



BANK OF CANADA
BANQUE DU CANADA

Staff working paper / Document de travail du personnel—2026-40

Last updated: October 16, 2026

Artificial Intelligence and Financial Tail Risks

Javier Ojea-Ferreiro

Financial Stability Department
Bank of Canada
jojaferreiro@bank-banque-canada.ca

Andrea Ugolini

Department of Economics, Management and
Statistics (DEMS)
University of Milano-Bicocca
andrea.ugolini@unimib.it


Juan C. Reboredo

Department of Economics
University of Santiago de Compostela
juancarlos.reboredo@usc.es

Bank of Canada staff research is produced independently from the Bank's Governing Council and may support or challenge prevailing views. The views expressed in this paper are solely those of the authors and may differ from official Bank of Canada positions. No responsibility for them should be attributed to the Bank.

DOI: <https://doi.org/10.34989/swp-2026-40> | ISSN 1701-9397

© 2026 Bank of Canada



Artificial Intelligence and Financial Tail Risks

Javier Ojea-Ferreiro¹, Andrea Ugolini² and Juan C. Reboredo³

¹Bank of Canada, jojeaferreiro@bank-banque-canada.ca

²Department of Economics, Management and Statistics (DEMS), University of Milano-Bicocca, Italy, andrea.ugolini@unimib.it

³Department of Economics, University of Santiago de Compostela, Spain, juancarlos.reboredo@usc.es

October 2026

Artificial intelligence (AI) has a double-edged relationship with financial stability: it may improve efficiency and risk management, while also creating new vulnerabilities and amplification channels. This paper studies whether market-perceived exposure to AI is reflected in the tails of financial-sector returns, and whether this relationship depends on macroeconomic conditions. We develop a framework that decomposes financial-sector tail outcomes across contemporaneous AI and macroeconomic scenarios, and identifies whether the AI-related episodes are primarily associated with extreme gains or losses. Using daily equity returns for financial sectors in Canada, the United States, Europe, the United Kingdom, and China, we document substantial heterogeneity. Financial-sector-specific scenarios dominate tail outcomes across markets, while the contribution of AI-related scenarios to return asymmetry differs substantially across countries. AI-sector scenarios contribute to downside asymmetry in Europe, China, and especially the United Kingdom, but make only a limited contribution in Canada and have the smallest tail weights in the United States. China is the only market exhibiting positive average skewness, driven by the financial and systemic scenarios rather than by AI-related outcomes. Overall, AI-related market movements are neither uniformly stabilizing nor destabilizing; their financial-stability interpretation depends on the market considered and on whether AI-sector tail events occur in isolation or jointly with broader macroeconomic movements.

JEL Classification: G01, G21, G32, C58

Keywords: Artificial intelligence, systemic risk, joint tail risk, financial stability

We would like to express our gratitude to Angel León, Thibaut Duprey, Radoslav Raykov and Robert Hill for their valuable comments and suggestions. We would also like to thank seminar participants at the Bank of Canada (RTM) and C.r.e.d.i.t. 2026 conference in Venice. The mistakes are, of course, our own. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Bank of Canada.

1. Introduction

Artificial Intelligence (AI) is driving a disruptive transformation of the financial system. This transformation may improve financial-sector performance by enhancing the efficiency and quality of financial decision-making, but it may also introduce or amplify vulnerabilities.¹ Its overall implications for financial stability are therefore ambiguous. This raises the question of whether AI-related developments are associated with the extreme gains and losses of the financial sector, and whether this association depends on the broader macroeconomic conditions in which they occur.

AI-related vulnerabilities may propagate into systemic risk through several transmission channels. FSB (2024) groups these channels into four broad categories: third-party dependencies and service-provider concentration, market correlations, cyber risk, and model risk related to data quality and governance. First, financial institutions may become increasingly dependent on a small number of providers of cloud infrastructure, specialized hardware and data services, creating common operational points of failure and increasing the potential for disruptions at key service providers to propagate across the financial system (FSB 2024; Cecchetti et al. 2025).² Second, the use of similar AI models, data sources, and automated decision rules may contribute to a convergence in how market participants process information, perceive risk, and manage exposures. This convergence can encourage more correlated and procyclical behaviour, thereby amplifying herding, liquidity pressures, and asset-price vulnerabilities (Danielsson et al. 2022; Aldasoro et al. 2025; Leitner et al. 2024). Third, AI may strengthen the capabilities of malicious actors by reducing the time and cost required to identify and exploit vulnerabilities in financial institutions' digital infrastructure, increasing cyber risk, fraud, scams, and disinformation (Danielsson and Uthemann 2025; FSB 2024; Lopez 2026).³ Fourth, AI systems may increase model-risk and governance concerns when their outputs are difficult to explain, rely on sensitive or opaque data, or generate inaccurate responses such as hallucinations (Fernández 2019; Crisóstomo Ayala and Mykhalyuk 2026). These

¹AI-driven technologies promise substantial efficiency gains in trading (Chen and Ren 2022; Sheng et al. 2025; Lopez-Lira and Tang 2026), portfolio optimization (Gu et al. 2020; Bianchi et al. 2021), fraud identification (Zheng et al. 2025), risk management (Liu and Pun 2022; Zhang et al. 2026; Hansen and Lee 2025; Aldasoro and Desai 2025), and information processing (Gambacorta et al. 2025; Fuster et al. 2019; Berg et al. 2020; Oehler and Horn 2024). More broadly, AI is increasingly viewed as a potential driver of productivity growth, innovation and structural transformation in financial services (Siddik et al. 2025; Han and Meng 2025; Lin et al. 2025; Ca'Zorzi et al. 2025; Rammer et al. 2022). For recent policy discussions emphasizing productivity and structural change channels, see the Bank of Canada's speeches by Deputy Governor Michelle Alexopoulos and Governor Tiff Macklem.

²Recently, Rishabh and Shreeti (2026) document that AI production is geographically concentrated across economies and supply-chain layers, reinforcing the relevance of concentration risks in the AI ecosystem.

³See recent IMF note, speech by the ECB representative Pedro Machado and interview to the ECB representative Frank Elderson.

risks may be particularly relevant in customer-facing applications, where AI can improve financial advice and client interaction (Oehler and Horn 2024), but where privacy failures, misleading recommendations, or unreliable outputs may damage financial firms' reputation (Weitzner 2024) and weaken client relationships (Choi et al. 2025). These channels suggest that AI can become relevant for systemic risk when its effects are common across institutions, amplified through markets, or interact with existing financial and macroeconomic vulnerabilities.

This paper asks what share of financial-sector extreme losses and gains is associated with contemporaneous movements in market-perceived AI exposure, and whether the contribution of these AI-related episodes to financial tail outcomes differs when they occur in isolation or jointly with broader macroeconomic conditions. This distinction is crucial as the interaction between AI-related dynamics and broader macroeconomic conditions is particularly important because AI can affect the financial system through the real economy as well as through financial markets. By increasing productivity, profitability, and resilience in some sectors, AI may improve firms' creditworthiness and reduce default risk; by contrast, sectors that are slower to adapt may face weaker profitability, lower valuations, or higher financial vulnerability. In parallel, the growing concentration of equity market capitalization in a small number of AI-intensive technology firms may blur the distinction between AI-specific performance and aggregate market performance, as broad equity indices become more sensitive to the valuation of AI-intensive firms (Potrykus 2024; Basele et al. 2025). This, in turn, can increase financial institutions' exposure to AI-related shocks through portfolio valuations, market sentiment, and broader risk appetite. Distinguishing AI-related scenarios from broader macroeconomic conditions is therefore necessary to assess whether financial-sector tail outcomes co-move with AI-specific developments, the real-economy environment, or the joint occurrence of both.

To assess how market-related AI movements shape the extreme losses and gains of the financial sector, we decompose the expected shortfall (ES) into four components defined according to whether financial tail outcomes occur alongside AI tail events, macroeconomic tail events, both jointly, or neither. This decomposition, built on Ojea-Ferreiro (2025), starts from realized financial tail events and identifies the economic environments associated with those outcomes, thereby revealing how an expected financial-sector tail risk is distributed across alternative AI and macroeconomic scenarios. We then examine whether the AI-related component of financial tail risk is primarily associated with extreme gains or extreme losses. To do so, we introduce a tail-sensitive measure of financial-sector asymmetry that can be decomposed across the same economic scenarios identified in the ES framework, allowing us to identify the economic environments that contribute to financial-sector asymmetry. To the best of our knowledge, this is the first

measure of skewness that can be decomposed across observable economic scenarios. Together, these measures provide a structured characterization of how financial tail outcomes relate to AI-sector tail episodes, macroeconomic conditions, and their interaction.

Market data provide a natural setting to answer these questions because asset prices summarize investors' real-time assessment of AI-related risks and opportunities. We therefore interpret the AI variable as market-perceived AI exposure, rather than as direct evidence of financial institutions' internal AI adoption or realized AI-related incidents. AI exposure is proxied by an AI-related equity index constructed from the holdings of the WisdomTree Artificial Intelligence and Innovation Fund, while broader macroeconomic conditions are captured by real-economy indices built from broad equity benchmarks after excluding financial and AI-related firms.

Empirical results for financial-sector indices in Canada, the United States, Europe, the United Kingdom, and China over the period December 2021 to December 2025 reveal substantial cross-country heterogeneity in the economic scenarios associated with extreme financial returns and their asymmetry. Financial-sector-specific conditions account for the largest share of tail outcomes in all five markets, although their relative importance and contributions to skewness differ considerably. In the United States, negative skewness is driven primarily by the financial and systemic scenarios, while the AI-sector scenario has the smallest weight in both tails. Europe exhibits pronounced downside asymmetry driven mainly by the financial scenario and, among the remaining scenarios, by macroeconomic conditions. In China, extreme gains are slightly larger than extreme losses, producing mild positive skewness driven by the financial and systemic scenarios, whereas the AI-sector scenario contributes negatively. In the United Kingdom, the financial scenario dominates both tails, but the AI-sector scenario makes the largest average contribution to negative skewness. Canada displays the most balanced ES decomposition across scenarios, although its downside asymmetry is driven predominantly by the financial and systemic scenarios. Overall, the evidence suggests that AI-related market movements should not be interpreted as uniformly stabilizing or destabilizing for the financial sector. Their financial-stability interpretation depends on the market considered and on whether they occur as AI-specific episodes or jointly with broader macroeconomic conditions.

These findings have several implications. For investors, they show that AI-related market movements are relevant not only for downside risk, but also for the balance between extreme gains and losses in financial-sector returns. Importantly, the prevalence of an AI-related scenario in the tails does not necessarily imply a similarly large contribution

to return asymmetry, which also depends on the magnitude of tail outcomes and their relation to the central region of the return distribution. This balance depends critically on the macroeconomic environment in which AI-related episodes occur. For supervisors and central banks, the framework provides a market-based monitoring tool that identifies which economic environments are associated with realized financial-sector tail events and whether they contribute primarily to downside or upside asymmetry. This reverse stress-testing perspective complements traditional approaches that begin with a scenario and evaluate its implications for the financial system. More broadly, the results emphasize that AI-related financial vulnerabilities should not be assessed in isolation from macroeconomic conditions, and that decomposing financial tail outcomes across observable economic scenarios can provide timely information about the financial system's exposure to AI-related risks and opportunities.

