The Economist

A SPECIAL ISSUE OF SUMMER READS

JULY 26TH–AUGUST 1ST 2025

The economics of superintelligence

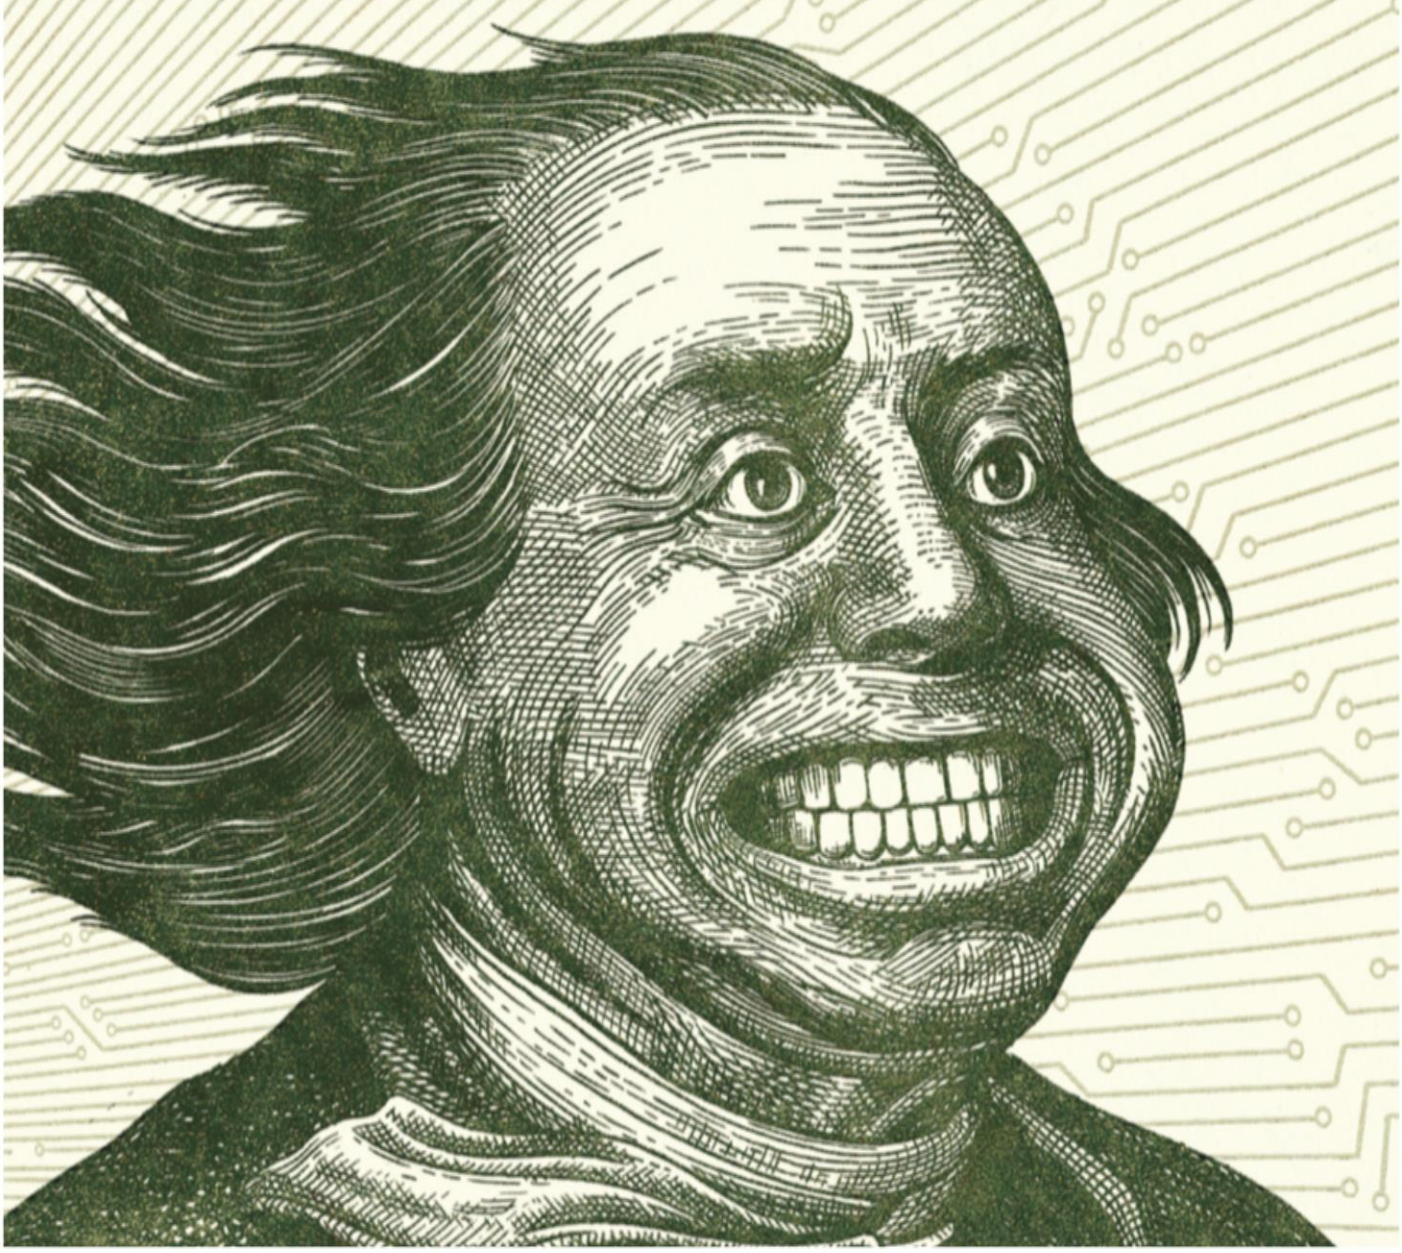

Business

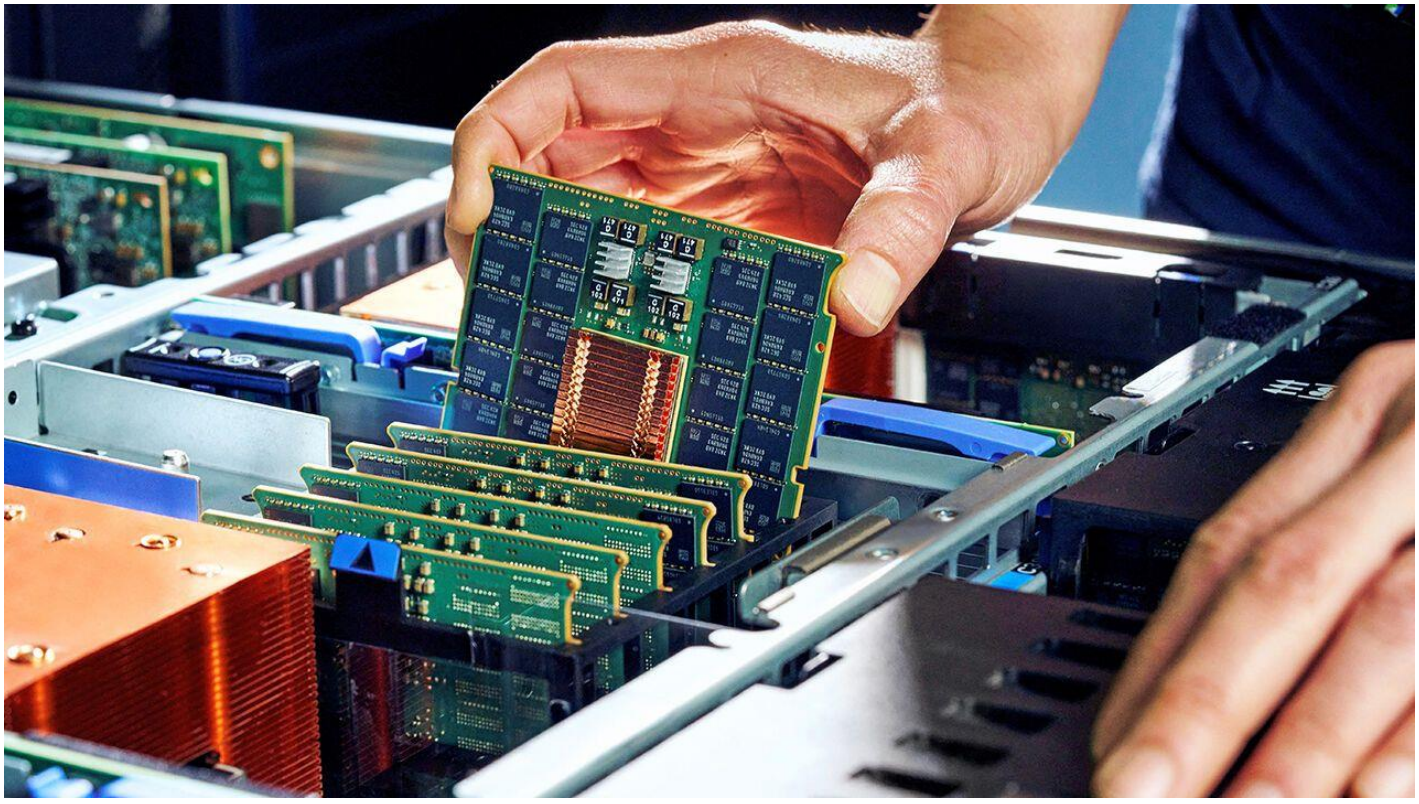

Photograph: Reuters

Donald Trump announced plans to relax regulations governing the development of artificial intelligence in America, including penalties for states with stern rules on deploying AI. Mr Trump aims to unify an increasingly fragmented legal landscape for AI among the 50 states, and to massively expand AI exports to America's allies. The president also signed an order intended to root out ideological bias in AI models, or "woke Marxist lunacy", as he described it.

