You could hand it a task, as you would to a skilled employee or contractor, and expect the job done competently. Current systems, for all their rapid progress, are not close to AGI; my own attempts to automate my job have been unsuccessful. However, once AGI appears, we should expect it to be able to perform physical work through robotics.

An AI system that can do every remote job can also operate a robot. Definitions of AGI focus on cognitive work, but there is nothing special about the physical world that confines intelligent agents to their data centers. Remote cognitive work is hard: doing all of it means mastering the real-time control, spatial reasoning, physical prediction, and continual learning required for such varied tasks as mechanical engineering and animation. A system with those skills has what it needs to operate machinery, if given actuators to work with.

Progress in AI is already pulling robotics forward. The general-purpose methods that cracked language and vision turn out to be good at manual manipulation too. Large language models can increasingly supply high-level control: today they run robot arms, fly drones, and program robot dogs. Hobbyists have even wired them into homemade robots. Meanwhile, the same training methods, applied to recordings of humans teleoperating machines, are producing the low-level dexterous control that humans perform instinctively; today's systems can fold laundry and assemble cardboard boxes, using ordinary cameras and simple grippers.

Hardware is not the bottleneck for robotic automation. Industrial arms have demonstrated precision and force beyond any human arm for over half a century. But, because they could neither see nor think, every motion had to be programmed in advance, at massive cost. The other work-around was a human operator: remote manipulators have handled radioactive material since the late 1940s, built underwater structures using remote-controlled submersibles, and performed surgery. That remote control was worthwhile only when a person could not be physically present. But it does show something we've known for decades: a machine under competent control can do skilled physical work. Once control can come from AIs rather than people, it becomes worth building general-purpose machinery for them.

Robots will not necessarily resemble humans. While the human hand is a remarkable instrument, robot hands are already fairly good, and certainly better than the split-hook prostheses—two rigid fingers on a cable with no sense of touch—used by amputees to farm, weld, and perform all manner of other tasks. With AI, simple actuators go far, and AGIs will have intelligence and patience to spare. Nor will the machines work at human workstations forever. Production demands dexterity today because every process was designed around human workers with hands. New possibilities open up when the workers are machines; just as today's programming agents do not type at keyboards as their human counterparts do, robotics tools need not be gripped and triggered by fingers—they could be mounted directly on a robotic arm instead.

Machine Labor Will Be Cheap

Automating labor is not free. But the computer chips and robotic actuators that replace workers are simply more capital goods for the economy to produce—and they're not particularly expensive ones. A humanoid robot, broadly speaking, uses parts similar to those of a car: metal, motors, batteries, electronics, and sensors. At automotive production volumes, it should cost tens of thousands of dollars; Unitree already sells its child-sized G1 humanoid for around $13,500, and Tesla is reportedly aiming for $20,000 for its full-size Optimus. If a robot costs $30,000 and, working around the clock, can substitute for a single human worker who costs $30 an hour, it could pay for itself in six weeks.

AI cognitive labor is generally significantly cheaper than the humans it replaces. It is hard to price the future AI cognitive labor involved in an industrial explosion, because it does not exist yet. However, on the tasks AI systems can already do—like transcription, routine translation, or writing one-off scripts—they are typically far cheaper, per unit, than the people they replace, and the cost of a fixed level of AI capability falls rapidly year over year. Even if the first AGI arrived merely cost-competitive with human workers, within a year it would cost a fraction as much.

Input-Output Analysis

How fast would the economy grow once human labor is automated? At first blush, this looks hard to answer. Standard growth models assume that labor and capital substitute for each other, treating each as a single dollar-denominated aggregate—barristers and bricklayers combined into "labor," bulldozers and bridges into "capital." Such abstractions are poorly equipped for a world where labor is unnecessary and capital reproduces itself, and classic concepts like GDP can become profoundly misleading.

Input-output analysis helps forecast how an automated economy would grow. Input-output analysis takes a different approach, recording what each industry physically needs from others. The method was developed by Wassily Leontief, in the 1930s, to study the relationships between industrial sectors. In 1945, Leontief used his tables to project the United States' 1950 steel requirements to within a couple of percentage points. Military planners, concerned about war with the Soviet Union, took up the method to study how fast industry could remobilize.

Input-output analysis tracks what everyone buys from everyone else. The basic idea is simple: ask every business what it buys and from whom, then aggregate the responses into tables that track what each industry buys from others. An entry for aluminum smelting might record the electricity, ore, and machining required to produce aluminum, while an entry for aircraft manufacturing records how much aluminum is required. The Census Bureau runs this survey every five years, and the Bureau of Economic Analysis assembles the results. The tables take years to build, so the most recent set covers 2017. Companion tables record the equipment and structures that each industry holds, from machine tools to chip fabs. Entries are recorded in dollars, but they are detailed enough to closely follow the underlying kilowatt-hours of electricity, tons of steel, and other physical quantities. Most output is ultimately consumed; the rest is reinvested, either adding to physical capital or replacing what has worn out.

