

You're Gonna Need *a Bigger Boat*

By a slightly worried optimist

Insights

Last week, probably like many readers, I worked through two research pieces that jolted financial markets and beyond. The first report was "The Adolescence of Technology" written by Dario Amodei, the CEO of Anthropic. The second report "The 2028 Global Intelligence Crisis" written by Citrini Research & Alap Shah. Even though I thought I was prepared for what came next, both reports were enlightening and help put things into perspective (important to keep in mind that the Citrini report starts with the following sentence "What follows is a scenario, not a prediction"). A long time ago, my dad told me "better be ready than sorry". It's with this mindset that I strongly encourage everyone to read these reports. When I reflected on what I have been reading, I thought about the Steven Spielberg movie *Jaws* and the now famous sentence by Chief Brody "you're gonna need a bigger boat". It's not subtle. Neither is this moment for AI.

AI resembles the great engines of the Industrial Revolution, not because it moves steel or plows fields, but because it restructures society, labour, and power in similarly profound ways. Like the internal combustion engine, which multiplied human capability while simultaneously enabling exploitation, colonial extraction, and environmental damage, **AI is not inherently good or bad. It is a force amplifier.** It magnifies human intent, whether toward innovation or inequity, cooperation or concentration of power.

And as with the Industrial Revolution, the distribution of AI's gains will not happen organically; it will be contested. The middle class did not emerge automatically from mechanisation, it was forged through decades of struggle, bargaining, and political reform. AI, too, will generate windfalls and upheavals, but who benefits and who bears the costs will depend on choices made now, not on the technology itself.

Beneath the rhetoric sits an operational fact: adoption has moved from toy to tool. Inside enterprises, usage has surged and deepened, from casual prompts to embedded workflows. OpenAI's own telemetry reports an eight-fold rise in weekly enterprise messages in 2025 and, more tellingly, a 320-fold jump in "reasoning" tokens associated with multistep tasks. This is the tell that systems are moving from the periphery of knowledge work into its centre.

The consultants' spreadsheets are therefore unsurprising: one widely cited estimate puts annual value creation from generative AI at \$2.6-44 trillion, concentrated in customer operations, software engineering, marketing and R&D. Yet the macro literature is cooler; the Organisation for Economic Co-operation and Development (OECD) and International

Monetary Fund (IMF) emphasise that in the near-term AI may lift labour-productivity growth by only tenth of a percentage point per year, with adoption frictions and regulatory constraints pacing the gains. Translation: micro benefits arrive well before the macro data crown them.

That is the setup. The uncomfortable part is what comes next: second round effects that start in offices and leak into household balance sheets, local tax bases and politics. The knife cuts first where routine cognitive work is thickest.

The U.S. Bureau of Labor Statistics embeds AI exposure in its 2025-33 projections, with the heaviest risk in computer, legal, business and financial occupations, a reminder that task content, not industry label, defines the frontier.

Investment bank economics now temper the 2023 headlines into a baseline of roughly 6-7% displacement during the transition with a temporary unemployment hump of around 0.5 percentage point, benign only if redeployment works in practice as well as it does in models.

The International Labour Organization's 2025 refinement adds distributional texture: clerical roles, disproportionately held by women, sit at the top of the risk gradient; in high income economies, roughly 9.6% of women's jobs fall into the highest exposure category versus 3.5% for men. Exposure is not unemployment; widespread job transformation is the base case.

The frontier is no longer confined to screens. If even part of the robotics roadmap lands on schedule, **2026 could be the year "AI meets atoms."** Tesla's Optimus targets external sales after internal pilots; Agility's Digit is already clocking paid shifts in logistics; Boston Dynamics' electric Atlas is moving from demonstrations toward partner fleets.

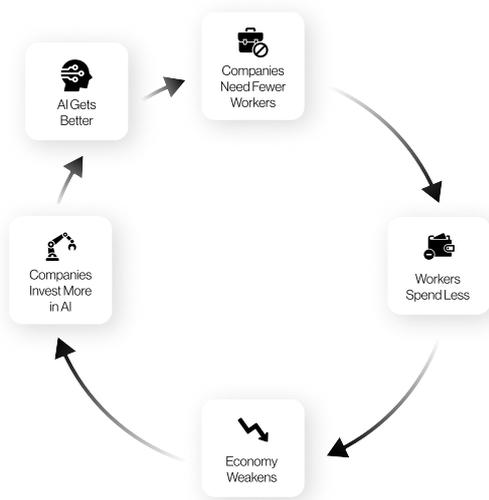


Schedules slip in robotics as a rule, but they still matter for wage setting and capital allocation. Warehouses and lines do not require science fiction generality to rewrite rosters. Economic output generated by AI boosts headline GDP figures and this “Ghost GDP” lens matters because it ties the firm and the household in a coherent loop.

