Good AI, bad for banks?

Stefan Kerbl[1]

The rapid advancement of Artificial Intelligence (AI) is set to reshape the financial system. We synthesize the assessments of five reputable institutions – the IMF, BIS, FSB, BoE, and ECB – on how AI may pose risks to financial stability. The most frequently highlighted risks include (i) model opacity, (ii) herding behavior, (iii) cyber risks, and (iv) supplier concentration. We explore four scenarios based on varying speeds and capabilities of AI development. In a moderate scenario, market manipulation via deep fakes emerges as a key concern, while in a fast-paced scenario, credit risks stemming from layoffs and firm failures take precedence. We address the implications of artificial general intelligence and superintelligence (AGI, SI) and highlight the negative wealth effects in a scenario where AI fails to commercialize. Finally, we empirically examine AI's impact on banks by (i) using Nvidia stock as a proxy for AI advancements and (ii) conducting an event study on banks' stock price reactions to major AI breakthroughs. With regards to the US banks, we do not find statistically significant effects with either approach. For European banks, we find a positive albeit small reaction to Nvidia stock returns. This effect is more pronounced for large European banks, while for the smallest, we find negative effects. The event study finds insignificant responses to events for the largest European banks, but significant negative responses for the small European banks.

JEL classification: G14, G28, O33.
Keywords: artificial intelligence, technological change and financial stability, event study, cyber risk.

AI technology is poised to transform the financial sector. Given the financial sector's reliance on AI-compatible capital, labor and (non-physical) output —combined with its prevalence of large firms—it is unsurprising that finance is among the most AI-exposed industries (Aldasoro et al., 2024a; Felten et al., 2021). But internal adoption is only a small part of the effect of AI on finance. For once, the financial sector also plays a key role in funding AI development, given the capital and R&D expenditures required. Moreover, AI will impact banks' traditional debtors, including firms, households, and governments.

For these reasons—and others we will explore—reputable financial stability bodies are assessing the implications of AI's rapid development for the sector. In this paper, we summarize their views (section 1), present a scenario-structured analysis of key risks (section 2), and provide an empirical evaluation (section 3). While AI promises many benefits for banking and beyond, the nature of our discipline compels us to focus on its potential risks to financial stability.

[1] Oesterreichische Nationalbank, Financial Stability and Macroprudential Supervision Division, stefan.kerbl@oenb.at;
Opinions expressed by the authors of studies do not necessarily reflect the official viewpoint of the OeNB or the Eurosystem.
We thank Sebastian Rötzer, Christian Wipf, Stefan Schellerer, Franz Mohr, Philip Imhof and Martin Lackner for valuable suggestions and ideas.

1. View of bodies concerned with financial stability

In this section, we distill recent perspectives from major financial stability institutions. These include, in chronological order, the Bank of England (Bank of England & Financial Conduct Authority, 2022), the ECB (Leitner et al., 2024), the BIS (Aldasoro et al., 2024b), the IMF (IMF, 2024b), and the FSB (Financial Stability Board, 2024) [2]. A cross-reading of these contributions yields the following key observations:

1. **Consensus on AI's transformative potential and risks** – There is broad agreement on AI's potential to reshape finance, the risks it poses, and the need for financial stability bodies to address the topic.
2. **Little or no categorization of risks** – No standard has emerged on how to categorize risks emerging from AI to financial stability. Most do not group risks but list them without a particular order. An exception here is the BIS categorizing risks based on financial intermediation, insurance, asset management, and payments.
3. **Commonly identified risks** – Despite differences in focus, figure 1 summarizes the key risks highlighted across the reports. Model opacity, herding, and cyber risk emerge as core concerns in all contributions. We discuss these risks in more detail below.
4. **Limited consideration of AI's indirect effects** – The primary focus of these reports is AI adoption within the financial sector. Discussions on indirect AI-driven effects – such as transformations in non-financial industries – are largely confined to cyber risks and misinformation. Again, the BIS report is an exception, examining how AI-driven shifts in legal, non-financial economic activity might affect finance.[3] Notably, no study among the five explores the financial implications if AI fails to deliver expected returns. We argue that these indirect effects deserve more attention, given the financial sector's deep integration into the broader economy – a point we develop further in section 2.

[2] This is not to say that these are the only contributions, see e.g. Finance Watch (2025). They are not even the only contributions of these organizations – see the early contribution of the FSB 2017, or several further contributions by the BIS with varying focus (Aldasoro et al, 2024a, Aldasoro et al, 2024c, Gambacorta et al, 2024, Araujo et al, 2024). The ones we synthesize are all fairly recent, are of comparable length and come from respected bodies with a history of being at the forefront of the international debate on many issues of financial stability.

[3] This does not say that these are blind spots at these institutions, see e.g. other publications by the IMF (2023) and BIS (Aldasoro et al 2024a).

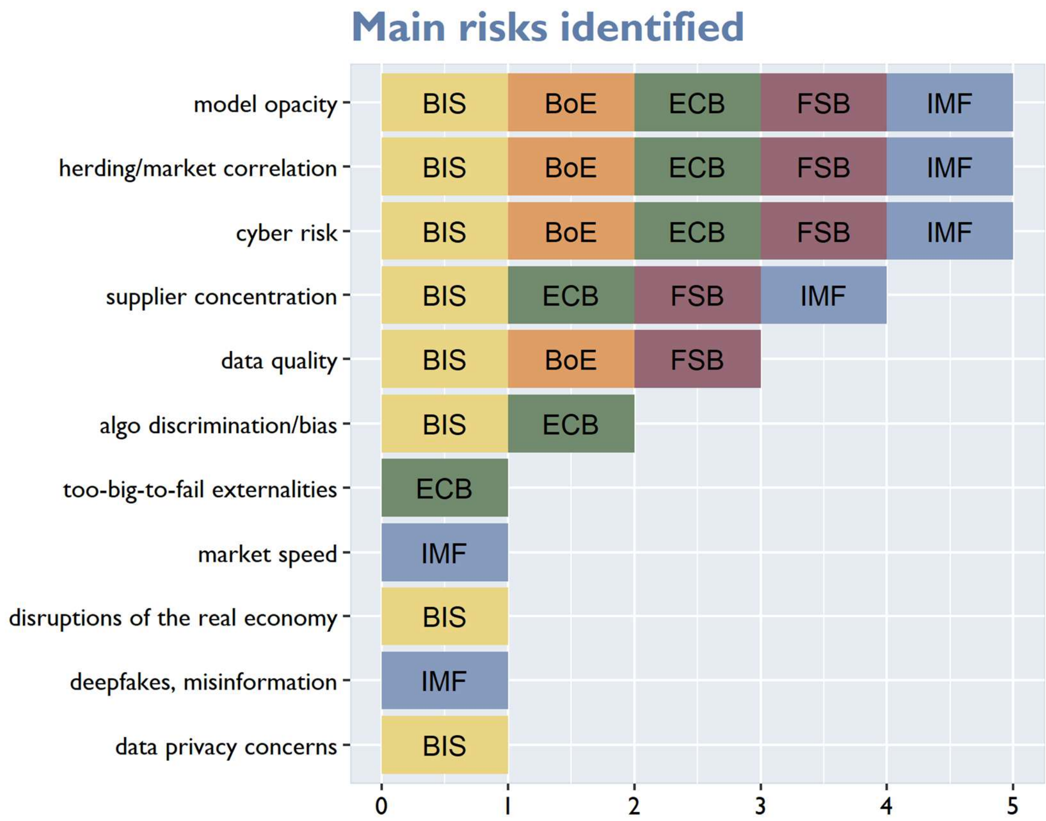

Figure 1: Overview of main risks by international bodies.

