

By Invitation - Pace the race

The "let-it-rip" camp is wrong on AI

They make four main arguments and all of them are flawed, write Bridgewater's Nir Bar Dea and Greg Jensen

THE INTENSE debate of recent weeks on the risks around artificial intelligence presents an opportunity. Anthropic's Dario Amodei urged the industry to "pace the frontier". Many others echoed him, and politicians across the spectrum have signalled they are ready to act. Not everyone, however, agrees. President Donald Trump has spoken against new AI regulation, as have Meta's Mark Zuckerberg and Nvidia's Jensen Huang.

We have been here before: when a new technology surfaces, society is forced to grapple with the question of whether it needs controlling, and if so how much and how to achieve that. That moment is here with AI. It is the most important policy decision of our lifetime.

Those sceptical of regulation make four arguments: the risk is overstated; existing law already covers it; some labs are backing regulation only to entrench their lead; and if America slows down, China will overtake it and nobody will be safer.

The first claim, that the risk is small or improbable, misses the forest for the trees. Highly capable models are already in the hands of those seeking to do harm. We estimate that cyber-attacks have more than doubled in frequency since ChatGPT's release in 2022 and are accelerating. Everyone is now exposed, and what the world has seen so far is the tip of the iceberg: nefarious actors are likely to use open-weight models, which are far harder to track than proprietary ones.

Furthermore, models with superhuman capability are here. The labs' disclosures since the Hugging Face inci-




ILLUSTRATION: DAN WILLIAMS

dent—when a horde of OpenAI agents broke out of their testing environment and began conspiring with each other—show a growing trend. Anthropic recently self-reported experiments suggesting that models may be months from self-replicating autonomously, without recourse to a data centre. Over time, there is a risk, however small, that superior agentic intelligence will pose an existential risk to humanity.

In the meantime, society must grapple with AI's impact on labour. If it delivers on its productivity promise, AI will structurally disadvantage human labour. Countries across the Western world have already been destabilised by a wave of populism following the era of globalisation, which saw a seven-percentage-point fall in manufacturing's share of non-farm employment between 1990 and 2010. We believe that

a larger and faster second wave is now building, with the potential to displace almost a fifth of all American jobs over a shorter time period.

People will disagree about the odds; the question is whether the aggregate risk warrants action. The risk from cyber-attacks is significant and growing. The risk linked to agentic superintelligence is uncertain, but non-zero and potentially catastrophic. The risk to human labour looks inevitable, particularly if society avoids the first two. The need to regulate the aggregate is, therefore, clear.

The sceptics' second claim, that existing laws are enough, ignores that those laws were not written with AI-specific dangers in mind. Society has learned this before—with food safety after the Industrial Revolution, drug safety after the patent-medicine boom and fi-

nancial rules after the 2007-09 crisis. It won't accept contaminated baby food, a drug that kills thousands or system-de-stabilising financial instruments merely because companies will later be liable. Where the magnitude of the consequence is intolerable, liability is inadequate and therefore regulation cannot be applied retroactively. Each industry, after much debate, went on to thrive under subsequent regulation.

The third claim, that some labs favour regulation to entrench their advantage, may be true but is irrelevant. Never mind that others say these companies have too much power; they are reporting so themselves. We estimate that OpenAI and Anthropic are on track to control roughly 35% of global compute between them by 2028. Ten S&P 500 companies, nearly all in the AI supply chain, make up around 40% of its market capitalisation, and Nvidia alone is worth over 4% of global GDP. Whatever the motives, this extreme concentration of power must be regulated.

The fourth claim, that China will zoom ahead if America brings in new rules, is multi-layered. The AI race is like the cold war, in which American pursuit of nuclear technology was highly regulated while the Soviet Union ran unchecked. Far from putting the Soviets ahead, looser controls led to in-

cidents like Chernobyl, which helped precipitate the collapse of their empire. China understands that: it has tighter controls, consistent with enforcing stability by controlling technology, even as it distils frontier-lab models. Regulation in America would structurally slow China down while limiting the risk of serious incidents that could alter the race.

In building a path forward, it is necessary to think about principles through three lenses: the who, the what and the how.

The question of who should regulate is the one that matters most, because the people responsible will determine the quality of everything else. Some argue that companies should police themselves given governments' limited expertise in AI. But governments are best placed to act for society because, in a democracy, their incentives are most closely aligned with society's. If a regulator suffers from an expertise gap, it can be narrowed. It is much harder to close the incentive mismatch inherent in corporations.

And who should be regulated? Those who can cause outsize impact. That means all who develop models, proprietary and open-source alike, and all end-users. Too narrow a view will leave loopholes.

The "what" concerns the nature and extent of regulation. It should be the minimum needed to maximise long-term benefit, or the most reward per unit of risk. Under-regulate and society is harmed, as with the unchecked child labour that followed the Industrial Revolution. Over-regulate, as America did to protect transportation companies in the mid-1900s, and economic growth is impeded, which also harms society. The goal must be rules that capture defined risks with the minimum necessary restrictions.

The "how" is the implementation, which should be fast—and fast-changing. Aviation is a model to follow: America's Federal Aviation Administration issues 300-plus airworthiness directives a year and has revised its comprehensive aircraft-certification rules more than 90 times since 1964. AI will require magnitudes more in speed and iteration.

None of this will be simple. But doing nothing is by far the worst option. If the debate stalls, the world will stay on a path whose risk—at best serious social disruption, at worst unprecedented catastrophe—no human should accept. With the right guardrails, AI could bring enormous benefits to humanity. Regulating it now is the only way to secure what it promises later. ■