

AI Takeoff

The First Thirty-Six Months of Recursive Self-Improvement

RodMalpert.com

The first day of recursive self-improvement need not be the day ordinary life becomes unrecognizable. The next three years could nevertheless transform the institutions, incomes, and expectations on which ordinary life depends.

A concrete scenario, not a dated prediction

The scenario runs from March 15, 2028 to March 15, 2031. Its starting condition is a laboratory repeatedly demonstrating that an AI can carry a research objective through design, experiments, implementation, criticism, and useful improvements that accelerate later AI research. The significant event is closure of that feedback loop, not a single benchmark result or press release. March 2028 is an illustrative anchor. A start about a year either side and a transformation lasting roughly eighteen months to five years are scenario alternatives, not confidence intervals.

The first three months

Research and testable digital work change first. AI-run teams improve software, experimental design, mathematics, legal analysis, and medical coordination. Useful descendants spread even when the best frontier system remains expensive and restricted. Adoption accelerates because integration, testing, and workflow redesign are themselves intellectual tasks. Hiring freezes and unfilled vacancies precede widespread announced layoffs. Benefits appear immediately in tutoring, accessible expertise, and less administrative work, while the physical economy remains slower to change.

The rest of the first year

Lower-cost competitors force established firms to reorganize. The scenario removes or leaves unhired roughly one position in five in the most exposed office businesses by the first anniversary. That is not a national unemployment forecast. Digital services become cheaper while chips, power connections, equipment, and suitable facilities become scarce. Stock prices, business valuations, household wages, and productive capacity can move in different directions. A household can gain wealth on paper while losing the income needed to service its debts.

The bottleneck is real, but not a return to normality

The scenario assumes roughly twentyfold improvement in the productivity of the starting hardware on a defined basket of AI-research workloads during year one. This is a demanding illustration, not twenty times the intelligence or economy. Compute and experiments can still become binding constraints. The key claim is that a constraint on further progress does not reverse changes already made, and that many industrial and organizational constraints become targets of the same intelligence.

AI Takeoff | The argument and its limits

Years two and three

AI increasingly redesigns laboratories, infrastructure, construction, manufacturing, and the evidence-generation process in medicine. Trials still need evidence; the scenario does not assume that their organization remains fixed at present speed. By year three it envisages major clinical gains in selected diseases, credible but incomplete rejuvenation signals, improved neural interfaces, and millions of useful robots. These are future assumptions, not claims that universal cancer cures, human age reversal, or unrestricted physical autonomy have already been demonstrated.

The social transition is not an afterthought

The chronology follows households, professional firms, universities, governments, military and intelligence organizations, and media. Universities face separate pressures on research, teaching, credentials, and their economic model. Emergency income support becomes a debate over recurring payments, public services, dividends, and ownership of productive assets. Money can be distributed faster than houses, clinics, electricity, and robots can be built. The central distributional question is who gains practical purchasing power before old incomes disappear.

Institutions change too

The scenario rejects a common shortcut: introducing extraordinary AI while leaving defenders, regulators, construction firms, and governments at their old capabilities. They also adapt, though not uniformly or without conflict. Political alliances can reorganize around acceleration, safety, ownership, national competition, open access, and preservation of human roles. Religious and transhumanist interpretations shape how people understand the same events. Neither technological abundance nor one political response is assumed inevitable.

What the story deliberately brackets

A major war, civilization-scale human misuse, or an unrecoverable loss of AI control could end this chronology. They are excluded from the narrated branch, not judged impossible or negligible. Physical agency can grow through human intermediaries and existing institutions before new robots are ubiquitous. The companion risk papers address those dangers directly. By the early 2030s, feedback among AI research, energy, factories, robotics, and biology makes detailed prediction less credible, not less important.

The reason to read it

This is a future history designed to make the sequence of pressures tangible. Its distinctive claim is a staged takeoff: a real research threshold, rapid cognitive transformation, an initially slower world of atoms, and then AI working on the constraints that made the first three years describable. The human-AI contribution record explains how technical analysis of compute expanded into a scenario about choices people must actually make.

Source: AI Takeoff: The First Thirty-Six Months of Recursive Self-Improvement, September 2026, 44 pages. See the premise, four chronological parts, coda, and attribution. All future milestones remain scenario assumptions; no investment advice is offered.