Our analysis adds to three strands of literature. First, it provides market-based evidence to complement the growing policy and academic literature on AI and financial stability, which has so far been largely conceptual or theoretical (Danielsson et al. 2022; Aldasoro et al. 2025; Aldasoro and Desai 2025; Leitner et al. 2024; FSB 2024). This literature identifies channels through which AI may create or amplify financial vulnerabilities, but provides limited evidence on whether AI-related episodes are embedded in realized financial-sector extremes. Our approach addresses this gap by using market prices to measure whether investors' perceived exposure to AI-related risks and opportunities appears in the tails of financial-sector returns. Second, the paper contributes to empirical work on AI and financial markets. Existing studies examine the effects of AI adoption or AI-related technological developments on firm performance, valuation, risk taking, and market outcomes (Bertomeu et al. 2025; Gambacorta et al. 2025; Ca'Zorzi et al. 2025; Sheng et al. 2025; Chen and Ren 2022; Fuster et al. 2019). These studies provide important evidence on the economic effects of AI, but they generally focus on average outcomes, firm-level exposures, or event-based responses. We instead focus on the shape of the financial-sector return distribution, asking whether AI-related episodes are associated with extreme losses, extreme gains, and the asymmetry between them. Third, the paper contributes methodologically by extending the decomposition of the Expected Shortfall from Ojea-Ferreiro (2025) to a measure of financial-sector asymmetry that admits the same state decomposition. This extension allows us to move beyond the size of extreme losses and gains and to track how AI and macroeconomic scenarios are represented in the shape of the return distribution. The resulting framework provides a market-based monitoring tool for assessing whether the AI-related component of financial-sector asymmetry is concentrated in the upside or downside tail.

The remainder of the paper is organized as follows. Section 2 presents the methodological framework. Section 3 describes the data and index construction. Section 4 reports the empirical results. Finally, Section 5 concludes the paper.

2. Methodology

In the following, we describe the methodological framework we use to characterize how AI-sector tail events shape extreme outcomes in financial sector returns. The framework integrates two complementary measures. The first is the decomposition of the ES according to whether tail realizations occur during AI-sector stress, macroeconomic stress, both jointly, or neither. The second is a tail-sensitive skewness index that we introduce, the Bowley-type Expected Shortfall (BES), which summarizes the balance between upside and downside tail exposure within a bounded and economically interpretable measure of return asymmetry.

2.1. Expected Shortfall decomposition

A central challenge in assessing AI-related tail risk in financial institutions is distinguishing AI-specific shocks from broader macroeconomic conditions. To that end, we built on Ojea-Ferreiro (2025) to decompose the value of the ES according to the contemporaneous states of the AI sector and the real economy. Unlike CoVaR-type measures, which condition directly on stress of a conditioning variable, our decomposition first conditions on realized tail outcomes of the financial institution and then identifies the combinations of AI-sector and macroeconomic conditions associated with those outcomes. This approach provides a direct characterization of the economic states under which extreme financial losses materialize.

Let $r_{i,t}$ denote the financial index returns i and a distress event as

$$R_t^L = \{r_{i,t} \leq Q_{i,t}(\alpha)\}$$

where $Q_{i,t}(\alpha)$ denotes the α -quantile of the return distribution, corresponding to the Value-at-Risk at level α . The superscript L denotes events defined relative to the lower tail of the corresponding return distribution. Hence, the Expected Shortfall (ES) at confidence level α is defined as

$$ES_{i,t} = E(r_{i,t} | R_t^L) = \frac{1}{\alpha} \int_{-\infty}^{Q_{i,t}(\alpha)} r_{i,t} f_{i,t}(r_{i,t}) dr_{i,t},$$

This definition does not impose independence between financial returns and the AI and macroeconomic returns. Rather, their dependence is accounted for when R_t^L is

partitioned according to the contemporaneous tail states of the AI and macroeconomic indices. Specifically, define

$$A_t^L = \{r_{a,t} \leq Q_{a,t}(\alpha)\}, \quad M_t^L = \{r_{m,t} \leq Q_{m,t}(\alpha)\},$$

where $r_{a,t}$ and $r_{m,t}$ denote the returns of the AI and macroeconomic indices. The overbar denotes the corresponding no-distress event, e.g. $\bar{A}_t^L = \{r_{a,t} > Q_{a,t}(\alpha)\}$

Conditioning on these distress states yields the following decomposition of Expected Shortfall:

$$\begin{aligned} ES_{i,t} &= E\left[E(r_{i,t} \mid R_t^L, r_{a,t}, r_{m,t}) \mid R_t^L\right] \\ &\quad \text{financial} \\ &= \overbrace{E(r_{i,t} \mid R_t^L \cap \bar{A}_t^L \cap \bar{M}_t^L)P(\bar{A}_t^L \cap \bar{M}_t^L \mid R_t^L)}^{\text{AI-sector}} \\ &\quad + \overbrace{E(r_{i,t} \mid R_t^L \cap A_t^L \cap \bar{M}_t^L)P(A_t^L \cap \bar{M}_t^L \mid R_t^L)}^{\text{macroeconomic}} \\ &\quad + \overbrace{E(r_{i,t} \mid R_t^L \cap \bar{A}_t^L \cap M_t^L)P(\bar{A}_t^L \cap M_t^L \mid R_t^L)}^{\text{systemic}} \\ &\quad + \overbrace{E(r_{i,t} \mid R_t^L \cap A_t^L \cap M_t^L)P(A_t^L \cap M_t^L \mid R_t^L)}. \end{aligned} \tag{1}$$

Figure 1 illustrates the ES decomposition. The top-left panel (1A) reports the empirical distribution of the financial index returns. The vertical red line marks the stress threshold, with observations beyond it classified as tail realizations. The top-right panel plots AI and real economy index returns, together with their marginal distributions. Observations associated with financial tail events are highlighted in orange. These points represent sectoral returns conditional on the financial index being in its lower tail. The bottom panel (1B) restricts attention to this conditional sample and partitions the observations into four mutually exclusive regions depending on whether the AI and real economy indices lie in their respective tails. The bottom-left panel depicts the resulting four-component decomposition of the expected shortfall, with each component associated with a distinct joint tail configuration of the AI and real economy indices.

Hence, ES is split into four mutually exclusive components corresponding to:

- (i) Financial: losses when neither AI index nor the macroeconomic index is simultaneously in distress. These events capture financial-sector tail realizations that are not associated with concurrent AI-sector or macroeconomic tail events.

$$E(r_{i,t} | R_t^L \cap \bar{A}_t^L \cap \bar{M}_t^L) P(\bar{A}_t^L \cap \bar{M}_t^L | R_t^L) = \frac{1}{\alpha} \int_{-\infty}^{Q_{i,t}(\alpha)} r_{i,t} f(r_{i,t}, r_{a,t} > Q_{a,t}(\alpha), r_{m,t} > Q_{m,t}(\alpha)) dr_{i,t}$$

- (ii) AI-sector: losses coinciding with distress in the AI index only. These events reflect losses associated with AI-related or technology-sector shocks in the absence of a general downturn.

$$E(r_{i,t} | R_t^L \cap A_t^L \cap \bar{M}_t^L) P(A_t^L \cap \bar{M}_t^L | R_t^L) = \frac{1}{\alpha} \int_{-\infty}^{Q_{i,t}(\alpha)} r_{i,t} f(r_{i,t}, r_{a,t} \leq Q_{a,t}(\alpha), r_{m,t} > Q_{m,t}(\alpha)) dr_{i,t}$$

- (iii) Macroeconomic: losses coinciding with distress in the real economy. These events capture exposure to broad macroeconomic stress without AI-specific turmoil.

$$E(r_{i,t} | R_t^L \cap \bar{A}_t^L \cap M_t^L) P(\bar{A}_t^L \cap M_t^L | R_t^L) = \frac{1}{\alpha} \int_{-\infty}^{Q_{i,t}(\alpha)} r_{i,t} f(r_{i,t}, r_{a,t} > Q_{a,t}(\alpha), r_{m,t} \leq Q_{m,t}(\alpha)) dr_{i,t}$$


- (iv) Systemic: losses occurring when simultaneous stress materialized in both the macroeconomic indicator and the AI index. These “perfect storm” states represent episodes of joint macroeconomic and AI-sector turmoil.

$$E(r_{i,t} | R_t^L \cap A_t^L \cap M_t^L) P(A_t^L \cap M_t^L | R_t^L) = \frac{1}{\alpha} \int_{-\infty}^{Q_{i,t}(\alpha)} r_{i,t} f(r_{i,t}, r_{a,t} \leq Q_{a,t}(\alpha), r_{m,t} \leq Q_{m,t}(\alpha)) dr_{i,t}$$


These components have clear economic interpretations. For instance, a large AI-sector component indicates that shocks in the AI index can generate substantial losses even without the presence of a broader economic downturn. A dominant systemic component suggests that tail losses are concentrated in episodes where macroeconomic and AI stresses coincide, reflecting limited diversification of extreme outcomes under these conditions. In contrast, a dominant financial component suggests that tail realizations are primarily associated with conditions within the financial sector rather than with concurrent AI-sector or macroeconomic distress.

Conceptually, this decomposition of ES proposed by Ojea-Ferreiro (2025) resembles a reverse stress-testing exercise. Rather than imposing a stress scenario *ex ante* and evaluating its implications for financial institutions, we begin with realized tail losses and then identifies the combinations of AI-sector and macroeconomic conditions associated with those outcomes. This approach provides a direct characterization of the economic states under which extreme financial losses materialize.

This decomposition naturally extends to upper-tail realizations, i.e., $R_t^U = \{r_{i,t} \geq Q_{i,t}(1 - \alpha)\}$. In this case, the decomposition characterizes how extreme gains co-move



A. Returns of the financial index (left) and the joint distribution of AI and real economy index returns (right). Orange bars indicate observations occurring when the financial index is in its left tail.



B. Expected Shortfall allocation (left), distinguishing four regions in the tail of the financial index returns based on the conditional behavior of the AI and real economy index (right).

FIGURE 1. Concept of the Expected Shortfall decomposition.

with AI-sector and macroeconomic conditions, i.e.,

$$A_t^U = \{r_{a,t} \geq Q_{a,t}(1 - \alpha)\}, \quad M_t^U = \{r_{m,t} \geq Q_{m,t}(1 - \alpha)\}.$$

The right-tail decomposition therefore complements the downside analysis by distinguishing whether unusually strong performance is associated with AI-driven growth, broad macroeconomic expansions, or their interaction. Comparing left- and right-tail decompositions provides a comprehensive view of the asymmetry of institutions' exposure to AI-related developments.

2.2. Skewness index

Artificial intelligence may affect not only the magnitude but also the asymmetry of financial tail outcomes. To capture this asymmetry, we introduce a tail-sensitive skewness measure that captures whether AI-related exposures disproportionately contribute to extreme gains or to extreme losses, that is, whether AI exposure is associated with downside-dominated risk, upside-dominated gains, or a more balanced tail structure.

An appropriate skewness measure for this setting should satisfy three properties. First, it should be bounded to facilitate comparisons across institutions and over time. Second, it should admit a direct economic interpretation in terms of expected upside gains versus downside losses. Third, it should remain sensitive to tail behavior, as AI effects are likely to manifest disproportionately in the extremes of the return distribution. A natural starting point is the skewness coefficient proposed by Bowley (1920).

