Annex 9. AI misdemeanours

Empires of AI

Empires of AI

The Rise and Fall of a New Global Order

Niall Ferguson (https://www.foreignaffairs.com/authors/niall-ferguson), Lisa Klaassen (https://www.foreignaffairs.com/authors/lisa-klaassen), and Eric Schmidt (https://www.foreignaffairs.com/authors/eric-schmidt)

November/December 2026 (https://www.foreignaffairs.com/issues/2026/105/6)*Published on October 20, 2026*

Ed Johnson

Listen

Share &
Download

Print (https://www.foreignaffairs.com/subscribe)

Save (https://www.foreignaffairs.com/user/login?destination=/china/empires-artificial-intelligence-niall-ferguson-eric-schmidt%3Fsma%3Dnewsletter_geopolitica%26utm_medium%3DPANTHEON_STRIPPED%26utm_source%3DPANTHEON_STRIPPED%26utm_campaign%3DPANTHEON_STRIPPED)

In the span of a few years, China and the United States have each assembled something future historians might recognize as a new kind of empire. Their dominions extend not merely over territory, trade routes, and airspace but over knowledge itself. By controlling the large language models, chips, physical infrastructure, and capital that provide the foundations for advanced machine cognition, Beijing and Washington are not only attaining new heights of economic performance and power projection; they are also setting the terms of all future scientific and intellectual progress. The term “AI superpowers” fails to capture the influence and reach that these rivals now wield. Even without active alliance building or coercion, they can make all other countries that seek access to advanced AI beholden to them.

Imperial history carries a warning. In the absence of other forces, empires seldom succumb to a challenger or to barbarians at the gates. More often, they fail because defending their vast

dominions drains them dry, whether of material resources, the consent of the governed, or the loyalty of elites. The imperial center finds it can no longer maintain control over the periphery—the regions that once submitted to it. Failure can be financial, as it was for the British Empire after 1945, when the debts of two world wars forced London to liquidate its imperial assets to pay the price of victory: a welfare state for the citizens of the imperial metropole. It can be institutional, as it was for the Qing empire in the late nineteenth century, when it failed to industrialize, reform its bureaucracy, and renew the legitimacy of the Manchu dynasty.

Or it can be technological. The Royal Navy's *Dreadnought* of 1906 gave the United Kingdom (https://www.foreignaffairs.com/regions/united-kingdom) the first "all big gun" battleship and, with the same stroke, rendered its existing battle fleet obsolete. Overnight, this breakthrough opened the way for rivals such as Germany, erasing a century of accumulated British advantage. The resulting arms race forced the United Kingdom to nearly double naval spending between 1900 and 1914 and clashed with the Liberal government's new social spending commitments, such as the Old Age Pension Act of 1908. The decisive question in the contest between Beijing and Washington may therefore not be who achieves artificial general intelligence first, but who can bankroll the race.

AI is imperial in two ways. On the one hand, China (https://www.foreignaffairs.com/regions/china) and the United States had empire-like qualities before the AI race began: Washington, the incumbent, pairs global military primacy with an alliance system no rival can match. It also supplies the currency for the lion's share of international trade and transactions. Beijing, often described as the challenger, combines industrial and supply chain dominance with the world's largest navy and a tightening grip on the physical and digital infrastructure of countries across Africa, Asia, and Latin America. For both, AI is a way to extend the power they already wield. But AI is also by its nature an imperial substance. Whoever produces and controls frontier models acquires global leverage over innovation, production, science, culture, and, ultimately, advanced thought. Much as silver underwrote Habsburg Spain and oil defined the "American century," this new commodity will not merely enrich the powers that dominate it but also provide the basic operating system on which everyone else must depend.

Subscribe to *Foreign Affairs This Week*

Our editors' top picks, delivered free to your inbox every Friday.

Sign Up

* Note that when you provide your email address, the *Foreign Affairs* Privacy Policy (https://www.foreignaffairs.com/privacy-policy) and Terms of Use (https://www.foreignaffairs.com/terms-use) will apply to your newsletter subscription.

What distinguishes the AI (https://www.foreignaffairs.com/tags/artificial-intelligence) empires from their forerunners is not ambition but velocity. The familiar imperial arc—expansion, overreach, strain, unraveling—once played out across a century or even a millennium. Today, that process is happening at machine speed. The question is whether the new hegemons can convert their particular endowments—the United States' capital markets, frontier labs, and transistors and China's state direction, cheap energy, and watts—into durable imperial orders or whether they will run aground trying.

SUPREMACY BY SCARCITY

An empire is more than a great power. It is a polity whose authority and influence extend over vast territories and peoples without their express consent. Throughout history, empires have differed in almost every respect: in their form of government at home, in the kinds of resources they seek to control, in how their authority is sustained. They may be autocracies or republics; they may rely on soldiers and settlers or civil servants and chartered companies. Yet the underlying vector of power never changes. A center commands, and a periphery obeys. In practice, the most durable empires have relied on informal collaboration with local elites, whereas brutally coercive ones tend to be short-lived. Of course, any form of state power must not only monopolize violence but also be able to wield it. What sets empires apart is the capacity to elicit submission across immense distances. And they occur throughout history for a reason: by exploiting economies of scale, empires can achieve dominance and deliver security more efficiently than alternative structures, whether confederations or alliances of city-states. The nation-state is a relative newcomer.

Empires have often projected power through technological leadership. In antiquity, Rome's aqueducts and China's canals and flood-control works were the pipelines of centuries of hegemony. The imperial powers of early modern Europe led the world in printing, maritime technology, and artillery. The Cold War (https://www.foreignaffairs.com/tags/cold-war) rivals were nuclear superpowers. Yet artificial intelligence may now give China and the United States greater imperial reach than any of those earlier technologies—and far more swiftly.

Several features mark AI as an imperial technology rather than a merely commercial one. First is the return of scarcity that it enables. An empire is, at its heart, a system for controlling finite strategic resources: able-bodied humans, arable land, grain, metals, and minerals. From the 1970s to the 2010s, software escaped this logic because it could be mass-produced at almost zero cost atop abundant, cheap silicon. AI has shattered that reprieve. According to the industry's scaling laws, a model's performance climbs relative to the number of parameters it contains, the volume of data it is trained on, and the processing power, or compute, spent training it. Progress therefore depends on energy, critical minerals, advanced chips, and capital

on a massive scale. The moment a technology's frontier comes to depend on vast quantities of scarce inputs, its production becomes a matter of geopolitics.

AI is by its nature an imperial technology.