Below, we summarize the four primary risks identified across the reports. The counterarguments provided under "*Discussion*" are not meant to dismiss these risks but to critically examine their prominence among the most relevant concerns.

1. **Model Opacity** encompasses issues such as reduced predictability in banks' internal models (e.g., risk modeling, pricing, or asset allocation) and the lack of explainability in AI-generated outputs.

Discussion: While it is true that the mechanical operations even in mid-sized neural networks becomes hard to trace, recent advances in Large Language Models, esp. their reasoning capabilities and chain-of-thought technology, address the lack of explainability (Wei et al. 2022). Yet, even advanced models show pockets of surprising behavior and unexpected weak performance[4]. Would it be in the self-interest of financial institutions to

[4] E.g. the failure of GPT-4 to accurately answer "which number is larger 9.9 or 9.11", or the failure of 2025 image generators to produce "a wine glass filled to the brim". For more serious behavior see emergent misalignment in Betley et al (2025).

avoid opaque models when more interpretable alternatives exist? Certainly, supervisors have a role in emphasizing explainability over short-term prediction performance.

2. **Herding and Market Correlation.** If financial institutions rely on the same AI systems for asset allocation and investment decisions, systemic common exposures could arise. In turn, this would amplify market bubbles and trigger correlated responses to market corrections, exacerbating financial instability. The BIS highlights the 1987 U.S. stock market crash, where portfolio insurance strategies contributed to the largest single-day drop in U.S. stock history. Shin (2010) discusses modern hedging strategies with similar destabilizing effects. As a counter position, Shiller (2010) and Scharfstein & Stein (1990) emphasize the psychological and social dimensions of herding.

Discussion: Algorithmic trading and technical analysis have been widespread for decades (Gehrig & Menkhoff, 2006), so have common data sources for risk models and asset allocation, like Blackrock's Aladdin. Financial institutions have long used portfolio insurance strategies and stop-loss mechanisms, sometimes with unintended externalities for market stability. Whether AI will exacerbate these risks remains uncertain. Investors' pre-existing portfolio positions and risk appetites differ, leading to diverse decision-making even when using similar AI tools. More advanced AI models can be expected to pay greater attention to those differences, resulting in a variation in recommendations.

Additionally, a model-driven trade, like any trade, influences market prices, making identical subsequent trades less attractive. Given these factors – and the inherently social and psychological nature of herding – it is not evident that AI will significantly worsen the problem.

3. **Cyber risk** encompasses attacks where the attacker was supported by AI and attacks targeting AI-driven financial systems. One example of the former is using AI-generated code to create malware capable of infiltrating a bank's infrastructure. An example of the latter is jailbreaking an AI-powered customer service assistant to extract confidential client data (Wei et al., 2024). A third category of attacks, partly also filed under the umbrella of cyber risk, is "social hacking" where AI exploits human psychology and perpetrators use AI to personalize their phishing attacks ("spear phishing") to gain

unauthorized access to information or systems. We will address this last facet in section 2.1.

Discussion: The overlap between AI risk and cyber risk is sometimes overstated. Many cyber risks exist independently of AI, just as AI-related risks do not always involve cyber security threats. That said, AI-driven cyber risks will continue to evolve (Zou et al., 2023) and increase in importance with the wider deployment of AI in financial institutions. At least, AI will not only empower attackers but also enhance defensive cyber security capabilities.

4. **Supplier concentration** captures the risk that the AI software or the hardware required to train or host AI models might be in the hands of an oligopoly.[5] Such dependence creates risks due to limited substitutability if suppliers fail to deliver, as well as potential cost pressures arising from market power. Sitaraman & Narechania (2024) conclude that in semiconductor manufacturing, and to a lesser extent, cloud-infrastructure provision oligopolistic structures are particularly prevalent.

Discussion: The risks of IT system dependency are a central focus of the EU's Digital Operational Resilience Act, and reliance on AI providers is likely to grow. However, the number of AI model developers remains large and is expanding. As the IMF notes, the share of open-source AI models has risen from approximately 30% in 2021 to 80% in 2024. While an oligopoly in AI development is undesirable from an economic competition standpoint (Sitaraman & Narechania, 2024), its implications for financial stability are less clear. If all banks face higher AI-related costs due to supplier concentration, consumers would be worse off but not financial stability. Higher fixed costs may favor large banks with economies of scale, leading to greater financial sector concentration – what the ECB refers to as the "too-big-to-fail" externalities of AI. Our empirical insights from section 3 support this concern.

2. Risk grouping and AI scenarios with relevance for finance

[5] Note that higher supplier concentration would also fuel cyber risk and herding, but we classified both risks separately.

After reviewing the five papers, two approaches to grouping these risks emerges. The cross-section differentiation distinguishes first between internal AI use (ie. use by the financial sector entities) and external AI use. Second, internal use could be further divided according whether the risk exists in isolation (idiosyncratic) or only in conjunction with other financial entities (systemic), see figure 2. External use could be further divided into direct effects of external AI use and indirect, e.g. via credit events of debtors.