Let $Q_{i,t}(0.25)$, $Q_{i,t}(0.50)$, and $Q_{i,t}(0.75)$ denote the first quartile, median, and third quartile of the return distribution of financial institution i , i.e. $r_{i,t}$. The Bowley coefficient is defined as

$$B = \frac{(Q_{i,t}(0.75) - Q_{i,t}(0.50)) - (Q_{i,t}(0.50) - Q_{i,t}(0.25))}{Q_{i,t}(0.75) - Q_{i,t}(0.25)}.$$

The coefficient is bounded between -1 and 1 . Positive values indicate relatively stronger upper-tail outcomes, negative values indicate stronger lower-tail outcomes, and zero corresponds to a symmetric distribution.⁴

However, a limitation of quantile-based skewness measures is that they depend only on selected quantile locations and therefore remain relatively insensitive to the magnitude of realizations beyond those thresholds. To incorporate tail magnitudes directly, we replace quantile distances with conditional expectations computed over percentile regions of

⁴Hinkley (1975) generalises this idea to arbitrary tail quantiles. Ghysels et al. (2016) provide a review of skewness measures.

equal probability mass. Let us define

$$\mu_{i,t}^L = E(r_{i,t} | R_t^L), \quad \mu_{i,t}^C = E(r_{i,t} | R_t^C), \quad \mu_{i,t}^U = E(r_{i,t} | R_t^U),$$

where $R_t^C = \{Q_{i,t}(0.5 - \frac{\alpha}{2}) \leq r_{i,t} \leq Q_{i,t}(0.5 + \frac{\alpha}{2})\}$. These quantities correspond to the average lower-tail, central-region, and upper-tail realizations of the distribution, respectively, as illustrated in the upper panel of Figure 2. Then, we can define the Bowley-type Expected Shortfall (*BES*) skewness index as

$$BES_{i,t} = \frac{(\mu_{i,t}^U - \mu_{i,t}^C) - (\mu_{i,t}^C - \mu_{i,t}^L)}{\mu_{i,t}^U - \mu_{i,t}^L} = \frac{\mu_{i,t}^U + \mu_{i,t}^L - 2\mu_{i,t}^C}{\mu_{i,t}^U - \mu_{i,t}^L}. \quad (2)$$

The index satisfies $-1 < BES_{i,t} < 1$. Positive values indicate distributions with relatively stronger upper-tail realizations, whereas negative values indicate greater downside tail intensity.

To identify the economic states of skewness, we apply the ES decomposition to the numerator of $BES_{i,t}$. The denominator is kept fixed as a common normalizing term to preserve comparability across states. Each conditional mean admits its own ES decomposition, allowing skewness to be traced back to the economic environments associated with lower-tail, central-region, and upper-tail realizations.

For instance, we can decompose $\mu_{i,t}^L$ into four sections following the law of total expectation shown in Eq. (1), i.e.

$$\mu_{i,t}^L = \sum_{j=1}^4 P(S_{j,t} | R_t^L) \mu_{i,t}^L(S_{j,t}),$$

where $\mu_{i,t}^L(S_{j,t}) = E(r_{i,t} | R_t^L \cap S_{j,t})$ and $S_{j,t}$ indexes one of the four ES economic states at time t , corresponding to the different combinations of AI-sector and macroeconomic distress. Similarly, we could do the same type of decomposition with $\mu_{i,t}^U$ and $\mu_{i,t}^C$ using the corresponding central (C) and upper (U) regions from their return distributions. Importantly, AI and macroeconomic states are defined relative to the region of the financial return distribution under consideration; for example, AI extremes correspond to negative AI realizations in the lower tail and to positive AI realizations in the upper tail.


The skewness index in Eq. (2) can therefore be rewritten as the sum of state-specific contributions:

$$BES_{i,t} = \sum_{j=1}^4 \frac{\left(P(S_{j,t} | R_t^U) \mu_{i,t}^U(S_{j,t}) + P(S_{j,t} | R_t^L) \mu_{i,t}^L(S_{j,t}) - 2P(S_{j,t} | R_t^C) \mu_{i,t}^C(S_{j,t}) \right)}{\mu_{i,t}^U - \mu_{i,t}^L}. \quad (3)$$


Each component measures how a given economic state contributes to the asymmetry

of financial returns, with normalization fixed at the aggregate level. This decomposition admits a direct economic interpretation. A positive AI-sector component indicates that AI-related states contribute more strongly to extreme gains than to extreme losses. Conversely, a negative systemic component implies that joint AI and macroeconomic distress disproportionately amplifies downside tail outcomes. The decomposition therefore reveals not only whether financial returns are asymmetric, but also which economic environments generate that asymmetry.

The lower panel of Figure 2 illustrates how the skewness index can be partitioned across the four ES states, highlighting which economic environments contribute positively or negatively to distributional asymmetry.



A. Areas to compute the Bowley-type skewness index using the conditional expectation.



B. Decomposition of the areas to compute the skewness index based on the realizations of the real economy and AI index.

FIGURE 2. Bowley-type skewness index based on conditional expectations.

2.3. Implementation of the ES decomposition

The ES decomposition requires evaluating tail losses conditional on the joint behaviour of financial, AI and macroeconomic returns. A copula framework is particularly well suited for this purpose. First, it allows the marginal distributions and dependence structure to be modelled separately, accommodating the heteroskedasticity, asymmetries and fat tails characteristic of financial returns observed in each market while allowing dependence to depart from linear correlation. This flexibility is especially relevant in our setting because the decomposition in Eq.(1) is driven by tail dependence rather than average co-movement (Embrechts et al. 2002). Second, the copula framework yields a tractable representation of the conditional distributions entering Eq. (1), reducing each component of the ES to a one-dimensional integral and avoiding the multidimensional integrations required by alternative multivariate approaches.

To implement this framework, we model the joint return distribution of financial sector index, a real-economy index, and an AI index. We follow the standard ARMA-GARCH model used in the literature (Girardi and Ergün 2013; Ojea-Ferreiro and Reboredo 2022; Ojea-Ferreiro et al. 2024; Reboredo and Ugolini 2015; Reboredo et al. 2016; Reboredo and Ugolini 2016) for the marginal distribution, where skewed epsilon Skew-t innovations (Arellano-Valle et al. 2005; Gómez et al. 2007) capture heavy tails and evolving asymmetries following GAS dynamics (Delle Monache et al. 2024) in individual return series. Dependence across markets is modeled through a Skewed Student-t copula with cDCC dynamics (Christoffersen et al. 2012; Christoffersen and Langlois 2013; Christoffersen et al. 2018; Ojea-Ferreiro 2026). The resulting framework combines autocorrelation and conditional heteroskedasticity in the marginals with skewed and heavy-tailed innovations, while the copula captures time-varying dependence, asymmetric dependence patterns, and tail dependence across markets. Details of the marginal model, copula specification, and estimation procedure are reported in Appendices A.1-A.4.

Under this representation, the marginal cumulative distributions of financial, AI, and macroeconomic returns, denoted by $F_i(\dots)$, $F_a(\dots)$, and $F_m(\dots)$, are linked through a copula function $C(\dots)$. By Sklar's theorem (Sklar 1959),

$$F(r_{i,t}, r_{a,t}, r_{m,t}) = C(q, u, v),$$

where $q = F_i(r_{i,t})$, $u = F_a(r_{a,t})$ and $v = F_m(r_{m,t})$. The associated joint density can be written as

$$f(r_{i,t}, r_{a,t}, r_{m,t}) = c(q, u, v) f_i(r_{i,t}) f_a(r_{a,t}) f_m(r_{m,t}), \quad (4)$$

where c denotes the copula density and $f_i(\dots)$, $f_a(\dots)$ and $f_m(\dots)$ are the marginal

distribution. Eq. (4) makes explicit the separation between marginal behavior and dependence. For the ES decomposition, the key implication of this representation is that the conditional probabilities characterizing each AI-macroeconomic scenario can be expressed directly through conditional copulas. Following Joe (1997),

$$C_{a,m|i}(u, v | q) = \frac{\partial C(q, u, v)}{\partial q}. \quad (5)$$

The conditional copulas of the individual markets, $C_{a|i}(u | q)$ and $C_{m|i}(v | q)$, are obtained by setting $v = 1$ or $u = 1$.

Eq. (1) partitions Expected Shortfall according to whether the AI sector and the real economy are in distress. The corresponding scenario probabilities are obtained directly from the conditional copulas above. To derive the copula representation, we first write Expected Shortfall as

$$E(r_{i,t} | R_t^L) = \frac{1}{\alpha} \int_0^\alpha F_i^{-1}(q) dq,$$

where $F_i^{-1}(q) = r_{i,t}$ and $dq = dF_i(r_{i,t}) = f_i(r_{i,t})dr_{i,t}$. Using the scenario probabilities implied by Eq. (5), the unit weight inside the integral can be decomposed according to the four scenarios in Eq. (1). Hence,

$$E(r_{i,t} | R_t^L) = \frac{1}{\alpha} \int_0^\alpha F_i^{-1}(q) \left\{ C_{m,a|i}(\alpha, \alpha | q) + \right. \quad (6)$$

$$\left[1 - C_{m|i}(\alpha | q) - C_{a|i}(\alpha | q) + C_{m,a|i}(\alpha, \alpha | q) \right] \quad (7)$$

$$+ \left[C_{a|i}(\alpha | q) - C_{m,a|i}(\alpha, \alpha | q) \right] \quad (8)$$

$$\left. + \left[C_{m|i}(\alpha | q) - C_{m,a|i}(\alpha, \alpha | q) \right] \right\} dq. \quad (9)$$

The four terms correspond to the systemic, financial, AI-sector and macroeconomic scenarios in Eq. (1). Since these scenarios are mutually exclusive and exhaustive, the four probabilities sum to one for every quantile level q . The decomposition therefore assigns each tail loss to one of four possible scenarios depending on whether the AI sector and the real economy are in distress. By linearity of the integral, Expected Shortfall can be obtained as the sum of four components, each computed by weighting tail losses with the corresponding scenario probability. Appendix A.5 shows the detailed derivation of the formula.

These components distinguish whether financial-sector tail losses occur in isolation or together with distress in the AI sector, the real economy, or both. The relative importance of each component therefore depends on how likely financial losses are to coincide with adverse outcomes in the other two markets. The copula framework provides a flexible way to capture these dependence patterns and their implications for systemic risk. Appendix A.6 extends the same representation to the upper-tail and central-region decompositions used in the construction of the skewness index.

3. Data

This paper uses daily total return data over the period from December 2021 to December 2025 for three groups of equity indices covering Canada, Europe, the United States, the United Kingdom, and China: a market-based proxy for artificial intelligence exposure, the regional financial sector indices, and a set of custom-built regional indices representing the non-financial, non-AI real economy. All series are denominated in local currency.⁵ Exposure to artificial intelligence is proxied using the WisdomTree Artificial Intelligence and Innovation Fund (RIC: WTAI.K), which we use as a market-based proxy for AI exposure, capturing firms with substantial AI-related activities. This index is used both as a standalone measure of AI-related equity performance and as a benchmark for excluding AI-exposed firms from the construction of the real economy index. The financial sector is represented in each region by the Datastream Financials sectoral indices (DS Financials). The sample starts on December 9, 2021, corresponding to the inception of the WisdomTree Artificial Intelligence and Innovation Fund, and ends in December 2025.

To measure macroeconomic performance, we construct a custom regional real economy index for Canada, the United States, Europe, the United Kingdom and China. For each region, we begin with the monthly constituent lists of the main equity benchmark indices: the S&P/TSX Composite Index for Canada, the S&P 500 Index for the United States, the STOXX Europe 600 Index for Europe, the FTSE 100 Index for the United Kingdom, and the Shanghai Stock Exchange Composite Index for China. From these monthly compositions, we exclude (i) all firms classified in the financial sector and (ii) all firms that are constituents of the AI index. The remaining firms are interpreted as representing the non-financial, non-AI real economy. Daily equity returns of these firms are aggregated into a regional real economy index using market-capitalization weights, where weights are computed at the monthly frequency based on each firm's relative market capitalization within the filtered index. The resulting real economy indices vary over time to reflect changes in index composition and relative firm size. All data are obtained from

⁵The analysis is performed in the local currency to prevent distortions in return distributions arising from exchange rate movements (see Ojea-Ferreiro 2020).

Datastream.