Not on the same page

A report surfaced that OpenAI and SoftBank are at odds over the direction of the Stargate project, a $500bn AI investment that was announced in January with the support of Mr Trump. The companies are struggling to co-ordinate the construction of data centres. Countering the narrative, OpenAI said it had entered into an agreement with Oracle, another Stargate partner, to develop 4.5 gigawatts of data-centre capacity, though it did not say where or how it would be funded.

Mr Trump described Jerome Powell as a "numbskull" and said he would be out of his job as chairman of the Federal Reserve within eight months. The president is waging a public war against the Fed for not cutting interest rates fast enough, though he said recently that he would not sack Mr Powell. Adding to the pressure, Scott Bessent, the treasury secretary, called on the central bank to "conduct an exhaustive internal review of its non-monetary policy operations", accusing it of "mission creep".

The European Central Bank kept interest rates on hold, following a blistering pace of cuts over the past year. The headline annual rate of inflation is 2% in the euro area, which is the ECB's medium-term target. The currency bloc's inflation rate is lower than that of America, Britain and Japan.

Alphabet's quarterly net profit rose by 19%, year on year, and revenue from Google's search and advertising business grew by 12%. That helped alleviate concerns from investors that AI chatbots are eating into its core search business, for now.

Meanwhile, an AI system developed by Google's DeepMind was awarded the "gold" standard at the International Mathematical Olympiad, an annual event where pre-university students compete to solve six exceptionally difficult maths problems. It was the first time an AI reached the top standard at the competition. OpenAI said it had notched up a similar score to DeepMind, though it did not officially enter the contest.

Chevron sealed its $53bn acquisition of Hess, after winning a lengthy legal dispute that had delayed the takeover. Hess owns a big stake in Guyana's fast-growing offshore oilfields, but ExxonMobil had challenged the takeover, claiming it held the rights to negotiate the stake with Hess. Exxon said it disagreed with the ruling by the Paris-based International Chamber of Commerce, but respected the arbitration process. Earlier, the Federal Trade Commission lifted its ban on John Hess, the chief executive of Hess, joining Chevron's board of directors.

UniCredit, one of Italy's biggest banks, withdrew its hostile takeover bid for Banco BPM, a smaller rival, and criticised the government for trying to block the deal. The European Commission has also chastised the government for meddling in the takeover process.

Elon Musk warned of a "few rough quarters" ahead for Tesla, as the carmaker reported a big drop in sales and net profit. Mr Musk pointed to the loss of incentives in America to buy electric vehicles as one source of Tesla's troubles. A $7,500 federal tax credit for EV purchases ends in September. The company did not update its outlook, saying it was "difficult to measure the impacts of shifting global trade."

General Motors' profit plunged in the second quarter, as it took a $1.1bn charge related to the cost of tariffs imposed on cars imported from its factories in South Korea, Mexico and elsewhere. The

company said it expects those costs to abate in coming months. Meanwhile, Stellantis recorded a net loss for the first half of the year, mostly because of charges related to its business, but also because of a €300m ($350m) hit incurred from tariffs.

Coca-Cola registered another decline in North American sales of its trademark drink, as health-conscious Americans seek alternatives to soda. The company is introducing a new version of Coca-Cola made with cane sugar, a change that was trailed by Mr Trump a week ago. Some people think cane sugar is healthier than high-fructose corn syrup, which has been used to sweeten Coke for decades. Both versions will be available.

Duty free

Michael O'Leary, the boss of Ryanair, Europe's biggest airline, raised the possibility of switching the registration of new Boeing aircraft deliveries to Britain to ensure it avoids the cost of potential tariffs from the EU. Ryanair is based in Ireland, an EU member; the bloc is considering stiff duties on Boeing planes delivered to the region if it can't reach a trade deal with America.

Finance & economics
Crypto's big bang will revolutionise finance

The more useful stablecoins and tokens prove to be, the greater the risk

Illustration: George Wylesol

Among the strait-laced denizens of Wall Street, crypto's "use cases" are often discussed with a smirk. Veterans have seen it all before. Digital assets have come and gone, often in style, sending hype-prone investors in memecoins and NFTs on a ride. Their use as anything other than a tool for speculation and financial crime has been repeatedly found wanting.

Yet the latest wave of excitement is different. On July 18th President Donald Trump signed the GENIUS Act into law, providing stablecoins—crypto tokens backed by conventional (usually dollar) assets—with the regulatory certainty that insiders have long craved. The industry is booming; Wall Streeters are now scrambling to get involved. "Tokenisation" is also taking off: a rapidly growing

volume of assets trade on blockchains, representing stocks, money-market funds, and even private-equity stakes and debt.

As with any revolution, the insurgents are euphoric and the old guard concerned. Vlad Tenev, chief executive of Robinhood, a digital-assets broker, says the new tech can "lay the groundwork for crypto to become the backbone of the global financial system". Christine Lagarde, president of the European Central Bank, sees things a little differently. She worries that the rush of new stablecoins amounts to nothing less than "the privatisation of money".

Both appreciate the scale of the change at hand. The present moment holds the potential of something far more disruptive for mainstream markets than earlier crypto speculation. Whereas bitcoin and other cryptocurrencies promised to be digital gold, tokens are wrappers, or vehicles representing other assets. That may sound unimpressive, but some of the most transformative innovations in modern finance simply changed the way in which assets are packaged, sliced and reconstituted—the exchange traded fund (ETF), the eurodollar and securitised debt among them.

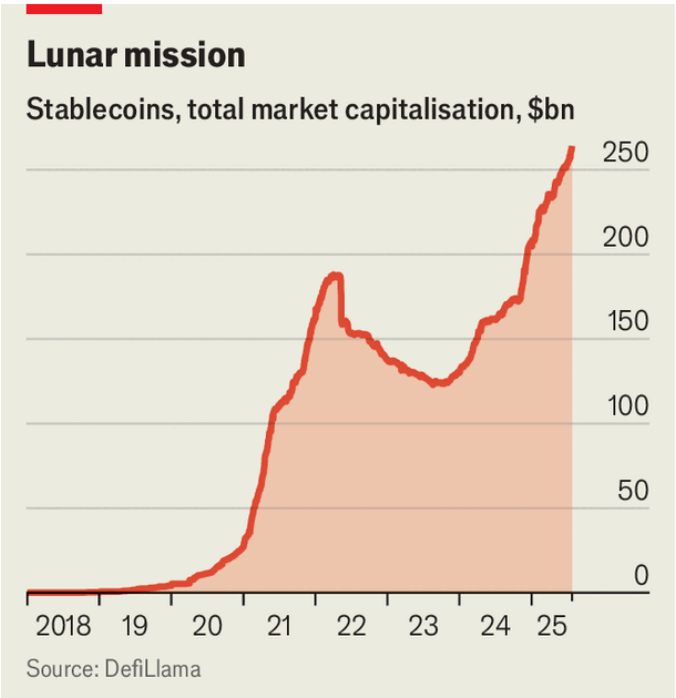

Chart: The Economist

Today there is $263bn in stablecoins in circulation, some 60% more than a year ago. Standard Chartered, a bank, expects the market to be worth $2trn in three years' time. Last month JPMorgan Chase, America's biggest bank, announced plans for a stablecoin-like product called JPMorgan

Deposit Token (JPMD), despite the long-held crypto scepticism of the firm's boss, Jamie Dimon. The market for tokenised assets is worth just $25bn but has more than doubled in size over the past year. On June 30th Robinhood launched over 200 new tokens for European investors, enabling them to trade American stocks and ETFs outside of ordinary trading hours.

Stablecoins allow for transactions that are cheap and fast, as ownership is registered instantaneously on digital ledgers, cutting out intermediaries who run traditional payment rails. This is especially valuable for cross-border transactions that are currently expensive and slow. Although stablecoins are now involved in less than 1% of financial transactions around the world, the GENIUS Act will provide a boost. It confirms stablecoins are not securities, and requires the coins to be fully backed by safe, liquid assets. Retail giants, including Amazon and Walmart, are reportedly considering their own coins. To consumers, these might work like a gift card, providing a balance to spend with the retailer, perhaps at lower prices. That would cut out firms such as Mastercard and Visa, which make a margin of 2% or so on sales they facilitate in America.