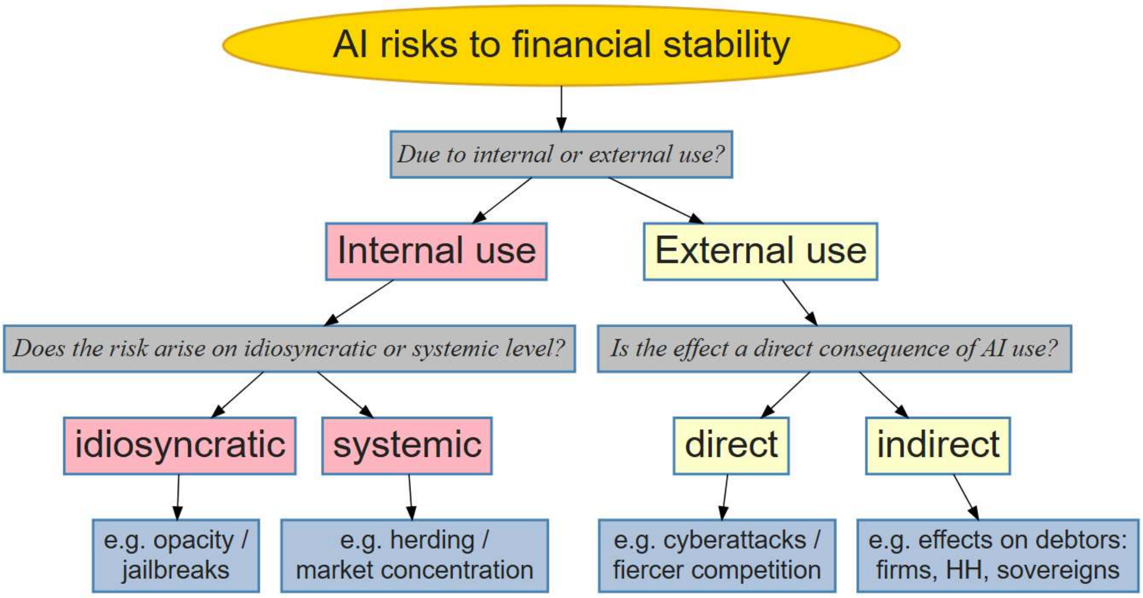

Figure 2:cross-section grouping of risks.

The second approach to group risks is a scenario-based approach, as recommended by Korinek & Suh (2025). Here, we structure risks according to four AI scenarios: 1. AI Moderate (our current state), 2. Fast and Powerful AI, 3. Artificial General Intelligence (AGI) and Superintelligence (SI) and 4. AI Winter.

2.1 Scenario 1: We already live in the AI moderate scenario.

History provides numerous examples where public comments by influential figures have triggered strong market reactions.[6] In this context, one of the most severe risks to financial markets is market manipulation via deepfakes. Even with current AI technology, it is possible to produce high-quality fake videos using limited hardware and labor. With greater computational

[6] E.g. the comments by Saudi National Bank, the top Credit Suisse shareholder, not to raise its stack in March 2023.

power and dedicated efforts of larger, skilled teams of experts, such fake content could become "information bombs" for financial markets. Consider, for example, a falsified recording of a private meeting in which a visibly distressed Chief Risk Officer informs top regulators of an unexpected portfolio loss at a major bank. Such a fabricated video could either be released publicly to go viral or – more insidiously – shared privately with select large shareholders. The latter approach would make it harder for both regulators and the targeted institution to trace the source of the market reaction. Other than market manipulation, social hacking is a serious concern. Already now, AI assists fraudsters in obtaining access to bank accounts or attract funds. With stronger AI, the manipulative potential increases and we might transition from social hacking to social engineering where AI exploits our psychological vulnerabilities in an ever more targeted way.

Mitigation: The best defense against market manipulation is strong capital and liquidity buffers at banks and non-bank financial institutions (NBFIs), ensuring they can withstand sudden market shocks. The expectation of a swift and credible response by a trusted regulatory body could discourage attempts at manipulation. Greater awareness of deepfake technology, along with AI-driven tools to detect and debunk fraudulent content, would be essential (Aldasoro et al 2024c).

2.2 Scenario 2: Fast and powerful AI

This scenario envisions a sudden and significant productivity shock – one in which AI capabilities rapidly advance, e.g. enabling autonomous agents to perform a broad range of cognitive tasks. If such an acceleration occurs within a short time frame, firms and labor markets may struggle to adapt, disrupting entire industries (Korinek, 2023). The scale of such potential transformation has been compared to the Industrial Revolution (Abis & Veldkamp, 2024).

In general, all productivity advances must lead to a distribution over three possible consequences:

1. Increased output with the same input (e.g., likely in healthcare, but unlikely for potato chip production where saturation is more prevalent).
2. Reduced input for the same output – either voluntarily (increased leisure) or involuntarily (unemployment and firm closures).
3. A shift in labor and capital to new sectors (structural transformation).

The faster AI technology advances and the more powerful it becomes, the more we will see layoffs and business failures, leading to surges in credit defaults in both corporate and household sectors. Furthermore, as the BIS points out, governments would face declining tax revenues while simultaneously incurring higher social security costs due to rising unemployment.

The impact of AI will not be uniform. It will vary across industries, skill levels, labor and capital, and national economies (Gambacorta et al., 2024; Makridis & Hickman, 2023; Korinek & Suh, 2024; Toner-Rodgers, 2024). Some societies will successfully channel the productivity shock into higher output and structural change, while others may struggle to maintain their economic stability. The resilience of a country's financial sector will largely depend on its capacity to adapt to this transformation.

While the IMF's AI Preparedness Index suggests that national rankings closely correlate with GDP per capita[7], Autonomous (2024) finds that EU banks currently lack a high degree of AI readiness. A key reason for this may be strict and ambiguous regulations, such as the EU AI Act and the General Data Protection Regulation (GDPR). These laws impose significant constraints on AI deployment across industries, slowing adoption. As a result, major tech firms – including Apple, Microsoft, OpenAI, Meta, and xAI – have delayed or indefinitely postponed rolling out certain AI features in the EU due to legal uncertainty. In a globalized economy, this regulatory burden places European firms at a comparative disadvantage – a gap that will only widen if AI advances rapidly. The consequences of such delays will be explored further in Section 3, where we empirically examine geographic differences in AI adoption.

Mitigation: Maintaining adequate capital buffers at banks and NBFIs will help absorb credit losses arising from AI-driven productivity shocks. But policymakers must also strike a balance: while AI safety regulations serve a legitimate purpose, overly restrictive frameworks can limit access to AI, hindering adaptation and productivity, and in the long run, potentially undermining financial stability.

[7] See IMF AI Preparedness Index under https://www.imf.org/external/datamapper/AI_PI@AIPI/ADVEC/EME/LIC. Spearman rank correlation ρ with GDP per capita is 0.89.

2.3 AGI / SI: Artificial General Intelligence or Superintelligence

There is no fundamental physical or theoretical limit preventing AI from surpassing human intelligence. Artificial General Intelligence (AGI) refers to AI systems that match or exceed human expert-level performance across all economically relevant cognitive tasks. Superintelligence (SI) represents an even more advanced scenario, where AI surpasses human intelligence by orders of magnitude – similar to the way humans surpass monkeys and monkeys surpass chickens. Both scenarios, while considered tail risks, have profound implications.

Chow et al. (2024) argue convincingly that if markets expected AGI or SI to emerge soon, interest rates would already be significantly higher today. The reasoning is straightforward: if such a transformative event were anticipated, consumption smoothing would lead to higher present-day consumption, driving up interest rates – regardless of whether AGI/SI is expected to be benevolent or malevolent.

Nevertheless, recent breakthroughs in AI have led experts to shorten their timelines for AGI (Dilmegani, 2025). If AGI were achieved, the historical analogy would not be the Industrial Revolution but rather the Manhattan Project, with increasing dominance of geopolitical over economic considerations (Aschenbrenner, 2024). A central issue would be the alignment problem – the challenge of ensuring that AGI and SI act in ways aligned with human values and goals.

