01

PART I · THE ATOMIC WORLD

Industrial Self-Replication Is a Policy Choice

Emerson Alden



The tool-and-die building at Ford's River Rouge plant, Dearborn, Michigan, 1941. Industrial expansion depended on the ability to manufacture the machines of mass production.
Photograph by Alfred T. Palmer / Library of Congress.

SUMMARY

Robotics will enable autonomous industrial self-replication, at least in principle. But amid rapid AI-enabled technological progress, the industrial base will be in a perpetual state of catch-up. Investors will hesitate to invest in processes or products that could be obsolete by the time they come online. Therefore, the industrial base will be the greatest bottleneck to realizing the gains of advanced AI. The U.S. should take action now to recompose the economy in preparation for this. This means prioritizing the security and onshoring of critical supply chains, providing certainty for investors, and creating new institutions to ensure that industrial self-replication benefits American citizens.

ONCE AI MODELS CAN PERFORM ANY TASK OF A REMOTE WORKER, COGNITIVE AUtomation will diffuse much faster than physical automation. But as compute becomes cheap enough that large amounts can be dedicated to high-fidelity robotics

SIMULATIONS, HARDWARE COMPANIES WILL BUILD THE INFRASTRUCTURE FOR MASS robotics production. The first stage of industrial buildout and recomposition will still rely heavily on humans. Only as robot production scales will robots increasingly build and operate the datacenters, power stations, mines, refineries, and factories required for further automation.

This begets an "industrial flywheel": a positive feedback loop in which robots manufacture and construct the infrastructure used for further robot production. Once established, the flywheel has the potential to spin astonishingly fast, possibly doubling industrial output every year.[1] But this assumes that the industrial economy is already composed around the sectors required for rapid self-replication—a far reach from the present-day U.S. economy built around software and financial services.

Furthermore, the actual speed of the industrial flywheel is heavily dependent on the rate at which output is reinvested into further growth. But reinvestment is largely downstream from expectations of future demand, as additional capacity is only financially justified insofar as it produces goods that people desire.

Therefore, reinvestment will depend on the multiplier that AGI applies to scientific discovery, and the degree to which these discoveries unlock new goods and services that people demand. But these are radically uncertain variables! Investors will be left forecasting demand for technologies that don't yet exist.

In hard tech, this problem is magnified, as investments come with hefty fixed costs and long expansion timelines. Consider that AGI's compression of 50 years of scientific discovery into five years, as Dario Amodei claims will occur, would be akin to 1920s industrialists completing their Model T production lines just in time for the Apollo program.[2] Such rapid progress will increase the value of waiting for more information for investors and pushes more investment into the hard-tech flywheel.

Notably, demand for construction and manufacturing robots is likely to be brutally cyclical. By the "accelerator effect," because demand for upstream equipment is derived from demand for the goods they produce, small swings in consumption produce massive swings in equipment orders. Orders for American machine tools historically fluctuated by as much as 90 percent year over year, and it was in one such trough, in the early 1980s, that the industry was offshored for good.[3]

Taken together, these effects present two ways in which the industrial flywheel is vulnerable: technically, the economy must first recompose towards the maximum-growth composition; and financially, the reinvestment rate depends on expected returns to an industrial base whose ultimate uses are radically uncertain. The industrial flywheel represents a powerful economic force with the potential to overshadow even the industrial revolution, but it is also a fragile tower of costly investments whose continued growth is subject to empirical unpredictability and political whims.

1 Damon Binder, "The AI Industrial Explosion—Part 1: Maximum Growth Rates with Current Production Methods," *Defenses in Depth*, May 4, 2026.

2 Dario Amodei, "Machines of Loving Grace," October 2024.

3 Brian Potter, "What Happened to the US Machine Tool Industry?" *Construction Physics*, January 18, 2024.

The immediate role of government is to improve the economy's starting position in anticipation of an industrial flywheel. This should take the form of supply chain security and industrial recomposition. Once the flywheel starts turning, market pressure should be sufficient to direct investment toward the general-purpose tools and inputs it demands.

From a national security perspective, we should err on the side of overbuilding general-purpose capacity. We face competition from nations that may be able to sustain a higher savings rate, ignore environmental costs, or simply start from larger preexisting industrial bases. Amidst this dynamic, a payoff matrix emerges: if the U.S. overbuilds and exceeds the demand that ultimately materializes, the cost is just symmetrical malinvestment: capital is wasted, but the flywheel is not ceded to our adversaries.

