the journal of PORTFOLIO management

QUANTITATIVE TOOLS FOR ASSET MANAGEMENT

PORTFOLIO MANAGEMENT RESEARCH

with. Intelligence

September 2026 | Volume 52 | Number 10

pm-research.com

Reframing Diversification as a Network Problem: Graph Neural Networks for Portfolio Management

Andrew Chin and Che Guan

ALLIANCEBERNSTEIN®

AllianceBernstein is a leading investment-management firm with $867 billion in client assets under management, as of 31 December 2025. With a unique combination of expertise across equities, fixed-income, alternatives and multi-asset strategies, we aim to deliver differentiated insights and distinctive solutions to advance investors' success. Across our global network, we're fully invested in delivering better outcomes for our worldwide clients, including institutional, high-net-worth and retail investors. By embracing innovation, we seek to address increasingly complex investing challenges and opportunities. And we pursue responsibility throughout the firm—from how we work and act to the solutions we design for clients.

Andrew Chin

Chief AI Officer

Andrew Chin is Chief AI Officer and a member of the firm's Operating Committee at AllianceBernstein (AB). In this role, he leads the firm's strategy to leverage AI in transforming the organization and driving better outcomes for clients and the firm. Previously, Andrew was the Head of Investment Solutions and Sciences, overseeing the research, management and strategic growth of the firm's asset allocation, data science, index and tax-management businesses. From 2022 to 2023, Andrew was the head of quantitative research and chief data scientist, developing and optimizing the quantitative research and data science infrastructure, capabilities and resources across the organization. As the firm's chief risk officer from 2009 to 2021, he led all aspects of risk management and built a global team to identify, manage and mitigate the various risks across the organization. Andrew has held various leadership roles in quantitative research, risk management and portfolio management in New York and London since joining the firm in 1997. Before joining AB, he spent three years as a project manager and business analyst in Global Investment Management at Bankers Trust.

Andrew is a Visiting Lecturer in Cornell's Master of Financial Engineering program, and serves on the Board of Directors of the Society of Quantitative Analysts and the editorial boards of *The Journal of Portfolio Management* and *The Journal of Financial Data Science*.

Andrew holds a BA in math and computer science, and an MBA in finance from Cornell University.

Che Guan

Principal Data Scientist

Che Guan is Principal Data Scientist at AB, where he actively engages in research and development in natural language processing. His expertise spans supervised fine-tuning of large language models, reinforcement learning alignment, graph neural networks, and generative and autoregressive AI models for time series forecasting. Che joined AB in 2021, following an accomplished career in data science and quantitative research. From 2018 to 2021, he served as Principal Data Scientist at Raymond James, where he specialized in portfolio rebalancing and optimization, multi-asset simulation, and robust asset allocation strategies. Prior to that, he worked as a quantitative strategist at a startup fund from 2017 to 2018, developing advanced algorithms for foreign exchange and digital assets. From 2013 to 2017, he was a founding member of J.P. Morgan Chase's Intelligent Solutions team, where he applied machine learning and big data technologies across various business lines. Che earned his MS in Data Science from the University of California, Berkeley, as well as an MS in Statistics and a PhD in Electrical Engineering from the University of Connecticut, with a focus on forecasting applications.

Reframing Diversification as a Network Problem: Graph Neural Networks for Portfolio Management

Andrew Chin and Che Guan

Andrew Chin
is the chief AI officer at AllianceBernstein LP in New York, NY and previously served as head of Investment Solutions and Sciences, chief risk officer and head of Quantitative Research for the firm.
andrew.chin@alliancebernstein.com

Che Guan
is principal data scientist at AllianceBernstein LP in Nashville, Tennessee.
che.guan@alliancebernstein.com

KEY FINDINGS

- Network-based representations enhance diversification analysis by making dependence, hidden concentrations, and shock transmissions more transparent.
- Many diversification failures reflect unobserved economic linkages rather than misestimated correlations, a distinction that networks make explicit.
- Portfolio managers can leverage Graph Neural Networks to operationalize proprietary graphs, making economic connectivity actionable for forecasting, portfolio construction, and risk assessment.