Mitigation: In this scenario financial stability considerations would be overshadowed by alignment and geopolitical risks. We argue that AGI/SI governance falls outside the primary mandate of financial stability institutions. However, regulators should remain aware of its potential economic and systemic impact.

2.4 AI Winter

An AI winter traditionally denotes a period in which artificial intelligence technologies fail to translate into substantial productivity gains. Acemoglu (2025) reviews the fraction of tasks which are impacted by AI productivity gains and estimates a mere 0.66% increase in total factor productivity (TFP) over 10 years. This is strongly at odds with predictions made by e.g. leaders in the AI industry (see e.g. Aschenbrenner, 2024). One reason for the difference is that Acemoglu's approach does not account for scientific progress enabled by AI that feeds back to all professions. On the other hand, AI industry leaders might be biased due to (i) a self-selection bias (conviction leads to work in those fields, not the other way around), (ii) the availability-bias (their work consists of research and coding where AI is evidently stronger than in other fields), and (iii) a confirmation bias arising from their task in convincing investors and acquiring venture capital.

In any case, a more moderate variant of the AI winter scenario still poses potential risks to financial stability: AI might continue to advance but struggle to achieve effective commercialization. Technological breakthroughs do not inherently guarantee economic returns. At the time of writing, the development cost of current models is not covered by the subscription revenue *that that model generates*.[8] This means that AI labs are betting on current development cost reducing future development cost sufficiently to generate long-term value. However, AI innovations may exhibit characteristics of a public good – when they are difficult to patent or easily replicable. In such cases, AI could generate broad societal value without offering sufficient private returns to investors.

Geopolitical dynamics may further intensify these vulnerabilities. As Aschenbrenner (2024) highlights, the trained model parameters can potentially be copied, sold illicitly, or disseminated widely. Importantly, this outcome is not mutually exclusive with more transformative scenarios, including those described under the terms 2. Fast and powerful AI and 3.AGI/SI.

For finance, such a development is highly relevant, as it brings about a sharp correction in the valuation of AI companies and their suppliers, ranging from hardware to energy. Knock-on effects on finance would be via direct holdings of equity or debt instruments and indirect effects such as the wealth channel. Figure 3 left-panel depicts the capital expenditure (capex) of selected US tech giants. For 2025, Google, Meta and Microsoft are expected to spend more than 200bn USD. These figures and their growth rates are staggering but are indicative only. They overestimate true spending by including all capex spending (incl. not AI-related), but they underestimate investment in AI as they exclude key players such as Amazon, xAI, IBM, and Oracle, as well as research and development expenditures.

The right-panel of figure 3 underlines the importance of the wealth channel, should AI not return their investment. It shows the growing importance of corporate equities in US household

[8] Precise data is lacking but the (negative) difference is thought to be large. As the AI-sceptic opinion by Zitron (2024) puts it, "OpenAI currently spends $2.35 to make $1".

financial wealth (in red) and the combined market capitalization of Oracle, Meta, Google, Nvidia, Amazon and Microsoft as percentage of total US household financial assets (in blue). Again, these numbers are indicative only. A substantial part of the market cap of these companies is unrelated to AI, and US households do not hold all of it, yet they do hold equities of other AI-related firms. Despite these caveats, economic data points to the importance of the wealth channel, when considering the scenario of AI failing to return investment.

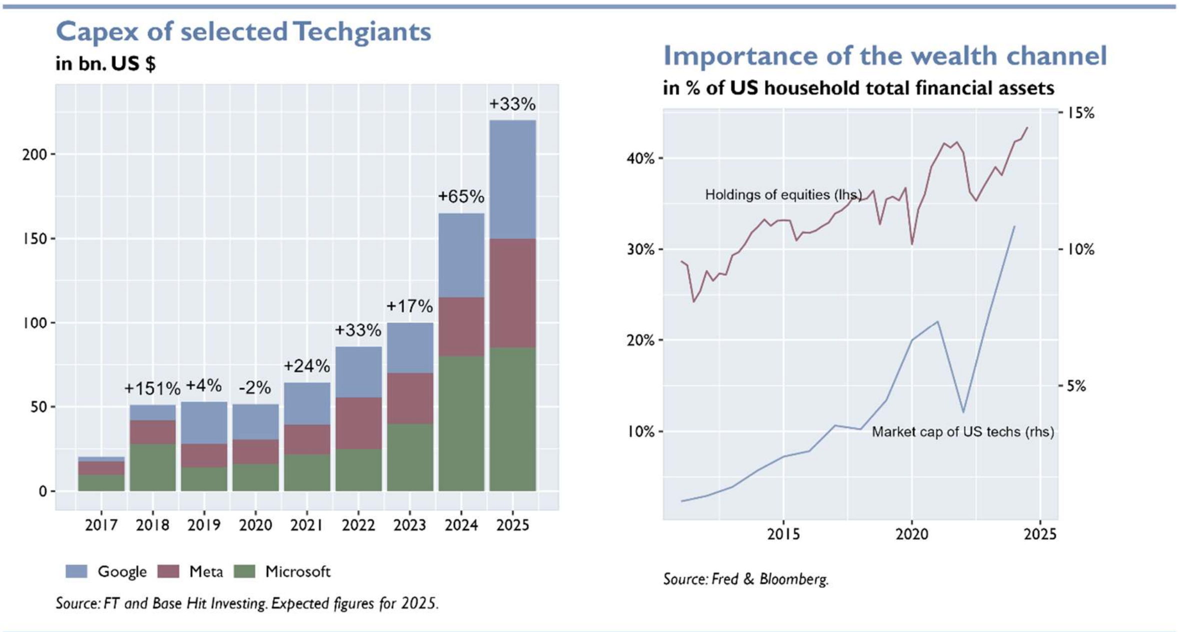

Figure 3: left-panel: Capex of selected tech giants. Data from Financial Times and Hit Investing. The year 2023 is available in both time series. We use FT data for this year as it is more recent. The alternative figure for 2023 is 13bn USD lower, 10bn USD thereof coming from lower capex of Microsoft. Right-panel: Households and Nonprofit Organizations directly and indirectly held corporate equities (lhs), sum of market cap of Oracle, Meta, Google, Nvidia, Amazon and Microsoft. Both as a percentage of US households total financial assets, Fred: TFAABSHNO.

Mitigation: As with previous scenarios, sufficient capital and liquidity buffers remain the most effective safeguard against systemic risks. Regulators should carefully monitor signs of speculative excess in AI investments to prevent financial instability.[9] Understanding the sources of AI financing is essential as the implications of an AI downturn will depend on whether risk-aware and

[9] See e.g. the 30bn USD valuation of a company that doesn't plan to release any products until it develops super intelligence (WSJ 2025).

resilient actors or weaker links within the financial system ultimately fund the AI boom. More clarity on this front should be a goal to financial stability institutions.

