
Managing Systemic Risks in Tech: Lessons from Finance



By François Cadelon and David Zuluaga Martínez , BCG; Peter Nathaniel , former Citibank and RBS; Theodoros Evgeniou and Ludo Van der Heyden , INSEAD

The financial sector's track record in risk management offers invaluable lessons for the tech industry.

Last month, the heads of seven major American AI companies emerged from the White House with an agreement on **“self-regulation”**. On the other side of the Atlantic, Europeans debate the long-awaited EU AI Act, the next major digital regulation following the EU’s Digital Services Act (DSA). The DSA is aimed at containing **“systemic risks”** from tech that include the “potentially rapid and wide dissemination of illegal content and of information” that is “incompatible with” large online platforms’ terms and conditions.

These are radically different approaches to address the AI challenge. The **risks posed by AI** have been debated for some time, including potentially systemic risks to political systems or public health due to **misinformation or disinformation** boosted by recommender systems and **deepfake technologies**. Striking the right balance between fostering innovation and ensuring safety is at the centre of the debate.

Given the speed of innovation, managing tech's systemic risks necessitates swift collaboration between regulators and the industry. Fortunately, there are lessons to be learned from other sectors without repeating costly mistakes – such as overreliance on self-regulation. The financial industry has spent decades, if not centuries, developing and refining mechanisms to contain, mitigate and respond to relatively similar risks. These efforts can provide a starting point for tech regulation.

Learning from finance

The financial sector has grappled with the phenomenon of **systemic risk** – understood as the risk that a shock to specific components of the financial system (say, individual banks) may have cascading effects that endanger the entire system. This is what happened in 2007-2008, when a shock in the US subprime mortgage lending space evolved into a global financial crisis. The repercussions extended well beyond finance, impacting global migration patterns and inequality within and across countries. The crisis was therefore “systemic” in yet another sense: a disruption within a single industry profoundly affected the entire “global system”. This is exactly the risk that **many fear** AI poses.

While tech and finance may both create systemic risks, they differ significantly in their approach to risk management. The tech sector, as a newcomer, would be wise to learn from the world of finance, given the similarities between AI and finance. Both sectors rely on opaque mathematical models built on large amounts of data and complex computations. More importantly, these models end up being used, in both industries, by executives with very limited understanding of the models themselves, while boards and regulators are distanced even further from the models they ought to govern. Similarities also exist regarding other risks, such as anti-money laundering concerns and the need to effectively monitor processes in handling so-called AI incidents.

Naturally, there are also great differences between tech and finance. While finance has faced challenges due to the knowledge gap between market players and regulators, this gap is even greater in the case of AI, and likely to worsen over time. Whereas financial regulators understand the specific types of risks that their industry faces, AI is rife with unknown unknowns that make risk management all the more challenging.

What can finance teach tech about managing risks?

Self-regulation is required, but insufficient

Hardly any attempts at self-regulation in tech have been successful (perhaps with the exception of the **Japanese gaming sector**). Even adequate risk management at the firm level may fail to address system-wide risks. Tech should embrace some form of external oversight to ensure what the finance world has come to accept: the role of regulators and independent third parties (like auditing firms) in ensuring and safeguarding the public interest and the firms' long-term **"social license"**.

Regulatory dialogue should involve the whole industry

Regulatory dialogue should largely take place at industry level and aim to strike a balance between keeping an industry innovative and competitive while protecting society. Too often the debate is about regulators sanctioning a particular "systemic" agent. However, true effectiveness lies in industry and government partnering to govern and manage systemic risks. Interestingly, such partnership has been more forthcoming in Canada and Scandinavia, which benefit from more collaborative and less individualistic cultures.

Tech needs 'nested' lines of defence

While self-regulation is insufficient, tech firms should nevertheless adopt strict risk management practices, with checks and balances and a governance structure not unlike that of banks. This essentially involves giving independent authority within the company to AI experts who can assess the appropriateness of deploying the technology in specific business cases. An **"AI watchdog board"** with real independence and teeth can enable companies that develop or use AI to define, implement and evolve rigorous internal risk management practices. Beyond individual firms, however, the tech industry needs to be regulated in each jurisdiction by appropriate agencies.

The models here are many, ranging from licensing requirements akin to those used in banking and pharmaceutical, to stricter corporate legal liabilities. Credible whistleblower processes and governance standards such as organisational structures, boards, disclosure requirements, contingency plans and transparency also need to be put in place. Of course, product safety requirements will continue to hold, but given the probabilistic nature

of AI systems, new processes, such as continuous **monitoring**, will need to be developed.

'Too big to fail' is also true in tech...

Similar to finance, certain tech actors – like Facebook or X (formerly Twitter) – are crucial to the entire system. Just as banks that are deemed domestically or globally “systemic” in terms of their importance face stricter regulatory oversight and liquidity requirements, so tech giants could be required to create redundancy for critical infrastructure, **explainability** standards for AI usage or mandatory stress tests and **red teaming**. Indeed, the DSA already imposes significantly stricter requirements on very large online platforms (defined as having more than 45 million unique monthly users in Europe).

