2026 Quarterly Brief – Rehn/Meidner model under pressure

Q1 2026: Intelligence as Infrastructure, a story about pivoting

Jonas Kjellstrand
www.kjellstrand.se
April 2026

This brief is prepared for senior decision-makers across government, industry, and investment. It reflects the author's independent analysis and is not written on behalf of any single client or institution. Distribution is limited to the author's advisory network and www.kjellstrand.se.

Table of Contents

Executive Summary

For two centuries, the wealth of nations was measured in steel, energy throughput, and the productivity of organized labour. That era is not winding down gradually. It is being replaced — rapidly, measurably, and with consequences that most leaders have not yet absorbed.

This brief points out that AI has crossed a decisive threshold. It is no longer an emerging technology arriving inside the existing economy. It is becoming the operating logic of the next one — a new

industry whose output is machine intelligence, whose capital base is compute, energy, and data, and whose downstream effects are already restructuring labour markets, industrial geography, and the architecture of national competitiveness. The evidence now supports this claim with a weight that scenario planning and speculative forecasting never could.

Three developments converge to make the first half of 2026 a pivotal moment.

First, the physical infrastructure of machine intelligence is being built at wartime pace. Hyperscaler capital expenditure on AI infrastructure exceeded $200 billion in 2025, and the economic centre of gravity is shifting from training — building capability — to inference — monetising it. This distinction reshapes the competitive map. The long-run winners will not necessarily be those who train the largest models, but those who can serve intelligence most cheaply, most reliably, and closest to demand. That is fundamentally an energy story, a geography story, and a chip architecture story — and it plays directly to advantages that Northern Europe, and Sweden in particular, hold in unusual combination.

Second, Europe is converting policy ambition into concrete industrial infrastructure — and concrete regulatory enforcement. The EU AI Act reaches its most consequential compliance date on 2 August 2026, with penalties exceeding even GDPR's ceiling. Simultaneously, the EuroHPC network has selected nineteen AI Factory consortia, and the InvestAI facility is mobilising €20 billion for up to five AI Gigafactories. The risk, documented in this brief, is structural: Europe may build prestigious training assets while systematically underbuilding the commercial inference capacity on which a broad-based AI economy actually depends. The gap between EU and US private AI investment — 12% versus 66% of global venture capital — tells the story plainly.

Third, the labour market evidence has hardened from hypothesis to measurement. Stanford Digital Economy Lab research using high-frequency ADP payroll data now documents a 13–16% relative decline in employment for workers aged 22–25 in AI-exposed occupations, and roughly 20% for early-career software developers since late 2022. These are not projections. They are observed distortions, falling disproportionately on the youngest and most educated cohort of workers, in precisely the occupations that used to absorb them. The canary is no longer singing. It is silent.

Against this backdrop, the old industrial paradigm is reaching exhaustion in parallel rather than in sequence. Euro-area industrial production hit its lowest seasonally adjusted level in over a year in January 2026. European machinery output is down 3% since 2019, chemicals and automotive down 14%, basic metals down 12%. BASF is exploring a US listing for a major division. ThyssenKrupp is eliminating 40% of its workforce. These are not cyclical fluctuations. They are symptoms of a developmental model that has been optimised to its limits.

Sweden occupies an unusual position within this transition. Abundant low-carbon electricity, a 97% data centre tax rebate, natural cooling, the highest unicorn density per capita outside Silicon Valley, and a cultural DNA forged by late but ferocious industrialisation — pragmatic adaptation, comfort with creative destruction, export orientation from birth — give it both the input stack and the entrepreneurial reflexes for an intelligence economy. Major commitments from EdgeConneX, Mistral AI, EcoDataCenter, and US based Hyperscalers confirm that global capital has noticed. The

open question, explored in the central sections of this brief, is whether Sweden becomes an architect of the AI ecosystem or merely a well-run host for foreign-owned assets-slippery sloping Sweden into an energy prone third world nation.

This brief is written for leaders who sense that the ground is shifting beneath familiar institutions but want the structural analysis to act on it — not the hype, and not the reassurance. The argument proceeds from infrastructure to regulation, from regulation to industrial geography, from geography to labour markets, from labour markets to the exhaustion of the old paradigm, and from there to what Swedish dynamism looks like when it pivots from exporting commodities to exporting intelligence.