3. Empirical insights

In this section we study market moves to understand the effects of AI on banks, or at least the markets' expectation of the effect. How do advances in AI change the market's view of banks' future earnings? Are there differences between European and US banks? Are there differences between smaller and larger banks?

To this end, we use daily stock market data starting with 1st January 2018.[10] To quantify advances in AI we use changes to the Nvidia stock in section 3.1 and an event study in section 3.2. On the left-hand side we use daily returns of individual stocks of US and European banks. The first approach relies on the assumption that Nvidia stock (after controlling for confounders) is a workable proxy for AI advances and that the market gauges the further impact without delay. The second approach relaxes these assumptions but relies on the assumption that the chosen events are market surprises with regards to AI advancement and that its implications are transmitted through the market within several weeks.

3.1 Nvidia proxy

Using Nvidia stock returns as proxy we need to be mindful of possible confounders. First, we account for general market sentiment or economic news. Second, the performance of crypto assets might be another confounder, influencing both banks as representatives of the 'legacy' financial system and Nvidia stock, whose GPUs are used for mining. Initially as a robustness check only, we also control for Fama-French factors (Fama French, 1993) to account for outperformance of small versus big companies, and of high book/market versus low book/market companies. As this changed the results considerably, we chose to always include these controls.[11]

Thus, we estimate the model,

$$\text{stock}_{i,t} = \beta_i \, nvidia_t + X_{t,i}\, \gamma_i + \alpha_i + u_{i,t} \qquad (1)$$

[10] The data comes from investing.com and yahoo finance. The data ranges to 14th February 2025 for the European market and to 31st December 2024 for the US market.

[11] This is a relevant finding on its own and an example of research considerably changed by careful robustness checks. We also use the Fama-French Factors in the event study.

with $stock_{i,t}$ is the daily return of bank' equity i at time t, nvidia_t the return of Nvidia's stock, $\boldsymbol{X}_{t,i}$ the vector of control variables encompassing the Fama-French three factors (separated into US and Europe) and the daily bitcoin return. β_i is the parameter of interest.

We run regression (1) on a large sample of bank stocks. We start with the largest 100 banks in the US and exclude those that are not publicly traded and those who are foreign subsidiaries. For Europe we start with all banks in the EBA transparency exercise, add the largest UK banks, and again filter for being publicly traded. Table 3 in the annex provides a list. In addition to the geographical breakdown, we classify banks with total assets lower than the 75th quantile (around 400bn EUR) as "small", and the others as "large". After running separate regressions for the 75 US and 54 European banks, we calculate a weighed mean over β_i and an according standard error.[12]

	Europe	US
total		
	0.021 ***	−0.0020
	(0.0035)	(0.0038)
small		
	0.008 **	0.0015
	(0.0037)	(0.0033)
large		
	0.024 ***	−0.0029
	(0.0042)	(0.0047)

Table 1: Regression results: Weighted mean of β and robust standard errors. *, **, and *** indicate significance at the 10%, 5%, and 1% level respectively

Table 1 presents the weighted means of coefficients for Europe and the US together with robust standard errors in parentheses. While the US estimates are insignificant for all size categories, the estimates for European banks are positive and statistically significant. Economically, the coefficients are small in both cases. A 1% increase in our AI proxy translates to a 0.021% increase

[12] Note that this is standard procedure to calculate sector betas. We use heteroskedasticity and autocorrelation robust standard errors. Inspecting individual regressions, we did not observe severe issues with autocorrelated residuals.

in European bank (weighted) equities after controlling for Fama-French factors and changes in bitcoin.

Interestingly, we find substantially larger effects for big banks compared to small banks in Europe. To further analyze bank-level heterogeneity we regress our bank-specific coefficients on (logarithmic) bank size. As table 2 reports, we indeed find that larger banks are more positively affected by our AI proxy. This finding is significant both for the total sample of banks and for Europe separately, but not for US banks separately, though the sign of the coefficient still points in the same direction.

	Intercept		log(TotalAssets)		n	R^2
total						
	-0.0661	***	0.0059	***	129	8.59
	(0.0234)		(0.0021)			
Europe						
	-0.091	**	0.0087	**	50	16.66
	(0.0401)		(0.0036)			
US						
	-0.0436	*	0.0035		75	2.91
	(0.0261)		(0.0022)			

Table 2: Coefficients regressing the response to Nvidia stock of individual bank stocks on banking size. Robust standard errors in parenthesis, *, **,
and *** indicate significance at the 10%, 5%, and 1% level respectively.

As the intercept is negative, the estimated effect of our AI proxy on the smallest banks is negative too. Based on our results banks smaller than 73bn EUR total assets (total sample) or 35bn EUR (European sample) are expected to see a decline in their share price with an increase of our AI proxy.

3.2 Event study

Using Nvidia as proxy in section 3.1 relies on the assumption that the Nvidia stock change (net of general market and bitcoin development) is a good indicator of AI news. This section presents another approach that does not rely on this assumption. However, we assume here that product releases are news to the market and that the market incorporates their implications within a given timeframe – an assumption that is not necessary in the approach taken in section 3.1.

We survey a small panel of experts with the question which events in the past five years were the most positive AI surprises or unexpected breakthroughs. The panel quickly converged on four main events: the release of GPT-3.5 in November 2022, the release of GPT-4 in March 2023, the preview of o1 in September 2024 and the release of DeepSeek-R1 in January 2025.

The technical procedure follows the established literature of event studies (MacKinlay, 1997). First, we use data on individual bank equities to construct four indexes (Europe small, Europe large, US small and US large). Second, for each point in time, we estimate the relation between the general market (Fama-French factors) and our bank indexes in the three months before the event. We then use these relations to calculate the expected returns to bank stocks given the general market move in the two weeks following the event. The differences between these expected returns (based on the previously observed co-movement of the general market and bank indexes) and the realized bank stock returns are the abnormal returns (MacKinlay, 1997). Formally,

$$AR_{j,t} = R_{j,t} - \mathbb{E}\left(R_{j,t} \middle| market_{j,t}\right)$$

where $AR_{j,t}, R_{j,t}$ and $\mathbb{E}\left(R_{j,t} \middle| market_{j,t}\right)$ are the abnormal, realized and expected returns respectively. We calculate the abnormal returns as cumulative returns from 1 to 15 trading days after the event.

Figure 4 depicts the abnormal returns for each of the four indices and for each of the four events[13]. The grey shaded areas are the 5%-95% bands of all time points in our data sample. We see that the events mainly did not trigger a significant market reaction. Only after the release of GPT4 small European banks showed significantly negative abnormal returns and small US banks significantly positive abnormal returns.