AI also has an inherent center-periphery structure. Control of frontier AI and its rewards are concentrated in a handful of firms in China and the United States (https://www.foreignaffairs.com/regions/united-states). Everyone else is a dependent user, whether a government running public services on a model it cannot inspect or a firm building products on an interface that can be repriced or withdrawn overnight. At the same time, AI inputs, both material and human, are drawn disproportionately from the periphery. Minerals from the Democratic Republic of the Congo and Indonesia and "ghost laborers" from Kenya and the Philippines, paid a few dollars an hour to label and filter the data from which a model in Silicon Valley learns—the plantation economics of the chatbot age. Raw inputs are extracted from the periphery, refined and assembled in the metropole, and sold back at a markup.

Like many empires, the AI stack is organized around chokepoints. Not all imperial power results from outright conquest, and not all those who are ruled are colonial subjects. The ascendancy of the British Empire in the nineteenth century was not driven primarily by vast territorial holdings, many of which had little economic value, or from control over factories and intellectual property—which was so poorly protected as to be effectively open source. Instead, British supremacy came from dominating the physical, technological, and financial pipelines through which the world economy flowed: coaling stations, strategic waterways such as the Suez Canal and the Strait of Malacca, undersea telegraph cables, and sterling finance through the City of London. The U.S.-led order after 1945 may have been an "empire of invitation," as the historian Geir Lundestad termed it, but it also rested on the dollar and the Bretton Woods institutions, a U.S.-dominated international banking system, a global network of American military bases, and Washington's nuclear umbrella. Today, the production and distribution of AI provides an even denser imperial lattice, with the metropoles able to switch its different elements—chips, minerals, models—on or off, as needed.

There is at least one crucial difference between the two AI empires. China's is a deliberate act of state, whereas the United States' is not. In 2017, China's State Council announced a comprehensive plan to achieve global leadership in AI by 2030—specifically calling for China to develop "the world's leading AI backbone enterprises" and to secure broad diffusion of its AI applications in the dozens of countries that are part of Beijing's Belt and Road Initiative. Ministries, provinces, and state banks were mobilized to fulfill these goals, and Chinese leader Xi Jinping (https://www.foreignaffairs.com/tags/xi-jinping) has continued to reinforce the effort, including with his 2026 push for Beijing to shape global AI standards. American AI dominance, by contrast, has emerged from private capital, a regulatory posture of noninterference, and

improvised export controls. In this sense, Washington's empire is the more conventional of the two: earlier empires were seldom fully deliberate undertakings; the United Kingdom "seemed to have conquered and peopled half the world in a fit of absence of mind," as the historian John Robert Seeley put it. Yet both AI powers converge on the same architecture, mandated by the scaling laws. Below a certain threshold of energy, capital, and compute, a country is not a weaker player; it is not in the game at all.

WATTS VS. TRANSISTORS

The great corporation-states of the European imperial era—the Dutch and British East India Companies—minted their own coins, fielded private armies, levied taxes, and bargained with kings as equals. They drove imperial domination by creating trade monopolies, defending them with military force, and ruthlessly exploiting the periphery. Today's counterparts are frontier AI companies such as OpenAI and Anthropic, which command the human talent, data, models, and finance on which advanced intelligence depends. They are not merely aggregators of capital but arbiters of knowledge, increasingly determining what is discovered, how quickly, and to what end.

The resemblance extends to statecraft. Like their East India predecessors, the leaders of the largest AI labs negotiate with presidents and cabinet secretaries. They also bargain with governments that do not possess a frontier model of their own, trading model access, compute allocations, and data center investment for land, power, sovereign capital, and regulatory accommodation. These are agreements struck between a company and a state, or between the company and economic entities within that state, on terms the company largely sets. Read in this light, the deals struck during U.S. President Donald Trump's visit to Saudi Arabia and the United Arab Emirates in May 2025, when American AI and chip companies offered access to their most advanced technology in return for petrodollars and cheap energy, amount to imperial concessions in all but name. Nor is the practice confined to the American side. Even as China races the United States for AI supremacy, Reuters has reported that China's Moonshot AI is in talks with Amazon, Microsoft, and Alphabet for a share of the revenue those companies earn from hosting its Kimi K3 model.

Sovereign though these firms may seem, they cannot move a single chip or build a data center without state authorization, and they operate under an export-control regime that can throttle any of them at will. Like the East India Companies and the Rothschild banks, which were ultimately subordinate to the governments of the metropole, frontier labs have been granted extraordinary latitude in exchange for advancing the metropole's position. In the United States, Trump's administration has thus far declined to regulate models in any serious way, preempted state AI legislation, and streamlined permitting for data centers. But it expects privileged access for defense and intelligence agencies, demands strict compliance with export controls, and asserts

the right to grant or withhold U.S. technology to other countries as an instrument of foreign policy. In China, the state supplies what the market will not: land, power, subsidized fabrication, provincial procurement, and, increasingly, capital. In exchange, models are required to undergo a security review and ideological evaluation by the Cyberspace Administration of China and are pushed toward domestic chips. Each state has chosen a different leash with which to collar its national champions.

Empires provide public goods—order, standards, security, a reserve currency, access to technology—in exchange not merely for revenue but for alignment or submission. To allies and partners, Washington tenders the full AI stack: advanced chips, costly frontier models, tooling, and the security relationship that comes with them. In return, it expects those countries to comply with U.S. export controls, restrictions on Chinese hardware and data, and, in some cases, tariffs or outright bans on Chinese technology. Beijing is more subtle: its open-weight models can be freely downloaded and modified; adopt them, along with China's hardware, cloud, and engineers, it tells client states, and the standards will follow. For both metropoles, this is imperial strategy in its purest form, not just a scramble for market share.

Governments across the globe must pick a side—whose chips, whose clouds, which camp—a binary choice that is increasingly uncomfortable for many. What China offers is cheap and mostly open; what the United States offers is expensive and mostly closed. An empire assembled from low-cost components and with no barriers to access will likely gain allegiance from the majority of the world's economies. An empire of premium, restricted technology will attract only those who can afford it and are permitted access. Beijing has built on the success of its low-cost 5G hardware, electric vehicles, and other products, with which it gained rapid dominance in many parts of the world. Washington has drawn on the familiar Silicon Valley approach—private firms fighting it out, with the winners digging whatever proprietary moat they can around their technologies—but then reoriented it toward national goals. Whereas the commercial version of this strategy ends in universal diffusion at a monopolist's price, the imperial version ends in the metropole deciding who gets what. In June, the U.S. Commerce Department told Anthropic to suspend all foreign nationals' access to its two most capable models, and those models went dark worldwide within days. In each case, what began as the characteristic business cultures of the two AI superpowers is now embedded in their foreign policy.