ABSTRACT

Diversification and portfolio analysis are traditionally framed through covariance, correlation, and factor exposures. While foundational, these tools summarize co-movement without revealing the structure through which information and risk propagate, particularly during periods of market stress. This article reframes diversification as a network problem. By viewing the market as an interconnected system of assets, network-based representations make dependence, hidden concentration, and shock transmission more transparent. Building on this intuition, the authors focus on Graph Neural Networks (GNNs) as a practical framework for learning from market connectivity. Through intuitive examples and portfolio-relevant applications, the article shows how GNNs formalize reasoning that portfolio managers already apply informally, making diversification analysis more transparent, repeatable, and scalable in modern asset management.

Diversification is one of the oldest and most widely taught principles in asset management. From Markowitz onward, the central idea has been that combining imperfectly correlated assets can reduce portfolio risk without sacrificing expected return. In practice, this principle is traditionally implemented through covariance matrices, correlation estimates, and factor models, all of which summarize how asset returns move together on average.

While powerful, these representations are typically derived from empirical and statistical relationships. They quantify co-movement but do not explain why assets move together or the underlying structure that links the movements. For example, a portfolio diversified across technology, industrials, energy, and consumer sectors may appear well balanced under traditional risk metrics yet still be highly exposed to a single underlying economic driver. As a result, investors often find that portfolios which appear diversified on paper become highly correlated in periods of stress, when diversification matters most. This gap between statistical diversification and structural diversification is a recurring challenge in portfolio construction.

Network-based representations provide a complementary way to think about diversification. Instead of viewing correlation as a table of pairwise numbers, a network representation treats assets as nodes and relationships as connections. This shift makes structure visible. Clusters of tightly connected assets, pathways through which shocks propagate, and concentrations of exposure that are difficult to see in a covariance matrix become immediately apparent when viewed as a network.

The purpose of this article is to explain how this network-based approach can enhance the understanding, practice, and application of diversification and portfolio analysis. Rather than presenting empirical results, the authors focus on providing a conceptual framework and practical applications. In doing so, they show how graph-based representations connect naturally to traditional covariance-based approaches, while offering additional insight into co-movement, risk concentration, and spillover effects that are central to contemporary asset management.

This market as a graph concept, representing financial markets as networks, has developed over several decades, spanning portfolio theory and systemic risk research. While early work used network representations primarily for visualization and diagnosis of dependence structure, recent advances in graph-based learning and neural networks, like Graph Neural Networks (GNNs), treat the market explicitly as a graph that can be used for forecasting and portfolio decision-making.

The subsequent sections reframe diversification through a network lens, discuss how graph-based representations address the shortcomings of classical techniques, highlight methods to construct market graphs, describe how GNNs can improve investment insights, and finally, provide practical examples to enable portfolio managers to enhance decision-making using these tools.

DIVERSIFICATION THROUGH A NETWORK LENS

Traditional diversification relies on covariance matrices. Each entry measures the degree to which two assets move together, and portfolio risk is computed by aggregating these pairwise relationships. While mathematically sound, this representation may be difficult to interpret and provides little intuition about the underlying structure of dependencies in a portfolio. In addition, some of the relationships found in the empirical data may be spurious.

A network representation reframes the same information in a more intuitive way. Assets are represented as nodes, and relationships between them are represented as connections. Stronger relationships correspond to stronger connections. The resulting network provides a visual and conceptual map of how assets are linked, with the edges suggesting economic linkages among the nodes. Viewed as a network, a seemingly sector-diversified portfolio may reveal a dense cluster of firms connected through shared suppliers, capital investment cycles, or technological dependencies, despite belonging to different sectors.

This perspective highlights an important distinction: diversification is not about the number of assets held, but about the number of independent economic and fundamental themes represented in a portfolio. Two assets may appear diversified because they belong to different sectors or regions, yet be tightly connected through shared suppliers, common ownership, or sensitivity to the same macroeconomic shock. Conversely, assets within the same sector may provide meaningful diversification if they occupy different positions in the network of economic relationships.