How likely is the combined realization across all four events for each index? For this purpose, we compute the p-values for each event and produce a combined p-value (for each index and time

[13] Note that our data for the US currently ends 2024-12. Therefore, the deepseek event is not available for the US as of now.

horizon) under the null that these p-values are uniformly distributed. Values below <0.01 are marked with * in figure 4. Combining the market realizations, the positive response of small US bank vanishes, as their response is negative in the other events. However, the significance of the reaction of small European banks increases as all four events triggered negative abnormal returns. We find significant negative abnormal returns (p<0.01) for small European banks at the horizons 8, 11, 13 and 14 trading days after the event, and also for large European banks at 8 trading days afterwards. At other horizons, small European banks still show significant negative abnormal returns at the 5% or 10% significant level.

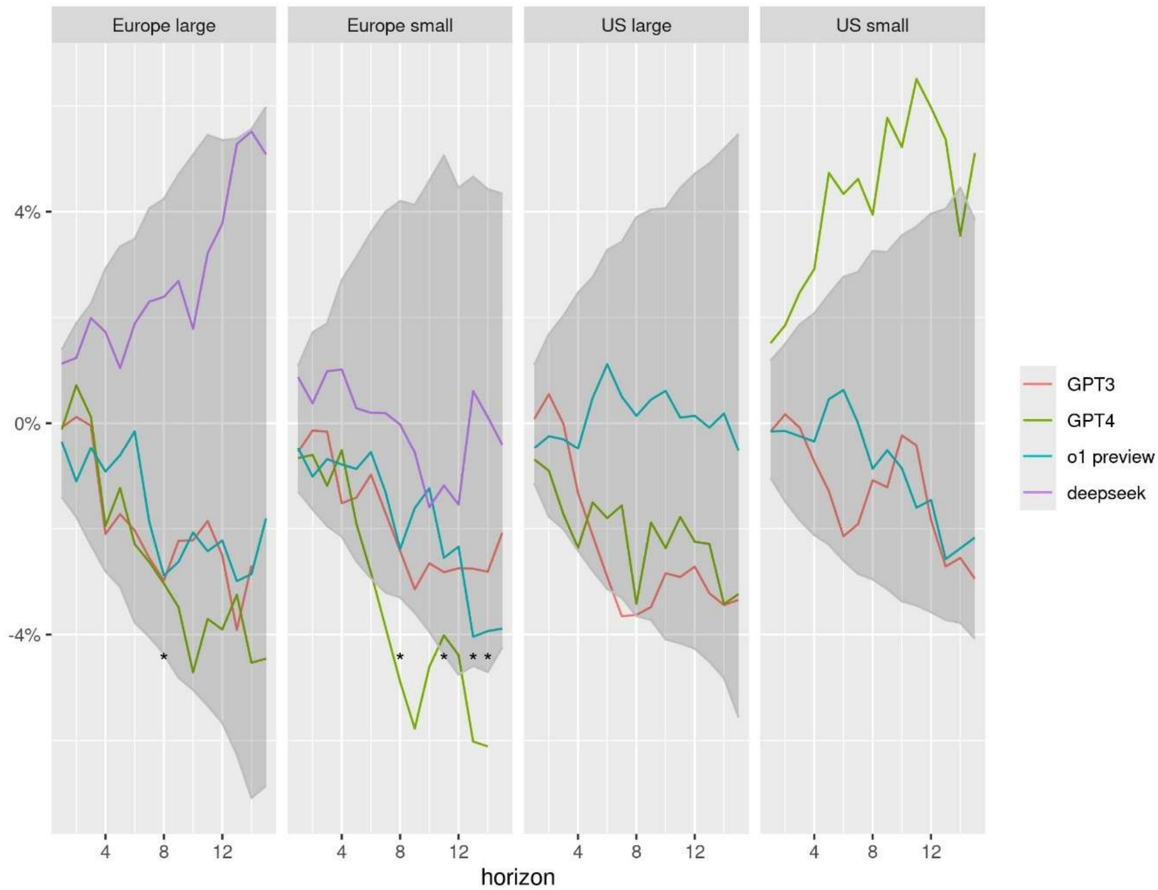

Figure 4: Abnormal returns in the event study. Grey shaded areas are the 5%-95% bands of all time points in our data sample. * mark combined levels of the four events beyond the 1%-significance level.

What do we make of the empirical findings?

We used Nvidia stock returns as proxy in section 3.1 and an event study with positive AI-surprises in section 3.2. With regards to the US banks, we do not find statistically significant effects with either approach. For European banks, the results are mixed. In section 3.1 we find a positive albeit small reaction to Nvidia stock returns. This effect is more pronounced for large European banks, while for the smallest, we find negative effects. Section 3.2 finds insignificant responses to events for the largest European banks, but significant negative responses for the smaller European banks. We conclude that the market too is confronted with uncertainty concerning the effects of AI on banks. For now, there does not seem to be a clear expectation on the effects on both side of the Atlantic. For Europe, though, there is evidence that large banks may benefit from AI advancement, while for smaller banks the opposite is true. One reason for this finding might be that Europe has more miles to go in terms of cost efficiency (where AI might help) and in terms of banking sector consolidation than the US.

4. Conclusion

As Aldasoro et al (2024b) note, finance has historically been among the earliest adopters of many new technologies, like the telegram, the computer and the internet. AI is no exception. Given its reliance on data-heavy information processing, financial intermediation is particularly well-suited to AI-driven enhancements. However, the risks associated with AI extend far beyond internal applications. A comprehensive approach to financial stability must also account for adversarial AI use, the economic disruptions AI may impose on borrowers, and the potential fallout from a speculative AI-driven market bubble.

International bodies tasked with financial stability stress the risks of model opacity, herding, cyber risk and supplier concentration. Our approach categorizes risks according to four scenarios. In the moderate scenario, which is essentially currently, deepfakes might deliberately destabilize markets. The more powerful AI becomes and the faster it does so, indirect effects on banks' debtors will dominate. We argue that a society adaptive to the new technology is in a better position to channel the productivity growth into more output or into a shift of output towards other sectors. Ultimately, such a stance will benefit financial stability. In the unlikely case that AI fails to commercialize, we must expect direct losses on finance (although the lack of data inhibits quantitative estimates) as well as indirect effects e.g. via the wealth channel.

In conclusion, there is no such thing as *the one* AI risk to financial stability. As we have seen, AI impacts finance through many channels which are best described as new sources of well-known risks to banks and NBFIs, such as market risk, credit risk, operational risk and liquidity risk. The fundamental safeguards of financial resilience – adequate capital buffers, ample liquidity, and a robust resolution framework – will remain critical. However, as AI adoption expands, the systemic nature of these risks may intensify, with a growing share of the financial system exposed to simultaneous AI-driven disruptions. Recognizing and addressing these evolving dynamics will be essential to ensuring financial stability in an AI-powered future.