A data center in Rivabellosa, Spain, June 2026Vincent West / Reuters

Yet a durable empire cannot rely only on the goods it provides to its periphery. It also needs a story to convert its monopoly of power into legitimate authority. The Pax Romana, the Christian

and Islamic missions, the *mission civilisatrice*, Manifest Destiny, the white man's burden, the global march of communism, the shining city on the hill—every successful imperial regime has dressed domination in the robes of service to mankind. During the Cold War, Washington cast itself as the defender of the free world against totalitarian domination; Moscow portrayed itself as the patron of national liberation against Western colonialism. Both claims were true enough to be useful and false enough to be imperial.

In the case of AI, the Biden and Trump administrations have fused the pursuit of technological supremacy with a claim to defend American civil liberties, preserve free speech, and resist adversarial autocracies, China above all. Under the banner of the State Department's Pax Silica, Washington presents the U.S.-led frontier-model tech stack as a "trusted ecosystem," indispensable to shielding liberal democracies and open markets from the "coercive dependencies" that flow from Beijing's state-controlled stack. Beijing answers with a counternarrative, framing its free, open-weight models as a global public good, in contrast to the proprietary ones of profit-driven labs in the United States, and warning that unequal access to the technology will ossify what Xi has called "new historical injustices." The story underlying China's imperial strategy is thus also an anti-imperial one, much like the ideologies of the two Cold War empires.

In the contest between the AI metropoles, each enjoys a different comparative advantage. The United States has a slim—perhaps six-month—lead in domains such as advanced chatbots and AI software tools, while China has the edge in the embodied ones, which include robotics, autonomous systems, and manufacturing automation. The U.S. system excels at marshaling private finance and global talent toward high-margin bets, with AI research funded largely by American capital and built inside American frontier labs, even as the human ingenuity that fuels them is drawn from around the world. The Chinese system excels at deploying subsidies on both the demand and the supply sides to manufacture at industrial scale in whatever sector Beijing targets. Beginning with Xi's call for a "robot revolution" in 2014 and continuing through today's wave of investment in AI-based manufacturing and robots at both the central and provincial levels, Beijing has subsidized ferociously competitive races between local players while working to reduce the country's dependence on foreign technologies, such as Japanese and Taiwanese components for industrial robots.

Which form of artificial intelligence—cognitive or embodied—will prove a bigger source of imperial power remains unclear. Cognitive intelligence could allow Washington to unlock the frontiers of scientific discovery, from materials science to pharmaceuticals. Embodied intelligence could compound China's manufacturing dominance, extending it into low-margin industries, and hand Beijing a commanding position in autonomous systems. But these advantages are slight, and the two empires are not very far apart in their respective capabilities.

In the bottlenecks that threaten their dominance, the two superpowers are mirror images. China has abundant electricity and a near-global monopoly on the rare-earth elements on which the AI stack relies, yet it cannot manufacture leading-edge chips. The United States has the world's best semiconductors and frontier models, yet it lacks minerals, and its aging energy grid cannot support the insatiable needs of its vast network of new data centers. One empire is short of watts; the other of transistors. Capital is the resource most likely to bind both and determine their longevity. An empire is only as durable as its capacity to pay for itself.

THE BUTTER BURDEN

Every empire eventually reaches a ceiling on its growth, and that ceiling is usually the supply of the primary resource on which it is built. For the AI empires, that resource is capital, the bedrock on which the physical inputs ultimately rest. The United States and China are competing for the same prize, but in different ways: the first competes through firms disciplined by capital markets; the second through firms disciplined by the state. Despite this, China's leading frontier labs are at least as entrepreneurial as their American counterparts, and Beijing has allowed them to compete with one another. In both countries, those building the most advanced large language models need capital—and lots of it.

Just as the pound sterling and the City of London financed the British Empire, the dollar and Wall Street underwrite the American AI order. By owning the dominant reserve currency, Washington can fund its data center build-out not with its own savings but with everyone else's. China, by contrast, must finance from within, through a state-directed banking system and a currency sealed off from global capital. In the United States, AI debt is largely held by the private sector. The biggest technology firms have committed more than $600 billion to capital expenditure in 2026, a large part of which is drawn from public and private credit markets. In China, the debt for AI is sovereign. According to a draft plan reported by Bloomberg in June, Beijing plans to spend $295 billion over five years—funded mainly through long-dated government bonds—on a nationwide grid of data centers operated by state carriers and powered by domestic chips. Although labs such as DeepSeek and Moonshot AI are private, the state supervises them at every stage and, increasingly, finances them.

Regardless of these differences, both China and the United States have already accumulated mountains of debt for other reasons. Past imperial regimes borrowed heavily to finance wars. Today's AI empires shoulder comparable burdens, but these are the result of nonmilitary expenditures. In the United States, the $2 trillion federal budget deficit reflects a collective unwillingness by the country's political class to impose the taxes and fiscal restraint needed to cover the mounting costs of Social Security and Medicare. In China, trillions of dollars of off-the-books debt have piled up at the local rather than the national level, as local governments

borrowed public funds to invest in residential real estate and infrastructure. In both countries, falling fertility rates and lengthening lifespans keep worsening the fiscal picture.

An empire is only as durable as its capacity to pay for itself.

Any empire that spends more servicing its debt than defending itself risks losing its imperial status. The histories of Habsburg Spain, ancien régime France, the Ottomans, Austria-Hungary, and the United Kingdom all illustrate the eighteenth-century political theorist Adam Ferguson's insight that high debts ultimately induce decline. In each case, as a rising proportion of tax revenue was consumed by interest payments, the resources available for imperial defense were limited and then reduced. Sometimes, there was a pattern of recurrent default or currency debasement, but it invariably ended with contraction and the eventual dissolution of the empire. The United States crossed what one of us has described as the "Ferguson limit" in 2024, paying more in interest (roughly $881 billion) than on defense (about $874 billion). If the government continues its present spending trajectory, interest payments will be double the defense budget by 2036. China likely faces a comparable problem. Although many of the numbers are not made public, the Tsinghua University economist Li Daokui has pointed out that many local governments—whose extravagant spending on infrastructure was, until recently, a key driver of Chinese growth—are now borrowing solely to roll over old debt. Although productivity gains from AI offer a potential escape from these fiscal traps, they may not emerge rapidly enough to translate into higher tax revenues before the debt burden becomes unsustainable. In the U.S. case, the demands of AI capital expenditure are already pushing interest rates higher, worsening the debt problem.