The window, for Sweden, is not permanently open. It is in fact closing. What follows is the case for understanding why — and acting accordingly.

The new [AI] industry is no longer arriving-It is here in full force.

The most consequential mistake a leader can make in 2026 is to treat artificial intelligence as an important technology arriving inside the existing economy. That framing was defensible two years ago. It is not defensible now. AI is not arriving inside the existing economy. It is replacing the operating logic of the existing economy. The distinction is not semantic. It determines whether organisations adapt or are rendered structurally obsolete.

For two centuries, industrial civilisation was built around the scaling of energy, machines, logistics, and standardised labour. The wealth of nations depended on how efficiently societies could organise physical production and administrative coordination. AI changes the equation because it scales something more foundational: cognition itself. Once intelligence becomes reproducible, distributable, and continuously improving, every sector — from education to law, medicine to manufacturing, public administration to scientific research — enters a fundamentally different developmental path.

What is emerging is not a new application layer. It is a new industry whose output is machine intelligence, whose capital base is compute, chips, energy, data, and models, and whose downstream effects will restructure every other industry. What electrification did to manufacturing and the internet did to communication, AI is doing to intelligence itself. The difference is that this transformation is moving faster, scaling harder, and concentrating value more aggressively than either predecessor-it is now in full force.

Current state of AI: the infrastructure phase

AI has crossed from experimentation into new industrial infrastructure. The clearest evidence is not in any single model release but in the scale of capital now committed to the physical substrate of intelligence: chips, data centres, energy contracts, and networking. Global capital expenditure on AI

infrastructure exceeded $200 billion in 2025 across the major hyperscalers alone, and 2026 commitments are running higher still. This is not speculative investment. It is the construction of a new utility layer-intelligence as a service.

The model layer itself has matured in ways that are easy to underestimate from outside. Frontier models now reason, plan, use tools, write and execute code, and operate with multimodal fluency across text, image, audio, and video. Agentic AI — systems that can take sequences of actions toward goals with limited human supervision — is no longer a research concept. It is shipping in products. The implication is that the next wave of economic impact will come not from smarter chatbots, building your own large language model, but from software agents that can perform entire workflows: research, analysis, drafting, scheduling, procurement, compliance, customer service, coding. The shift from 'tool' to 'agent' is the shift from augmenting human effort to partially replacing it.

A second structural development is the emergence of inference as the dominant economic workload. The public discourse still focuses on training: who has the biggest model, who has the most GPUs for pre-training runs-this part of AI is already in a way over. But the economic value of AI increasingly sits in inference — serving models at scale, cheaply, reliably, and close to demand. Training builds capability. Inference monetises it. This distinction matters enormously for infrastructure strategy, because it means the long-run winners in AI compute will not necessarily be those who train the largest models, but those who can serve intelligence most efficiently and at the lowest marginal cost. That is an energy story, a chip architecture story, and a geography story — and it plays directly to Nordic and Swedish structural advantages-it is easier sending bits thru fiber than subsidising building adjacent countries electricity infrastructures.

AI policy becomes industrial policy: the EU AI Act at enforcement

On 2 August 2026, the EU AI Act reaches its most consequential enforcement date. The full framework for high-risk AI systems listed in Annex III — covering biometrics, critical infrastructure, education, employment, law enforcement, migration, and democratic processes — becomes binding. Transparency obligations under Article 50 take effect simultaneously, requiring disclosure of AI interactions, labelling of synthetic content, and deepfake identification. The Commission's enforcement powers over general-purpose AI model providers also activate on the same date, following a one-year adjustment period since obligations first applied in August 2025.

What makes this moment structurally significant is not the rules themselves but their combined effect. Compliance, documentation, transparency, conformity assessment, human oversight, risk management, post-market monitoring, and incident reporting are no longer legal peripherals. They are becoming core components of the industrial process of building and deploying AI in Europe. Organisations that have not completed conformity assessments, finalised technical documentation, affixed CE marking, and registered in the EU database by August 2026 face penalties of up to €35 million or 7% of global annual turnover — exceeding even the GDPR's penalty ceiling.