In Table 1, we report the main descriptive statistics for the daily log-returns of the AI index and of the regional financial (F) and real economy non-financial (NF) indices over the sample period. Average returns are small and, for most series, positive, ranging between -0.01% (China financials) and 0.05% (Canada non-financials), consistent with the well-documented difficulty of distinguishing the drift component at the daily frequency. The AI index displays the highest unconditional volatility (1.93%), followed by the Chinese financial index (1.66%), whereas the non-financial real economy indices are systematically less volatile than their financial counterparts within each region. Most return series are negatively skewed, with the notable exception of the AI index (approximately symmetric) and the Chinese financial index (0.71), indicating a higher propensity for large negative realizations in the remaining markets. Furthermore, all series exhibit kurtosis well above the Gaussian benchmark of three, ranging from 4.84 (Europe non-financials) to 13.08 (China non-financials), which points to fat-tailed distributions. Consistent with this evidence, the 10% left Value-at-Risk (VaR) and Expected Shortfall (ES) are largest in absolute value for the AI index (-2.52% and -3.58% , respectively) and for the Chinese financial index (-1.74% and -2.76%), confirming their greater exposure to extreme downside movements; by contrast, the UK and Canadian non-financial indices display the most contained tail losses. Turning to the upside, the 10% right VaR and ES are again largest for the AI index (2.27% and 3.24% , respectively) and for the Chinese financial index (1.72% and 3.26%), pointing to a comparable exposure to extreme positive realizations; notably, for both indices the right-tail risk is of the same order of magnitude as the left-tail one, in line with the approximately symmetric distribution of the AI index and the positively skewed distribution of the Chinese financial index, whereas for the remaining series the right tail is systematically smaller than the left, consistent with the negative skewness documented above.

Finally, the Pearson correlation coefficients reported in the lower panel reveal a pronounced co-movement between the AI index and the financial sectors of developed markets, with ρ_{AI} peaking at 74.25% for the United States and declining markedly for the emerging Chinese market (12.77%). Likewise, the correlation between the financial and the real economy indices within each region ($\rho_{NF,F}$) is substantial in developed markets, exceeding 88% in the United States, but weak in China (3.10%), which may reflect the more segmented nature of the Chinese equity market.

TABLE 1. Summary statistics for daily log-returns of different equity indices

	AI	US-NF	US-F	CA-NF	CA-F	EU-NF	EU-F	UK-NF	UK-F	CN-NF	CN-F
Mean (%)	0.01	0.03	0.04	0.05	0.04	0.01	0.01	0.04	0.00	0.01	-0.01
Volatility (%)	1.93	0.90	1.23	0.77	1.21	0.88	0.83	0.74	0.91	1.10	1.66
Skewness	0.00	-0.14	-0.09	-0.32	-0.30	-0.32	-0.71	-0.85	-0.13	-0.45	0.71
Kurtosis	5.31	8.43	7.51	6.04	5.90	4.84	9.63	8.87	6.87	13.08	9.17
Max. downturn (%)	-7.67	-5.85	-7.50	-4.62	-5.98	-4.21	-6.44	-4.95	-5.07	-8.39	-10.34
Max. upturn (%)	12.30	6.45	8.20	4.15	7.01	4.68	5.20	3.49	4.56	8.41	9.95
1st quartile (%)	-1.06	-0.44	-0.62	-0.38	-0.57	-0.47	-0.40	-0.33	-0.47	-0.50	-0.85
3rd quartile (%)	1.23	0.55	0.73	0.51	0.78	0.57	0.50	0.44	0.49	0.58	0.66
10% left VaR (%)	-2.52	-0.99	-1.42	-0.87	-1.41	-1.07	-0.90	-0.75	-1.08	-1.13	-1.74
10% left ES (%)	-3.58	-1.60	-2.29	-1.41	-2.25	-1.67	-1.58	-1.38	-1.68	-2.02	-2.76
10% right VaR (%)	2.27	1.04	1.37	0.93	1.45	1.03	0.88	0.83	0.94	1.16	1.72
10% right ES (%)	3.24	1.59	2.14	1.36	2.06	1.48	1.34	1.27	1.56	1.87	3.26
ρ_{AI} (%)		70.53	74.25	52.71	69.85	30.66	43.51	21.52	46.34	14.49	12.77
$\rho_{NF,F}$ (%)			88.48		64.44		66.03		7.75		3.10

Notes. This table presents summary statistics for daily log returns in the corresponding local denomination for AI, Non-Financial (NF), and Financial (F) assets over the sample period December 2021 to December 2025 from the five regions under analysis (the United States, Canada, Europe, United Kingdom, and China). For each asset category, we report the average returns, volatility, skewness, kurtosis, maximum downturn and upturn, value-at-risk (VaR), and expected shortfall (ES) for the left and right sides of the return distribution, both computed at the 10% tail probability. ρ indicates the correlation between AI and the different equity returns, or between the Financial and Non-Financial equity returns from the same region.

4. Results

Figures 3–12 report the decomposition of financial-sector tail outcomes and return asymmetry for the United States, Europe, China, the United Kingdom, and Canada. For each market, the upper chart of the ES figure reports the average financial-sector return within the lowest and highest 5 percent of the conditional return distribution, shown in red and green, respectively. The central and bottom charts decompose the corresponding left- and right-tail ES across the financial, AI-sector, macroeconomic, and systemic scenarios. The subsequent figure reports the Bowley-type Expected Shortfall index and its decomposition across the same scenarios. Positive contributions increase right skewness, whereas negative contributions increase downside asymmetry. For each market, we first compare the magnitude and scenario composition of extreme losses and gains and then examine how they translate into scenario-specific contributions to return asymmetry.

4.1. United States

In the United States, the left-tail ES is slightly larger in absolute terms than the right-tail ES, averaging -2.55 percent and 2.47 percent, respectively. This modest imbalance between extreme losses and gains is reflected in the predominantly negative skewness index shown by the red line in Figure 4. The AI-sector scenario has the smallest weight in the ES decomposition, accounting for approximately 7 percent of both tails.

The bottom charts of Figure 3 nevertheless reveal some differences in scenario composition. The systemic scenario accounts for 31 percent of the left-tail ES, compared with 28 percent of the right-tail ES. Conversely, the financial scenario is more prevalent in the right tail, with a weight of 32 percent, than in the left tail, where it accounts for 30 percent.

The skewness index is slightly positive during the first half of 2022 but negative throughout most of the remaining sample. On average, the financial scenario makes the largest contribution to the negative skewness of U.S. financial-sector returns, closely followed by the systemic scenario. During the episodes of positive skewness, the systemic and macroeconomic scenarios provide the main positive contributions. Thus, the relatively small difference between average extreme losses and gains masks meaningful variation over time and across scenarios.




FIGURE 3. US ES decomposition




FIGURE 4. US skewness decomposition

4.2. Europe

Europe exhibits a more pronounced imbalance between its two tails. The left-tail ES averages -2.44 percent, compared with 2.06 percent for the right-tail ES. The difference is even more evident at their most extreme values: the left-tail ES reaches a minimum of

–10.5 percent, whereas the right-tail ES reaches a maximum of 7.03 percent. Consistent with this stronger downside tail, the skewness index in Figure 6 is predominantly negative and averages –0.148 over the sample.

The ES decomposition in Figure 5 is dominated by the financial scenario, whose average weight is larger in the right tail than in the left tail, 59.5 percent versus 52.1 percent. The other three scenarios are more prevalent in the left tail. The macroeconomic scenario accounts for 23.2 percent of the left-tail ES and 19.9 percent of the right-tail ES, while the corresponding weights of the AI-sector scenario are 14.3 percent and 11.8 percent. The systemic scenario has the smallest weight in both tails, accounting for 10.4 percent of the left-tail ES and 8.8 percent of the right-tail ES.

All four scenarios contribute negatively to average skewness. The financial scenario makes the largest contribution, averaging –0.071 and accounting for approximately half of the aggregate Bowley index. The macroeconomic scenario provides the second-largest negative contribution, averaging –0.037, followed by the AI-sector and systemic scenarios at –0.026 and –0.015, respectively. European financial-sector returns therefore exhibit persistent downside asymmetry driven primarily by the financial scenario. Among the remaining scenarios, macroeconomic conditions are the most important source of negative skewness, while the AI-sector and systemic scenarios play smaller roles.




FIGURE 5. EU ES decomposition




FIGURE 6. EU skewness decomposition

4.3. China

China is the only market in which the average right-tail ES exceeds the absolute value of the average left-tail ES. The right-tail ES averages 2.49 percent, compared with -2.36 percent in the left tail. Similarly, its maximum of 5.65 percent slightly exceeds the absolute value of the left-tail minimum of -5.33 percent. This modest predominance of extreme gains is reflected in the slightly positive skewness index in Figure 8, which averages 0.021 over the sample.

The ES decomposition in Figure 7 is strongly dominated by the financial scenario. Its average weight increases from 81.2 percent in the left tail to 87.7 percent in the right tail. The AI-sector scenario is the second-largest component in the left tail, accounting for 10.8 percent, but its weight falls to 3.5 percent in the right tail. The macroeconomic scenario likewise has a larger weight in the left tail than in the right tail, 7.0 percent versus 4.3 percent. By contrast, the systemic scenario accounts for 4.6 percent of the right-tail ES but only 1.0 percent of the left-tail ES.

The positive average skewness is driven by the financial scenario, which contributes 0.050, and, to a lesser extent, by the systemic scenario, which contributes 0.019. These contributions are partly offset by the AI-sector and macroeconomic scenarios, whose average contributions are -0.035 and -0.013 , respectively. The negative AI-sector contribution is consistent with its substantially larger presence in the left tail, although its precise value also reflects tail magnitudes and the central-region decomposition.

Chinese financial-sector returns therefore exhibit mild positive asymmetry generated primarily by the financial scenario. The systemic scenario reinforces this upside asymmetry, whereas the AI-sector scenario provides the largest negative offset. Although the financial scenario dominates both tails, the decomposition shows that the aggregate positive skewness does not originate from AI-related tail outcomes.



FIGURE 7. China ES decomposition




FIGURE 8. China skewness decomposition

4.4. United Kingdom

The United Kingdom displays a clear predominance of downside tail risk. The left-tail ES averages -2.28 percent, compared with 1.97 percent for the right-tail ES. At their most extreme values, the left-tail ES reaches a minimum of -6.11 percent, while the right-tail ES reaches a maximum of 5.94 percent. This imbalance is reflected in the predominantly

negative skewness index in Figure 10, which averages -0.131 over the sample.

The financial scenario dominates both tails in Figure 9, but its weight is larger in the right tail than in the left tail, 71.0 percent versus 61.3 percent. The AI-sector scenario is the second-largest component and is more prevalent in the left tail, where it accounts for 22.6 percent, than in the right tail, where it accounts for 15.9 percent. The macroeconomic scenario displays a similar pattern, with weights of 11.8 percent and 5.5 percent in the left and right tails, respectively. Conversely, the systemic scenario accounts for 7.6 percent of the right-tail ES but only 4.2 percent of the left-tail ES.

The AI-sector scenario makes the largest average contribution to negative skewness, at -0.052 , closely followed by the financial scenario at -0.050 . The macroeconomic contribution is also negative and averages -0.041 . In contrast, the systemic scenario contributes positively, averaging 0.013 , and therefore partially offsets the downside asymmetry generated by the other scenarios.

The United Kingdom consequently differs from Europe in the relative importance of the AI-sector scenario. Although the financial scenario dominates the ES decomposition, the AI-sector scenario is the largest average contributor to negative skewness. Its greater presence in the left tail, together with the larger magnitude of extreme losses, is consistent with this contribution. The systemic scenario displays the opposite pattern, with a larger right-tail weight and a positive contribution to return asymmetry.



FIGURE 9. UK ES decomposition




FIGURE 10. UK skewness decomposition

4.5. Canada

Canada also exhibits stronger downside than upside tail outcomes. The left-tail ES averages -1.85 percent, compared with 1.57 percent for the right-tail ES. The difference is more pronounced at their most extreme values, with a left-tail minimum of -6.77 percent and a right-tail maximum of 4.19 percent. The skewness index is correspondingly negative

throughout most of the sample and averages -0.143 .