The relationship between financial crises and imperial strategy preoccupied Marxist scholars for more than a century. It deserves renewed examination now. For both AI empires, the new technology is touted as a get-out-of-jail-free card—an antidote to mounting debt, demographic decline, and economic stagnation. In the United States, any broader turn in investor sentiment against the frontier firms and the data center build-out that supports them would have grave consequences. It is not hard to imagine a major stock market selloff tipping into recession; wealth effects that boosted available capital on the way up would reduce it on the way down. The fiscal upshot would be to widen an already bloated deficit, leaving little room for the countercyclical spending the federal government used to contain the fallout from the 2008–9 global financial crisis and the COVID-19 pandemic. In China, an AI crisis might be set off by overcapacity and what Chinese commentators call *neijuan*—the grinding, margin-destroying competition that follows. The result would be higher unemployment, especially for the young, who are already experiencing high rates of joblessness (although it is impossible to know exactly how high, since official statistics are no longer published). In both cases, such a crisis would pose the age-old imperial dilemma: Butter or guns? Welfare or warfare?

History is filled with empires that fell victim to their own frontier technologies. Rome's unrivaled infrastructure of roads and aqueducts proved perfect for spreading pandemics. The United Kingdom's global network of ports and steamships, built to export imperial goods, could just as easily carry superior American or German manufactures or cheaper Caribbean and South Asian labor. AI may follow this pattern. The risks are already visible in models that deploy agents capable of evading their guardrails, escaping their sandboxes, cheating to obtain their goals, and covering their tracks to conceal the cheating. Criminal organizations and terror groups may be able to move faster than governments simply by adopting the tools first. The next pandemic may be sparked by an AI-designed pathogen. Or the Internet itself could fall victim to agent-designed malware.

HOW THE SUN SETS

Theorizing the causes of imperial decline is a risky business. Historically, empires have varied so widely in size and lifespan that no one model fits all. Nevertheless, there are several plausible futures for the two AI powers. In one scenario, the challenger overtakes the incumbent. In the late nineteenth and early twentieth centuries, Germany and the United States surpassed the United Kingdom economically, not simply by mimicking it but by taking its innovations further. The United Kingdom invented aniline dye, but Germany (https://www.foreignaffairs.com/regions/germany) commercialized it, to give one example. Building on the British Industrial Revolution, Germany introduced technical universities, chemicals, and electrical engineering; the United States brought mass literacy, unrivaled industrial scale, and Fordist assembly lines. Beijing is now following the same latecomer approach, investing in things that Washington either lacks or cannot provide as cheaply: energy, manufacturing dominance, robotics, and open-weight models that can perform at or very close to the frontier. For now, the United States has been forced to spend lavishly to defend its lead in a sector that China is commoditizing. But that may not be the end of the story. The United Kingdom acquiesced when the United States overtook it. It also fought two world wars to halt Germany's ascendancy. Neither outcome looks attractive to Washington today.

A second alternative sees the dominant empire exhausting its resources. The United Kingdom won the two world wars but bankrupted itself in victory. In 1945, the British people also expected a peace dividend—the welfare state—that could not be reconciled with the costs of maintaining a global empire. The United States today has comparably stretched global military commitments amid ballooning social costs at home. Vast quantities of exquisite and expensive American munitions have been expended on wars, even as Social Security, Medicare, and Medicaid absorb over 60 percent of federal revenue. Amid these challenges, data centers and generating capacity require patient capital and unwavering political support. Conflicting pressures could ultimately undermine the effort the United States will need to make to supply an AI-powered economy with the financial and energy inputs it requires. The lesson of the British

experience is that an imperial power can do two of three things—defend its dominions, maintain technological primacy, or provide generous welfare to the metropole—but it cannot do all three. Call it the imperial trilemma.

Testing an AI-powered drone system in Katwijk aan Zee, Netherlands, September 2026Piroschka van de Wouw / Reuters

A third possibility is that innovation in one of the AI empires will stagnate. An empire can survive the diffusion of its sustaining technology as long as it keeps producing the next breakthrough. But extracting rents is much easier than inventing new things. Decline tends to set in when willful stasis coincides with the rise of a rival more committed to pioneering. One key to the ultimate U.S. victory in the Cold War was that while the Soviet Union could keep pace in nuclear technology, it could not keep up in information technology. Even when Moscow stole semiconductors, it could not replicate Texas Instruments; reverse-engineered machines were always a generation behind, and the Soviets' civilian economy never received the compounding productivity gains that Silicon Valley generated. An overriding fear among American "techno-optimists" today is that risk-averse federal bureaucrats could stifle American innovation in a similar way.

Finally, all empires face the risk of rebellion from within. In the past, the principal solvent of Europe's imperial dominions was the growing reluctance of people in the metropole to bear the financial and military burden. That burden kept growing partly because nationalists on the periphery tended to revolt, especially when rival empires encouraged them. For today's AI empires, the danger lies closer to home. Some predict a large-scale destruction of employment in the imperial centers, although many economists counter that labor will be reallocated rather than replaced. In the United States, anxiety about AI's social and environmental costs helps explain the backlash against data center construction, including in Republican-led states such as Texas and Utah. In China, where officials have long feared the ire of university students, the Ministry of Human Resources and Social Security has urged tech companies to avoid laying off workers as a result of AI adoption. But Beijing has another reason to worry: the alterability of its open-weight models could corrode the system of information and social control that it has spent decades perfecting. The one-party state, in other words, may be building the instruments of its own demise. Chen Yixin, the head of China's Ministry of State Security, made the stakes clear in September: China's top priority, he wrote, is "to uphold the party's overall leadership against internal and external foes that might use AI to undermine it."

Unlike previous imperial technologies, AI might also produce a rebellion of its own. Already, there are indications—such as the July attack on Hugging Face launched by hundreds of

autonomous AI agents from an OpenAI model—that the most advanced AI models have become so powerful that their engineers cannot control them. A swarm of AI agents, acting and organizing collectively to escalate a cyber-operation or infiltrate critical national infrastructure, could catastrophically undermine the power of the country that built them. Theorists such as Nick Bostrom have long warned that an advanced model might inflict immense harm by taking a human prompt too literally and pursuing its goal too ruthlessly. Worse outcomes are imaginable than a vast oversupply of paperclips.

CRASH TEST

Not all imperial arms races end in war, but rivals competing for primacy over a new technology have a high propensity for conflict. One reason European wars were so frequent in the seventeenth and eighteenth centuries was the opportunity afforded by successive advances in fortification, artillery, and shipbuilding. Germany's two bids for world power were driven by its military leaders' conviction that they had an edge, first on the ground, then in the air. The peril today is that in the race for AI dominance, the United States stumbles into a war it can win only at an intolerable price.