Tokenised assets are a digital copy of another asset, whether that is a fund, a share in a company or a bundle of commodities. Like stablecoins, they can make financial transactions faster and easier, particularly ones involving less liquid assets. Some offerings are gimmicky. Why tokenise individual stocks? Doing so may enable 24-hour trading, since the exchanges on which the shares are listed do not need to be open, but the advantages of that are questionable. And marginal trading costs are already very low, or even zero, for many retail investors.

Token effort

A lot of offerings are less gimmicky, however. Consider money-market funds, which invest in Treasury bills. A tokenised version could double as a form of payment. The tokens are, like stablecoins, backed by safe assets, and can be swapped seamlessly on blockchains. They are also an investment that beats bank interest rates. The average American savings account offers a rate of less than 0.6%; many money-market funds offer yields of 4%. BlackRock's tokenised money-market fund, the largest, is now worth over $2bn. "One day, I expect tokenised funds will become as familiar to investors as ETFs," wrote Larry Fink, the firm's boss, in a recent letter to investors.

This will prove disruptive for incumbents. Banks may be trying to get involved with the new digital wrappers, but they are doing so in part because they are aware tokens are a threat. A combination of stablecoins and tokenised money-market funds could, in time, make bank deposits a less attractive product. The American Bankers Association notes that if banks lost about 10% of their $19trn in retail deposits—their cheapest form of funding—it would raise their average funding cost from 2.03% to 2.27%. Although total deposits, including commercial accounts, would not be reduced, bank margins would be squeezed.

The new assets may also prove disruptive for the broader financial system. Holders of Robinhood's new stock tokens, for example, do not actually own the underlying securities. Technically, they own a derivative that tracks the value of the asset, including any dividends the company pays, rather than the stock itself. Thus they do not gain the voting rights usually conveyed by stock ownership. And if the issuer of the tokens goes bankrupt, its customers would find themselves in a difficult legal situation, competing with the collapsed firm's other creditors over who should take possession of the underlying assets. Something similar has happened with Linqto, a fintech startup that filed for bankruptcy earlier this month. The company had offered shares in private firms through special-purpose vehicles. Buyers are now unclear whether they own the assets they believed they possessed.

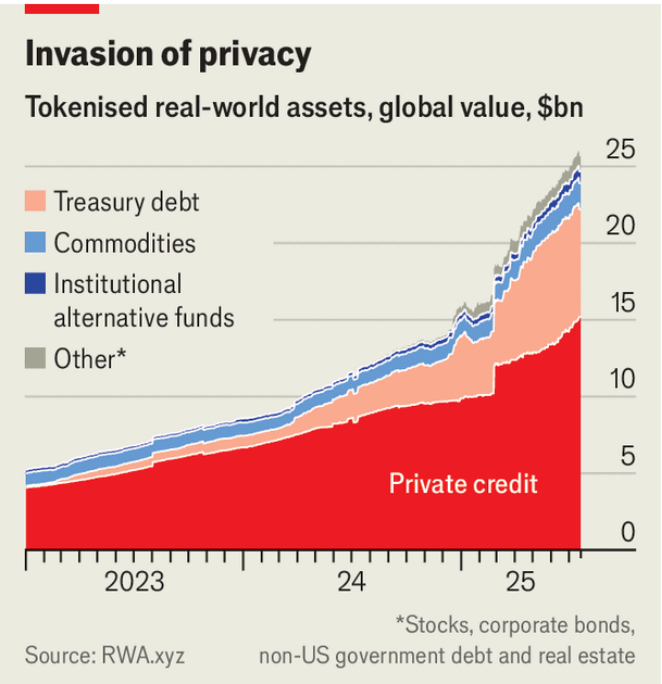

Chart: The Economist

It is one of the greatest opportunities for tokenisation that presents the greatest difficulty for regulators. Pairing illiquid private assets with easily exchanged tokens opens a cloistered market to millions of retail investors, who have trillions of dollars of capital to allocate. They could buy slivers of the most exciting private companies, currently beyond their reach. This raises questions. Agencies such as the Securities and Exchange Commission (SEC) have far more sway over publicly listed firms than private ones, which is what makes the former suitable for retail investment. Tokens representing private shares would turn once-private stakes into assets that could be traded as easily as an ETF. But whereas the issuers of an ETF promise to provide intraday liquidity by buying and selling the underlying assets,

the providers of tokens do not. At a large enough scale, tokens would in effect turn private firms into public ones, without any of the disclosure requirements normally required.

Even pro-crypto regulators want to mark clear lines in the sand. Hester Peirce, an SEC commissioner known as “crypto mom” for her digital-friendly approach, emphasised in a statement on July 9th that tokens ought not to be used to skirt securities laws. “Tokenised securities are still securities,” she wrote. As such, disclosure rules for companies issuing securities will be enforced, regardless of whether those come wrapped in new crypto packaging. Although that makes sense in theory, a plethora of new assets with novel structures means that, in practice, watchdogs will be endlessly playing catch-up.

So there is a paradox. If stablecoins are to be truly useful, they will also be truly disruptive. The more attractive tokenised assets are to brokers, customers, investors, merchants and other financial firms, the more they will change finance, in ways both welcome and worrying. Whatever the balance between the two, one thing is already clear: the view that crypto has not produced any innovations of note can be consigned to the past.

Want higher pay? Stay in your job

America's cooling labour market is bad news for those who move about

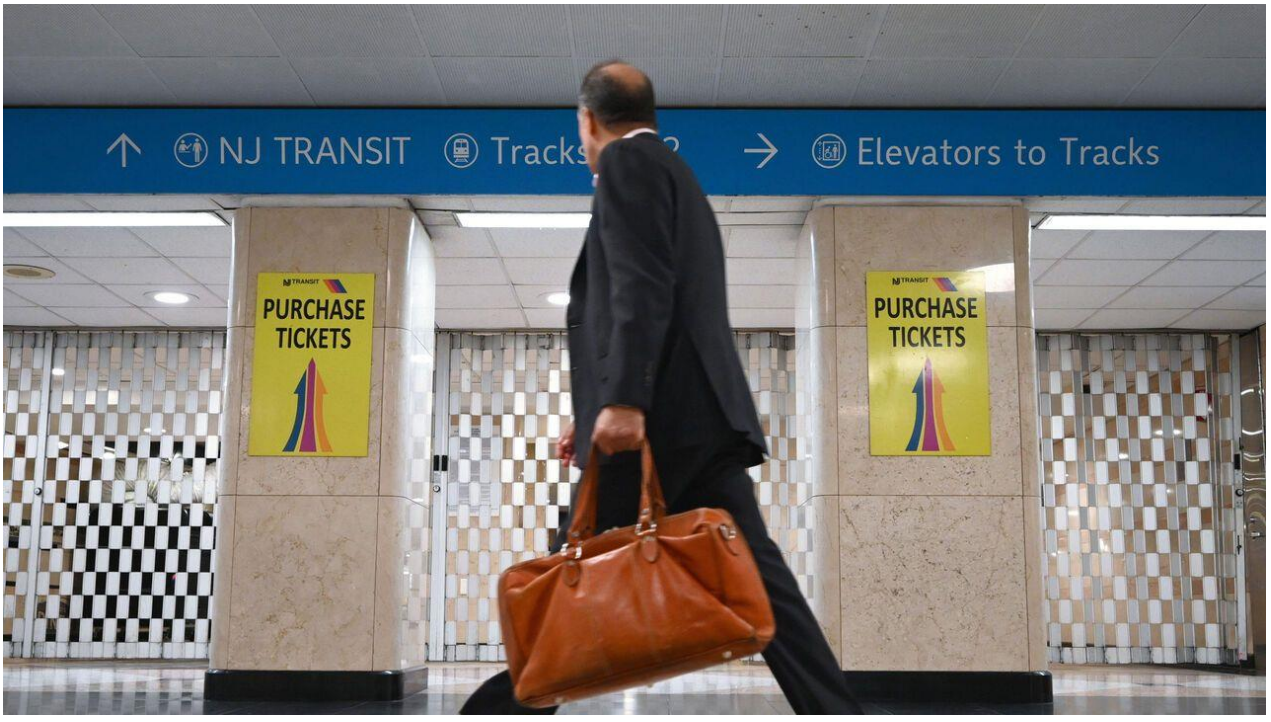

Photograph: Alamy

For years, America's job market has rewarded the footloose. The surest route to a higher salary, the advice goes, is to string together a series of one- or two-year stints, each paying a bit better than the last. Career gurus on TikTok set videos of their own salary progression to jaunty pop beats, cloaking online bragging as guidance for the uninitiated. On Reddit, posters debate just how little time in a role a job-hopper can get away with before future employers might start to fret about disloyalty. (A year or so is the consensus, though a brave few argue for six months.)