The strategic implication is that AI policy in Europe is no longer merely about ethics or permissible conduct. It is an operating system for market access. This creates a dual dynamic. For well-prepared European firms and sovereigns, the AI Act becomes a competitive moat: a compliance infrastructure that raises barriers to entry for less regulated competitors and creates premium market positioning for trustworthy AI. For those unprepared, it becomes an execution risk of the first order. An applied AI study of 106 enterprise AI systems found that 40% had unclear risk classifications as recently as early 2026. The gap between regulatory ambition and operational readiness is real and widening.

The contrast with the United States is sharpening. The US continues to lack a comprehensive federal AI framework, relying instead on a fragmented patchwork of state laws, executive orders, and sectoral enforcement. The EU's approach is more systematic and ex ante; the American model is more permissive, more innovation-first, and more willing to let tort risk and market dynamics shape outcomes. Neither approach is obviously superior. But they are producing different industrial incentive structures, and those incentive structures will increasingly determine where AI companies build, where they deploy, and where value accrues.

EU AI Factories and Gigafactories: from ambition to infrastructure

Europe's infrastructure answer to the AI race is now materialising physically. The European Commission identified AI Factories as a strategic priority in 2024, and the network has expanded rapidly since. By late 2025, the EuroHPC Joint Undertaking had selected nineteen AI Factory consortia across more than twenty member states, with at least nine new AI-optimised supercomputers being procured and deployed, more than tripling existing EuroHPC AI computing capacity. Sweden is among the countries selected in the first wave, placing it directly inside the physical substrate of Europe's AI buildout-remember sending bits beats building electricity wires.

The next step is the move from factories to gigafactories. In January 2026, the EU Council adopted the amended EuroHPC regulation enabling the creation of up to five AI Gigafactories — large-scale facilities each integrating over 100,000 (a number sized by US microchip producers) advanced AI chips, energy-efficient data centres, and AI-driven automation. The InvestAI initiative mobilises a €20 billion facility to support these gigafactories, backed by a memorandum of understanding between the European Commission and the European Investment Bank. The Commission received 76 expressions of interest across 16 member states and 60 proposed sites. Cities from Vienna to Prague to Mora la Nova in Spain have formally submitted bids, alongside consortia led by Deutsche Telekom and Brookfield, Scaleway's AION consortium, and national public-private partnerships from Portugal to the Netherlands.

The ambition is real, but so are the structural risks. A recent policy brief from Interface analysed the first thirteen AI Factories and found that partnership consortia are primarily composed of research institutions rather than commercial actors — a significant mismatch given that global AI-specific compute infrastructure is overwhelmingly held by the private sector. The brief argues that AI Factories are suited to supporting research and training medium-sized models but are insufficient to

boost commercial AI innovation at scale across the EU. If gigafactories replicate this pattern, Europe may build impressive assets without achieving the speed, openness, or market dynamism required to challenge [US based] hyperscalers.

There is a further issue of workload design. Policy discussions still over-index on training capacity, but competitive advantage increasingly depends on inference infrastructure. If Europe builds symbolic training prestige while underbuilding commercial inference capacity, it may succeed politically without creating a broad-based AI economy. The gap between EU and US private investment tells the story plainly: between 2023 and mid-2025, private venture capital investment in AI startups was 66% in the United States compared with 12% in Europe. The US produced more than 50% of all significant AI models in 2024. American and Chinese firms together operate over 90% of AI-specialised data centres globally, while European companies operate six. The gigafactory programme is necessary. It is not sufficient.