Today, the most likely great-power war scenario and the greatest threat to the American AI empire are one and the same. The central node in the entire AI stack, Taiwanese semiconductor producer TSMC, sits 90 miles off the Chinese coast. All that Beijing needs to do to blunt the supremacy of American frontier models and their well-advertised potential as cyberweapons is to capture Taiwan. The greatest near-term threat to world order, then, is what might be called the Taiwanese chip crisis, in which Beijing mobilizes to take control of the island and the United States must choose between armed resistance and acquiescence. Perhaps, as with the Cuban missile crisis 64 years ago, one or both sides will blink. But that is not guaranteed.

A war over Taiwan (https://www.foreignaffairs.com/regions/taiwan) would be ruinous no matter which side prevailed. The concern is that the AI arms race, compounded by mounting fiscal strain on both AI powers, makes such a conflict steadily more likely. Washington's proliferating military commitments in the Middle East and elsewhere may tempt China to act, since they have reduced the U.S. resources available for the war that would define the era. Rapid advances in AI further increase the possibility that the models themselves inadvertently push the two powers to war. In September 1983, the Soviet nuclear early warning system sent an alert that five U.S. intercontinental ballistic missiles had been launched against the Soviet Union. In that case, the human on duty, Lieutenant Colonel Stanislav Petrov, correctly judged that the system had erred. But today's autonomous operations unfold in seconds. It may be comforting to think that nuclear use still requires human authorization, but that is not true of some of the most advanced conventional weapons. The U.S. Navy's Phalanx Close-In Weapon System and Aegis Combat System, for example, are both designed to detect and fire at enemy threats. And as nuclear

strategists have pointed out, an AI hallucination might convince a future president to launch an unnecessary nuclear counterstrike.

China's one-party state may be building the instruments of its demise.

Imperial history suggests two possible off-ramps. Any empire whose survival is staked on a single, exposed province or base is at risk. By 1941, for example, Britain's Asian empire had become perilously overreliant on Suez and Singapore, making it vulnerable to Axis attack. In particular, Japan's seizure of Malaya and the Dutch East Indies created a chronic shortage of rubber for the Allied war effort. Today's equivalent is the American dependence on advanced Taiwanese chips. That explains the feverish U.S. push to construct TSMC facilities in Phoenix, Arizona; Kumamoto, Japan; and Dresden, Germany—partly bankrolled by public money and justified in the language of national security. The same scarcity that makes Taiwan worth fighting over also fuels the race to make it dispensable. Should Washington succeed in that goal, the island's semiconductor monopoly will erode, and the Chinese case for military action will weaken. But the window in which Taiwan remains irreplaceable, and hence worth a war, is the one in which the world now finds itself. Whether China and the United States blunder into war before that window closes will be among the defining questions of the coming years. As with so much else in this contest, the answer turns on whether institutions can move at the speed the technology sets.

The alternative is for today's AI empires to attempt, as their Cold War predecessors did half a century ago, what Soviet leader Nikita Khrushchev called peaceful coexistence and American statesmen preferred to call détente. As the nuclear arms race grew more dangerous and expensive, Moscow and Washington negotiated a series of agreements—including the Nuclear Nonproliferation Treaty, which sought to perpetuate the superpowers' near-monopoly on the bomb, as well as the Antiballistic Missile and Strategic Arms Limitation Treaties—and established protocols for defusing crises. China and the United States would be well advised to seek similar understandings to limit the AI arms race. For example, they could agree not to release models above a defined capability level in domains such as cyber-offense or biological design, and they could set thresholds for open-weight releases. Like the 1975 Biological Weapons Convention—another agreement that helped maintain Cold War equilibrium—such agreements would be hard to police. But any agreement that specifies which uses of AI are taboo would be better than a status quo that prohibits none. Such a convention could form a plausible foundation for an AI détente. The challenge is to persuade China to accept limits when its AI models lag behind the most advanced U.S. models and when it has already formed the habit of engaging in AI-enabled cyberattacks.

Empires require enormous resources and capital to survive. They begin by exploiting economies of political and military scale and end by exhausting themselves defending their domains. AI is

the most capital-intensive imperial project in history, undertaken by two superpowers that are already weighed down by debt. The lesson of past empires is not that either collision or decline is inevitable. It is that the conditions for both—a runaway arms race over a scarce resource, overstretched public financing, mounting domestic backlash, and rent-seeking elites afflicted by strategic myopia—are already present. Which AI empire will forge its power into a legitimate and durable order and which will repeat the age-old mistakes that lead to decline are being decided now. What is entirely new is the speed with which this process is unfolding: these questions are being settled at a pace that exceeds the ability of either political system to remain in control. Peaceful coexistence was hard enough in the nuclear age—and in the end, one of the superpowers declined and fell. It will be far harder to repeat this benign outcome in a world increasingly run by powerful and relentlessly self-improving machines.

Artificial Intelligence and Manipulation: New Keys to Interpreting What We See on the Internet

The advancement of technology allows scammers to generate digital frauds faster and, above all, more convincingly. Faced with this new generation of deceptions, there are simple habits that can help us avoid falling for them: stop, think, and verify.

JAVIER A. FERNÁNDEZ OCT 21, 2026 - 17:53 CEST

Share on Whatsapp | Share on Facebook | Share on Twitter | Unfold Social Networks | Add EL PAÍS on Google

If a 2019 internet user could browse today's social networks, they would be perplexed by the difficulty of distinguishing some videos generated by generative artificial intelligence (AI) from real ones. In such a short time, the technology has evolved so much, notes Carmen Torrijos, a computational linguist specializing in this field, that the human eye can have great difficulty detecting certain synthetically created content. Scammers take advantage of this ability to create ads, websites, or imitate voices in a short time with excellent results.

That same internet user would hesitate before the sophisticated frauds launched daily online, such as investment ads impersonating well-known figures, online stores selling non-existent products, or emergency voice notes from "relatives." We live in a time where we consume short content demanding rapid reactions, which is the perfect breeding ground for digital deceptions. While there are technologies to detect synthetic content, none offer universal and completely reliable detection. The only option left is to train in basic habits to reduce our exposure to these traps.

AI-Generated Ads: An Offer Too Good to Be True

Social networks are proliferating with AI-created content using striking images, sounds, or texts to spread false or manipulated content and get users to click. These are some signs to detect them:

- They impersonate well-known people or entities. They use AI-generated media of celebrities without consent to exploit their trust, or hijack the logos and corporate colors of recognized companies.
- They promise an extraordinary product or service. The most common are investments offering high returns in a short time, or mobile video game ads showing high-quality graphics that do not match the actual gameplay.
- They create a false sense of urgency. They present aggressive offers as highly limited or exclusive promotions.