Compared with the other markets, the Canadian ES decomposition in Figure 11 is more evenly distributed across scenarios. The financial scenario remains the largest component, but its average weights are similar in the left and right tails, at 41.2 percent and 40.1 percent, respectively. The macroeconomic scenario is more prevalent in the right tail, where it accounts for 26.0 percent, than in the left tail, where it accounts for 22.5 percent. The AI-sector scenario also has a larger right-tail weight, 18.9 percent compared with 15.5 percent. The systemic scenario displays the opposite pattern, accounting for 20.8 percent of the left-tail ES but only 15.0 percent of the right-tail ES.

Despite the comparatively balanced ES allocation, the scenario contributions to skewness differ substantially. The financial scenario makes by far the largest negative contribution, averaging -0.092 and accounting for nearly two-thirds of the aggregate index. The systemic scenario provides the second-largest negative contribution, averaging -0.043 . The macroeconomic and AI-sector contributions are much smaller, averaging -0.007 and -0.001 , respectively.

Canadian downside asymmetry is therefore driven primarily by the financial and systemic scenarios. The greater prevalence of the systemic scenario in the left tail is consistent with its negative contribution to skewness. By contrast, the macroeconomic and AI-sector scenarios have meaningful weights in both tails and are somewhat more prevalent in the right tail, but their average contributions to skewness are close to zero. The Canadian results thus illustrate that a scenario may account for a meaningful share of tail outcomes without making a similarly large contribution to return asymmetry.




FIGURE 11. Canada ES decomposition




FIGURE 12. Canada skewness decomposition

4.6. Cross-country evidence

The results reveal substantial cross-country heterogeneity in both the composition of financial-sector tail outcomes and their asymmetry. The financial scenario is the largest component of both tails in all five markets, although its importance varies considerably. It is particularly dominant in China, Europe, and the United Kingdom, whereas the ES

decomposition is more evenly distributed across scenarios in the United States and Canada. The AI-sector scenario is most prominent in the United Kingdom, while macroeconomic conditions play a comparatively larger role in Europe and Canada. The systemic scenario is particularly relevant in the United States and Canada but accounts for a considerably smaller share of tail outcomes in Europe, China, and the United Kingdom.

The skewness decomposition further shows that the prevalence of a scenario in the two tails does not necessarily determine its contribution to return asymmetry. The financial scenario is the principal source of negative skewness in the United States, Europe, and Canada and the main source of positive skewness in China. The United Kingdom is distinctive because the AI-sector scenario makes the largest average contribution to downside asymmetry, although the financial scenario continues to dominate both tails. More generally, AI-sector scenarios contribute negatively to average skewness in Europe, China, and the United Kingdom, while their contribution is comparatively limited in Canada. Systemic scenarios also differ across markets, contributing to downside asymmetry in the United States and Canada but partially offsetting it in China and the United Kingdom. Several mechanisms may help explain these cross-country differences. First, the transmission of AI-sector movements to financial returns is likely to depend on the degree of market integration between AI-related firms, the real economy, and the domestic financial sector. The relatively large systemic contribution in the United States and Canada may reflect the stronger co-movement of their financial sectors with both the AI index and broader market conditions. By contrast, the dominance of the financial scenario in China is consistent with weaker transmission through the market-based channel captured by our framework. Second, differences in financial structure, sectoral composition, and exposure to internationally traded technology firms may affect whether AI-related movements are reflected primarily in financial-sector returns, broader macroeconomic conditions, or their interaction. Third, the comparatively large AI-sector contribution in the United Kingdom may reflect the financial sector's extensive use of AI and growing reliance on external technology providers.⁶ Common exposures to providers of models, data, and cloud infrastructure could make financial firms' valuations more sensitive to AI-specific developments, including when those developments occur independently of broader macroeconomic tail events. This interpretation is consistent with evidence of widespread AI use and concentrated third-party dependencies in UK financial services, but it remains suggestive because our analysis does not directly measure firms' AI adoption or technology-provider exposures. Overall, the results suggest that the financial-stability implications of AI-related tail events depend not only on whether these events occur in isolation or jointly with broader macroeconomic movements, but also on the financial structure and market linkages through which AI-sector movements are transmitted.

⁶See for instance this survey from Bank of England.

5. Conclusion

This paper develops a framework to assess how artificial intelligence (AI) exposure shapes the tail behavior of financial-sector returns. Combining an Expected Shortfall (ES) decomposition with a tail-sensitive skewness index, the approach decomposes extreme losses and gains according to the contemporaneous tail scenarios of the AI sector and macroeconomic conditions. The ES decomposition identifies the economic environments associated with realized financial tail outcomes, while the skewness decomposition summarizes how these environments contribute asymmetrically to downside and upside risk. Together, the framework provides a structured and interpretable way to link AI-sector dynamics, macroeconomic conditions, and financial tail risk.

The empirical analysis uncovers substantial cross-country heterogeneity in how market-perceived AI exposure is associated with financial tail risk. In the United States, extreme losses are slightly larger than extreme gains, and the resulting negative skewness is driven primarily by the financial and systemic scenarios. The AI-sector scenario has the smallest weight in both tails, accounting for approximately 7 percent of each. Europe displays more pronounced downside asymmetry, driven mainly by the financial scenario and, among the remaining scenarios, by macroeconomic conditions. The AI-sector scenario contributes negatively to skewness but plays a comparatively smaller role. China exhibits a distinct pattern: extreme gains are slightly larger than extreme losses, producing mild positive skewness. This upside asymmetry originates primarily in the financial and systemic scenarios, whereas the AI-sector scenario makes the largest negative contribution. In the United Kingdom, the financial scenario dominates both tails, but the AI-sector scenario makes the largest average contribution to negative skewness, closely followed by the financial scenario. Finally, Canada presents the most balanced ES decomposition across scenarios, although its negative skewness is driven predominantly by the financial and systemic scenarios, with the AI-sector contribution remaining close to zero.

In sum, AI-related tail events are neither uniformly associated with greater downside risk nor with stronger upside outcomes. The AI-sector scenario contributes to negative skewness in Europe, China, and the United Kingdom, with its largest contribution observed in the United Kingdom. Its contribution is comparatively limited in Canada, while its weight in both tails is smallest in the United States. The systemic scenario is also heterogeneous across markets: it contributes to negative skewness in the United States and Canada but supports more positive asymmetry in China and the United Kingdom. Across all five markets, however, financial-sector-specific conditions remain the principal determinant of tail outcomes and return asymmetry. These findings indicate that the financial-stability implications of market-perceived AI exposure depend on the structure

of each financial market and on whether AI-sector tail events occur in isolation or jointly with broader macroeconomic movements.

Several limitations should be acknowledged. First, AI exposure is proxied using a market-based equity index, which may not capture all dimensions of AI adoption, particularly non-listed activities or firm-level implementation intensity. Second, the analysis focuses on equity returns of financial institutions and does not directly address other channels of risk transmission, such as credit risk, funding markets, or operational risk. Third, while the framework identifies associations between financial tail outcomes and economic scenarios, it is not designed to establish causal effects of AI adoption on financial stability. Finally, the sample period coincides with an unusual macro-financial environment, which may influence the observed tail dynamics.

Future research could extend the framework in several directions. Firm-level measures of AI adoption could be incorporated to distinguish between direct exposure and broader sectoral effects. The ES and skewness decompositions could also be extended to other measures of financial risk, such as downside semivariance. From a policy perspective, integrating the framework with stress-testing exercises or supervisory data could help assess whether AI adoption alters the resilience of financial institutions under adverse macroeconomic scenarios. More broadly, the approach offers a flexible template for studying how emerging technologies shape not only the magnitude but also the asymmetry of financial risk.

References

- Aldasoro, I. and Desai, A. (2025). Money talks: Ai agents for cash management in payment systems. Technical report, Bank of Canada.
- Aldasoro, I., Gambacorta, L., Korinek, A., Shreeti, V., and Stein, M. (2025). Intelligent financial system: How ai is transforming finance. *Journal of Financial Stability*, 81:101472.
- Arellano-Valle, R. B., Gómez, H. W., and Quintana, F. A. (2005). Statistical inference for a general class of asymmetric distributions. *Journal of Statistical Planning and Inference*, 128(2):427–443.
- Basele, R. B., Phillips, P. C., and Shi, S. (2025). Speculative bubbles in the recent ai boom: Nasdaq and the magnificent seven. *Journal of Time Series Analysis*.
- Berg, T., Burg, V., Gombović, A., and Puri, M. (2020). On the rise of fintechs: Credit scoring using digital footprints. *The Review of Financial Studies*, 33(7):2845–2897.
- Bertomeu, J., Lin, Y., Liu, Y., and Ni, Z. (2025). The impact of generative ai on information processing: Evidence from the ban of chatgpt in italy. *Journal of Accounting and Economics*, 80(1):101782.
- Bianchi, D., Büchner, M., and Tamoni, A. (2021). Bond risk premiums with machine learning. *The Review of Financial Studies*, 34(2):1046–1089.
- Bowley, A. (1920). Elements of statistics (charles scribner's sons, new york).
- Ca'Zorzi, M., Manu, A.-S., and Lopardo, G. (2025). Verba volant, transcripta manent: what corporate earnings calls reveal about the ai stock rally. Technical report, European Central Bank.
- Cecchetti, S. G., Lumsdaine, R. L., Peltonen, T., and Sánchez Serrano, A. (2025). Artificial intelligence and systemic risk. *ESRB: Advisory Scientific Committee Reports*, 16.
- Chen, R. and Ren, J. (2022). Do ai-powered mutual funds perform better? *Finance Research Letters*, 47:102616.
- Choi, J. J., Huang, D., Yang, Z., and Zhang, Q. (2025). How good is ai at twisting arms? experiments in debt collection. *NBER Working Paper*, (w33669).
- Christoffersen, P., Errunza, V., Jacobs, K., and Langlois, H. (2012). Is the potential for international diversification disappearing? a dynamic copula approach. *The Review of financial studies*, 25(12):3711–3751.
- Christoffersen, P., Jacobs, K., Jin, X., and Langlois, H. (2018). Dynamic dependence and diversification in corporate credit. *Review of Finance*, 22(2):521–560.
- Christoffersen, P. and Langlois, H. (2013). The joint dynamics of equity market factors. *Journal of Financial and Quantitative Analysis*, 48(5):1371–1404.
- Crisóstomo Ayala, R. and Mykhalyuk, D. (2026). Large language models and stock investing: Is the human factor required? *Documentos de trabajo (CNMV)*, (96):1–39.
- Danielsson, J., Macrae, R., and Uthemann, A. (2022). Artificial intelligence and systemic risk. *Journal of Banking & Finance*, 140:106290.
- Danielsson, J. and Uthemann, A. (2025). Artificial intelligence and financial crises. *Journal of Financial Stability*, 80:101453.
- Delle Monache, D., De Polis, A., and Petrella, I. (2024). Modeling and forecasting macroeconomic downside risk. *Journal of Business & Economic Statistics*, 42(3):1010–1025.
- Demarta, S. and McNeil, A. J. (2005). The t copula and related copulas. *International Statistical Review*, 73(1):111–129.
- Embrechts, P., McNeil, A., Straumann, D., and Dempster, M. (2002). Risk management: value at