Once robots can replace human labor, the economy will no longer be bound by a fixed workforce. It can produce every input it needs, including the "workers" themselves, so output can be fed back into building more capacity, and growth compounds. John von Neumann worked out how to compute the maximum rate at which such an economy could expand. Combining his analysis with the production data in the tables tells us how fast an autonomous industrial sector could grow.

Input-output analysis measures economic growth in physical output, not value. This approach to forecasting growth tracks how much stuff can be produced. I make no attempt to convert this into GDP, because an industrial explosion would upend the prices any such conversion relies on: when machines produce everything, including more machines, goods become exponentially cheap. Physical output is also the better guide to what is at stake. Military power, human employment, and material abundance all depend on how much gets built, rather than on its dollar price.

How Fast Could an Autonomous Economy Grow?

Using the US government's 2017 input-output tables, which track 402 industries, I find that a fully automated economy using US production methods could double its output roughly every year; input-output tables for other advanced economies give comparable results. The robots and computer chips needed to automate production represent only a small fraction of total output; even a tenfold increase in their cost would not impact the growth rate significantly. Construction lags can be incorporated too; even with these delays included, the economy still doubles in well under two years.

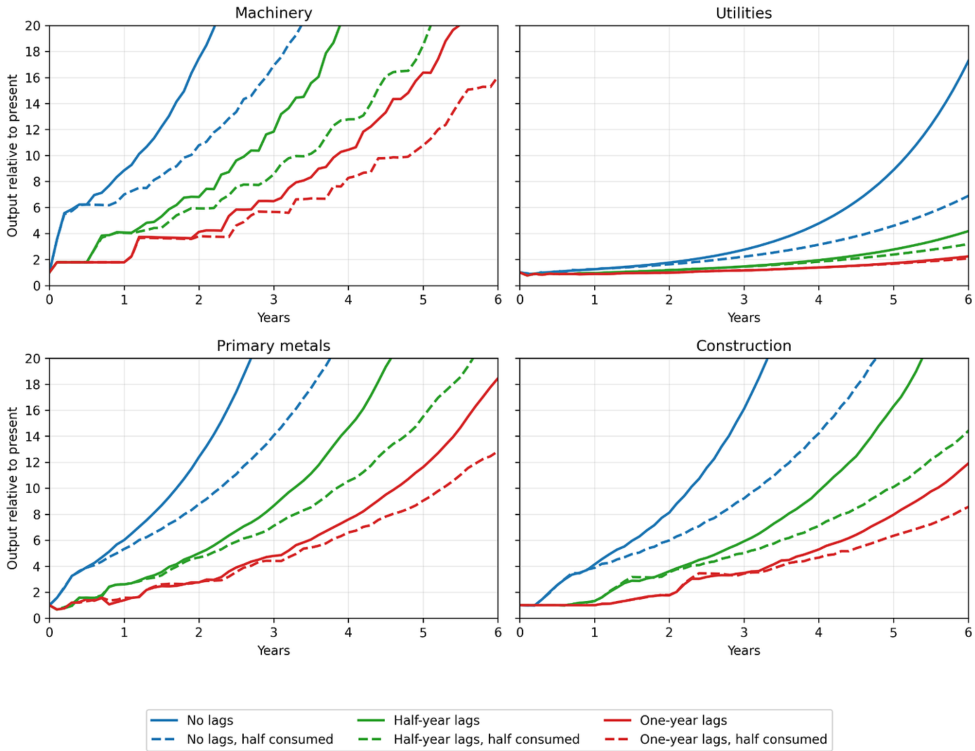

Projected physical output of four key industrial sectors in the years after AGI arrives, relative to today's level. Each scenario assumes that new capacity takes zero, six months, or a year to build before it starts producing output. Solid lines hold human consumption at today's level and reinvest all output above that level into industry; dashed lines split that output 50-50 between industry and human consumption. A real buildout would differ somewhat from these projections: it would be slower initially as AI capabilities diffuse throughout the economy, but accelerate later as production methods adapt to robot workers. Source: author's calculations from the 2017 US input-output tables, as described in Part 2 of the author's memo series.

Why does a fully automated economy accelerate so much faster than the current one? The main difference is that an automated economy can reinvest far more of what it produces. Building more factories today is not useful; there would be no workers to staff them nor to use the extra machines and tools produced by them. Thus, we consume about four-fifths of what we produce, and most reinvestment goes to replacing worn-out capital rather than adding to it. AGI would remove this obstacle, allowing robot workers to be manufactured along with other physical goods and thus breaking the link between human labor and physical production.

The pace of growth is then set by how fast the machine economy can copy itself, which input-output analysis shows us is fast.