From a conceptual perspective, networks can uncover important insights from the underlying market structure:

- Hidden concentration—A portfolio that appears diversified by sector or region weights may still be concentrated in a tightly connected cluster of assets.
- Shock propagation—When an adverse event affects one asset, the impact often spreads along specific pathways rather than uniformly across the market.
- Fragility—Portfolios dominated by highly connected nodes tend to be more fragile during stress, while portfolios with exposure to peripheral nodes may be more resilient.

Importantly, network representations do not replace covariance-based analysis. Rather, they provide a complementary lens that emphasizes structure and connectivity, helping investors reason about diversification in economic rather than purely statistical terms.

GRAPHS ADDRESS THE LIMITATIONS OF EXISTING MODELS

Modern portfolio theory and quantitative asset management have been built on a sequence of increasingly sophisticated models, with each innovation delivering genuine insight and practical value. However, when viewed through the lens of diversification and asset return co-movements, they share a common limitation: they struggle to represent structure.

LIMITATIONS OF CLASSICAL LINEAR ECONOMETRIC MODELS

Linear econometric models, including mean–variance optimization and factor models, formalize dependence by quantifying how assets move together on average. These models are typically estimated with historical data and they remain foundational for portfolio construction and risk management.

Yet covariance-based representations are inherently abstract. They summarize pairwise relationships from empirical movements. Two portfolios may exhibit similar aggregate risk metrics while having very different underlying dependency structures. In practice, this often manifests as portfolios that appear diversified ex ante but behave as a single concentrated position during periods of market stress.

More fundamentally, linear econometric models assume that relationships among assets are stable, symmetric, and well captured by variance and correlations. In real markets however, dependence evolves over time. Supply-chain disruptions, ownership concentration, regulatory actions, and narrative-driven sentiment create channels of co-movement that are difficult to capture with static covariance estimates or linear factor exposures alone.

As a result, classical models are more effective at measuring diversification after the fact, but less effective at explaining why diversification breaks down or how co-movement is likely to change as conditions evolve.

LIMITATIONS OF TEMPORAL-ONLY DEEP LEARNING MODELS

Deep learning models such as recurrent neural networks and long short-term memory networks address some of the limitations of linear models. By learning non-linear patterns directly from data, they improve short-horizon forecasting and adapt more flexibly to changing regimes. In many applications, they outperform classical econometric approaches.

However, many deep learning models used in finance remain fundamentally temporal. They process each asset's time series independently, learning from its own history but not explicitly from its relationships with other assets. Cross-asset information enters these models only indirectly, typically through shared market regimes or common features created by the researcher.

This design choice limits their ability to model how information and risk propagate across assets. A temporal model may detect momentum or mean reversion in an individual security, but it does not inherently capture how shocks travel across a set of assets. As a result, temporal-only models often react to changes in co-movement rather than anticipate them.

From a diversification perspective, this is a critical shortcoming. Portfolio risk is not driven solely by the behavior of individual assets over time, but by how those assets interact with one another at each point in time.

GRAPHS ADD STRUCTURE

What classical econometrics and temporal-only deep learning have in common is that they treat structure implicitly. Relationships among assets are summarized statistically or absorbed into model parameters. In practical terms, fundamental or economic perspectives from portfolio managers regarding market dynamics and transmission pathways are difficult to embed into the underlying traditional models.

This is precisely where network-based representations can add value. By making relationships explicit, networks provide a way to reason about diversification in structural terms, capturing many of the economic linkages that portfolio managers use in their analysis and research.

FROM NETWORK REPRESENTATIONS TO LEARNING ON GRAPHS

Network-based representations of asset return co-movements are powerful descriptive tools by making patterns of dependence visible in a way that covariance matrices do not. Clusters of tightly connected assets, bridges between groups, and concentrations of exposure that are difficult to infer from tables of correlations become immediately apparent when viewed as a network. These perspectives are critical for both diversification and portfolio construction.

However, descriptive network representations have an important limitation. They are inherently static and diagnostic. They summarize relationships over a given window, but they do not, by themselves, answer forward-looking questions that portfolio managers care about. They do not forecast returns, anticipate shifts in co-movement, or quantify how shocks are likely to propagate as market conditions evolve. As a result, they are most often used to explain outcomes after the fact rather than to inform decisions in real time.