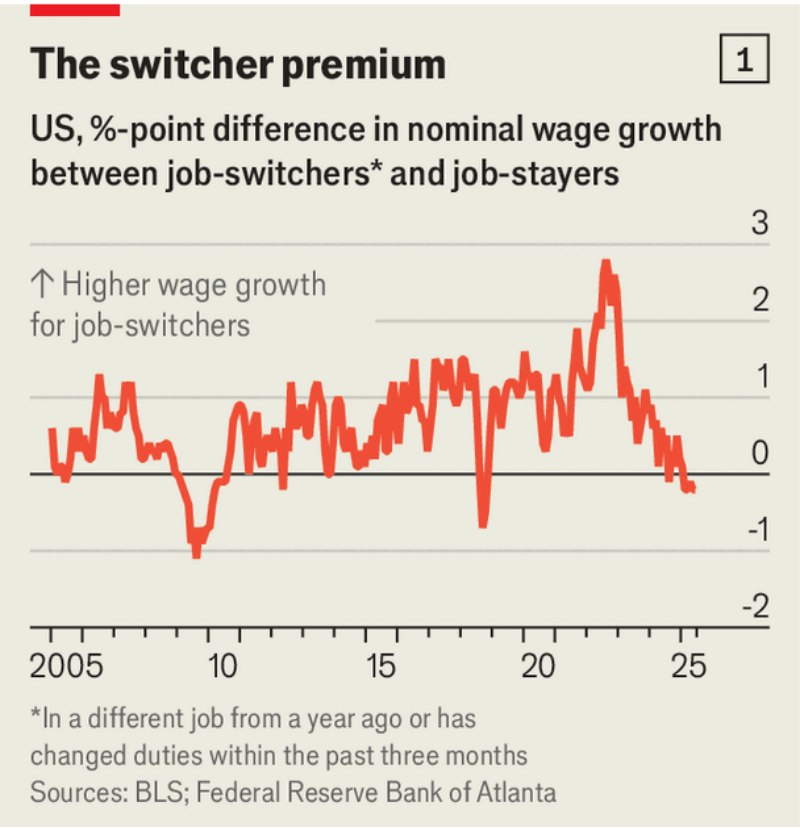

Chart: The Economist

Now things are changing. For the first time in 15 years, barring blips in 2012 and 2018, wage growth for "job-stayers", those who have stuck with their employer, is running faster than for "job-switchers", those who jumped ship, according to data from the Federal Reserve Bank of Atlanta (see chart 1). Hunkering down and impressing the boss may now be a better bet than polishing your CV or responding to those LinkedIn messages from head-hunters.

The woes of job-hoppers are a particularly visible sign of a wider trend: the softening of America's once rock-solid labour market. Another sign is the rising unemployment rate for fresh university graduates, who are struggling to persuade anyone to take a punt on them.

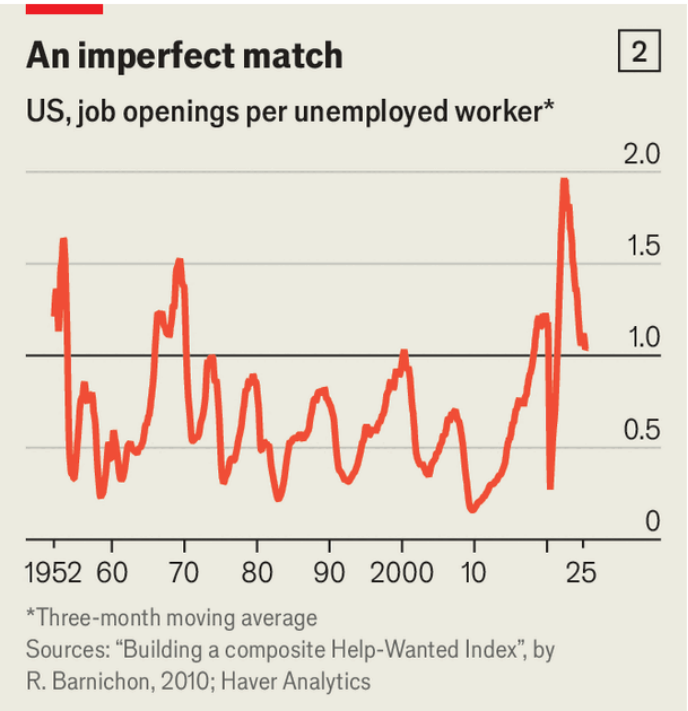

Chart: The Economist

There have been weak patches in the jobs market over the past few years. Some firms, like the big consultancies, vastly over-hired during the post-pandemic boom, and have since spent years painfully paring back. Coders have complained about poor prospects for a while; first, high interest rates hurt tech-firm valuations, then new artificial-intelligence models turned out to be particularly adept at programming.

But a broad jobs slowdown is a fresh development. Even as the stimulus-fuelled excesses of the post-pandemic years cooled, America's labour market remained impressively robust. For month after month, new payroll data would show that hundreds of thousands of jobs were being added to the economy. Now that motor may be slowing. June's payroll numbers looked strong on the surface—some 147,000 jobs were added across the economy, beating expectations—but half of those were in government, mostly teaching. Private-sector employment, a better measure of the underlying state of the economy, was disappointingly weak.

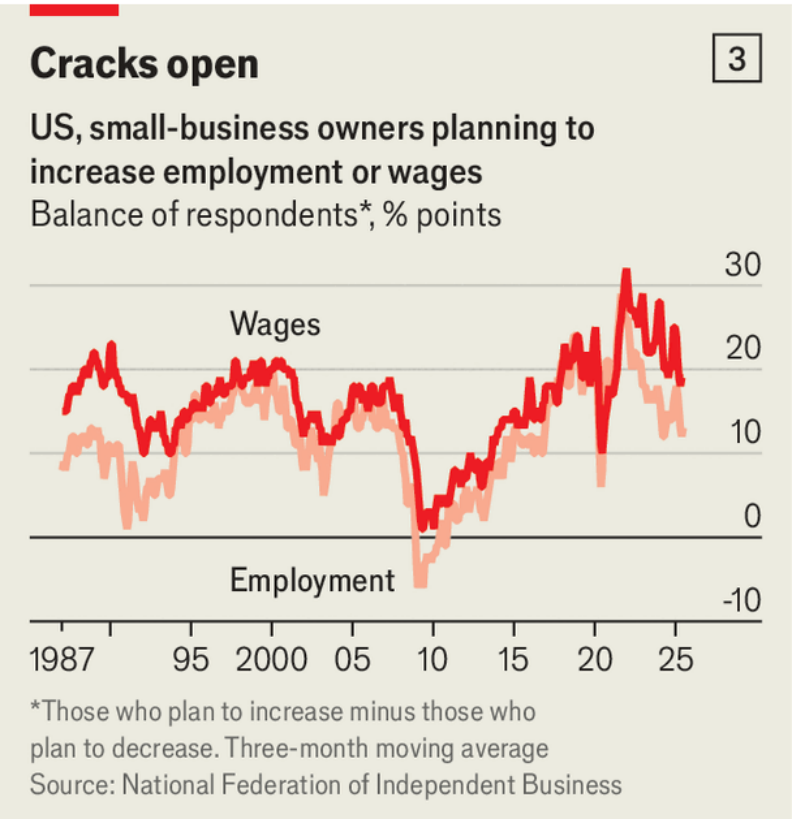

Chart: The Economist

In recent years there have been more job openings than people on the unemployment rolls, suggesting high demand for workers. Yet this ratio of vacancies to unemployment, which economists watch as a measure of labour-market health, has recently fallen back to one, meaning there are about as many openings as job-seekers (see chart 2). And the share of small-business owners who are planning to raise hiring or salaries is at its lowest since the pandemic (see chart 3).

What does all this mean for America's would-be job-hoppers? Once labour markets start cracking, they have a habit of deteriorating further. The Federal Reserve may, in due course, come to the rescue by lowering interest rates—but only after it can be persuaded that tariffs will not set off another round of inflation. If ever there was a time to redecorate your office and invest in a comfier chair, now might be it. You could be there for a while.