Sweden's comparative advantage: energy, talent, and startup density

Sweden occupies an unusually strong position within the European AI transition — stronger, arguably, than its political class has yet fully internalised. The country combines characteristics that are becoming scarce in the AI era: abundant low-carbon electricity from a generation mix split roughly 75% hydropower and 25% wind, and a cold climate that enables natural free cooling. An electricity tax rebate reduces levies to 0.6 öre per kilowatt-hour for qualifying industrial loads — though, regrettably, data centre loads are not yet eligible. Long-term power-purchase agreements, now a double-edged sword, allow hyperscalers to lock in energy costs for fifteen years, insulating them from the gas-linked price volatility that afflicts Germany, the Netherlands, and Southern Europe. European energy prices average roughly three times those of the United States and China; Sweden is one of the very few European countries where this gap narrows to a competitive differential rather than a structural handicap — a place that exports bits globally, which is a lot easier, rather than building new grid/wire infrastructures to pass on energy. This is why properly structured data centres in Sweden should enjoy energy costs comparable to those of heavy manufacturing – it is comparable to when we built smelting furnaces over 100 years ago only to radipdly understand we nneded to capture the whole value chain making steel into cars.

The capital markets have noticed. EdgeConneX announced in February 2026 a new data centre site in Skellefteå [the real value in the bankrupt Northvolt site] with potential capacity of up to one gigawatt, powered primarily by renewable energy — which upon completion would be one of the largest facilities in Europe. Mistral AI partnered with EcoDataCenter in February 2026 to develop a new AI-focused facility in Sweden for large-scale model training and inference, leveraging Swedish renewable energy and EcoDataCenter's sustainable operations. CapMan Infra acquired three EcoDataCenter sites in 2025, signalling private-equity appetite for Nordic AI infrastructure. The Sweden data centre market is forecast to grow from 0.47 GW installed capacity in 2025 to 0.71 GW by 2031, a 7.4% compound annual growth rate driven by hyperscale AI workloads.

But Sweden's advantage extends well beyond energy and cooling. Stockholm has the highest number of unicorns per capita of any city in the world outside Silicon Valley. Sweden is home to fourteen unicorn startups with a collective ecosystem valuation exceeding €239 billion — a figure that has doubled in the past five years. The AI-native cohort is accelerating: Lovable, the AI vibe-coding platform, reached unicorn status in July 2025 just eight months after launch, raising $200 million at a $1.8 billion valuation. Legora, automating tasks for lawyers, raised $150 million at a $1.8 billion valuation. Neko Health, co-founded by Spotify's Daniel Ek, raised $260 million for preventive AI health scans. These sit alongside Spotify, Klarna, Einride, and the broader ecosystem shaped by generative AI, climate tech, and deep engineering talent from KTH and the Stockholm School of Economics-all these companies are consumers of AI compute and also a sign of new industry time to come.

The strategic question is whether Sweden will content itself with being a well-run host for foreign-owned data centre assets, or whether it will become an architect of the new industry AI ecosystem: a country that links green power, compute infrastructure, universities, applied research, public procurement, domestic venture capital, and startup formation into a self-reinforcing system. The infrastructure investments are arriving. The startup culture is proven. What remains underdeveloped is the connective tissue: research-to-commercialisation pathways, public-sector AI procurement that creates domestic demand, targeted risk capital for AI-native firms, and an explicit new AI industrial strategy that converts Sweden's enabling conditions into owned capabilities rather than hosted capacity-it is a current battle between the old and new.

European sovereignty in information technology

Europe's sovereignty question has moved from abstraction to operational urgency. The numbers are unambiguous. Between 2000 and 2025, the number of European firms among the world's top 100 by market capitalisation fell from more than forty to eighteen. None of the European companies founded after 2000 has sustained a market capitalisation above €100 billion. The EU's share of global GDP fell from approximately 22% to 14% over the same period. Productivity growth halved. American firms dominate cloud, AI, and platform markets. Chinese firms increasingly challenge in hardware, manufacturing AI, and frontier models. Europe retains niches in pharmaceuticals, industrial automation, and green process engineering, but has struggled to translate laboratory excellence into scalable global platforms.

This is the deeper significance of the AI Factory and Gigafactory agenda. A sovereign technology position does not require full autarky. But it does require bargaining power: the capacity to build, host, regulate, and deploy critical digital infrastructure without total dependence on foreign hyperscalers or non-European supply chains. If Europe succeeds, regulation and industrial policy will reinforce each other. If it fails, Europe risks becoming a sophisticated supervisory layer atop value chains whose economic rents accrue in Washington, Shenzhen, and Moscow.