This distinction between understanding structure and operationalizing it motivates the move from network representation to learning on graphs.

Graph-based learning treats the network not merely as a visualization or summary statistic, but as an input to the modeling process itself. Instead of asking only how assets have moved together in the past, these methods ask how information embedded in the network can improve forecasts, portfolio construction, and risk assessment going forward. GNNs are the most prominent class of models built for this purpose. Broadly, while networks help investors see dependence, GNNs help them act on it.

At a high level, the distinction can be stated simply. Traditional models learn from asset-specific histories and, sometimes, from global factors. Network representations reveal how assets are connected. GNNs combine these ideas by allowing models to learn from both an asset's own history and the evolving state of its economic neighbors. In doing so, they operationalize the same intuition that network visualizations provide, but in a form that can be embedded directly into predictive and decision-making workflows.

This perspective is particularly natural in finance. Information rarely arrives uniformly across the market. It often appears first in one part of the network and then diffuses outward. A disruption at a supplier affects customers before it appears in their prices. A liquidity shock at a large investor leads to selling pressure across holdings that may otherwise seem unrelated. A regulatory or geopolitical event reshapes relationships across sectors and regions. Network representations help explain these patterns. GNNs allow models to learn from them to make better predictions.

It is important to emphasize that learning on graphs builds upon network-based intuition. An investor does not need GNNs to understand that diversification can fail when assets are tightly connected. But once that insight is accepted, GNNs offer a systematic way to incorporate connectivity into forecasting, portfolio construction, and risk analysis at scale.

THE MARKET AS GRAPH: PRACTICAL CONSTRUCTION CHOICES

Adopting a market as a graph perspective requires an explicit decision about how relationships among assets are represented. Unlike physical or social networks, where connections are often directly observed, financial networks are largely inferred. The graph is therefore not necessarily an objective truth about the market, but a model that reflects the researcher's beliefs about how information, risk, and returns propagate.

For practitioners, the key question centers on constructing a graph that is consistent with their investment philosophies and leverages their unique edge into the markets. Different graph constructions emphasize different aspects of market structure and serve different investment objectives. In a financial graph, assets are represented as nodes, and edges represent relationships that cause returns to move together or risk to propagate. These relationships may be direct or indirect, stable or time-varying, and symmetric or directional. Importantly, edges are not limited to correlations. They can represent economic dependence, common exposure, or channels through which shocks travel.

From a portfolio perspective, a financial graph provides a structured view of diversification. Highly connected nodes represent concentrations of shared risk, while loosely connected nodes may offer more independent sources of return. Changes in the graph over time often signal shifts in market regime, even before they are fully reflected in aggregate risk measures or market prices.

In practice, most applications fall into three broad categories: graphs based on economic relationships, graphs implied by market data, and graphs derived from information and narratives.

Economic Graphs: Encoding Fundamental Relationships

Economic graphs are built from observable relationships that link firms through real economic activity. Common examples include supply chains, ownership structures, sector classifications, and governance links.

Supply-chain graphs are particularly intuitive. A disruption at an upstream supplier often affects downstream firms. Representing these relationships explicitly allows investors to reason about second- and third-order effects that are difficult to capture with correlation alone. Ownership-based graphs capture a different mechanism: common investors can create co-movement through liquidity needs and rebalancing, even when firms are economically unrelated.

Economic graphs tend to be relatively stable and interpretable, which makes them attractive for institutional use. Their limitation is that they are often incomplete. Not all economically relevant relationships are observable, and their importance can vary across market conditions.

Market-Implied Graphs: Learning from Co-Movement

Market-implied graphs infer relationships directly from return data. The most common approach is to construct graphs from rolling correlations or other measures of dependence. These graphs reflect how assets have moved together in recent history or during certain regimes and are closely tied to traditional covariance-based portfolio analysis.

Viewed through a network lens, correlation-based graphs provide immediate insight into diversification. Assets that appear distinct in a portfolio may form tightly connected clusters, indicating shared exposure. During periods of market stress, these clusters often expand and merge, revealing why diversification breaks down when it is most needed.