Why 24/7 trading is a bad idea

There are advantages to the old-fashioned working day

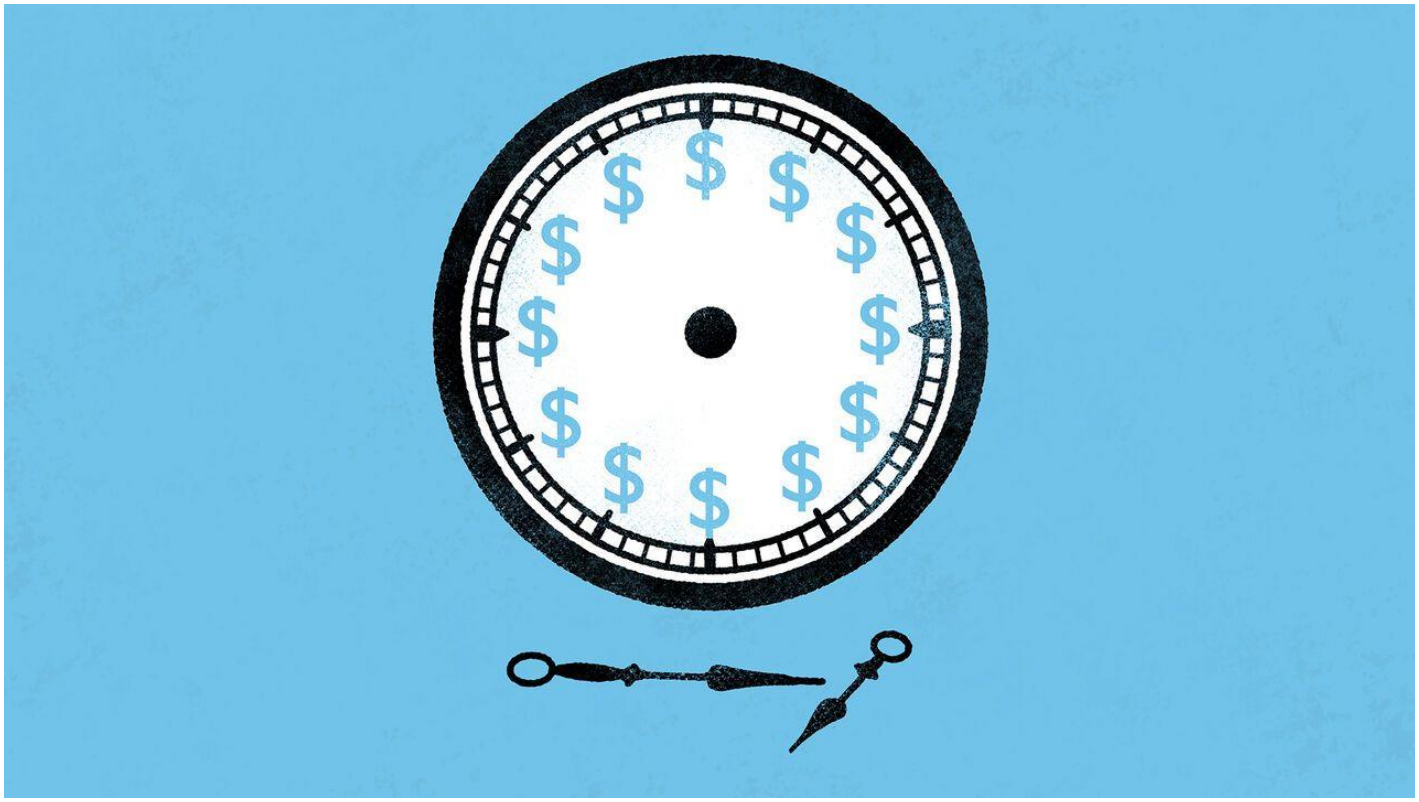

Illustration: Satoshi Kambayashi

Stock exchanges are the quaintest corners of modern finance. At other financial institutions, the typical trading floor features people in t-shirts sitting at ergonomic keyboards, sipping herbal tea and reviewing computer code. Enter New York's bourse, meanwhile, and you might as well have been through a time warp. Tense-looking people bustle everywhere sporting headsets and relics known as neckties. Everyone looks as if they shout a lot. As in a cattle market, opening and closing bells are rung to mark either end of the day's trading session, at 9.30am and 4pm.

This last bit is the most anachronistic, and not because of the bells. Today's markets, after all, have long ceased to respect the working day. Currencies, American Treasury bonds and crypto assets all trade around the clock.

Although investors the world over want to buy and sell American stocks, standard trading at both the New York Stock Exchange (NYSE) and its digital cousin, the Nasdaq, begins close to Hong Kongers' bedtime and ends while Californians are having lunch. Pre- and post-market sessions facilitate some transactions from New York's 4am and until 8pm. But few marketmakers operate during these, and liquidity can evaporate. Stock exchanges in Europe and Asia keep similar core hours to American ones, adjusted by time zone.

Now change is afoot. Both the NYSE and Nasdaq have applied for regulatory permission to extend their trading sessions, to 22 and 24 hours a day respectively. On July 20th the Financial Times reported that the London Stock Exchange is considering something similar. The transition is coming fast: Nasdaq expects to be open all night from the second half of 2026. So are investors, traders and financial plumbers ready for another market that never sleeps?

Some certainly are. Stock exchanges might not yet trade overnight, but plenty of alternative platforms do. In May 2023 Robinhood, a digital retail broker, began to allow all-night trading of 43 popular single stocks, before expanding this to many more. It joined Charles Schwab and Interactive Brokers, two rivals already facilitating some overnight trading; this year, on July 21st, Schwab said it would start offering the service for 1,100 securities. For decades institutional investors have had "dark pools"—venues for off-exchange stock trading, which can stay open all night. It is through these pools that the retail platforms often execute their clients' overnight trades.

The problem with such trading is that price discovery can be fraught with difficulty. In fact, this is partly why institutional investors like dark pools: their lighter reporting requirements, compared with exchanges, allow big orders to be executed without alerting the wider market beforehand, which would move the price. Professionals taking the other side of these trades accept the risks and know how to navigate them. Amateurs, getting a worse price than they might have done in daylight, often do not.

Overnight exchanges would increase transparency, but might not improve the second drawback of the small hours: low liquidity. Even during the day trading is concentrated around the opening and closing auctions (held just before the respective bells), with much lower volumes at other times. The phenomenon is self-fulfilling, since high auction volumes allow for better price discovery, which means professionals often prefer to place their trades during them.

Regulators could compel marketmakers to offer quotes throughout the night, as they currently do throughout the day. However, thin liquidity decreases the proportion of buy and sell orders that can be netted off, forcing marketmakers to hold positions on their books for longer. The result is more risk,

more capital that must be held against it and therefore higher trading costs. Even then, prices would remain more volatile than during the day.

And that is before considering the logistical nightmare that 24-hour exchanges would entail. The witching hours are currently when all manner of dull, but vital, post-trade processes take place, from settlement and valuation to the reconciliation of mistakes. Once trading is non-stop, there will be no pause for the financial pipes to clear—nor for traders to rest in the knowledge that the market is resting with them, so there is no need to refresh their screens. In today's always-on world, stock exchanges' limited opening hours might seem old-fashioned. But get ready to miss them once they're gone.

What economics can teach foreign-policy types

Hegemons should care about even puny countries

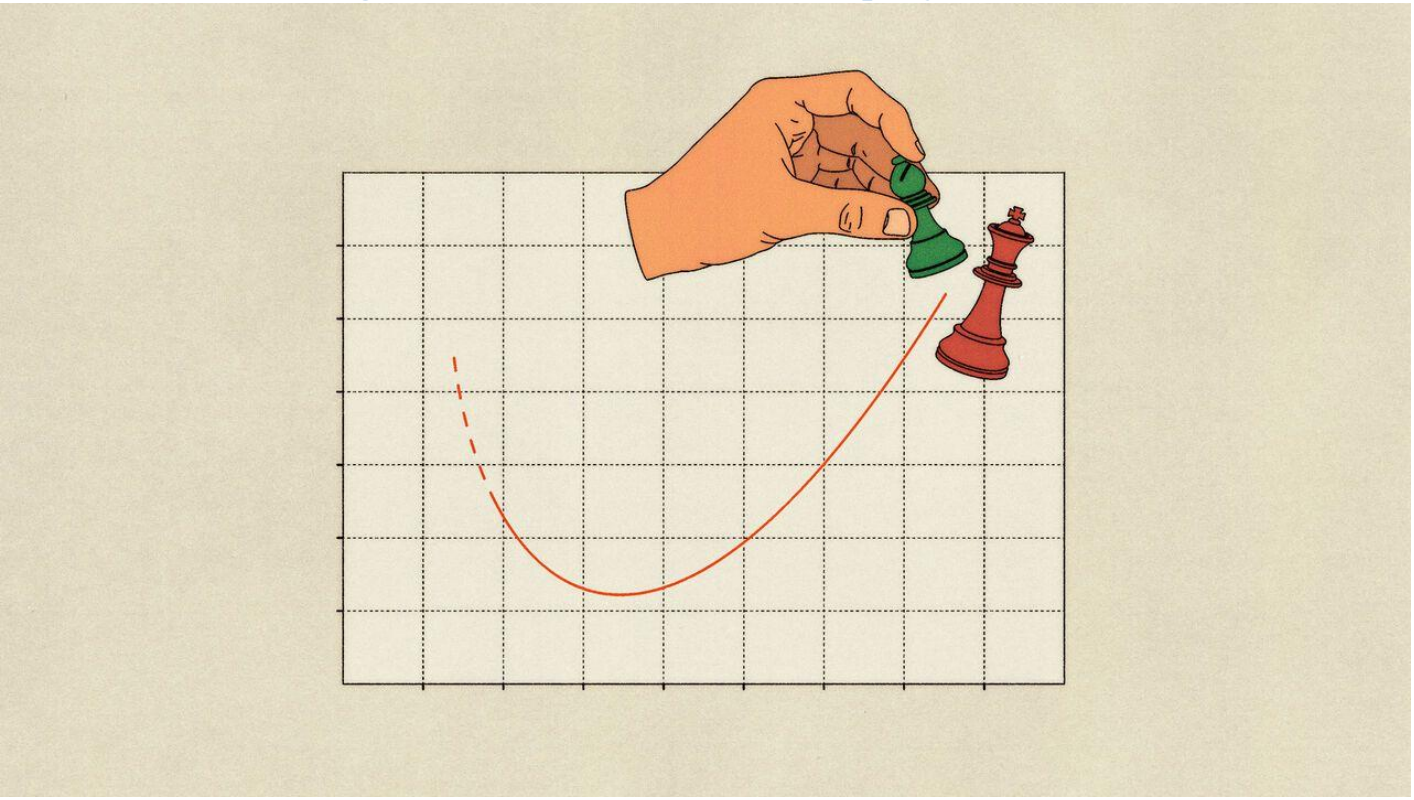

Illustration: Álvaro Bernis

Markets work best when many companies vie for customers' favour. They work badly when a few firms dominate, carving up sales between them. Economists therefore need a measure of whether markets are competitive or concentrated.