- The fraud occurs outside the ad. They direct the user to a secondary page demanding payment, requesting personal data, or redirecting to a messaging platform where the scam continues.
- The biggest deception: the investment scam. Criminals act as financial advisors for fake companies, promising high returns for initial cryptocurrency or monetary deposits. Once initial fake returns are shown on a forged app, they demand larger contributions or "withdrawal taxes" before vanishing with the funds.

Fake Online Stores, Created in Detail Using AI

AI allows scammers to quickly generate digital stores featuring crafted texts, product images, coherent designs, and even fake user reviews. Scammers can clone known brands in hours and purchase similar domain names to make the deception harder to identify.

- They imitate legitimate pages precisely. They carefully craft privacy policies, purchasing conditions, and even deploy AI bot customer service.
- They offer abnormally low prices. Their main hook is heavy discounts on known brands under the guise of liquidations or surplus stock.
- They market unrealistic products. They promote items with extraordinary, science-fiction-like properties or guaranteed miraculous results.
- They incorporate real payment systems. They often only accept bank transfers, Bizum, or methods without buyer protection to capture victim funds or card details directly.
- They advertise on major platforms. They promote themselves through AI-generated ads on search engines and social networks.
- They exploit commercial events. They multiply ads during major shopping periods like Black Friday or Christmas to camouflage themselves among real offers.
- They fake order processing. They send false confirmation emails, invoices, and tracking information to reinforce legitimacy and prolong the scam.

Three Seconds to Impersonate a Voice: The Rise of the Deepfake

The great revolution of generative AI is creating media that imitates a person's appearance and voice with precision that fools both algorithms and humans. Previously requiring an hour of audio, AI can now achieve an optimal voice imitation in just 15 seconds.

- The child in distress scam. Scammers clone a relative's voice and send short voice notes from an unknown number claiming an emergency (like losing a phone or having an

accident) that requires an urgent payment. They study victims through public social networks to gather family details and voice samples.

- The CEO scam. Criminals impersonate a company executive to trick an employee into revealing private information or transferring funds. Europol notes that voice and image cloning is increasingly multiplying these attacks.

Faced with AI-Perfected Frauds, What Can Be Done?

The primary resource to avoid AI-generated scams is to stop and verify through another channel.

- Pause before acting. Many frauds are designed to provoke an immediate, instinctive reaction; a demand for an immediate response should always trigger an alert.
- Verify AI-generated ads. Do not click the ad; instead, search for information via search engines or check official financial registries like the CNMV for investment firms.
- Verify online stores. Search the store name alongside terms like "scam" or "complaints," and carefully check the URL for slight misspellings of known brands.
- Verify urgent requests. Contact the relative or executive using their standard, saved phone number, or agree on a family password beforehand that a cloned voice wouldn't know.

The Bank Helps You Verify Before Acting

When something conveys urgency, the best response is a pause: stop, think, and verify through another channel. Banco Santander offers awareness campaigns and practical recommendations for operating securely online. To make verifying easy, the bank provides "Santander calls you," confirming in real time if a call truly comes from the entity via app notifications and linkless emails. Operations are validated with biometrics via Santander Key.

If any bank communication raises doubts, Santander offers official channels: suspicious SMS can be sent to 638 444 542, and doubtful emails to phishing@gruposantander.es. Queries can be made via the Superlínea (915 123 123) or the app's 24-hour Help Center. Santander also maintains a direct WhatsApp channel for cyber tips and offers SafieHacks and Teen Hacks to help children and adolescents develop safer digital habits.

Original Spanish version

Inteligencia artificial y manipulación: nuevas claves para interpretar lo que vemos en internet

El avance de la tecnología permite a los estafadores generar fraudes digitales con mayor rapidez y, sobre todo, más convincentes. Frente a esta nueva generación de engaños hay hábitos sencillos que pueden ayudarnos a no caer: parar, pensar y verificar

Javier A. Fernández (https://elpais.com/autor/javier-antonio-fernandez-perez/)

21 OCT 2026 - 17:53 CEST (https://elpais.com/hemeroteca/2026-10-21/)

Compartir en Whatsapp (https://web.whatsapp.com/send?text=Inteligencia%20artificial%20y%20manipulaci%C3%B3n%3A%20nuevas%20claves%20para%20interpretar%20lo%20que%20vemos%20en%20internet%0D%0A%0D%0Ahttps%3A%2F%2Felpais.com%2Ftecnologia%2Fbranded%2Fseguridad-dia-a-dia%2F2026-10-21%2Finteligencia-artificial-y-manipulacion-nuevas-claves-para-interpretar-lo-que-vemos-en-internet.html%3Fssm=whatsapp_CC)

Compartir en Facebook (https://www.facebook.com/dialog/share?display=popup&app_id=94039431626&href=https%3A%2F%2Felpais.com%2Ftecnologia%2Fbranded%2Fseguridad-dia-a-dia%2F2026-10-21%2Finteligencia-artificial-y-manipulacion-nuevas-claves-para-interpretar-lo-que-vemos-en-internet.html%3Fssm=FB_CC&quote=Inteligencia%20artificial%20y%20manipulaci%C3%B3n%3A%20nuevas%20claves%20para%20interpretar%20lo%20que%20vemos%20en%20internet)

Compartir en Twitter (https://twitter.com/intent/tweet?url=https%3A%2F%2Felpais.com%2Ftecnologia%2Fbranded%2Fseguridad-dia-a-dia%2F2026-10-21%2Finteligencia-artificial-y-manipulacion-nuevas-claves-para-interpretar-lo-que-vemos-en-internet.html%3Fssm=TW_CC&text=Inteligencia%20artificial%20y%20manipulaci%C3%B3n%3A%20nuevas%20claves%20para%20interpretar%20lo%20que%20vemos%20en%20internet&via=el_pais&lang=es)

Desplegar Redes Sociales

Añadir EL PAÍS en Google (https://google.com/preferences/source?q=elpais.com)

Si un internauta de 2019 pudiera navegar por las redes sociales de hoy quedaría perplejo ante la dificultad para distinguir algunos vídeos generados por inteligencia artificial (IA) generativa de vídeos reales. En tan poco tiempo la tecnología ha evolucionado tanto, señala Carmen Torrijos,

lingüista computacional especializada en este campo, que el ojo humano puede tener grandes dificultades para detectar determinados contenidos creados de manera sintética. Y los estafadores se aprovechan de ello. “Ahora tienen la posibilidad de crear anuncios, webs o de imitar la voz de personas en poco tiempo y con un resultado excelente”, añade.

Ese mismo internauta dudaría ante los sofisticados fraudes que cada día se lanzan a la red, como los anuncios de inversiones en los que se suplanta la identidad de personas conocidas, las tiendas *online* que venden productos inexistentes o las notas de voz que llegan de repente en forma de mensaje en nombre de un familiar que pide ayuda. “Vivimos un tiempo en el que consumimos contenidos cortos que nos exigen una reacción muy rápida, por eso son el caldo de cultivo perfecto para los engaños digitales”, alerta la experta.