Market-implied graphs adapt quickly to changing conditions, but they are also noisy and sensitive to choice of lookback period. While researchers have decades of market data, markets are non-stationary, and the various regimes over time challenge their ability to create robust models over different cycles. For this reason, these graphs are best viewed as complements to economic graphs rather than replacements.

Information Graphs: Capturing Narratives from Unstructured Data

With more advanced natural language processing techniques, investors now have the ability to systemically capture insights from significantly more data sources. Information graphs can be constructed from unstructured data sources such as news articles, earnings calls, and regulatory filings in multi-modal formats like text, audio, and video. These graphs capture relationships formed through shared narratives, events, or themes rather than direct economic ties.

For example, firms mentioned together in news coverage often experience correlated price movements as investors process information collectively from the content. Regulatory actions, geopolitical developments, or technological themes can link assets across sectors and regions in ways that are not visible in fundamentals or correlations alone.

Information graphs are particularly useful for understanding short-term co-movement and event-driven risk, especially when synthesizing news articles. However, these graphs can also be used to capture longer-term relationships using data sources that inform deeper ties amongst assets, like strategic alliances or M&A activity.

Choosing and Combining Graphs in Practice

No single graph captures all relevant aspects of market structure, and portfolio managers often combine multiple graph types, either explicitly or implicitly.

Economic graphs provide stability and interpretability. Market-implied graphs capture evolving co-movement. Information graphs reflect narrative-driven dynamics.

Ultimately, graph construction is a modeling choice that should be aligned with the portfolio manager's investment process. The goal is not to engineer the most complex graph, but to represent the relationships that matter most for the decision-making framework and reflect the skills of the investment team. For example, a healthcare investment team may use a graph to capture their proprietary insights into the sector's drug pipelines and competitive forces.

GRAPH NEURAL NETWORKS IN FINANCIAL MARKETS

GNNs extend network-based views of financial markets from description to decision support. Where network representations help visualize structure, GNNs allow market connectivity itself to influence forecasts, portfolio construction, and risk assessment. Their defining feature is that relationships among assets are treated as direct inputs to modeling rather than as background context.

An important practical capability of modern GNN architectures is their ability to model time and structure jointly. Financial markets evolve along two dimensions: assets change over time, and they interact with one another at each point in time. GNNs are designed to learn this time and cross-asset structure in a non-linear manner, allowing complex patterns of co-movement across multiple assets to emerge from data rather than being imposed through linear assumptions.

In addition, many contemporary GNN architectures incorporate attention mechanisms, enabling models to weight relationships dynamically and to integrate heterogeneous sources of information. This makes it possible to combine traditional market data with alternative data and to represent markets with different types of nodes, such as individual stocks and sectors, connected by multiple types of relationships. These relationships can be aggregated hierarchically, reflecting how portfolio managers reason across securities, industries, and themes.

Learning from Neighbors: Message Passing as Economic Intuition

At the core of a GNN is a simple idea: assets influence one another. Information does not arrive uniformly across the market, and shocks rarely remain isolated. GNNs capture this by allowing each asset's representation to be informed not only by its own characteristics, but also by the state of related assets.

In practice, this process resembles how investors reason about markets. When news affects a supplier, portfolio managers immediately consider the implications for customers. When volatility increases in one sector, they reassess risk across related industries. GNNs formalize this reasoning by allowing information to flow across connections in the graph, so that changes in one part of the market influence expectations elsewhere.

The result is a set of asset representations that reflect both individual behavior and broader market context. Compared to models that treat assets independently, these representations tend to be more stable and less sensitive to noise, particularly in cross-sectional applications.

Combining History and Structure

Financial markets are shaped by both time and relationships. Momentum, mean reversion, and volatility clustering are inherently temporal phenomena.

Spillovers, contagion, and diversification failures are inherently structural, affecting assets across the market at the same time. Many traditional models focus on one dimension at the expense of the other.

From a practical standpoint, this distinction is crucial. A sudden price move in a peripheral asset may carry little information for the rest of the portfolio. The same move in a highly connected asset may signal a regime change. GNNs are explicitly designed to capture this nuance, allowing them to incorporate non-linear co-movement across assets in a way that is difficult to achieve with purely temporal or purely cross-sectional models.