Consider how a white-collar shock would surface in credit: not as a wave, but as a map. At the national level, U.S. mortgage delinquencies ended 2025 near historic lows, but the New York Fed already notes deterioration concentrated in lower income areas and metros with falling home prices, classic weak points when labour income growth stalls.

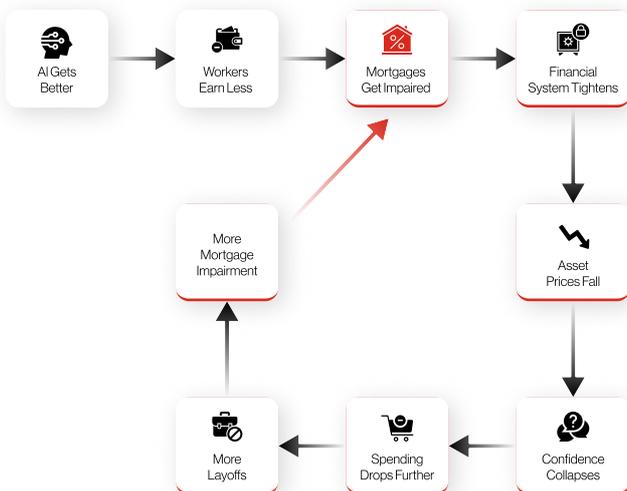
Private monitors see the same split screen: flat national rates masking a rise in early and serious delinquencies across many cities. Mechanically, if AI compresses entry ladders and slows progression, where task automation bites first, households that looked solid on paper discover fixed mortgage costs growing heavy. The current-to-30-day transition rate, measuring change in delinquencies, becomes the canary. Even with bankbook delinquencies still low by historical standards, direction matters.

Without Mortgage Concerns



Source: CitiriniResearch

With Mortgage Concerns

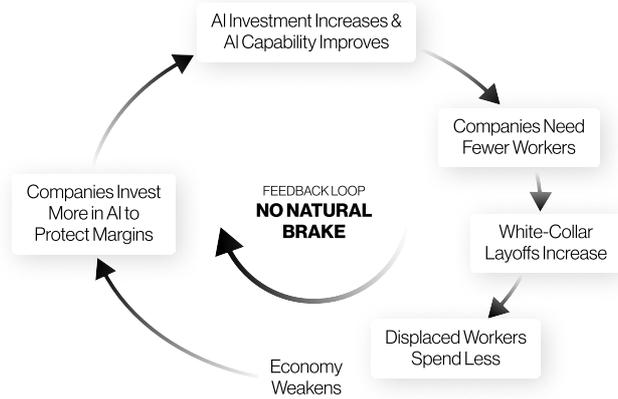


Source: CitiriniResearch

Consumption tells a similar story. If productivity rises faster than wages—and the early gains mostly go to capital and a small group of power users—spending splits into a K-shape: luxury and bargain goods hold up, while the middle gets squeezed. OECD and IMF work points to distribution as the key issue. Who gets the gains matters more for demand than how big the gains are. That fits Citirini's scenario. Companies reinvest profits into compute and data centres, while households either rebuild savings or fall behind on payments.

AI-Driven Economic Feedback Loop

(No Natural Brake)



Source: CitiriniResearch

This isn't a forecast. It's a coherent way to think about what can happen when micro-level progress runs ahead of the macro economy.

Commercial real estate and municipal finance sit in the blast radius. Organisations that replace layers of junior staff with AI mediated workflows do not simply lower costs; they shrink footprints and the service ecosystems around them. That lowers municipal tax bases and pressures downtowns already adjusting to hybrid work.

Studies mapping generative AI's value concentration in finance and tech hubs add a twist: the places best positioned to harness the upside may see the sharpest near-term adjustments if spending shifts from city centre floors to exurban server farms.

Markets are already trading the split screen. On one side sits a tidy P&L: productivity up, turnaround times down, compliance automated, client agents always on. Citi's analysis of the sector is explicit that AI can lift banking profits by raising internal productivity and improving the client experience, even as it warns about governance burdens and spread compression as always on agents intensify price transparency.

Strategically, established companies are up against newer cloud-native competitors who can more easily build agentic automation directly into their systems. It is perfectly possible for banks to get better while borrowers get bumpier, a combination that flatters quarterly numbers even as it raises the system's sensitivity to white collar income volatility.