The true growth rate of an automated economy is likely even higher. Our estimate of this maximum growth rate uses 2017 American production methods; therefore, it's conservative. Those methods were designed around human workers at human wages and are a poor fit for a robot economy; letting the production methods change makes it grow considerably faster. Assuming 2017 methods also ignores the efficiencies that come with producing at greater volume. However, it also neglects the depletion that makes raw materials costlier to extract, which I return to below.

Industrial explosions happen even without new science and technology. This assumption might seem strange to economists, because standard theory predicts that growth at the technological frontier requires new ideas. But that prediction assumes an economy in which the workforce cannot be manufactured; an economy that can build its own workers grows by building more of them and needs no new ideas to do so.

An economy growing at this maximum rate looks very different from today's. Its output is machinery, materials, energy, and factories rather than consumer goods and services, and its workforce is robotic. Table 1 shows its major sectors, by share of that workforce, of output, and of energy consumed. The workforce such an economy needs is overwhelmingly industrial: robotic arms on factory floors, self-driving excavators and cranes, and manipulators on mobile bases of all sorts and sizes. Automating physical production means automating construction, machining, and assembly; whether AI can substitute for doctors, lawyers, or babysitters is irrelevant to how fast this economy can grow.

Sector	Labor %	Output %	Energy %
Construction	35.1	24.0	30.5
Machinery	18.5	20.2	8.6
Fabricated metal products	9.6	7.3	6.1
Electrical equipment and appliances	6.4	5.3	1.6
Computers and electronic products	5.9	6.5	2.1
Primary metals	3.5	7.9	13.2
Motor vehicles and parts	2.3	5.0	1.6
Professional and technical services	1.9	1.3	0.3
Nonmetallic mineral products	1.8	2.0	3.3
Plastics and rubber products	1.6	1.7	1.8
Aerospace and other transport equipment	1.5	2.2	0.7
Miscellaneous manufacturing	1.3	1.2	0.6
Chemicals	0.8	2.8	7.3
Mining	0.7	1.2	5.5
Utilities	0.4	1.2	3.7
Paper	0.4	0.8	1.9
Petroleum refining	0.1	1.2	3.6
All other sectors	8.1	8.2	7.7

The self-replicating economy at maximum growth, computed from the 2017 US input-output tables at their most detailed level, then aggregated to the government's 71 summary sectors for display. Columns give each sector's share of the human workforce being replaced (at today's staffing levels), of output (valued at 2017 prices), and of energy use (electricity and fuel purchases, in 2017 dollars).

In reality, some output would directly support human needs instead of further growth. The same machines, materials, energy, and robotic labor could make homes just as well as they could make factories. Robots could also cook, clean, and drive for people. However, an economy that devotes significant resources to human consumption could still grow quickly in absolute terms, according to my forecasts. For example, an economy reinvesting half its output would double every 20 months, and the stream of goods and services flowing out to people would double every 20 months as well. Reinvesting only half is arguably conservative—common sense suggests (and economic theory concurs) that countries and companies controlling the machines would reinvest most of their output to stay competitive with one another.

From Investment Boom to Unprecedented Growth

Let's pretend that AGI arrived today. What would happen? The factories and machines needed to automate production do not yet exist.

Input-output analysis can tell us how a realistic robotics build-out would go. The standard economist's approach is to imagine a planner assigning each industry's output either to consumption or to building new capacity, trading off consumption now against faster growth later, according to how much people value each. I have modeled this transition using the 2017 tables at a simplified 71-sector level, under a range of assumptions about those preferences.

AGI would create an enormous investment opportunity. In an AGI-enabled world, building robots will be an industry comparable to building cars, and an automated factory costs little more than an ordinary one while needing no workers to run it. Capital would flood in.

There would be significant competition for the robotics supply chain. Investors would start outbidding everyone else for machinery and construction workers. Resources that would otherwise be consumed, or spent replacing worn-out equipment, would be used to build new robots, automated factories, and computer chips. In the short run, the prices of physical goods would rise, not fall—and, because making them would suddenly pay better, idle and underused factories would ramp up production. Plausibly more human workers, not fewer, would crowd into factories while the boom lasted.

It would take several years for the economy's overall production to double. While the nascent autonomous sector itself would grow rapidly, the starting baseline is low: today's economy makes plenty of consumer goods and services, but it has few autonomous robots, and the machine tools needed to build more are themselves in short supply. Electricity generation would grow relatively slowly, since it is already a large part of the economy and requires significant capital to expand. Still, it would double in about four years, and double again in another two. Robot production would follow a similar schedule, from near zero today to tens of millions of units within a few years.

An industrial explosion would not create instant material abundance. Despite this tremendous growth, demand for physical goods would take time to satiate. Bringing the whole world to the living standards of a wealthy American would take more than 10 times today's physical output —and that is before counting the goods and services that cheap production would newly make accessible, from custom-built mansions to round-the-clock personal care.