The sovereignty question is therefore not primarily ideological. It is about whether Europe participates in the ownership structure of the intelligence economy or merely governs its edge conditions. The data suggest that this question is being answered in real time, and not necessarily in Europe's favour. The coming two years — the buildout of gigafactories, the enforcement of the AI Act, the scaling of inference infrastructure — will determine whether Europe's regulatory-industrial strategy produces an independent new AI industrial base or an elaborate governance framework around someone else's economy.

Brynjolfsson's canaries in the coal mine: the evidence hardens

Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen's research at the Stanford Digital Economy Lab has now moved the debate about AI and employment from speculation to measurement. Using high-frequency administrative payroll data from ADP — the largest payroll provider in the United States, covering millions of workers at thousands of firms — they document six facts that characterise labour market shifts following the widespread adoption of generative AI. The findings are no longer preliminary.

The headline result is a 13–16% relative decline in employment for early-career workers aged 22–25 in the occupations most exposed to AI, controlling for firm-level shocks. For young software developers specifically, employment has fallen by approximately 20% since late 2022. Employment for experienced workers in the same occupations has remained stable or increased. The adjustments are occurring primarily through employment rather than compensation: firms are not cutting wages; they are not hiring entry-level workers in exposed roles. And the employment declines are concentrated in occupations where AI automates rather than augments human labour. In occupations where AI complements rather than replaces human work, employment has held steady or grown. But overemphasis on AI-ifying the old industrial economy is devastating for a country like Sweden — not only because it pushes an often young workforce into occupational obsolescence, but because it triggers a political reflex to compensate through the income tax transfer model. Tax transfers do not solve this; they amplify it. For a country like Sweden, which has built its middle class on precisely that model, this feedback loop is destructive.

This is the canary in the coal mine.

Close to 70% of occupations in the lowest AI-exposure quintile saw rising early-career employment between October 2022 and September 2025, compared to less than half of occupations in the highest-exposure quintile. The pattern is consistent across multiple measures of AI exposure and robust to alternative explanations including the exclusion of technology firms and remote-work-amenable occupations.

The implication for policymakers, educators, and corporate leaders is now concrete rather than speculative. The AI revolution is not a future event to be modelled in scenarios. It is producing measurable labour market distortions today, and those distortions are falling disproportionately on

the youngest, most exposed cohort of workers. Education systems that continue to train junior professionals for roles that AI is already compressing are producing graduates into a structural headwind. Reskilling frameworks, income support mechanisms, and university curricula must now be designed around the reality that entry-level knowledge work is contracting in precisely the occupations that used to absorb the most educated young workers.

The decline of the old paradigm

General industrial economies are now confronting the limits of a developmental model that served them for a century and a half. The evidence is not anecdotal; it is structural. EU industrial production fell by 1.5% in the euro area and 1.6% in the EU in January 2026 compared to the previous month, reaching its lowest seasonally adjusted level in over a year. Year-on-year, euro area industrial output declined by 1.2%. Since 2019, European production of machinery has fallen by 3%, fabricated metals by 6%, non-metallic mineral products by 11%, basic metals by 12%, and both chemicals and automotive by 14%. German crude steel production fell to one of its lowest levels since the 2008 crisis. BASF launched a cost-cutting programme targeting over €2 billion in annual European savings and began exploring a US stock market listing for a major division. ThyssenKrupp announced plans to eliminate 11,000 jobs — 40% of its workforce.

These are not cyclical fluctuations. They are symptoms of a paradigm reaching exhaustion. The old model excelled at scaling physical production, coordinating labour, financing large fixed assets, and diffusing productivity through stable sectors and national champions. It was less well designed for a world in which advantage depends on intangible capital, digital infrastructure, hyperscale compute, and the speed with which new knowledge is converted into deployable systems. European production hubs are approximately one-third less effective than their American counterparts at converting density into productivity, according to IMF research. Internal barriers to trade — estimated at 44% ad-valorem for goods and 110% for services — suppress the ability of European clusters to achieve American-style scale economies.