Adapting to Changing Market Structure

One of the most challenging aspects of portfolio management is that relationships among assets are not static. Correlations rise and fall, supply chains are reconfigured, ownership structures change, and narratives evolve. Models that assume stable relationships can perform well in calm periods and fail abruptly during transitions.

GNNs address this challenge by operating on graphs that can evolve over time. As the structure of the network changes, so do the pathways through which information flows. This makes GNN-based signals particularly useful during periods of market stress, thematic rotations, or event-driven dislocations, when traditional diversification assumptions tend to break down.

GRAPH NEURAL NETWORKS FOR PORTFOLIO MANAGERS

The primary value of network-based thinking and GNNs lies not in methodological novelty, but in how they improve real investment decisions. For portfolio managers, the most relevant applications are those that enhance understanding of co-movement, improve diversification, and provide earlier or more stable signals during periods when traditional models struggle.

Importantly, for institutional adoption, the outputs of GNNs should fit naturally into existing workflows. These signals are most effective when used to augment existing processes and tools, refining cross-sectional decisions, improving diversification diagnostics, or providing early warnings of changing co-movement patterns.

These benefits need to be assessed against the challenges in implementing GNNs. These tools are complex and require investment teams to carefully design the graph and operationalize the network to augment investment decisions. In addition, GNNs are still models and likely will not reveal the true structure of the market. As a result, economic judgment and investment reasoning remain the key levers to create competitive advantage amongst portfolio managers.

Ultimately, the effectiveness of a GNN depends critically on how the underlying graph is constructed and how its outputs are interpreted and applied. In this sense, GNNs are best viewed as tools that amplify network-based intuition rather than as black-box predictors.

Understanding and Anticipating Return Co-Movements

One of the most immediate benefits of GNNs is improved understanding of how and why assets move together. They can be used to make cross-sectional rankings and relative-value decisions rather than point forecasts. Portfolio managers are often less concerned with predicting exact returns than with identifying which assets are likely to outperform or underperform their peers. GNN-based signals tend to be more stable in this context because they incorporate information from related assets rather than relying solely on noisy, asset-specific histories.

For example, if several firms within a tightly connected network begin to exhibit weakness, a GNN can propagate that information across the network, adjusting expectations for related assets before the co-movement is fully reflected in prices or market narratives. This allows portfolio managers to respond earlier to emerging themes or dislocations. Exhibit 1 highlights an example where a latent factor impacts multiple sectors.

Portfolio Construction and Network-Aware Diversification

Diversification is ultimately about managing dependence, not simply spreading capital across many assets or sectors. Network representations make this dependence visible, and GNNs allow it to be incorporated directly into portfolio construction.

A key insight from network analysis is that highly connected assets contribute disproportionately to portfolio risk, especially during stress periods. These assets often behave like latent factors, transmitting shocks across the portfolio. Conversely, assets that occupy more peripheral positions in the network may provide more independent sources of return, even if they appear similar along traditional dimensions such as sector or region. The result is a reduction in exposure to a tightly connected subnetwork that cuts across sector boundaries.

GNNs naturally learn and adapt to this structure. In practice, their outputs can be used to adjust position sizes to avoid hidden concentration, tilt portfolios away from tightly connected clusters, and identify diversification that is robust to regime shifts. Some implementations use GNNs directly within portfolio optimization frameworks,

EXHIBIT 1

When Sector Diversification Fails: Artificial Intelligence as a Network Shock

Consider an equity portfolio constructed to be well diversified across sectors. The portfolio holds technology firms developing artificial intelligence software, semiconductor manufacturers supplying advanced chips, industrial firms providing automation equipment, energy companies supporting data center infrastructure, and consumer firms adopting AI-driven pricing and logistics tools. From a traditional sector allocation perspective, exposures appear balanced. No single sector dominates portfolio risk, and historical correlations across these sectors are moderate.

Viewed through a network lens, however, a different picture emerges. Many of these holdings are tightly connected through a common dependence on artificial intelligence adoption. Semiconductor firms supply computing capacity to cloud providers. Cloud providers enable AI applications used by industrial and consumer firms. Energy companies support the physical infrastructure required to power data centers. Although these firms belong to different sectors, they form a dense cluster linked by shared exposure to the same technological and capital investment cycle.