- risk and beyond. *Correlation and dependence in risk management: properties and pitfalls*.
- Fernández, A. (2019). Artificial intelligence in financial services. *Banco de Espana Article*, 3:19.
- FSB (2024). The Financial Stability Implications of Artificial Intelligence. Technical report, Financial Stability Board .
- Fuster, A., Plosser, M., Schnabl, P., and Vickery, J. (2019). The role of technology in mortgage lending. *The Review of Financial Studies*, 32(5):1854–1899.
- Gambacorta, L., Sabatini, F., and Schiaffi, S. (2025). Artificial intelligence and relationship lending. *Bank of Italy Temi di Discussione (Working Paper) No*, 1476.
- Ghysels, E., Plazzi, A., and Valkanov, R. (2016). Why invest in emerging markets? the role of conditional return asymmetry. *The Journal of Finance*, 71(5):2145–2192.
- Girardi, G. and Ergün, A. T. (2013). Systemic risk measurement: Multivariate GARCH estimation of CoVaR. *Journal of Banking & Finance*, 37(8):3169–3180.
- Gómez, H. W., Torres, F. J., and Bolfarine, H. (2007). Large-sample inference for the epsilon-skew-t distribution. *Communications in Statistics—Theory and Methods*, 36(1):73–81.
- Gu, S., Kelly, B., and Xiu, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies*, 33(5):2223–2273.
- Han, W.-Z. and Meng, W. (2025). Does ai contribute to systemic risk reduction in non-financial corporations? *The Quarterly Review of Economics and Finance*, 100:101973.
- Hansen, A. L. and Lee, S. J. (2025). Financial stability implications of generative ai: Taming the animal spirits. Technical Report 2025-090, Board of Governors of the Federal Reserve System, Washington, DC.
- Hinkley, D. V. (1975). On power transformations to symmetry. *Biometrika*, 62(1):101–111.
- Joe, H. (1997). *Multivariate Models and Dependence Concepts*. Number 73 in Monographs on Statistics and Applied Probability. Chapman and Hall, London.
- Joe, H. and Xu, J. J. (1996). The estimation method of inference functions for margins for multivariate models. Technical report, Department of Statistics, University of British Columbia.
- Leitner, G., Singh, J., van der Kraaij, A., and Zsámboki, B. (2024). The rise of artificial intelligence: Benefits and risks for financial stability. In *Financial Stability Review*. European Central Bank. Special Feature B, May 2024 issue.
- Lin, X., Zhang, S. S., and Zachariadis, M. (2025). Open data and api adoption of us banks. *Journal of Financial Intermediation*, 63:101162.
- Liu, R. and Pun, C. S. (2022). Machine-learning-enhanced systemic risk measure: A two-step supervised learning approach. *Journal of Banking Finance*, 136:106416.
- Lopez, J. C. (2026). Responsible use of artificial intelligence in financial market infrastructures. *Journal of Financial Market Infrastructures*, 13(2):1–14.
- Lopez-Lira, A. and Tang, Y. (2026). Can chatgpt forecast stock price movements? return predictability and large language models. *Journal of Financial Economics*. Forthcoming.
- Lucas, A., Schwaab, B., and Zhang, X. (2014). Conditional euro area sovereign default risk. *Journal of Business & Economic Statistics*, 32(2):271–284.
- Lucas, A., Schwaab, B., and Zhang, X. (2017). Modeling financial sector joint tail risk in the euro area. *Journal of Applied Econometrics*, 32(1):171–191.
- Mudholkar, G. S. and Hutson, A. D. (2000). The epsilon-skew-normal distribution for analyzing near-normal data. *Journal of statistical planning and inference*, 83(2):291–309.

- Oehler, A. and Horn, M. (2024). Does chatgpt provide better advice than robo-advisors? *Finance Research Letters*, 60:104898.
- Oh, D. H. and Patton, A. J. (2018). Time-varying systemic risk: Evidence from a dynamic copula model of CDS spreads. *Journal of Business & Economic Statistics*, 36(2):181–195.
- Oh, D. H. and Patton, A. J. (2023). Dynamic factor copula models with estimated cluster assignments. *Journal of Econometrics*, 237(2):105374.
- Ojea-Ferreiro, J. (2020). Disentangling the role of the exchange rate in oil-related scenarios for the european stock market. *Energy Economics*, 89:104776.
- Ojea-Ferreiro, J. (2025). A Market-Based Approach to Reverse Stress Testing the Financial System. Technical report, Bank of Canada.
- Ojea-Ferreiro, J. (2026). Tail market linkage between canadian banks and non-bank financial intermediaries. *Finance Research Letters*, page 109820.
- Ojea-Ferreiro, J. and Reboredo, J. C. (2022). Exchange rates and the global transmission of equity market shocks. *Economic Modelling*, 114:105914.
- Ojea-Ferreiro, J., Reboredo, J. C., and Ugolini, A. (2024). Systemic risk effects of climate transition on financial stability. *International Review of Financial Analysis*, 96:103722.
- Potrykus, M. (2024). Dot-com and ai bubbles: Can data from the past be helpful to match the price bubble euphoria phase using dynamic time warping? *Finance Research Letters*, 67:105799.
- Rammer, C., Fernández, G. P., and Czarnitzki, D. (2022). Artificial intelligence and industrial innovation: Evidence from german firm-level data. *Research Policy*, 51(7):None.
- Reboredo, J. C., Rivera-Castro, M. A., and Ugolini, A. (2016). Downside and upside risk spillovers between exchange rates and stock prices. *Journal of Banking & Finance*, 62:76–96.
- Reboredo, J. C. and Ugolini, A. (2015). Systemic risk in european sovereign debt markets: A covar-copula approach. *Journal of International Money and Finance*, 51:214–244.
- Reboredo, J. C. and Ugolini, A. (2016). Systemic risk of spanish listed banks: a vine copula covar approach. *Spanish Journal of Finance and Accounting/Revista Española de Financiación y Contabilidad*, 45(1):1–31.
- Rishabh, K. and Shreeti, V. (2026). The geography of ai firms. Technical report, Bank for International Settlements.
- Sheng, J., Sun, Z., Yang, B., and Zhang, A. L. (2025). Generative ai and asset management. *Review of Financial Studies*, forthcoming.
- Siddik, A. B., Li, Y., Du, A. M., and Migliavacca, M. (2025). Fueling financial development: The crucial role of generative ai financing across nations. *Finance Research Letters*, 72:106519.
- Sklar, M. (1959). Fonctions de répartition à n dimensions et leurs marges. In *Annales de l'ISUP*, volume 8, pages 229–231.
- Smith, M. S. (2023). Implicit copulas: An overview. *Econometrics and Statistics*, 28:81–104.
- Weitzner, G. (2024). Reputational algorithm aversion. Papers 2402.15418, arXiv.org.
- Zhang, J., Luo, L., and Yang, J. W. (2026). Does artificial intelligence mitigate climate change exposure? *Journal of Banking & Finance*, 183:107623.
- Zheng, M., Zhang, L., Tripe, D., and Zhang, Y. (2025). Can artificial intelligence mitigate green-washed green credit? evidence from loan contracts of chinese listed firms. *International Review of Financial Analysis*, page 104948.

Appendix A. Appendix

A.1 Marginal behaviour

The marginal densities of equity returns is characterized by an ARMA(p,q)-GARCH(d,m) model, i.e.,

$$\begin{aligned} r_{i,t} &= \mu_{i,t} + \epsilon_{i,t}, \\ \epsilon_{i,t} &= \sigma_{i,t} \varepsilon_{i,t}, \\ \mu_{i,t} &= \phi_0 + \sum_{j=1}^p \phi_j r_{i,t-j} + \sum_{h=1}^q \varphi_h \epsilon_{i,t-h}, \\ \sigma_{i,t}^2 &= \omega + \sum_{k=1}^m \alpha_k \epsilon_{i,t-k}^2 + \sum_{h=1}^d \beta_h \sigma_{i,t-h}^2, \end{aligned}$$

where ϕ_j and ψ_k are the parameters of the AR and MA components of the marginal model, ω , α_k , β_k are the components of the GARCH. The standardized innovation ε_i follow a Skew-t (Arellano-Valle et al. 2005; Gómez et al. 2007) distribution, which captures the skewness and excess of kurtosis that we might find in financial returns. The density function is

$$f(\varepsilon_{i,t}; \eta_{i,t}, \nu_i) = \begin{cases} c(\nu_i) \left(1 + \frac{1}{\nu_i} \left(\frac{\varepsilon_{i,t}}{1+\eta_{i,t}} \right)^2 \right)^{-\frac{\nu_i+1}{2}} & \text{for } \varepsilon_{i,t} < 0 \\ c(\nu_i) \left(1 + \frac{1}{\nu_i} \left(\frac{\varepsilon_{i,t}}{1-\eta_{i,t}} \right)^2 \right)^{-\frac{\nu_i+1}{2}} & \text{for } \varepsilon_{i,t} > 0 \end{cases},$$

where $c(\nu_i) = \frac{\Gamma(\frac{\nu_i+1}{2})}{(\nu_i \pi)^{1/2} \Gamma(\nu_i/2)}$, the number of degrees of freedom ν_i must be higher than 3 and the parameter of asymmetry $\eta_{i,t}$ could take a value between -1 and 1. This distribution converges to the Gaussian distribution when $\eta_{i,t} = 0$ and $\nu_i \rightarrow \infty$ if $\eta_{i,t} \neq 0$ we obtained the epsilon-Skew-Gaussian distribution of Mudholkar and Hutson (2000) and the symmetric Student-t when the number of degrees of freedom are finite and $\eta_{i,t} = 0$.

To analyse changes in the shape of the skewness over time, we allow the skewness parameter of the distribution to evolve dynamically. Following the Generalized Autoregressive Score (GAS) framework proposed by Delle Monache et al. (2024), the transformed skewness parameter is updated using the scaled score of the conditional log-likelihood. Specifically,

$$\tau_{i,t+1} = \omega_i + \alpha_i s_{i,t} + \beta_i \tau_{i,t}, \quad (\text{A1})$$

$$\eta_{i,t} = \frac{1 - \exp(-\tau_{i,t})}{1 + \exp(-\tau_{i,t})}, \quad (\text{A2})$$

where $s_{i,t}$ denotes the scaled score of the conditional log-likelihood with respect to the skewness parameter. The transformation maps $\tau_{i,t} \in \mathbb{R}$ into $\eta_{i,t} \in (-1, 1)$, thereby guaranteeing that the skewness parameter remains in its admissible domain. In contrast to Delle Monache et al. (2024), who employ the inverse hyperbolic tangent transformation, we use the alternative mapping above, which we found to provide greater numerical stability while preserving the bounded support of the skewness parameter.

A.2 Skewed Student-t Distribution and Copula Properties

The multivariate Skewed Student-t distribution discussed in Demarta and McNeil (2005), denoted by $ST(\mu, P, \lambda, \nu)$, has density

$$f_X(x) = c \frac{K_{\frac{\nu+N}{2}} \left(\sqrt{(\nu + d(x)) \lambda' P^{-1} \lambda} \right) \exp([x - \mu]' P^{-1} \lambda)}{\left(\sqrt{(\nu + d(x)) \lambda' P^{-1} \lambda} \right)^{-\frac{\nu+N}{2}} \left(1 + \frac{d(x)}{\nu} \right)^{\frac{\nu+N}{2}}}, \quad (\text{A3})$$

where

$$c = \frac{2^{\frac{2-(\nu+N)}{2}}}{\Gamma\left(\frac{\nu}{2}\right) (\pi\nu)^{\frac{N}{2}} |P|^{1/2}}, \quad d(x) = (x - \mu)' P^{-1} (x - \mu),$$

and $K_a(\cdot)$ denotes the modified Bessel function of the second kind. The skewness vector λ controls asymmetry in the joint distribution, while ν governs tail thickness. A finite variance requires $\nu > 4$.

The flexibility induced by the skewness parameters is reflected in a wide range of dependence structures. Figure A1 illustrates how different values of λ_1 and λ_2 generate asymmetric tail dependence even for fixed correlation ρ and degrees of freedom ν .