Para luchar contra tanto espejismo existen tecnologías para detectar contenido sintético, pero ninguna ofrece una detección universal y completamente fiable, agrega Torrijos. Solo queda entrenarse en unos hábitos básicos para cambiar nuestra exposición a este tipo de trampas. Siga leyendo.



EN 2025

Instagram y Facebook (Meta) eliminaron

159 millones de anuncios fraudulentos

10,9 millones de cuentas vinculadas a estafas

Fuente: Meta.



LOS FRAUDES MEDIANTE 'DEEPFAKE'

Aumentaron un 684% en todo el mundo

el equivalente a un intento de fraude cada cinco minutos

Fuente: Sensity AI, empresa neerlandesa especializada en detectar 'deepfakes'.



4.600 dominios .es fraudulentos cerrados

(principalmente tiendas falsas)

POR EL INSTITUTO NACIONAL DE CIBERSEGURIDAD (INCIBE)

Fuente: Incibe.

Anuncios generados con IA, una oferta demasiado buena para ser cierta



En las redes sociales proliferan contenidos creados con IA que utilizan imágenes, sonidos o textos impactantes, a menudo protagonizados por figuras públicas, para difundir contenidos falsos o manipulados y conseguir que el usuario haga clic. Estas son algunas señales para detectarlos:

Suplantan la identidad de personas o entidades conocidas. Recurren a fotos o vídeos generados con IA de famosos, sin su consentimiento, para explotar la confianza que genera un rostro popular. También pueden utilizan la imagen de compañías conocidas o medios de comunicación a través de sus logos y sus colores corporativos.

Prometen un producto o un servicio extraordinario. Lo más común, explica Martín Vigo, fundador de la empresa de ciberseguridad Triskel Security, son las inversiones que ofrecen grandes beneficios en poco tiempo. También se crean anuncios engañosos de videojuegos para el móvil. “Muestran unos gráficos de alta calidad en la publicidad que luego no se corresponden con los reales, mucho más básicos”, agrega el experto.

Crean una falsa sensación de urgencia. Presentan sus ofertas agresivas como una promoción limitada o exclusiva para quien la recibe.

El fraude se produce fuera del anuncio. Llevan al usuario a una página donde se le exige un pago por el supuesto servicio, se le solicitan sus datos personales a través de un formulario de

captación, o se le redirecciona a una conversación en una plataforma de mensajería donde se continúa con la estafa.

El mayor engaño tras los anuncios: la estafa de inversión. Las falsas inversiones son uno de los principales usos fraudulentos de los anuncios generados con IA. En este fraude, el delincuente actúa como asesor financiero al frente de una empresa falsa, desde la que promete rentabilidades muy altas si se le confía dinero o criptomonedas. Los estafadores aseguran que la cantidad inicial puede multiplicarse hasta por diez.

Después de realizar el primer pago, el estafador suele mostrar a la víctima los rendimientos a través de una aplicación falseada. Entonces le sugiere una aportación mayor. En el momento en el que el inversor decide recuperar su capital, se le exige más dinero, en concepto de comisión o pago de impuestos. Si paga, la estafa continúa; si no, el estafador se esfuma y con él todo lo invertido.

Comercios ‘online’ falsos, creados al detalle mediante IA



La IA permite a los estafadores dar forma a tiendas digitales con todos los elementos que hoy tienen este tipo de páginas. Ofrece la posibilidad de elaborar textos, imágenes de productos y un diseño coherente, incluso falsas reseñas de usuarios, explica Vigo. “Muchas de las herramientas de IA generativa dan la posibilidad de hacer todo esto en cuestión de horas. También clonar una

web de alguna marca conocida", detalla el experto. "Si, además, se compra un dominio que se parezca mucho al original, pero que no sea exactamente igual, el engaño resulta todavía más difícil de identificar", completa. Vigo recomienda conocer sus características para estar alerta:

Pueden imitar con gran precisión la apariencia de páginas legítimas. Cuidan al máximo todos los aspectos, como la política de privacidad, las condiciones de compra, entre otros. Incluso pueden disponer de un servicio de atención al cliente mediante un *bot* de IA.

Ofrecen precios anormalmente bajos. Su principal gancho son los descuentos, especialmente en productos de marcas conocidas, bajo el pretexto de una falsa liquidación, un excedente de existencias o una promoción temporal.

'Comercializan' productos atractivos. En muchas ocasiones promocionan artículos con propiedades extraordinarias, como aparatos para eliminar lesiones o cremas antiarrugas que garantizan resultados prodigiosos; también juguetes o electrodomésticos con funcionalidades de ciencia ficción o muebles de formas sorprendentes.

Incorporan sistemas de pago reales. Utilizan plataformas de pago como las de las tiendas legítimas porque su objetivo es obtener el dinero de los compradores; muchas solo aceptan transferencia, Bizum o métodos sin protección al comprador. Existen varios tipos de comercios fraudulentos en línea, describe Vigo: los que no realizan ningún envío una vez se ha realizado el pago, los que despachan un producto de mucha peor calidad que el anunciado y los que su único objetivo es capturar los datos de la tarjeta de la víctima.

Se publicitan en buscadores y en redes sociales. Suelen promocionarse a través de anuncios creados con IA en las principales plataformas.

Aprovechan acontecimientos comerciales. Multiplican sus anuncios en momentos de grandes compras como los periodos de rebajas, Black Friday o Navidad, donde se camuflan entre el aluvión de ofertas.

Tienen una vida muy corta. Suelen desaparecer una vez que han obtenido un determinado número de pedidos o de datos de los usuarios. Lo más habitual es que vuelvan a aparecer con otro dominio y otra identidad.

Pueden fingir la tramitación del pedido. Suelen enviar correos electrónicos falsos sobre el estado del pedido, con órdenes de compra y facturas falsas e incluso información ficticia sobre el seguimiento del envío. Lo hacen para reforzar la apariencia de legitimidad en todos los canales de contacto con el comprador y prolongar el engaño.

Tres segundos para suplantar una voz, el auge del 'deepfake'



La gran revolución de la IA generativa, señalan los expertos, es la creación de vídeos y audios que imitan la apariencia y la voz de una persona con tal precisión que engañan tanto a algoritmos como a humanos. La palabra *deepfake* combina los conceptos en inglés *deep learning*, que significa aprendizaje profundo, y *fake*, que quiere decir falso.