References

Abis, S., & Veldkamp, L. (2024). *The changing economics of knowledge production*. The Review of Financial Studies, 37(1). https://doi.org/10.1093/rfs/hhad059

Acemoglu, D. (2025). *The simple macroeconomics of AI***,** Economic Policy, Volume 40, Issue 121, January 2025, Pages 13–58, https://doi.org/10.1093/epolic/eiae042

Aldasoro, I., Doerr, S., Gambacorta, L., & Rees, D. (2024a). *The impact of artificial intelligence on output and inflation* (BIS Working Paper No. 1179). Bank for International Settlements. https://www.bis.org/publ/work1179.pdf

Aldasoro, I., Gambacorta, L., Korinek, A., Shreeti, V., & Stein, M. (2024b). *Intelligent financial system: How AI is transforming finance* (BIS Working Papers No. 1194). Bank for International Settlements. https://www.bis.org/publ/work1194.htm

Aldasoro, I., Doerr, S., Gambacorta, L., Notra, S., Oliviero, T., & Whyte, D. (2024c). *Generative artificial intelligence and cyber security in central banking* (BIS Papers No. 145). Bank for International Settlements. https://www.bis.org/publ/bppdf/bispap145.pdf

Araujo, D., Doerr, S., Gambacorta, L., & Tissot, B. (2024). *Artificial intelligence in central banking* (BIS Bulletins No. 84). Bank for International Settlements. https://www.bis.org/publ/bisbull84.pdf

Autonomous. 2024. *Banks' AI Readiness Index*, European Retail Banks, 11 December 2024

Aschenbrenner, L. (2024). *Situational awareness: The decade ahead***.** Retrieved June 2024 from https://situational-awareness.ai/.

Bank of England & Financial Conduct Authority. (2022). *AI Public-Private Forum final report*. Bank of England.

Betley, J., Tan, D., Warncke, N., Sztyber-Betley, A., Bao, X., Soto, M., Labenz, N., & Evans, O. (2025). *Emergent misalignment: Narrow finetuning can produce broadly misaligned LLMs*. arXiv. https://arxiv.org/abs/2502.17424

Dilmegani, C. (2025). *When will singularity happen? 995 experts' opinions on AGI.* AI Multiple. Retrieved March 10, 2025, from https://research.aimultiple.com/artificial-general-intelligence-singularity-timing/

Eurostat. (2025). *Usage of AI technologies increasing in EU enterprises*. European Commission. https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20250123-3

Fama, E. F. & French, K. R. (1993). "Common risk factors in the returns on stocks and bonds". Journal of Financial Economics. 33: 3–56. CiteSeerX 10.1.1.139.5892. doi:10.1016/0304-405X(93)90023-5.

Fernández, P. (2009). *Betas used by professors: A survey with 2,500 answers* (Working Paper No. WP-822). IESE Business School, University of Navarra.

Felten, E., Raj, M., & Seamans, R. (2021). *Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses*. Strategic Management Journal, 42(12), 2195–2217. https://doi.org/10.1002/smj.3286

Finance Watch. (2025). *Artificial intelligence in finance report*. Finance Watch. March 2025. Link.

Financial Stability Board. (2017). *Market developments and financial stability implications of artificial intelligence (AI) and machine learning (ML).* https://www.fsb.org/wp-content/uploads/P011117.pdf

Financial Stability Board. (2024). *The financial stability implications of artificial intelligence*. November 14**.**

Hoffmann, J., Borgeaud, S., Mensch, A., Buchatskaya, E., Cai, T., Rutherford, E., ... & Sifre, L. (2022). Training Compute-Optimal Large Language Models. Advances in Neural Information Processing Systems, 35, 12776-12788.

IIF (Institute of International Finance) & Ernst & Young. 2024. "Annual Survey Report on AI/ML Use in Financial Services." Washington, DC.

IMF. (2023). Artificial Intelligence What AI means for economics. *Finance & Development*, December 2023. International Monetary Fund. Link.

Gambacorta, L., Qiu, H., Shan, S., & Rees, D. M. (2024). *Generative AI and labour productivity: A field experiment on coding* (BIS Working Paper No. 1208). Bank for International Settlements. https://www.bis.org/publ/work1208.pdf

Gehrig, Menkhoff (2006) Extended Evidence on the Use of Technical Analysis in the Foreign Exchange Market. *International Journal of Finance & Economics*, 11(4), 327-338.

IMF. (2023). *Artificial intelligence What AI means for economics.* In Finance & Development, December 2023. International Monetary Fund.

IMF. (2024a). *Cyber risk: A growing concern for macro-financial stability*. In Global Financial Stability Report, April 2024 (Chapter 3). International Monetary Fund.

IMF. (2024b). *Advances in artificial intelligence: Implications for capital market activities*. In Global Financial Stability Report, October 2024 (Chapter 3). International Monetary Fund.

Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J., & Amodei, D. (2020). Scaling Laws for Neural Language Models. arXiv:2001.08361.

Korinek, A. (2023). Scenario Planning for an AGI Future. *Finance & Development*, December 2023. International Monetary Fund. Link.

Korinek, A., & Suh, D. (2024). Scenarios for the transition to AGI (NBER Working Paper No. 32255). National Bureau of Economic Research. https://www.nber.org/papers/w32255

Leitner, G., Singh, J., van der Kraaij, A., & Zsámboki, B. (2024). *The rise of artificial intelligence: Benefits and risks for financial stability*. European Central Bank. Link.

MacKinlay, A. C. (1997). *Event studies in economics and finance*. *Journal of Economic Literature, 35*(1), 13–39.

Makridis, C., & Hickman, L. (2023). *Beyond skills: The role of personality traits in shaping economic returns amid technological change*. SSRN. https://ssrn.com/abstract=4365176

Scharfstein, D. and Stein, C. (1990), *Herd behaviour and investment,* The American economic review, 80(3): 465-479, June.

Shiller, Robert J., 2000, *Irrational Exuberance*, Princeton, NJ: Princeton University Press.

Shin, H. S. (2010). *Risk and Liquidity*. Oxford University Press.

Sitaraman, G. and Narechania, T. N. (2024). *An Antimonopoly Approach to Governing Artificial Intelligence*. 43 Yale Law and Policy Review 95 (2024), Vanderbilt Law Research Paper No. 24-8.

Toner-Rodgers, A. (2024). *Artificial Intelligence*, Scientific Discovery, and Product Innovation. arXiv:2412.17866.

WSJ. (2025). *This scientist left OpenAI last year. His startup is already worth $30 billion. The Wall Street Journal*. March 5 2025. https://www.wsj.com/tech/ai/ai-safe-superintelligence-startup-ilya-sutskever-openai-2335259b

Wei, A., Haghtalab, N., & Steinhardt, J. (2024). *Jailbroken: How Does LLM Safety Training Fail?* Advances in Neural Information Processing Systems, 36.