An equivalent stochastic representation of the Skewed Student-t distribution is given by

$$X = \mu + \lambda W + \sqrt{W} Z, \quad (\text{A4})$$

where $W \sim IG(\nu/2, \nu/2)$ and $Z \sim N(0, R)$ are independent. Conditionally on W ,


$$X|W \sim N(\mu + W\lambda, WR),$$

which implies that the density and distribution functions can be written as scale mixtures of Gaussian distributions:

$$f_X(x) = \int_0^\infty \phi_{X|W}(x|w) f_W(w) dw, \quad (\text{A5})$$


$$F_X(x) = \int_0^\infty \Phi_{X|W}(x|w) f_W(w) dw. \quad (\text{A6})$$

FIGURE A1. Skewed Student-t copula density for different values of λ_1 and λ_2 , with $\rho = 0.8$ and $\nu = 5$.



This representation is particularly convenient for simulation and for illustrating dependence patterns. Figure A2 shows simulated observations from the Skewed Student-t copula under uniform and Gaussian marginals, highlighting asymmetric tail dependence that cannot be captured by symmetric elliptical copulas.

FIGURE A2. Simulation from a Skewed Student-t copula



The Skewed Student-t copula is an implicit copula (Smith 2023), defined as the ratio between the joint and marginal densities. Conditional copulas can be derived from Eq. (A6). In particular, the joint conditional copula is given by

$$C(u, v|q) = \int_0^\infty \Phi_{X_1, X_2|X_3, W} \left(F_{X_1}^{-1}(u), F_{X_2}^{-1}(v) \right) f_{W|X_3}(w) dw. \quad (\text{A7})$$

These conditional copulas constitute the key building block for the ES decomposition introduced in Section 2.3, whose derivation is reported in Appendix A.5.

A.3 Estimation process.

Estimation is implemented using a two-step procedure (Joe and Xu 1996). In a first step, the marginal distributions are estimated. In a second step, the dependence structure is estimated using the pseudo-probability integral transforms of the marginal distributions. Dependence between returns is modeled using a Skewed Student-t copula. The Skewed Student-t distribution has been widely employed in financial and economic applications (Lucas et al. 2014; Lucas et al. 2017; Oh and Patton 2023; Oh and Patton 2018; Ojea-Ferreiro 2025) due to its flexibility in capturing asymmetric tail dependence.

The Skewed Student-t copula is an implicit copula (Smith 2023), as it does not admit a closed-form expression. Instead, it is defined as the ratio between the joint and marginal

densities. Specifically, the copula density is given by

$$c_{ST}(u, v, q) = \frac{f_{ST}(F_X^{-1}(u), F_Y^{-1}(v), F_Z^{-1}(q))}{f_X(F_X^{-1}(u)) f_Y(F_Y^{-1}(v)) f_Z(F_Z^{-1}(q))},$$

where $f_{ST}(x, y, z)$ denotes the trivariate Skewed Student-t density with parameters $\mu = [0, 0, 0]'$,

$$P_t = \begin{bmatrix} 1 & \rho_{12,t} & \rho_{13,t} \\ \rho_{12,t} & 1 & \rho_{23,t} \\ \rho_{13,t} & \rho_{23,t} & 1 \end{bmatrix},$$

$\lambda = [\lambda_1, \lambda_2, \lambda_3]'$, and degrees of freedom ν . The marginal densities $f_X(\cdot)$, $f_Y(\cdot)$ and $f_Z(\cdot)$ are univariate Skewed Student-t distributions with $\mu = 0$, $P = 1$, skewness parameters λ_1 , λ_2 and λ_3 , respectively, and common ν .

This specification allows for flexible and asymmetric dependence patterns while preserving tractability in the construction of conditional distributions required for the risk measures analysed in this study. Additional properties of the Skewed Student-t distribution, as well as details on simulation, on the derivation of the conditional copulas, are provided in Appendix A.2.

A.4 Corrected Dynamic Conditional Correlation (cDCC) dynamics

We model the dependence between artificial intelligence (AI), financial sector, and real economy returns using a copula-based corrected Dynamic Conditional Correlation (cDCC) framework. Let $z_{i,t} = F^{-1}(u_{i,t})$ denote the inverse cumulative distribution function of the skewed Student-t distribution, mapping the probability integral transform variables $u_{i,t}$ into skewed-t innovations. Let $\mathbf{z}_t = (z_{1,t}, \dots, z_{N,t})'$. The time-varying dependence is captured through the conditional correlation matrix of $\mathbf{z}_t$, denoted $\mathbf{R}_t$. The cDCC dynamics are specified via the auxiliary matrix $\mathbf{Q}_t$:

$$\mathbf{Q}_t = (1 - a - b) \bar{\mathbf{Q}} + a \tilde{\mathbf{z}}_{t-1} \tilde{\mathbf{z}}_{t-1}' + b \mathbf{Q}_{t-1},$$

where $\tilde{\mathbf{z}}_t$ denotes the variance-corrected innovations as in the cDCC framework, $\bar{\mathbf{Q}}$ is the unconditional correlation matrix of $\tilde{\mathbf{z}}_t$, and the parameters satisfy $a \geq 0$, $b \geq 0$, and $a + b < 1$. The correction ensures that the unconditional moments of the correlation process are properly centered around $\bar{\mathbf{Q}}$.

The conditional correlation matrix is obtained by normalization:

$$\mathbf{R}_t = \text{diag}(\mathbf{Q}_t)^{-1/2} \mathbf{Q}_t \text{diag}(\mathbf{Q}_t)^{-1/2}.$$

The resulting $\mathbf{R}_t$ governs the dependence structure of the copula, allowing for time-

varying correlations across AI returns, financial sector returns, and real economy returns. This framework combines flexible marginal modeling with a parsimonious and consistent specification of dynamic dependence.

A.5 Proof of ES decomposition for lower tail

This appendix derives the copula representation of the lower-tail ES decomposition stated in Eqs. (6)–(9), using the copula density in Eq. (4) and the conditional copulas in Eq. (5) introduced in Section 2.3.

Eq. (1), together with the four integrals of Section 2.1, defines each component over the joint distribution of the three returns. The copula reduces every one of them to a single integral over the quantile level of the financial return. We show this explicitly for the systemic component; the other three follow from the same chain. Throughout, all distributions are conditional on the information set available at $t - 1$ and are therefore time-varying; the time subscript on F_i , C and their conditional counterparts is omitted to keep the notation compact.

Starting from the systemic component as written in Section 2.1,

$$\begin{aligned}
& E\left(r_{i,t} \mid R_t^L \cap A_t^L \cap M_t^L\right) P\left(A_t^L \cap M_t^L \mid R_t^L\right) \\
&= \frac{1}{\alpha} \int_{-\infty}^{Q_{i,t}(\alpha)} r_{i,t} f\left(r_{i,t}, r_{a,t} \leq Q_{a,t}(\alpha), r_{m,t} \leq Q_{m,t}(\alpha)\right) dr_{i,t} \\
&= \frac{1}{\alpha} \int_{-\infty}^{Q_{i,t}(\alpha)} r_{i,t} P\left(r_{a,t} \leq Q_{a,t}(\alpha), r_{m,t} \leq Q_{m,t}(\alpha) \mid r_{i,t}\right) f_i(r_{i,t}) dr_{i,t} \quad (\text{A8}) \\
&= \frac{1}{\alpha} \int_0^\alpha F_i^{-1}(q) P\left(r_{a,t} \leq Q_{a,t}(\alpha), r_{m,t} \leq Q_{m,t}(\alpha) \mid F_i(r_{i,t}) = q\right) dq \\
&= \frac{1}{\alpha} \int_0^\alpha F_i^{-1}(q) C_{a,m|i}(\alpha, \alpha \mid q) dq.
\end{aligned}$$

The four equalities correspond to the following steps. The first is the definition given in Section 2.1: the loss is integrated over the distress region of institution i , while the joint density restricts the integration to the episodes in which the AI sector and the real economy are in distress as well. The second factorises that joint density into the marginal density of the financial return and the conditional probability that the other two markets are in distress given that return; it is the product rule, and no assumption is involved. The third is the change of variable $q = F_i(r_{i,t})$, which does two things at once: it maps the distress region $(-\infty, Q_{i,t}(\alpha)]$ onto $(0, \alpha]$, and it absorbs the marginal density, since $dq = f_i(r_{i,t}) dr_{i,t}$. This is why no density survives in the last line: the only trace left of the marginal distribution of the financial return is its quantile function $F_i^{-1}(q)$. The fourth recognises the conditional probability as a conditional copula: because $F_a(Q_{a,t}(\alpha)) = F_m(Q_{m,t}(\alpha)) = \alpha$, the distress events of the AI and macroeconomic indices are $\{u \leq \alpha\}$ and $\{v \leq \alpha\}$, so that probability

is exactly $C_{a,m|i}(\alpha, \alpha | q)$ of Eq. (5).

Read downwards, Eq. (A8) replaces a trivariate integration problem with a univariate one. Read upwards, it says that each tail loss $F_i^{-1}(q)$ is weighted by the probability that it occurs jointly with distress in the other two markets. The AI and macroeconomic returns are never integrated out explicitly: as in Reboredo and Ugolini (2015), they enter only through the copula evaluated at their quantile levels.

The other three components follow from the very same chain, because only the event entering the third line changes; the change of variable is unaffected. Their conditional probabilities are obtained from $C_{a,m|i}$ by inclusion–exclusion. Writing $\pi^s(q)$ for the conditional probability of state s given q ,

$$\begin{aligned}\pi^{\text{financial}}(q) &= P(\bar{A}_t^L \cap \bar{M}_t^L | q) = 1 - C_{a|i}(\alpha | q) - C_{m|i}(\alpha | q) + C_{a,m|i}(\alpha, \alpha | q), \\ \pi^{\text{AI-sector}}(q) &= P(A_t^L \cap \bar{M}_t^L | q) = C_{a|i}(\alpha | q) - C_{a,m|i}(\alpha, \alpha | q), \\ \pi^{\text{macroeconomic}}(q) &= P(\bar{A}_t^L \cap M_t^L | q) = C_{m|i}(\alpha | q) - C_{a,m|i}(\alpha, \alpha | q), \\ \pi^{\text{systemic}}(q) &= P(A_t^L \cap M_t^L | q) = C_{a,m|i}(\alpha, \alpha | q),\end{aligned}\tag{A9}$$

where the second and third lines subtract, from the conditional probability of one index being in distress, the event in which both are, and the first is the conditional survival probability of the pair. Substituting each $\pi^s(q)$ into the last line of Eq. (A8) gives the four components of Eq. (1) in a single expression,

$$ES_{i,t}^{s,L} = \frac{1}{\alpha} \int_0^\alpha F_i^{-1}(q) \pi^s(q) dq, \quad s \in \{\text{financial}, \text{AI-sector}, \text{macroeconomic}, \text{systemic}\},\tag{A10}$$

following the same order as the four terms of Eq. (1). The losses are the same in all four components: what differs is only the probability mass attached to them.

Since the four states partition the sample space, the weights are non-negative and add up to one at every quantile level $q \in (0, 1)$,

$$\pi^{\text{financial}}(q) + \pi^{\text{AI-sector}}(q) + \pi^{\text{macroeconomic}}(q) + \pi^{\text{systemic}}(q) = 1,\tag{A11}$$

so that $\sum_s ES_{i,t}^{s,L} = \alpha^{-1} \int_0^\alpha F_i^{-1}(q) dq = ES_{i,t}^L$. The decomposition is therefore exact for any copula: no tail loss is left unallocated, and the copula governs how Expected Shortfall is distributed across states, not its level.