But the cost of underbuilding will be much greater should AI-driven demand exceed expectations. Once robot production becomes self-sustaining elsewhere, it will become very difficult to buy back the years of early compounding; the size and quality of the defense industrial base will become decisive factors for national security.

The goal should be to maximize cheap early deployment without becoming structurally dependent on inputs that may cease to be available once other countries develop their own flywheels. But early recomposition faces a trilemma. First, prioritizing robot deployment requires allowing the import of cheap foreign robots and using them to rebuild the industrial base, at the cost of a domestic industry for robotics hardware. Second, prioritizing robot production requires allowing the import of cheap foreign components, at the cost of a domestic components industry. And third, prioritizing production of the full hardware stack means restricting component imports, at the cost of the cheap deployment that drives the industrial flywheel.

One possible solution is allowing the import of foreign component inputs in the short run, but with a strict countdown. This approach would involve credibly signaling demand for domestic robots and robot components down the line (e.g., with targeted statutory tariffs that go into effect four years after passing), but allowing for U.S. manufacturers to continue importing cheaper foreign tools in the meantime. During this window, the growing industry could be supported on the supply side, enabling them to scale during the four-year period and seamlessly sub in for the foreign incumbents when import restrictions hit.

Supporting the growing industry requires minimizing uncertainty for builders. Apart from statutory tariffs, this means addressing the accelerator effect that has taught U.S. investors to avoid upstream manufacturing equipment. This can be solved with a capacity market, in which the state pays the fixed cost of readiness rather than bulk-ordering units. Certified U.S. robot manufacturers bid for the annual fee they will accept to keep a production line warm: equipment maintained, workforce retained, output deliverable on call at pre-agreed prices. Readiness would be verified by random test orders with substantial penalties for non-performance. This has precedent in much of the electric grid: power generators earn annual payments, set by auction, to stand by for peak demand.[4] Such a mechanism retains production lines through troughs, mitigating the worst of the accelerator effect without bulk orders of robots destined to sit unused or obsolesce.

Defense procurement is a necessary demand driver that can serve to mitigate uncertainty and keep the flywheel spinning. But our classic defense procurement model breaks down in the face of rapid technological progress. The Department of War's traditional acquisition structure is "characterized by rigid, sequential processes" and has "proven inadequate in adapting to evolving threats and integrating emerging innovation."[5] The capacity to pivot rapidly as uncertainty resolves ought to be prioritized over attempts to predict the precise shape of future warfare and the particular products that enable it.

New portfolios should be created within the Department of War to focus on buying adaptive capacity from companies that have committed to redirecting production for military capabilities when called upon. The less certain the future product, the more our procurement systems should be willing to pay general-purpose industrial capacity with the ability to specialize later. In a period of AI-enabled rapid technological growth, we may know with high confidence that we will need enormous quantities of advanced military hardware, while having very low confidence about what that hardware will be.

The early-stage Civil Reserve Manufacturing Network (CRMN) is a step in the right direction. CRMN, authorized in the FY2026 NDAA, is intended to partner with commercial manufacturers that can rapidly convert existing facilities to defense production.[6] This includes direction to assess a "factory-as-a-service" model built around scalable manufacturing capacity that can be rapidly reconfigured. This is an excellent initiative that will only grow in importance as the industrial flywheel kicks off.

A more ambitious expansion of the CRMN could periodically give participating firms unfamiliar specifications, evaluating them on a number of metrics, including time to reach production at scale. Furthermore, it should increase the scope of its funding to cover the adaptation of existing factories to new products, alongside being extended to cover new expansions rather than just existing facilities.

4 Congressional Research Service, *PJM's Electric Capacity Market: Background and Current Issues*, R48553 (2025).

5 U.S. Government Accountability Office, *Defense Acquisition Reform: Persistent Challenges Require New Iterative Approaches to Delivering Capability with Speed*, GAO-25-108528 (2025).

6 National Defense Authorization Act for Fiscal Year 2026, Pub. L. No. 119-60, § 1841, 139 Stat. 1254 (2025).

ABOUT THE AUTHOR

Emerson Alden is a research fellow at the Foundation for American Innovation, where he focuses on American competitiveness in robotics. Before joining FAI, Emerson was a DC Seasonal Fellow with GovAI. Emerson holds a BS in mathematics-economics from Furman University.



From *Futurecraft: The Day After AGI*, Foundation for American Innovation, 2026.
Read or listen online at futurecraft.thefai.org/essays/industrial-self-replication-is-a-policy-choice.

© 2026 by the Foundation for American Innovation. All rights reserved.