Politics will arrive late and then all at once. Democracies do not excel at regulating technologies that evolve faster than legislative calendars.

Even as initial gains from AI are modest—registering productivity improvements only in small increments—the disparity between AI winners and the broader population can quickly become pronounced. This widening perception gap risks amplifying populist sentiment across the political spectrum, as individuals and groups who feel left behind may push back against the pace or terms of technological change. Such pressures often lead to policy responses that aim to slow or regulate the spread of AI, particularly through measures focused on strengthening labour protections and data rights. As a result, the momentum of AI adoption is not only shaped by technical capability but also by the social and political responses to its uneven impact.

On the right via reshoring and industrial policy. The governance problem is compounded by agentic AI. As systems are allowed to act, initiating payments, modifying code, orchestrating procurement, the question becomes less about model accuracy and more about autonomy plus access. Citi's "do-it-for-me" analysis reads like a set of pre-mortems: enforce least privilege access, require human in the loop on high-risk actions, log everything. Better to turn safety into boring compliance than discover it as headline risk.

There is also a monetary policy temptation to resist. The late 1990s offered a rare alignment: an ICT driven productivity surge, ample labour supply and deepening globalisation that allowed Greenspan to lean dovish without stoking inflation.

Today's surface rhymes—strongest productivity prints of the expansion, capex booms in software, data centres and information processing kit—are real enough.

Yet the backdrop is different: ageing demographics, tighter immigration and smaller global labour pools are headwinds, not tailwinds; deglobalisation and tariff uncertainty push against goods disinflation; and cuts to public R&D hobble long lag channels through which technology raises total factor productivity.

Even if AI delivers a genuine boom, it may not buy the same inflation slack Greenspan enjoyed. Caution is not anti-tech; it is arithmetic.

What, then, should we watch to separate signal from narrative?

- First, watch the hiring ladder. Entry-level roles—analysts, paralegals, junior coders—are the backbone of white-collar training. If campus hiring shrinks two years in a row, it's hard to keep calling it a "temporary blip." And that matters, because a big chunk of white-collar work is already squarely in AI's early strike zone.
- Second, look past the national averages as they can hide a lot. Look instead at the city-by-city metrics—think high debt-to-income, newer loans, and prime borrowers, specifically in tech-heavy cities as leading indicators. Those pockets are usually the first to crack.
- Third, look at how deep AI runs inside companies. When custom GPT workflows, API usage, and agent permissions start showing up everywhere, you know it's moved past experiments and become part of the core system.
- Fourth, robots—watch orders, not demos. Pilot projects and trials at places like Toyota or Hyundai are interesting, but they're just signals. The real warning signs show up when companies start placing real purchase orders. Demos hint at curiosity; orders mean commitment.
- Fifth, pay attention to how fast policy turns into action. That means real safety rules for AI agents, procurement rules that demand audit trails, and reskilling programs with actual money behind them.

These signals decide whether we surf or swallow water.

Investors and risk officers can draw a straightforward map from here. Household credit is a neutral headline with a bearish tail. You need to hedge by looking at specific regions and job mixes and stop treating "prime" borrowers as bulletproof when income growth stalls. Commercial real estate and municipal deserve close attention where office usage and entry-level hiring are both slipping.

Equities can book short term margin expansion for early adopters, but the medium-term demand risk is real if wages lag. In operations, permissions and audit trails for agents are not hygiene; they are the control room.

If that is the diagnosis, what would a bigger boat look like from a policy standpoint?

- Start by favouring augmentation over substitution in the tax code: tie accelerated deductions for AI capex to verifiable redeployment and upskilling outcomes, not just to racks of GPUs. If firms want the break for agents and orchestration layers, show the task level skill transfer and wage paths for at risk cohorts.
- Next, codify agent safety so autonomy becomes a managed variable rather than a leap of faith: least privilege access, human in the loop for high-risk actions, tamper proof logs.
- Build early warning dashboards that blend labour and credit microdata with adoption depth, entry level hiring and internal progression, city-by-city level current to 30-day transitions, and the share of work routed through agents and APIs.
- Finally, fund apprenticeship grade training measured by what juniors can actually do (task transfer), in place of attendance hours or a glossy certificate.

Not all sectors face AI disruption at the same speed. Early on, HALO (heavy assets, low obsolescence) businesses have more breathing room. AI is hitting tasks before physical assets.