The most famous measure was invented by Albert Hirschman in 1945. It starts by summing the squares of each supplier's market share. That makes it sensitive both to the number of suppliers and inequalities between them. A staple of antitrust investigations, his index, with minor modifications, crops up in courtrooms and classrooms. Almost every economics student comes across it.

Few will know that Hirschman invented the index to measure something else: the economic power wielded not by firms but by countries. It appeared in a book examining trade as a source

of "political pressure and leverage". Hitler's Germany had exerted its influence over its neighbours in the 1930s not only through "diabolical cunning" but also the gravitational pull of its economy.

Hirschman rejected the naive belief that because trade is voluntary and mutually beneficial, it is geopolitically innocuous. Benefits can be mutual without being symmetrical. And if one country depends less on a relationship than its partner, it can extract concessions by threatening to walk away.

The German economist would not have been surprised by President Donald Trump's tariffs or indeed China's own attempts at economic coercion, which include export restrictions on vital inputs, such as rare earths. It is, therefore, a good time to revive the spirit of his inquiries, according to Christopher Clayton of Yale School of Management, Matteo Maggiori of Stanford University and Jesse Schreger of Columbia Business School. They are seeking to apply the modern toolkit of economics to geopolitics. The result is something they call "geoeconomics" (borrowing a term coined by Edward Luttwak, a historian and military strategist, in 1990). Whatever the subject is called, it is inescapable. Three years ago people asked the authors what geoeconomics was. Now they are asked not to go on about Mr Trump too much.

In their models, big economies—hegemons—can make demands of smaller ones by subjecting them to economic sticks and carrots. The losses a country suffers if it rejects such demands are a measure of the power the hegemon wields. Smaller countries can try to protect themselves in advance by decoupling and diversifying their economies. They might, however, overdo it, according to the authors' models. Economic networks, be they banking systems, industrial ecosystems, or global trade itself, increase in value the more people take part in them. If one country withdraws to protect itself, it makes that network less attractive to others. That might shift their calculus in favour of decoupling, too.

To dissuade countries from insulating themselves, a hegemon might promise not to exploit its power too much. It might say, "Do business with me. I'm not going to bully you like crazy later. I'm only going to bully you a little bit," as Mr Maggiori has put it. The hegemon might, for example, submit to international trade rules that put a ceiling on its tariffs. If the rules are credible, they can benefit the hegemon as well as everyone else. In these models, trade rules emerge not as the result of "globalist" planning, but as an act of enlightened self-interest on the part of a hegemon. The models make the nationalist case for multilateralism.

Their theory also aids measurement. In the spirit of Hirschman, they calculate their own indices of power, based on market shares and the ease with which imported inputs can be replaced. Their calculations bring out a vital implication of their theory: power does not grow in a straight line. It tends to shoot up as a hegemon's market share nears 100%. There is an enormous difference between claiming most of a market and claiming almost all of it. The same is true of hard-to-replace inputs. There is a big difference between having few substitutes and none or nearly none.

This result makes intuitive sense. If a hostile America provides 80% of a country's foreign financial services, the country's other providers would have to expand their sales by 400% to replace it (an increase from 20 to 100 equals growth of 400%). If America instead provides 90%, the alternative providers would have to expand by 900%. An increase in market share of ten percentage points makes a 500 percentage-point difference.

All this has practical implications. America and its allies have a large share of the world's financial services. Adding a country like Singapore to the coalition might not seem like much of a prize. But it would make a big difference; a small increase in a large market share yields disproportionate gains in power.

Bully for them

The same is true in reverse. Even small efforts to decouple from an abrasive hegemon can diminish its power by a surprising degree. America's rivals are, for example, experimenting with alternatives to the dollar in international finance. These alternatives do not have to match the dollar in order to erode America's financial clout. If they can capture even 10% of the market in small and medium-size countries, they can make a big difference, according to the authors. Going from 1% of the market to 10% can reduce American financial power as much as going from 10% to 50%, they argue. According to this logic, currencies like China's yuan can defang the dollar even if they never dethrone it.

Hirschman's book, in which he outlined his index, was largely forgotten by the economics profession until recently. Political scientists and other scholars now account for most of its citations, according to Messrs Clayton, Maggiori and Schreger. The three authors hope to revive interest in Hirschman's work among their fellow economists. The inventor of the best known index of concentration warrants a broader, more diversified readership.

Tecnology
Inside the top-secret labs that build America's nuclear weapons

To maintain the bombs, and build new ones, scientists are pushing the frontiers of physics

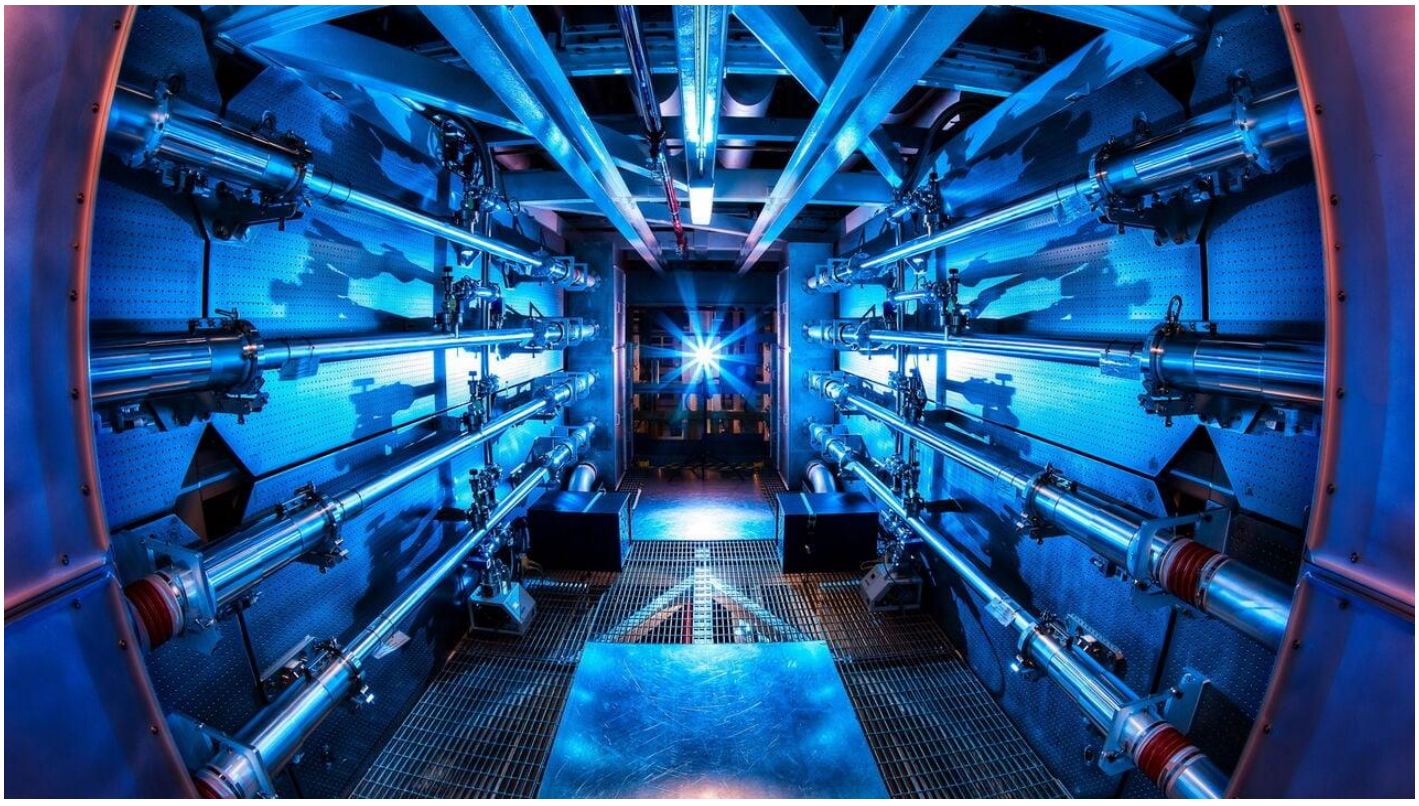

The National Ignition Facility's preamplifier module increases the laser energy as it travels to the target chamber Image: Damien Jemison/LLNL

Each experiment at the National Ignition Facility (NIF) in California—a "shot"—lasts just a few billionths of a second. A lot happens in that brief moment, however: 192 laser beams, totalling some 500trn watts, converge in the machine's target chamber and dump their energy onto a gold cylinder, which is just a few centimetres long. Inside the cylinder is a peppercorn-size diamond sphere filled with a mixture of deuterium and tritium, heavy isotopes of hydrogen.