... though not all risks are a function of size

As the industry becomes more interconnected, financial regulators have started to realise that size alone is an insufficient measure of risk. The recent collapses of Silicon Valley and Signature banks illustrate the point. While the contagion was rapidly avoided by regulators, it was clear that the failure of these institutions did pose significant risk to the system despite falling below the **size threshold** for the strictest scrutiny by the Federal Reserve. Much the same may be true for AI. For example, while LLMs may come from Big Tech, applications by smaller players across industries could pose major risks in specific domains, for example in critical infrastructure safety. A broader view of the tech system, considering sensitive applications within or by non-tech companies, is essential to effectively manage risk.

New global institutions and international coordination are paramount

Large tech companies operate globally and must adapt to diverse regulatory environments. As has been the case in finance, global cooperation is crucial to prevent “jurisdictional arbitrage” and properly coordinate responses to crises across governments. Some consistency and homogeneity of policies and their implementation within and across geographies and business models is necessary. For example, a safety net for the financial system in the event of a systemic crisis is to allow time (30 days in the case of the banking system) for G20 governments to coordinate their responses. Hence, those governments, through the **liquidity coverage ratio**, require all systemic

institutions to be able to survive for 30 days if the world comes to a standstill.

Ongoing innovation requires balancing regulatory stringency with sector profitability and competitiveness

Striking a balance between rigorous regulation and sector profitability is important to ensuring that there is continued investment in new technologies – including ways to make AI safer. For instance, stricter rules in EU banking have arguably impacted overall profitability compared to US banks. This asymmetry in a global financial market is simply not sustainable. It poses a risk that EU banks might not be able to efficiently recycle capital and fuel growth and stability of their countries, especially relative to their US competitors. A parallel situation regarding AI would generate strategic costs for lagging behind in technology development and could mirror the huge gap in profitability that affects the US and European bank sectors. This is not a call for weakening regulation, but for designing it in a thoughtful and more agile manner.

Learn with vigour, proceed swiftly and remain prudent

While the tech sector can learn valuable lessons from finance regarding industry-level oversight and international cooperation, there are also practices it should avoid emulating.

Tech demands faster regulatory processes

There's a notable difference in the speed of operations between tech and finance. Despite centuries of financial regulation, the quickest response time stands at 30 days. Most will agree that the response time for AI needs to be set to **one day at most** in serious crisis situations. This requires regulators and the industry to agree on rapid processes and protocols that finance doesn't even consider today. This should be approached with a balance of swift and gradual methods to avoid rocking the system and making the regulator a risk factor.

Tech likely requires a different engagement model

While large banks are sizable, the concentration of power is considerably higher in the tech space, particularly within the domain of AI. The system is poised to depend on a smaller number of behemoths that control critical IP

and resources underpinning advanced AI products. This, coupled with the gap of technical understanding vis-à-vis regulators, calls for more collaboration between large tech firms and regulators, as well as a greater commitment to the public interest duty by the former. Tech firms can help regulators design the right principles-based rather than rules-based regulatory framework that the rapidly evolving field of AI is likely to require.

Tech needs to remain continuously mindful of its unknown unknowns

Firms and regulators in finance can rely on quantitative risk models that leverage a wealth of historical data about previous crises. As noted earlier, finance has obtained a clearer sense of what a crisis looks like, even if potential root causes are not always identified. But matters are very different in the age of AI as there is no history to build on, nor data on past crises. Thus, any effort at replicating “[riskometers](#)” like those used in finance may overlook crucial sources of risk in the rapidly evolving tech landscape.

Collaborative learning is at the core of intelligence

Tech executives often advocate for self-regulation in order not to stifle innovation. However, effective and flexible regulation need not lead to stagnation, provided it avoids unnecessary complexity. Imperfect principles and rules that evolve and improve over time are undoubtedly much better than a complete absence of regulation.

If there is one lesson that the tech industry can learn from the financial sector, it is this: While it is not possible to eliminate or predict all risks, proactive and reactive regulations can co-exist harmoniously. Ultimately the key lies in continuously learning, adapting and improving. The recent advances in AI are built upon the power of (machine) learning, which is at the core of intelligence. It should come as no surprise to the AI and tech community that establishing deep learning processes might be the most crucial guiding principle for the regulation of technology as well.

Want more? [Register for our Tech Talk X](#) on 6 September entitled: *How to Regulate AI? Lessons from the Financial Sector*

Find article at

<https://knowledge.insead.edu/economics-finance/managing-systemic-risks-tech-lessons-finance>

About the author(s)

François Candelon is a Managing Director and Senior Partner at the Boston Consulting Group. He is also the Global Director of the BCG Henderson Institute.

David Zuluaga Martínez is a member of the Public Sector practice at Boston Consulting Group.

Peter Nathaniel conducts executive coaching and provides board development training and programmes for INSEAD. He is a former Adjunct Professor at INSEAD.

Theodoros Evgeniou is a Professor of Decision Sciences and Technology Management at INSEAD. He has been working on machine learning and AI for over 25 years.

Ludo Van der Heyden is the INSEAD Chaired Professor of Corporate Governance and Emeritus Professor of Technology and Operations Management. He is the founder of the INSEAD Corporate Governance Centre.

About the series

AI: Disruption and Adaptation

Delve deeper into developments in artificial intelligence, especially the disruptions across value chains. This series examines AI's impact on a range of sectors, including business consulting, education and the media. It also sizes up the regulatory and ethical questions tied to this game-changing technology.