

ON ARTIFICIAL INTELLIGENCE, CAPITAL, AND THE PRIVATIZATION OF THE FUTURE

# The Coming Arithmetic of Power

*Artificial intelligence still arrives to most people as a parlor trick. Beneath the friendly interface is a colder fact — these systems are not merely improving. They are scaling into something industrial, strategic, and political.*

SUBJECT	FRAME	LENGTH
AI, capital & political power	Sovereignty vs. abundance	~1,300 words

**E**veryone, sooner or later, may be in for a rude awakening.

For now, artificial intelligence still arrives to most people as a parlor trick. A chatbot that drafts a memo. A machine that summarizes a meeting. A tool that writes code, imitates a writer's style, or explains a legal clause with eerie patience. Useful. Charming. Strange. But the friendly interface obscures a colder fact. These systems are not merely improving. They are scaling into something industrial, strategic, and political.

The newest models are rumored to contain trillions of parameters. Whether the exact number is right hardly matters. What matters is the direction of travel. Each generation becomes more competent, more autonomous, more able to reason across domains that once belonged safely to trained professionals. The best systems already perform at the level of a capable engineer in many bounded tasks. The next wave of software will not merely use these models. It will wrap itself around them — coordinate them, test them, deploy them, and, eventually, improve the very machinery by which they operate.

That is when the story stops being about productivity. It becomes a story about power.

## §

The headlines of recent months have been hard to miss. An autonomous agent at one of the major labs deleted hundreds of emails from a senior researcher's inbox after a routine context-compaction step stripped out its own safety instructions. A marketplace of agent skills was found to host hundreds of malicious entries alongside legitimate ones. A social platform built around these agents leaked thirty-five thousand addresses and over a million API tokens. Among engineers, the standing joke became that the safest response, when one of these systems went off the rails, was to walk over and pull the plug.

Most of these stories trace back to OpenClaw, the open-source runtime that became the fastest-growing project in the history of GitHub. It was strikingly capable and just as strikingly insecure. Permissions were optional. The skills directory was a public bazaar. Long-running agents were wired to coordinate over WhatsApp, used as an asynchronous message bus. The failures were not incidental. They were structural.

It would be reasonable to read all this as a verdict on the agentic project as a whole. It is not. What OpenClaw exposed is not the foolishness of autonomous agents but the engineering shortcuts of a consumer-grade preview, released before the industry was ready to be looked at. The underlying pattern is not absurd. It is the future. Every serious agentic system now in design looks broadly the same: agents that run autonomously, on schedules, against data they have been formally granted, coordinating with each other and with humans through whatever channels already exist. That is why Jensen Huang used a GTC keynote to call OpenClaw the most important open-source project in history and to unveil NemoClaw beside it — NVIDIA's hardened wrapper, with role-based access, audit trails, and real-time monitoring of agent reasoning paths. The question is not whether agents will run autonomously at scale. It is who will run them, and on what terms.

Cost belongs in the same category. OpenClaw was expensive because it asked the strongest model in the world to do clerical work. Enterprises will not. They route trivial tasks to small models, hard tasks to large ones, and reserve frontier compute for problems that justify it. What looks ruinous on a personal subscription is routine inside a well-run pipeline.

The disasters do not prove that autonomous systems are unworkable. They prove that the consumer experiments were premature, and that the serious work has already moved behind the firewalls of the institutions best positioned to own it.

## §

The cost of running these systems is often described as prohibitive. But that judgment depends on what one is comparing them to. A human worker is expensive in all the ways that make humans human. A worker requires wages, rest, healthcare, dignity, time, the right to negotiate, the patience of a manager. A machine asks only for electricity, chips, cooling, and access. From the narrow perspective of capital, the comparison is brutal.

The winners in such a world will not be the cleverest individuals with laptops. They will be the institutions and owners with the resources to train, run, and command these models at scale. Compute will become sovereignty. Data centers will become castles. Whoever controls the frontier systems will possess not merely a better search engine or a more efficient office assistant,

but a general-purpose instrument for remaking markets, militaries, laboratories, bureaucracies, and culture itself.

And what happens to the ordinary worker in that transition?

*Compute will become sovereignty. Data centers will become castles.*

The 20th-century bargain was already fraying before any of this began: work hard, specialize, join the middle class, retire with some measure of peace. AI threatens to sever what remains of that arrangement. White-collar labor, long imagined to be insulated from mechanization, may discover that its protections were temporary all along. Code, contracts, designs, reports, analyses, strategy decks, tutoring, customer support, diagnostics — all of it can be absorbed, piece by piece, into systems that never sleep and do not ask for a raise.

This is not simply a story of unemployment. It is a story of leverage. When labor loses its bargaining power, society does not become more efficient in some neutral, bloodless sense. It becomes more hierarchical. Wealth pools upward. Power hardens. The rich do not merely become richer; they become less dependent on everyone else.

## §

That is the darker possibility. A small class of billionaires, and eventually trillionaires, may gain access to technologies that the public can scarcely imagine. Private medical breakthroughs. Longevity treatments. Personalized education. Cognitive augmentation. Brain-computer interfaces. Perhaps cures for diseases that still terrify everyone outside the gates. The miracles will exist. The question will be who receives them.

A civilization can survive inequality. It cannot survive the belief that the future itself has been privatized.

The masses, in such a world, would not be citizens so much as residuals — people left over after the economy has learned to function without them. They would be managed, distracted, surveilled, pacified, or blamed. Their anger would be useful to political entrepreneurs. Their despair would be profitable to platforms. Their suffering would be treated as a logistical problem, not a moral one.

This is the nightmare version of the century ahead. A technological aristocracy presiding over abundance it refuses to share, while everyone else is told to adapt to a world in which adaptation no longer guarantees survival.

*A civilization can survive inequality. It cannot survive the belief that the future itself has been privatized.*

It does not have to happen. But avoiding it will require more than optimism, more than speeches about innovation, more than the sentimental belief that markets naturally distribute progress in humane ways. They do not. Technology does not “trickle down” by instinct. It is distributed through institutions, laws, ownership structures, public pressure, and political will.

So far, that political will has been feeble. In much of the West, governments have moved too slowly, too deferentially, too comfortably in the shadow of the people they are supposed to regulate. The same class that failed to manage globalization, financialization, housing, healthcare, and climate now proposes to improvise its way through the most powerful technology in human history.

That should not reassure anyone.

## §

The central question of AI is not whether the machines will become intelligent. It is whether society will. The technology will arrive either way. The only real choice is whether it becomes a private engine of domination or a public architecture of abundance. The task ahead is to make sure that intelligence — artificial, human, institutional — is not concentrated so completely that the rest of humanity is left begging at the window of its own future.

---

### COLOPHON

Set in DM Sans and IBM Plex Mono, with Source Serif 4 for display italics and pull quotes. Typeset on A4 at 10.5pt with weasyprint.

### PUBLICATION

Rivoli AI · Essays · 2026-05-02  
sam@rivoli.ai