As the sphere absorbs the laser's energy, its outer layers rapidly ablate away. That creates a shock wave travelling at 300km per second that implodes the sphere's insides. As the atoms of deuterium and tritium are pushed together at billions of times atmospheric pressure, their temperatures exceeding 100m°C, they start fusing into helium, releasing vast amounts of energy.

This is the kit you need to be able to re-create a nuclear-weapon explosion without actually setting off a bomb. NIF was conceived in the 1990s, a few years after America decided to stop testing its nuclear arsenal in underground explosive tests. Without these tests, the people responsible for the country's nuclear deterrent still needed ways to guarantee the safety of their warheads as they sat in storage and, most important, instil confidence that they would perform as intended, if they were ever called upon.

The facilities that America's nuclear establishment developed to answer that challenge eventually included NIF, the world's most powerful laser, and El Capitan, its fastest and most capable supercomputer. Both have become central to a renewed mission for America's nuclear-weapons labs, as they upgrade their existing bombs and, for the first time in decades, design brand new ones.

Maintaining nuclear weapons takes an army of scientists and engineers. NIF is part of the Lawrence Livermore National Laboratory near San Francisco, set up in 1952 as a rival to the Los Alamos National Laboratory in New Mexico. It was at Los Alamos that the first nuclear bombs were built less than a decade before. "We were developing this advanced technology in a very classified environment," says Kim Budil, Livermore's boss. "It was really important to bring scientific rigour, peer review and competition to that technology race." The two labs purposefully pursue different designs for weapons and, though they sometimes collaborate, refer to each other as "competimates".

Testing under way at the Sandia Laboratory. Image: Sandia Labs

Livermore and Los Alamos design the "physics packages" in America's warheads, which is to say the nuclear bits of the nuclear bombs. A third institution, Sandia National Laboratories, adds the non-nuclear components (such as triggers, batteries, sensors and radiation-hardened electronics) and integrates the devices made by the two physics labs with the delivery systems (eg, missiles) that turn them into robust, deployable weapons. All told, the three labs of the National Nuclear Security Administration (NNSA) employ tens of thousands of scientists and engineers. All three granted The Economist rare access to their researchers and some of their facilities.

When Livermore opened, one of its primary goals was to accelerate the development of hydrogen, or thermonuclear, bombs. Unlike the fission bombs that had been developed in the Manhattan Project, which released energy by splitting atoms of heavy elements (uranium and plutonium), thermonuclear bombs were designed to release energy by fusing atoms of deuterium and tritium, some of the lightest in existence. (These bombs are called thermonuclear because they have two stages: first, a fission bomb made of plutonium which creates an intense burst of heat; that then ignites a second stage in which the fusion occurs.)

Thermonuclear technology opened the door to more powerful but also more compact weapons. In the 1950s, when the US Navy decided to create a sea-based nuclear deterrent, Livermore was assigned the task of miniaturising nuclear bombs so that they could be affixed to missiles that fit inside submarines. It took them less than four years to come up with Polaris, a missile system an order of magnitude smaller than anything that had come before and which Dr Budil proudly describes as "the single most important technology change in the history of nuclear weapons."

Small, compact thermonuclear devices became the workhorse of both the American and the Soviet nuclear arsenals as they were expanded during the cold war. Fortunately, none of these weapons was ever used in anger and, decades after being built, thousands remain in their stockpiles.

One of the biggest tasks occupying the scientists today at the Los Alamos, Livermore and Sandia labs is to keep a close watch on those warheads. "A nuclear weapon sitting on the shelf is sort of like a chemistry experiment cooking along year after year," says Dr Budil. "Things are changing. Radioactive materials decay over time. Polymer materials degrade."

Every year a few devices are taken apart and thoroughly examined. More extreme testing also happens. Microscopic samples of material are placed inside NIF's target chamber, where they can be imaged by X-rays while experiencing the equivalent of a nuclear blast. At Sandia, the Z machine is another way to approximate the core of a nuclear blast, but using intense electromagnetic fields rather than lasers.

At Los Alamos, by contrast, the non-nuclear parts of the weapons are blasted by shock waves from the conventional explosives that are used to initiate a nuclear bomb.

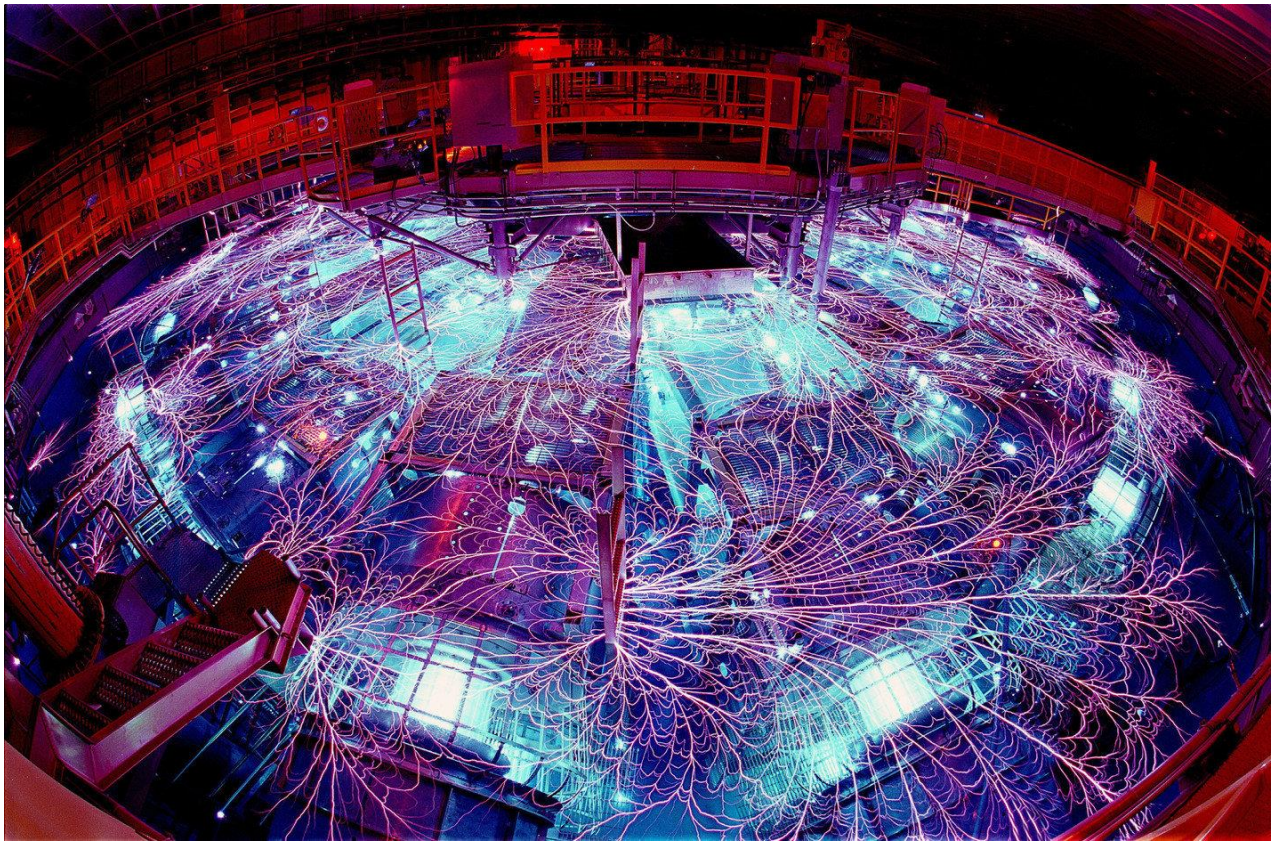

The Z machine uses electromagnetic fields to simulate a nuclear blast. Image: Randy Montoya/Sandia National Laboratories

All that experimental work is used to better understand the properties of materials that go into bombs. And, alongside the thousand or so full-scale nuclear-weapons tests carried out before 1992, the data are also used to build better computer simulations of nuclear blasts. These are now so good that Thom Mason, director of Los Alamos, reckons that scientists have a better understanding of how nuclear weapons work today than they did during the explosive-testing era. “The modern scientific tools really outstrip significantly anything that we had in the 1990s,” he says.

Number crunchers

Exactly how much better is demonstrated at Livermore’s computing centre, a few minutes’ walk from NIF. In January, scientists and government officials gathered there to unveil the NNSA’s latest (and

now the world's most powerful) supercomputer—El Capitan. This machine can run a quintillion (1018) floating-point operations (a measure of calculations) per second. That is around 100m times faster than a typical laptop, and makes it only the third ever exascale computer ("exa" being the measurement prefix for 1 followed by 18 zeros). Its roughly 90 refrigerator-size racks of processors are densely packed over the same space as a couple of tennis courts.

The supercomputer is part of the Advanced Simulation and Computing (ASC) programme, started in 1995, alongside NIF, as part of America's response to its moratorium on nuclear-weapons testing. One of its first goals, set for the turn of the millennium, was to assemble the hardware and software required to run a three-dimensional simulation of a weapon system.