The role of dependence is best seen in π^{systemic} . Averaging it over the tail returns the unconditional weight, $P(A_t^L \cap M_t^L | R_t^L) = \alpha^{-1} \int_0^\alpha \pi^{\text{systemic}}(q) dq$, which measures the share of the distress episodes of institution i in which the AI sector and the real economy are simultaneously in distress. If the financial return were independent of the other two markets, $\pi^{\text{systemic}}(q)$ would not depend on q and would coincide with the unconditional

probability of joint distress $P(A_t^L \cap M_t^L)$, which under full independence of the three markets would in turn reduce to α^2 . What makes the systemic component large is therefore the extent to which $\pi^{\text{systemic}}(q)$ rises as q approaches zero, that is, the extent to which joint AI and macroeconomic distress becomes more likely precisely at the deepest losses of institution i .

Taken together, these expressions make explicit how the copula representation reduces each ES component to a one-dimensional integral over the quantile level of the financial return, and how asymmetric dependence in the tails—driven by the copula skewness parameters λ and illustrated in Fig. A1—directly affects the allocation of Expected Shortfall across the different sources of systemic risk. For the Skewed Student-t copula adopted here, $C_{a,m|i}$ is not available in closed form but only as an integral over the mixing variable W (Eq. (A7)), so that each weight is itself an integral and every component is evaluated numerically as a double integral.

The corresponding decompositions for the upper tail and for the central region, which enter the skewness index in Eq. (2), are obtained from the same chain and are reported in Appendix A.6.

A.6 ES decomposition for the upper tail and the central region

The skewness index in Eq. (2) requires the partition of Eq. (1) to be applied not only to the lower tail of the return distribution, but also to its upper tail and to its central α region. Below we report both decompositions and their copula representations, which follow from the conditional copulas defined in Eq. (5).

A.6.1 Upper tail

An analogous decomposition can be obtained for upper-tail realizations, i.e. for the extreme gains of institution i . Let

$$R_t^U = \{r_{i,t} \geq Q_{i,t}(1 - \alpha)\}, \quad A_t^U = \{r_{a,t} \geq Q_{a,t}(1 - \alpha)\}, \quad M_t^U = \{r_{m,t} \geq Q_{m,t}(1 - \alpha)\},$$

denote the upper-tail events for the financial, AI and macroeconomic returns, respectively, where an overbar denotes the complement of the corresponding event, i.e. $\bar{A}_t^U = \{r_{a,t} < Q_{a,t}(1 - \alpha)\}$ and $\bar{M}_t^U = \{r_{m,t} < Q_{m,t}(1 - \alpha)\}$ indicate that the AI and macroeconomic indices are not in their respective upper tails. Applying the same partition to R_t^U yields

$$ES_{i,t}^U = E \left(\overbrace{r_{i,t} \mid R_t^U \cap \bar{A}_t^U \cap \bar{M}_t^U}^{\text{financial}} \right) P \left(\bar{A}_t^U \cap \bar{M}_t^U \mid R_t^U \right)$$

$$\begin{aligned}
& \overbrace{+ E\left(r_{i,t} \mid R_t^U \cap A_t^U \cap \bar{M}_t^U\right) P\left(A_t^U \cap \bar{M}_t^U \mid R_t^U\right)}^{\text{AI-sector}} \\
& \overbrace{+ E\left(r_{i,t} \mid R_t^U \cap \bar{A}_t^U \cap M_t^U\right) P\left(\bar{A}_t^U \cap M_t^U \mid R_t^U\right)}^{\text{macroeconomic}} \\
& \overbrace{+ E\left(r_{i,t} \mid R_t^U \cap A_t^U \cap M_t^U\right) P\left(A_t^U \cap M_t^U \mid R_t^U\right)}^{\text{systemic}}, \tag{A12}
\end{aligned}$$

whose copula formulation follows analogously, i.e.

$$\begin{aligned}
& \overbrace{E\left(r_{i,t} \mid R_t^U \cap \bar{A}_t^U \cap \bar{M}_t^U\right) P\left(\bar{A}_t^U \cap \bar{M}_t^U \mid R_t^U\right)}^{\text{financial}} \\
& \quad = \frac{1}{\alpha} \int_{1-\alpha}^1 F_i^{-1}(q) C_{a,m|i}(1-\alpha, 1-\alpha \mid q) dq, \\
& \overbrace{E\left(r_{i,t} \mid R_t^U \cap A_t^U \cap \bar{M}_t^U\right) P\left(A_t^U \cap \bar{M}_t^U \mid R_t^U\right)}^{\text{AI-sector}} \\
& \quad = \frac{1}{\alpha} \int_{1-\alpha}^1 F_i^{-1}(q) (C_{m|i}(1-\alpha \mid q) - C_{a,m|i}(1-\alpha, 1-\alpha \mid q)) dq, \\
& \overbrace{E\left(r_{i,t} \mid R_t^U \cap \bar{A}_t^U \cap M_t^U\right) P\left(\bar{A}_t^U \cap M_t^U \mid R_t^U\right)}^{\text{macroeconomic}} \\
& \quad = \frac{1}{\alpha} \int_{1-\alpha}^1 F_i^{-1}(q) (C_{a|i}(1-\alpha \mid q) - C_{a,m|i}(1-\alpha, 1-\alpha \mid q)) dq, \\
& \overbrace{E\left(r_{i,t} \mid R_t^U \cap A_t^U \cap M_t^U\right) P\left(A_t^U \cap M_t^U \mid R_t^U\right)}^{\text{systemic}} \\
& \quad = \frac{1}{\alpha} \int_{1-\alpha}^1 F_i^{-1}(q) (1 - C_{m|i}(1-\alpha \mid q) - C_{a|i}(1-\alpha \mid q) \\
& \quad \quad + C_{a,m|i}(1-\alpha, 1-\alpha \mid q)) dq.
\end{aligned}$$

Note that, since upper-tail events are the complements of the events entering the conditional copulas in Eq. (5), the roles of $C_{a|i}$ and $C_{m|i}$ are reversed relative to the lower-tail decomposition reported in Section ??: the financial component now collects the probability mass in which neither index exceeds its $(1-\alpha)$ quantile, whereas the systemic component is obtained by complementation.

A.6.2 Central region

The same partition can be applied to the central α region of the return distribution, which enters the skewness index in Eq. (2) as the benchmark against which the two tail allocations are evaluated. Let

$$L_\alpha = 0.5 - \frac{\alpha}{2}, \quad U_\alpha = 0.5 + \frac{\alpha}{2},$$

and define the central-region events as

$$R_t^C = \{Q_{i,t}(L_\alpha) \leq r_{i,t} \leq Q_{i,t}(U_\alpha)\}, \quad A_t^C = \{Q_{a,t}(L_\alpha) \leq r_{a,t} \leq Q_{a,t}(U_\alpha)\}, \quad M_t^C = \{Q_{m,t}(L_\alpha) \leq r_{m,t} \leq Q_{m,t}(U_\alpha)\},$$

where an overbar denotes the complement of the corresponding central-region event; for instance, $\bar{A}_t^C$ includes all AI-sector realizations outside the interval $[Q_{a,t}(L_\alpha), Q_{a,t}(U_\alpha)]$. Applying the same partition to the central region gives

$$\begin{aligned} \mu_{i,t}^C &= \overbrace{E(r_{i,t} | R_t^C \cap \bar{A}_t^C \cap \bar{M}_t^C) P(\bar{A}_t^C \cap \bar{M}_t^C | R_t^C)}^{\text{financial}} \\ &\quad + \overbrace{E(r_{i,t} | R_t^C \cap A_t^C \cap \bar{M}_t^C) P(A_t^C \cap \bar{M}_t^C | R_t^C)}^{\text{AI-sector}} \\ &\quad + \overbrace{E(r_{i,t} | R_t^C \cap \bar{A}_t^C \cap M_t^C) P(\bar{A}_t^C \cap M_t^C | R_t^C)}^{\text{macroeconomic}} \\ &\quad + \overbrace{E(r_{i,t} | R_t^C \cap A_t^C \cap M_t^C) P(A_t^C \cap M_t^C | R_t^C)}^{\text{systemic}}, \end{aligned} \tag{A13}$$

where $\mu_{i,t}^C = E(r_{i,t} | R_t^C)$ denotes the mean return of institution i over the central region, and the labels refer to the same partition applied to central realizations of the AI and macroeconomic indices, rather than to lower- or upper-tail distress events.

To ease notation, we define the conditional probabilities of the AI-sector and macroeconomic indices lying in their central regions as

$$\begin{aligned} \Delta_{a|i}(q) &= C_{a|i}(U_\alpha | q) - C_{a|i}(L_\alpha | q), \\ \Delta_{m|i}(q) &= C_{m|i}(U_\alpha | q) - C_{m|i}(L_\alpha | q), \\ \Delta_{a,m|i}(q) &= C_{a,m|i}(U_\alpha, U_\alpha | q) - C_{a,m|i}(L_\alpha, U_\alpha | q) \\ &\quad - C_{a,m|i}(U_\alpha, L_\alpha | q) + C_{a,m|i}(L_\alpha, L_\alpha | q), \end{aligned}$$

where $\Delta_{a|i}(q)$ and $\Delta_{m|i}(q)$ are the probabilities that the AI and the macroeconomic index

lie in their central regions, conditional on the financial return being at the quantile level q , and $\Delta_{a,m|i}(q)$ is the corresponding joint probability, obtained by inclusion–exclusion over the rectangle $[L_\alpha, U_\alpha]^2$. Since the central region of the financial return has probability $P(R_t^C) = U_\alpha - L_\alpha = \alpha$, the same normalization as in the tail decompositions applies, and the contribution of each state to $\mu_{i,t}^C$ can be written as

$$\begin{aligned}
& \overbrace{E\left(r_{i,t} \mid R_t^C \cap \bar{A}_t^C \cap \bar{M}_t^C\right) P\left(\bar{A}_t^C \cap \bar{M}_t^C \mid R_t^C\right)}^{\text{financial}} \\
&= \frac{1}{\alpha} \int_{L_\alpha}^{U_\alpha} F_i^{-1}(q) \left(1 - \Delta_{a|i}(q) - \Delta_{m|i}(q) + \Delta_{a,m|i}(q)\right) dq, \\
& \overbrace{E\left(r_{i,t} \mid R_t^C \cap A_t^C \cap \bar{M}_t^C\right) P\left(A_t^C \cap \bar{M}_t^C \mid R_t^C\right)}^{\text{AI-sector}} \\
&= \frac{1}{\alpha} \int_{L_\alpha}^{U_\alpha} F_i^{-1}(q) \left(\Delta_{a|i}(q) - \Delta_{a,m|i}(q)\right) dq, \\
& \overbrace{E\left(r_{i,t} \mid R_t^C \cap \bar{A}_t^C \cap M_t^C\right) P\left(\bar{A}_t^C \cap M_t^C \mid R_t^C\right)}^{\text{macroeconomic}} \\
&= \frac{1}{\alpha} \int_{L_\alpha}^{U_\alpha} F_i^{-1}(q) \left(\Delta_{m|i}(q) - \Delta_{a,m|i}(q)\right) dq, \\
& \overbrace{E\left(r_{i,t} \mid R_t^C \cap A_t^C \cap M_t^C\right) P\left(A_t^C \cap M_t^C \mid R_t^C\right)}^{\text{systemic}} \\
&= \frac{1}{\alpha} \int_{L_\alpha}^{U_\alpha} F_i^{-1}(q) \Delta_{a,m|i}(q) dq.
\end{aligned}$$

As in the lower-tail decomposition reported in Section ?? and in the upper-tail decomposition of Appendix A.5.1, the four components add up to $\mu_{i,t}^C$ by construction, so that the central-region allocation is directly comparable with the lower- and upper-tail allocations entering the skewness index in Eq. (2).

A.7 Extra Figures

Figure A3 reports the time-varying correlations between the AI, financial and real economy returns implied by the cDCC dynamics of Appendix A.4.




FIGURE A3. Time-varying correlation from a cDCC dynamics