Natural-resource constraints alone would not prevent an industrial explosion. The autonomous sector needs no human labor or inputs from the nonautomated economy, and there are minerals enough for economies far larger than today's. Mining is only about 1% of output, so even large cost increases do little to slow overall growth.

The one resource that would truly run short is oil; cheap reserves would be exhausted once the economy had grown severalfold. Coal and natural gas could power growth for far longer, but fossil fuel burning is ultimately limited by its climate effects rather than availability. An economy that declined to burn them would electrify, replacing fossil fuels with solar panels, batteries, and nuclear plants. At today's production costs, this would slow growth—with doubling times closer to two years than one—but not stop it.

Long before demand or minerals ran short, the assumption of frozen technology would fail. Within a few doublings, AI minds would vastly outnumber human ones, and whatever technology they built would transform production more radically still.

What Could Stop an AI Industrial Explosion?

US Admiral William D. Leahy assured President Truman that the atomic bomb "will never go off, and I speak as an expert in explosives." But the viability of a nuclear chain reaction was a question for physicists, not admirals, and the consequences of that physics were no less real for being counterintuitive. So it is with AGI.

Once physical production is automated, physical output can double every year. Objections to explosive growth that do not engage with this data miss the point. Cost disease is real: human services will become expensive relative to manufactured goods, and there are jobs, like babysitting or singing, that we may want humans to keep doing. Neither fact stops physical production from expanding on a scale we have never seen. New research ideas may become harder to find, but that too is beside the point, since the calculation assumes no research at all.

Nonautomatable production is a limited bottleneck on explosive growth. More serious reasons to expect slower growth come from the dynamics of the build-out itself. One possibility is that some small set of tasks cannot be automated. This changes the growth dynamics less than you might think, because the workers freed from every other task can supply the missing ones. If humans must still handle one task in 20, then production could grow 20-fold before it needed more workers. By the time the machine economy has grown 20-fold, we will probably have worked out how to automate those final tasks.

The time taken to integrate AGI may delay an industrial explosion but not stop it. Technological diffusion—the lag between a technology working and being used everywhere—is another potential slowdown, but at worst it delays the explosion rather than stopping it. No one knows whether AGI will arrive in 2030 or 2050; next to that, a few years of diffusion is a rounding error. In any case, the diffusion of AGI workers will look more like the arrival of skilled laborers in the economy than the spread of automobiles or trains. This is particularly true on the software side: generative AI has already been adopted faster than the personal computer or the internet, continuing several decades of ever-faster technological diffusion.

Explosive growth may be possible, but we may choose not to allow it. Regulations will undoubtedly slow deployment. But stopping it entirely is very hard, because an automated economy could be built almost anywhere and the profits would be enormous. If one US state regulates the robots away, the build-out moves next door and the permissive states collect the windfall. If the United States forswears the technology while China does not, China soon commands a vastly larger industrial base. Even if both agreed to stop, they would still have to prevent every other country from importing or developing the requisite robots and AI systems. And the temptation would never fade: a country that allowed the robots would be vastly richer within a few years. Someone would say yes.

Power Without People

War planners funded Leontief's input-output tables to learn how fast industry could be turned into arms. Military power still rests on industrial capacity, and an automated economy would outproduce any rival in every kind of weapon, from artillery shells to drones. It would need no defense workers to build the arsenal, and ever fewer soldiers to wield it. A state with a small population but an autonomous industrial base could field far more firepower than a much more populous state without one. Remove the need for labor, and many countries have raw materials enough for an enormous industrial base. And a two-year head start in a doubling economy means four times the production.

The same decoupling operates inside states. Governments have always needed their citizens as taxpayers and soldiers, and that need has quietly underwritten political accountability. Horses once worked fields and carried cavalry, before the engine made them surplus; Leontief famously expected a similar fate for human workers. Citizens who neither supply taxes nor serve as soldiers would lose the hold that they gained from the state's dependence on them. Rulers funded by oil wells rather than taxpayers can already afford to ignore their people, and often do. Full automation could make every state a petrostate.

While this may seem dire, none of it is inevitable. The same machines that build drones can build houses and hospitals, and within a decade or so of full automation, it would be possible to extend the living standards of a wealthy American to every human. Our biggest constraints will be political rather than physical—there is nothing in the machines themselves that needs to prefer an artillery shell to a dinner plate.

See things differently? AI Frontiers welcomes expert insights, thoughtful critiques, and fresh perspectives. Send us your pitch.

Damon Binder is a Senior Researcher at Coefficient Giving (formerly Open Philanthropy), where he investigates existential risks from biology and advanced AI. He previously worked at the Future of Humanity Institute, Oxford University, and has a PhD in physics from Princeton.