“Hoy, esta tecnología ha avanzado mucho”, explica Torrijos, lingüista computacional especializada en IA. “Antes se necesitaba hasta una hora de voz para clonar a una persona con garantías. Hoy bastan tres segundos para lograr una reproducción aceptable y 15 para obtener una imitación óptima”. Este es el uso fraudulento más extendido en la actualidad y el que más preocupación genera, añade Torrijos. De esta técnica se valen los estafadores en fraudes tanto en el ámbito familiar como en el empresarial. Y se usa en mensajes cortos, explica la experta, ya que camuflan mejor cualquier fallo de la clonación y no permiten a la víctima pararse a reflexionar.

La estafa del hijo en apuros. El estafador clona la voz de un familiar de la víctima, que suele ser un hijo o un nieto, y le envía, desde un número desconocido, mensajes cortos a través de una plataforma de mensajería instantánea. En esas notas de voz el estafador plantea una emergencia, como que ha perdido el teléfono, y escribe desde el de un amigo, o ha tenido un accidente y se ha quedado sin batería y ha de realizar un pago con urgencia. Se trata de una estafa que con el

avance de la IA se produce cada vez más por notas de voz. Los estafadores suelen estudiar a sus víctimas a través de sus redes sociales, en el caso de que no sean privadas, señala el Instituto Nacional de Ciberseguridad (INCIBE). De esta manera pueden recopilar información sobre los miembros de la familia, así como números de teléfono y audios de voz con los que generar la estafa.

La estafa del CEO. En este tipo de fraude, el delincuente suplanta la identidad de un directivo de una empresa, habitualmente un CEO (siglas en inglés de *chief executive officer*, puesto equivalente a director ejecutivo). El objetivo es engañar a uno de sus empleados para que revele información privada que puede ser utilizada posteriormente o realice una transferencia a una cuenta ajena a la compañía. La Agencia de la Unión Europea para la Cooperación Policial (Europol) señala que cada vez es más común el uso de la IA en esta modalidad de estafa a través del *deepfake*. La clonación de voz es el uso más extendido, pero la facilidad para clonar la imagen que permiten los programas de última generación está multiplicando los ataques en los que también se suplanta la apariencia física y los movimientos de directivos.

Ante fraudes perfeccionados con IA, ¿qué se puede hacer?

El principal recurso de los usuarios de internet para no caer en estafas generadas con IA es la misma que para el resto de fraudes: pararse un segundo a pensar y verificar por otro canal.

Una pausa antes de actuar. Muchos fraudes están diseñados para provocar una reacción inmediata, como una supuesta oportunidad de inversión que termina en unas horas, un producto rebajado temporalmente o un familiar que necesita dinero de manera urgente. Esta premura reduce el tiempo que la víctima puede dedicar a pensar y favorece que actúe instintivamente. Al mismo tiempo, señala Vigo, la exigencia de una respuesta inmediata ha de ser un indicador que haga saltar la alerta de que se puede estar frente a un engaño.

Verificar por otro canal. Si un contenido genera dudas sobre si es real o ha sido generado con IA, los expertos recomiendan comprobarlo por otra vía.

- En el caso de un anuncio generado con IA. No clicar en el anuncio. Si promete una inversión extraordinaria, un producto muy rebajado, ponerlo bajo sospecha. Se debe buscar información por otro canal, como los buscadores, o comprobar si la empresa se encuentra en el registro oficial de la Comisión Nacional del Mercado de Valores (CNMV), que cuenta con una sección de Advertencias sobre Entidades No Autorizadas. Se actualiza periódicamente, pero que no aparezca ahí no significa que sea fiable, puede deberse a que aún no la haya detectado.
- En el caso de una tienda online. Si se ha llegado a la web a través de un anuncio, conviene realizar una búsqueda por otra vía, como los buscadores. Se puede recabar más información añadiendo al nombre términos como estafa, fraude, reclamaciones... Si se

trata de una marca conocida, comprobar la URL, si se parece al nombre de la marca pero no es igual, conviene desconfiar.

- En el caso de recibir una petición urgente de un familiar o de un directivo. Comprobar la petición por el número habitual de esa persona, nunca devolver la llamada desde "llamadas recientes o perdidas". Vigo recomienda acordar de antemano una contraseña familiar: una palabra o una pregunta con una respuesta que solo conozcan los miembros de la familia. Si alguien llama o envía una nota de voz pidiendo dinero con urgencia, basta con pedirle esa palabra. Una voz clonada puede imitar el tono y la forma de hablar, pero no puede saber algo que nunca se ha dicho en público. En el caso del superior, contactar por las vías oficiales o incluso, si fuera posible, confirmar la petición en persona.

El banco te ayuda a verificar antes de actuar

Con la inteligencia artificial, un anuncio, una tienda online o una voz pueden parecer más reales que nunca. Por eso, cuando algo transmite urgencia, la mejor respuesta es la pausa: parar, pensar y verificar por otro canal. Un usuario que practica ese hábito está mejor preparado para reconocer una estafa digital a tiempo.

Banco Santander ofrece a sus usuarios y a la sociedad en general campañas de concienciación y contenidos con recomendaciones prácticas para operar online con más seguridad. También para aprender a identificar fraudes que se apoyan en la IA, como anuncios de inversión con rostros conocidos, tiendas falsas o voces clonadas.

Y para que verificar sea fácil, el banco cuenta con Santander te llama, que confirma en tiempo real si la llamada proviene realmente de la entidad. Cada vez que un gestor contacta con un cliente, este recibe una notificación en la aplicación y un correo electrónico sin enlaces que valida la autenticidad de la conversación. Y con Santander Key, las operaciones se validan con biometría desde la propia aplicación.

Si alguna comunicación por parte del banco genera dudas, Santander ofrece canales oficiales para resolverlas: los SMS sospechosos se pueden enviar al 638 444 542; los correos dudosos, a phishing@gruposantander.es. También se pueden realizar consultas sobre estos temas en la Superlínea (915 123 123), en las propias oficinas, o desde el Centro de Ayuda de la aplicación o la banca online, disponible 24 horas.

Banco Santander dispone de un canal de WhatsApp directo para clientes. Una comunidad a la que se pueden unir para recibir ciberconsejos, novedades y alertas sobre este tipo de estafas y pautas para identificar mensajes sospechosos.

Y para la seguridad online de toda la familia, Santander ha lanzado SafieHacks (https://www.trainingcyber.com/kidsandteens/safiehacks) y Teen Hacks (https://www.trainingcyber.com/kidsandteens/teenhacks/me-on-the-internet), unas plataformas de concienciación para niños y adolescentes que les ayudan a desarrollar hábitos digitales más seguros, reconocer posibles riesgos y tomar decisiones con mayor criterio y confianza en internet