Scientists overcame the enormous challenges using the parallel-computing architecture that was becoming possible at the time. This meant splitting up a simulation into small chunks that could be run simultaneously across the central-processing units (CPUs) and graphics-processing units (GPUs) found in high-end computers. It still took months to run a single simulation. "On El Capitan, we're now estimating we could be able to run upwards of 200 of those in a day," says Rob Neely, Livermore's associate director for weapon simulation and computing. And all that at much higher resolution too.

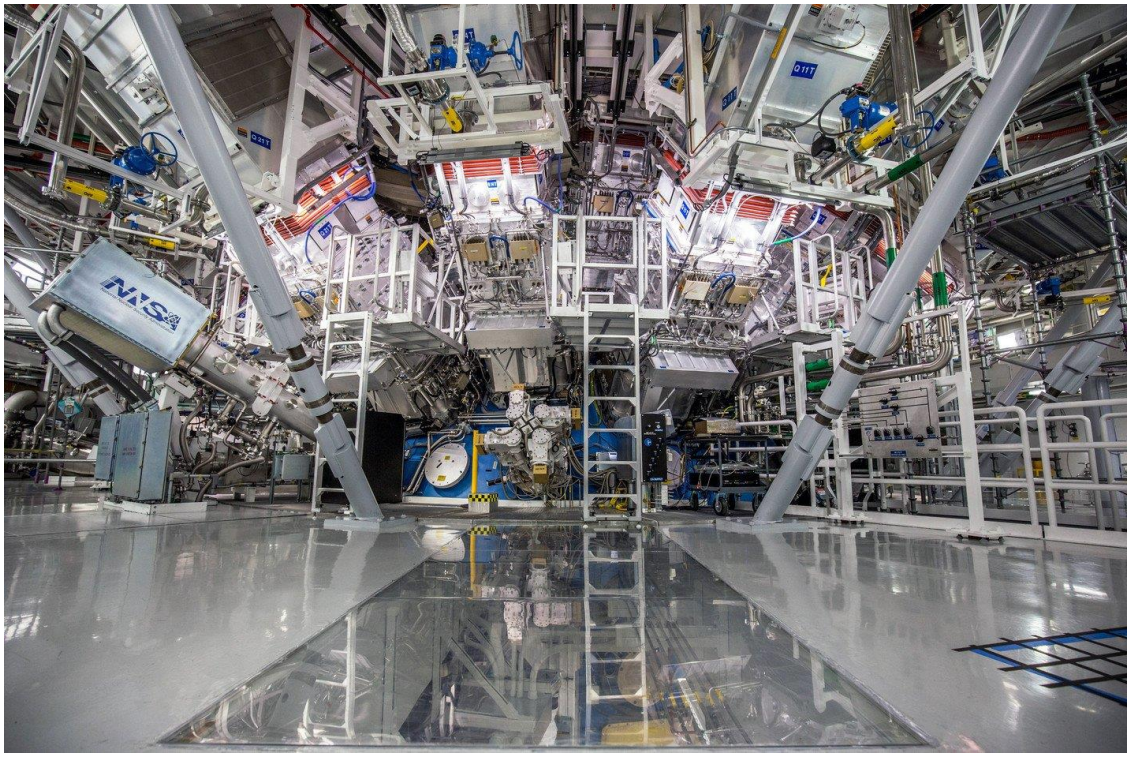

The target chamber of the National Ignition Facility, where fusion ignition has taken place. Image: Jason Laurea/LLNL

Look closer at the processors and something else becomes apparent. Instead of CPUs and GPUs, El Capitan uses specialised chips developed for Livermore by Advanced Micro Devices, a chip designer, called accelerated-processing units (APUs). Typically GPUs and CPUs will have their own storage and memory and the communication between them, known as the bus, can become a bottleneck to a system's speed. Each APU is, instead, a single piece of silicon with sections ("chiplets") that individually operate as CPUs or GPUs, allowing them to share memory and storage. "It's the only architecture in the world right now that we know of that's doing it this way," says Dr Neely.

The density and architecture of those APUs give El Capitan its edge over machines that might, on paper, have more raw computing power. At Los Alamos, the simulations are also being deployed for a new task—designing a new weapon from scratch. The W93, as it is called, will eventually be used on ballistic missiles deployed by the US Navy's new Columbia-class submarines. It is the first new weapon in the American nuclear arsenal since the 1980s and, with explosive tests off-limits, Los Alamos will need to run simulations from the very start of the design process. El Capitan will allow scientists to optimise the design, says Dr Neely.

The W93 is emblematic of the renewed energy at Los Alamos. "Our budget has roughly doubled over the past five or six years," says Dr Mason. That means thousands more scientists, modernised facilities and a restored ability to make plutonium pits, a core element of modern thermonuclear bombs. And, in contrast to many other areas of scientific research in America today, the budget for the NNSA is not expecting any cuts in federal funding.

All this is a response to what Dr Mason calls the "fourth age" of nuclear weapons. The first was the invention of nuclear bombs during the Manhattan Project; the second was the cold-war race to build up nuclear arsenals; and the third age was the period after the fall of the Soviet Union during which it was thought that nuclear deterrence would have a declining role in world affairs. The fourth nuclear age is a worrying time featuring the breakdown of arms control, Russia's threats of nuclear use, China's rapid build-up and tensions among other nuclear powers such as India and Pakistan. There is also uncertainty over new and would-be nuclear powers, and the risk that America's allies could develop their own nuclear weapons as they lose faith in its protective umbrella. "It's clear that deterrence is, once again, pretty important," says Dr Mason.

Though the primary purpose of the labs at Los Alamos and Livermore is never in doubt, their scientists are keen to point out that these facilities can do much more than national-security work. NIF, for example, is a leading laboratory in the attempt to create power from nuclear fusion.

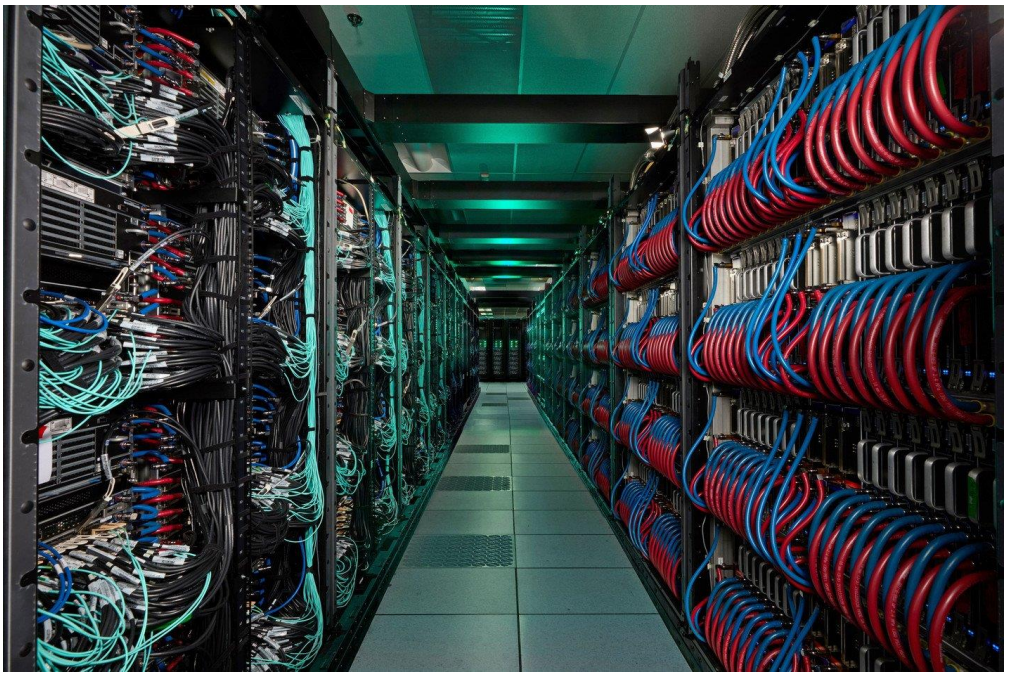

Top: El Capitan, based at the Lawrence Livermore National Laboratory is the world's most powerful computer. Bottom: B61-12 nuclear gravity bombs at Sandia. Image: Garry mcLeod/LLNL; Craig Fritz/Sandia National Labs

In December 2022 NIF made good on the "I" in its name and became the first site in the world to achieve ignition—releasing more energy from fusion than had been used to get it going. Since then the scientists there have achieved ignition on eight more occasions, gradually increasing the energy yielded each time.

Mark Herrmann, programme director for weapons physics at Livermore and a former director of NIF, is well aware that it will take a lot more work to turn these breakthroughs into a viable source of energy. For a start, the lasers themselves have to get a lot more energy-efficient and the fusion reactions would need to happen dozens of times per second (rather than just a dozen times per week). Although more engineering work is needed, says Dr Herrmann, "There are no scientific obstacles to those things happening."

Deterrence, undeterred

It's the weapons, though, that these labs exist for. And their terrifying power is never far from the minds and motivations of the scientists involved. When asked how he and his colleagues feel in their role developing nuclear bombs, Dr Mason points to the (albeit occasionally uneasy) geopolitical order that has been maintained as a result of people's fear of their power. "If the weapons we design are never used," he says, "we will have been successful."