The euro area's growth forecast for 2026 stands at 1.1%, close to its weak potential rate. Europe's working-age population will have shrunk in more than two-thirds of EU countries by 2050. Energy prices remain, on average, three to five times those in the United States. The closest the EU ever came to matching American productivity was in 1995, at 95%; the gap has widened since. What appears in many European economies as stagnation, institutional heaviness, and declining dynamism may therefore be more than cyclical weakness. It may reflect the exhaustion of an old developmental formula that AI is in the process of replacing.

This does not make existing institutions irrelevant. It raises the bar for their reinvention. Countries and companies that can combine industrial depth with AI-native infrastructure, entrepreneurial speed, abundant energy, and a culture of experimentation may still lead the next era. Those that merely digitise old structures will improve for a time, but they will not shape the new wealth architecture that is now emerging. The question for Europe is whether it has the political will to build a new industrial logic while the old one is still generating the tax revenues and employment that fund the transition. It is the rise of the New Swedish Dynamism.

Swedish dynamism in transition

Swedish dynamism refers to the paradox that has long puzzled outside observers: how a high-tax, consensus-driven Nordic welfare state consistently produces an outsized share of global entrepreneurs, unicorns, and transformative innovations.

The roots run deep. Sweden industrialised late but ferociously, compressing a century of British-style transformation into roughly fifty years between the 1870s and 1920s. That compression forged institutional habits — pragmatic adaptation, comfort with creative destruction, and a willingness to let old industries die so new ones can emerge — that persist today. Atlas, Ericsson, SKF, and ASEA were not accidents; they were expressions of a culture that paired current times engineering ambition with merchant pragmatism.

What makes the Swedish case theoretically interesting is that the welfare state didn't fully suppress this dynamism [not that it didn't try] — it may have amplified it. The Rehn-Meidner model of the postwar era deliberately used solidaristic wage policy to squeeze out inefficient firms while channelling labour toward productive ones. The safety net reduced the personal cost of failure, effectively socialising entrepreneurial risk. When Spotify, Klarna, and King emerged from Stockholm decades later, they did so in a society where founding a company didn't mean gambling your family's healthcare or your children's education or succumb to ancient family power structures or existing industrial traits.

There is also a cultural dimension. Sweden's combination of high interpersonal trust, flat hierarchies, widespread English fluency, and early digital infrastructure adoption created an unusually fertile environment for networked, technology-driven business models. The smallness of the domestic market forced an export orientation from day one — Swedish companies are born global because they have no choice.

But the old engine is losing power. The industrial model that built modern Sweden — steel, timber, automotive, telecom hardware — has been generating diminishing marginal returns for decades, optimised to its limits and increasingly unable to sustain its underlying profitability. The value captured per kilowatt-hour of energy consumed, per ton of raw material processed, per hour of labour deployed, has been in structural decline. Sweden, like much of the industrialized West, has been living increasingly on the institutional and financial capital accumulated during a golden century, not on the momentum of what comes next. This is not a crisis yet. But it is an unmistakable trajectory.

What is emerging now, still in its infancy, is a fundamentally different paradigm. The global economy is shifting from the production of physical goods to the production of intelligence — AI models, digital services, data infrastructure, and the ecosystems that surround them. In this new logic, value is not extracted from the ground but generated from compute, energy, talent, and ideas in combination. The economics are staggering: a kilowatt-hour routed through advanced AI compute can generate orders of magnitude more economic value than the same kilowatt-hour used to smelt ore, heat a paper mill or built electric car batteries.

Sweden is unusually well positioned for this pivot, and not by coincidence. The same structural advantages that defined the old dynamism — abundant low-cost energy, a highly educated population, deep engineering culture, political stability, and instinctive export orientation — map almost perfectly onto the input requirements of the intelligence economy. Add to that some of the cleanest electricity grids in Europe and a climate suited to cooling data centres, and Sweden holds infrastructure cards that most nations or large organisations simply cannot replicate — but would like to lay their hands on.

The historical parallels are instructive. Silicon Valley did not emerge from nothing; it grew from a specific confluence of research universities, defense spending, risk capital, and a culture that tolerated — even celebrated — failure. Israel's startup ecosystem was catalysed by military technology transfer, immigration, and government-backed venture funds. Shenzhen transformed from a fishing village into a hardware metropolis because policy, capital, and entrepreneurial energy converged at the right moment. In each case, a new generation of entrepreneurs, backed by new forms of capital and animated by new ideas, built something that the previous economic order could not have predicted and still can't.