Portfolio managers already recognize these connections in practice. When assessing diversification, they routinely consider questions such as whether industrial holdings rely on the same suppliers as technology firms, or whether consumer companies' margins depend on continued investment in data center infrastructure. These judgments, however, are often applied informally, vary across portfolios, and are difficult to document or scale consistently.

Now consider a shock to this network. A regulatory change restricts access to advanced chips, or capital spending on data centers slows sharply. Initial price weakness may appear in a small set of semiconductor or cloud infrastructure firms. A traditional time-series or factor-based model reacts asset by asset, adjusting exposures only after correlations across sectors increase and diversification breaks down visibly.

A network-based model makes this intuition explicit. A GNN propagates early weakness from the affected firms across connected holdings, systematically reflecting the same line of reasoning that portfolio managers apply informally. Because this process is embedded in the model, it is transparent, repeatable, and consistently applied across portfolios and market environments. As a result, the portfolio manager is alerted to rising concentration risk before it is fully captured by aggregate risk metrics.

The key insight is not that artificial intelligence is a technology story, but that it is a network shock impacting various sectors. Sector diversification alone obscures this dependence. A network-based perspective reveals it, and GNNs provide a disciplined framework for formalizing how experienced investors already think about interconnected markets.

while others use them to refine inputs such as expected returns or risk signals. In either case, the value comes from incorporating structural information that is absent from covariance estimates alone.

Risk Management, Contagion, and Stress Scenarios

Risk management is another area where network-based models are particularly valuable. Many of the most damaging portfolio losses arise not from individual asset moves, but from contagion, when shocks propagate rapidly across related positions.

Graph-based models are well-suited to capturing these dynamics by supporting stress testing based on realistic economic links, identifying prominent risk themes within portfolios, and conducting scenario analysis that reflects structural dependence.

When complemented with traditional risk measures, these GNNs can provide a more granular view of how risk is distributed and how it may spread under adverse conditions. Exhibit 2 examines how risk managers can leverage GNNs to improve stress tests during tail events.

Event-Driven and Thematic Investing

Event-driven and thematic strategies often hinge on understanding how information flows through the market. Corporate actions, regulatory changes, geopolitical developments, and technological shifts rarely affect assets in isolation. Their impact is mediated by networks of economic and informational relationships.

GNNs are particularly effective in these settings because they can integrate event information with network structure. For example, thematic narratives such as energy transition or artificial intelligence adoption can link firms across traditional sector boundaries.

For portfolio managers, this enables more coherent positioning around themes and events, reducing reliance on ad hoc judgment and improving consistency across portfolios.

EXHIBIT 2

Stress Testing as a Network Problem: Liquidity and Forced Selling

Traditional portfolio stress tests often apply shocks at the level of individual assets, sectors, or broad market factors. These exercises are useful, but they implicitly assume that shocks affect holdings independently or through the channels of covariance matrices. In practice, some of the most severe portfolio drawdowns arise not from isolated asset moves, but from forced selling that propagates through networks of ownership and liquidity.

Consider a multi-asset portfolio with exposures across equities, credit, and related derivatives. Several positions are held by the same group of large institutional investors or are commonly used as collateral. Under normal market conditions, these connections are not visible in standard risk reports, and historical correlations appear modest.

Portfolio managers are well aware of this dynamic. During periods of market stress, they routinely ask which holdings might be sold together if liquidity dries up or margin requirements rise. These assessments are typically qualitative, informed by experience, and applied on an ad-hoc basis under time pressure.

Now consider a liquidity shock. A sharp market move triggers margin calls or redemptions at a large investor. Selling pressure first appears in a small set of liquid positions. As losses accumulate, sales spread to other holdings that share common owners or funding channels, even if the underlying fundamentals are unrelated. Correlations rise rapidly, and diversification breaks down after the fact.

A network-based stress test makes this mechanism explicit by propagating the initial shock through these connections, identifying which positions are most likely to transmit stress and which clusters of assets are vulnerable to forced selling. This approach mirrors how experienced portfolio managers think about liquidity risk but applies that logic systematically and consistently.