Wei, J., Wang, X., Schuurmans, D., Bosma, M., Chi, E., Le, Q., & Zhou, D. (2022). *Chain of Thought Prompting Elicits Reasoning in Large Language Models.* arXiv preprint arXiv:2201.11903.

Zitron, E. (2024). OpenAI is a bad business. https://www.wheresyoured.at/oai-business/

Zou, A., Wang, Z., Carlini, N., Nasr, M., Kolter, J. Z., & Fredrikson, M. (2023). *Universal and Transferable Adversarial Attacks on Aligned Language Models. arXiv preprint arXiv:2307.15043.* Link.

Appendix – List of banks

European banks		US banks	
Name	Country	Name	City, State
ABN AMRO Bank N.V.	Netherlands	Ally Financial	Detroit, Michigan
AIB Group plc	Ireland	American Express	New York City
ALPHA SERVICES AND HOLDINGS S.A.	Greece	Ameriprise	Minneapolis, Minnesota
AS LHV Group	Estonia	Ameris Bancorp	Atlanta, Georgia
Akcinė bendrovė Šiaulių bankas	Lithuania	Associated Banc-Corp	Green Bay, Wisconsin
Arion banki hf	Iceland	Atlantic Union Bank	Richmond, Virginia
BANCA MEDIOLANUM S.P.A.	Italy	Axos Financial	Las Vegas, Nevada
BANCO BPM SOCIETA' PER AZIONI	Italy	BOK Financial Corporation	Tulsa, Oklahoma
BAWAG Group AG	Austria	Banc of California	Los Angeles, California
BNP Paribas	France	Bank of America	Charlotte, North Carolina
BPER Banca S.p.A.	Italy	Bank of Hawaii	Honolulu, Hawaii
Banca Monte dei Paschi di Siena S.p.A.	Italy	BankUnited	Miami Lakes, Florida
Banco Bilbao Vizcaya Argentaria, S.A.	Spain	Capital One	McLean, Virginia
Banco Comercial Português, SA	Portugal	Cathay Bank	Los Angeles, California
Banco Santander, S.A.	Spain	Charles Schwab Corporation	Westlake, Texas
Banco de Sabadell, S.A.	Spain	Citigroup	New York City
Bank Polska Kasa Opieki S.A.	Poland	Citizens Financial Group	Providence, Rhode Island
Bank of Ireland Group plc	Ireland	Columbia Bank	Tacoma, Washington
Bankinter, S.A.	Spain	Comerica	Dallas, Texas
Barclays PLC	United Kingdom	Commerce Bancshares	Kansas City, Missouri
COMMERZBANK AG	Germany	Cullen/Frost Bankers, Inc.	San Antonio, Texas
CaixaBank, S.A.	Spain	Customers Bancorp	Wyomissing, Pennsylvania
DEUTSCHE BANK AKTIENGESELLSCHAFT	Germany	Discover Financial	Riverwoods, Illinois
DNB BANK ASA	Norway	East West Bank	Pasadena, California
Danske Bank A/S	Denmark	Eastern Bank	Boston, Massachusetts
Deutsche Pfandbriefbank AG	Germany	FNB Corporation	Pittsburgh, Pennsylvania
Erste Group Bank AG	Austria	Fifth Third Bank	Cincinnati, Ohio
Eurobank Ergasias Services and Holdings S.A.	Greece	First BanCorp	San Juan, Puerto Rico
FINECOBANK S.P.A.	Italy	First Citizens BancShares	Raleigh, North Carolina
Groupe Crédit Agricole	France	First Hawaiian Bank	Honolulu, Hawaii
HSBC Holdings PLC	United Kingdom	First Horizon National Corporation	Memphis, Tennessee
ING Groep N.V.	Netherlands	First Interstate BancSystem	Billings, Montana
Intesa Sanpaolo S.p.A.	Italy	Fulton Financial Corporation	Lancaster, Pennsylvania
Jyske Bank A/S	Denmark	Glacier Bancorp	Kalispell, Montana
KBC Groep	Belgium	Goldman Sachs	New York City
Liechtensteinische Landesbank AG	Liechtenstein	Hancock Whitney	Gulfport, Mississippi
Lloyds Banking Group PLC	United Kingdom	Home BancShares	Conway, Arkansas
Mediobanca - Banca di Credito Finanziario S.p.A.	Italy	Huntington Bancshares	Columbus, Ohio
NatWest Group PLC	United Kingdom	Independent Bank	Rockland, Massachusetts
National Bank of Greece, S.A.	Greece	JPMorgan Chase	New York City
Nordea Bank Abp	Finland	John Deere Bank	Reno, Nevada
Nova Ljubljanska Banka d.d., Ljubljana	Slovenia	KeyCorp	Cleveland, Ohio
OTP-csoport	Hungary	M&T Bank	Buffalo, New York
Piraeus Financial Holdings	Greece	Morgan Stanley	New York City
Powszechna Kasa Oszczednosci Bank Polski S.A.	Poland	Northern Trust	Chicago, Illinois
Raiffeisen Bank International AG	Austria	Old National Bank	Evansville, Indiana
SPAREBANK 1 SR-BANK ASA	Norway	PNC Financial Services	Pittsburgh, Pennsylvania
Skandinaviska Enskilda Banken - gruppen	Sweden	Pinnacle Financial Partners	Nashville, Tennessee
Société Générale S.A.	France	Popular, Inc.	San Juan, Puerto Rico
Svenska Handelsbanken - gruppen	Sweden	Prosperity Bancshares	Houston, Texas
Swedbank - Grupp	Sweden	Provident Bank of New Jersey	Jersey City, New Jersey
UNICREDIT, SOCIETA' PER AZIONI	Italy	Raymond James Financial	St. Petersburg, Florida
Unicaja Banco, S.A.	Spain	Regions Financial Corporation	Birmingham, Alabama
Íslandsbanki hf.	Iceland	Simmons Bank	Pine Bluff, Arkansas
		SoFi	San Francisco, California
		South State Bank	Winter Haven, Florida
		State Street Corporation	Boston, Massachusetts
		Stifel	St. Louis, Missouri
		Synchrony Financial	Stamford, Connecticut
		Synovus	Columbus, Georgia
		Texas Capital Bank	Dallas, Texas
		The Bank of New York Mellon	New York City
		Truist Financial	Charlotte, North Carolina
		U.S. Bancorp	Minneapolis, Minnesota
		UMB Financial Corporation	Kansas City, Missouri
		United Bank (West Virginia)	Charleston, West Virginia
		United Community Bank	Greenville, South Carolina
		Valley Bank	Wayne, New Jersey
		WSFS Bank	Wilmington, Delaware
		WaFd Bank	Seattle, Washington
		Webster Bank	Stamford, Connecticut
		Wells Fargo	San Francisco, California
		WesBanco	Wheeling, West Virginia
		Western Alliance Bancorporation	Phoenix, Arizona
		Wintrust Financial	Rosemont, Illinois

Table 3: List of banks in the empirical section.