Sweden now stands at a comparable inflection point. A new generation of founders, engineers, and investors is beginning to build at the intersection of AI, energy infrastructure, and digital sovereignty — not merely adapting the old industrial playbook but writing an entirely new one. Ventures in sovereign AI compute, green data infrastructure, applied machine learning, and deep tech are appearing with increasing frequency, often connecting Sweden's legacy strengths to the frontiers of global technology competition.

This is a transition of scope, size, and direction. Scope, because it touches not just one sector but the foundational logic of how value is created. Size, because the global intelligence economy will dwarf the industrial economy in output within a generation. And direction, because Sweden's role shifts from exporting commodities and manufactured goods to exporting compute capacity, AI capability, and the institutional frameworks that make trustworthy digital infrastructure possible-all accessing and supply an ever-growing global market for intelligence using the BABIS acronym [BABIS, Buy American Build In Sweden].

Swedish dynamism, then, is not a finished story of industrial success. It is an unfinished story of reinvention — one in which the old foundations of trust, talent, energy, and institutional pragmatism are being repurposed for an era whose full contours we can only begin to discern. The question is not whether Sweden has the assets. It is whether it will move with the urgency and ambition the moment demands, rather than chasing short-term gains from the old industrial model.

A civilisational pivot and shift

The most important implication of AI is not that every company will become somewhat more efficient. It is that the foundations of value creation are pivoting from a world defined by labour, capital equipment, and linear process efficiency to one increasingly defined by software, models, compute, and continuously improving synthetic cognition. This shift has civilisational consequences.

In education, AI can make high-quality personalised learning available at scales that traditional teacher-student ratios never could. In medicine, it can tighten the loop between research, diagnosis, prevention, and treatment. In science, it can accelerate discovery by simulating possibilities before expensive real-world testing begins. In industry, it can collapse the distance between design, planning, production, and optimisation. Across sectors, AI reduces the marginal cost of applying intelligence to problems that were previously constrained by time, training, and institutional scarcity.

A world with abundant machine intelligence creates the possibility of better medicine, broader access to education, faster scientific progress, more adaptive institutions, and a reallocation of human effort away from repetition and toward judgement, care, creativity, leadership, and meaning. The ultimate significance of AI is not that it helps decaying old organisations do the same things faster. It is that it enlarges the set of things humanity can do at all — creating a completely new set of building blocks for the new industries.

But that outcome is not inevitable. It depends on choices made now: on how energy is allocated, on who controls compute, on whether regulatory frameworks enable or constrain innovation, on whether education systems adapt or fossilise, on whether political leaders treat AI as a new industry and a new layer of sovereignty or merely as a set of tools to be supervised. The choice is still open. The window is closing.

What leaders should embrace

The strategic challenge for current leaders is larger than adoption. Adding AI to today's workflows will improve competitiveness in the short term, but over time it will not measure against the organisations, sectors, and national systems that are rebuilt around intelligence abundance from the ground up. Today embracing the old is like pissing in your pants in below 40° – feels good to begin with but soon creeps up on you!

The right question is no longer "How can AI improve the current business?" It is "What forms of value creation become possible once intelligence is cheap, ubiquitous, and embedded everywhere?" Leaders who ask only the first question will optimise their way into obsolescence. Leaders who ask the second begin to see the outlines of a new industry and, with it, a new social contract.

AI is not merely augmenting the world that exists. It is helping bring into being a world organised around a new source of wealth, a new logic of production, and a higher baseline for the human condition. The task for leaders is not simply to adopt it. The task is to understand that history has opened, and to build accordingly.

Jonas Kjellstrand is a senior strategic adviser operating through UCC AB, based between Stockholm and Palo Alto. He advises at the intersection of European sovereign AI infrastructure, Nordic energy economics, government relations, and Silicon Valley technology partnerships. He holds an MBA from Stanford GSB (joint with Stockholm School of Economics) and business administration from MIT Sloan.