The value of this framework is not in predicting the precise path of a crisis, but in revealing structural fragilities before they materialize. By formalizing network-based reasoning, GNNs help transform stress testing from a static exercise into a dynamic assessment of how risk can spread through portfolios under adverse conditions.

When Network-Based Approaches Matter Most

It is important to recognize that network-based models may not be impactful in all environments. Their advantages are most pronounced when relationships among assets are changing rapidly or when co-movement dominates idiosyncratic behavior.

These tools are typically most useful during periods of market stress or elevated volatility. Examples may include sector rotations, liquidity-driven dislocations, and narrative- or event-driven markets. During calm periods, simpler models may perform comparably. The strength of GNNs lies in their ability to adapt when traditional diversification assumptions break down.

GOVERNANCE OF GRAPH NEURAL NETWORKS

For GNNs to be useful in institutional asset management, they must be interpretable, governable, and compatible with existing investment processes. Predictive performance alone is not sufficient. Portfolio managers and risk committees need to understand why signals change, how they relate to economic intuition, and how they should be used in decisions.

A key advantage of graph-based models is that their structure aligns naturally with how investors think about markets. Because GNNs operate on explicit networks of relationships, their outputs can often be linked back to identifiable connections among assets. This form of explanation is often more intuitive than explanations from purely time-series-based machine learning models.

From a governance perspective, the most important source of model risk in GNNs is the graph itself. Graph construction reflects assumptions about which relationships matter and how they are measured. Institutions, therefore, need to treat the graph as a model, subject to validation, stress testing, and ongoing monitoring. In particular, care is required to avoid overfitting, especially when relationships are inferred directly from noisy financial data.

Compounding the issue is that data limitations add another practical constraint. Economic, ownership, and information-based relationships are often incomplete or lagged, thus limiting the usefulness of graph-based tools. Relying on multiple perspectives on market structure and avoiding overdependence on any single data source can help mitigate this risk.

While GNNs can adapt to changing market conditions, they are not immune to regime shifts. Relationships among assets can change abruptly during periods of stress, policy shocks, or structural transitions. Continuous monitoring of model performance and network stability is required, particularly when diversification assumptions are most likely to fail. Feedback loops are thus essential in ensuring that the models are performing as expected, or whether further learning and fine-tuning are needed to improve model performance.

Finally, the role of human judgment is central to how GNNs are constructed, interpreted, and applied. GNNs should reflect existing investment processes and be used as decision support tools to augment the decision-making of portfolio managers. Clear documentation of how network-based insights are incorporated into portfolio decisions helps ensure accountability, trust, and consistent use.

CONCLUSION

Diversification and portfolio analysis have traditionally been taught and practiced through covariance matrices, correlations, and factor exposures. These tools remain essential, but they summarize dependence rather than explain it. In increasingly interconnected markets, this distinction matters. Asset return co-movements are shaped by economic linkages, market structure, and information flows that are difficult to interpret through covariance alone, particularly during periods of stress when diversification is most needed.

Network-based representations provide a natural extension of classical portfolio thinking. By making relationships explicit, they help investors see clusters of dependence and understand pathways of shock propagation and sources of hidden concentration. This network perspective aligns closely with how experienced portfolio managers reason about markets, even when it has not been formalized in their models.

The article shows that GNNs build directly on this network intuition and make it operational. Where network representations describe structure, GNNs learn from it. By treating the market as a graph and allowing information to propagate across economically meaningful connections, GNNs incorporate interdependence into forecasting, portfolio construction, and risk analysis in a way that traditional linear econometrics and temporal-only machine learning models cannot.

As with all quantitative tools, GNNs are not a panacea. Their effectiveness depends critically on thoughtful graph construction, careful governance, and integration into existing investment processes. This requires judgment, domain knowledge, and discipline. Used appropriately, however, GNNs offer a powerful and intuitive framework for incorporating network structure into portfolio construction. By grounding graph-based learning in the network intuition that underlies diversification, GNNs provide asset managers with a powerful framework to capture market connectivity and its impact on portfolio diversification.