Enterprise AI adoption is moving fast, but economy-wide productivity gains are still small. That gap puts near-term pressure on asset-light, services-heavy businesses doing routine cognitive work, while HALO sectors like utilities, transport assets, energy networks, towers, logistics hubs, and specialty industrials can use AI to boost margins without breaking their revenue model. Think predictive maintenance, outage forecasting, and asset health monitoring, not revenue replacement.

Where the work sits matters. The biggest exposure today is in office and professional roles such as computers, legal, business, and finance—not in physical infrastructure. AI automates screens and back offices before substations and steel. Revenue structure matters too. Regulated or contracted businesses are protected by physical bottlenecks and service agreements, not seat licenses or billable hours. The real capex threat comes later, when AI fully pairs with cheap, capable robotics. That's coming—but slowly. Robotics is moving from pilots to paid deployments, yet scale is still constrained by hardware limits and regulation. For HALO operators, that means a smoother transition: cost savings today, with automation risk arriving in waves rather than all at once.

There are risk contours investors should price. Second wave automation will chip away at labour content in inspection, warehousing and light assembly as mobile manipulators scale; data rich HALO verticals (say, automated terminals) will soon show robotics as a standard line item

alongside analytics. And because adoption in services depresses some white-collar incomes at the margin, demand side spillovers, office electricity loads, discretionary transport volumes, can soften before regulated resets catch up.

But these are timing issues, not existential threats. The same diffusion lag that tempers the national accounts gives HALO management teams a window to codify agent safety in control rooms, renegotiate reliability regimes with regulators and off takers around AI enabled performance, and redirect a slice of OPEX savings into low regret hardening of core assets. In a world where enterprise usage is sprinting while GDP registers by tenths, owning the infrastructure of the real economy, heavy assets with light obsolescence, is not a Luddite bet; it is a timing edge.

None of this requires doom. The evidence we do have, on adoption depth inside firms, on productivity ranges, on role by task exposure, on mortgage micro maps, argues for **urgency, not for despair**.

The path to "bigger boat" policy is not a mystery. Favour augmentation over substitution. Make safety routines dull and universal. Instrument the transition so you see fragility before it spills. Pay for measurable diffusion, not for slogans. If we do that, the water will still get rougher, but we will at least be sailing something built for turbulences. And unless policy, corporate practice, safety rails and social insurance scale with the wave, Brody's line will feel less like a quip and more like risk management.



Selected sources used throughout: OpenAI, *State of Enterprise AI 2025* (enterprise usage, "reasoning" tokens, structured workflows); McKinsey, *The economic potential of generative AI* (value pools; functional concentration); OECD & IMF (early phase productivity ranges; adoption frictions); BLS (AI exposure by occupation); Goldman Sachs (transitional displacement baseline); ILO 2025 (exposure gradients; gender asymmetry); TechCrunch & Humanoid.press (humanoid deployment timelines); New York Fed Household Debt and Credit (mortgage delinquency geography) and private monitors Cotality/ICE (metro level drift); Citi GPS, *AI in Finance and Agentic AI* (profit uplift; governance); Dario Amodei, "The Adolescence of Technology" (capability acceleration; adolescence framing); Citrini Research, "Global Intelligence Crisis" (Ghost GDP).

Disclaimer

This material is for the use of the recipient in accordance with the restrictions and/or limitations implemented by any applicable laws and regulations only. It is intended only for the recipient and may not be published, circulated, reproduced or distributed in whole or in part to any other person without the Bank's prior written consent. Unless otherwise indicated, the information is made available for informational purposes only, without considering the recipient's financial situation, investment objectives, risk tolerance, financial situation, or any other particular needs and should not be treated as legal or taxation advice.

The information is not and should not be construed as an offer or a solicitation to deal in any investment product or to enter into any legal relations. Any investment decision made based on the information provided is the sole responsibility of the client. The Bank disclaims any liability for any losses or damages resulting from the use of this information. The Bank assumes no responsibility for the way in which the client may choose to use or apply this information, or for any investment decision or

transaction that the client might undertake as a consequence. It is the client's own responsibility to ensure that this product is suitable for him or her and the client must make his or her own decision concerning this product. The client may also wish to obtain advice from other sources before making any decision.

Past performance is not indicative of future results. Any forecast on the economy, stock market, bond market and economic trends of the markets are not necessarily indicative of the future or likely performance of the product. Any investment involves risks, including the total loss of the invested capital.

For queries arising from, or in connection with this material, please contact the person who sent you this material.

This advertisement has not been reviewed by the Monetary Authority of Singapore.