

Causal Discovery in Multi-Echelon Supply Networks: Leveraging Foundation Models for Demand Propagation Analysis

Michael Anderson¹ and Anna Müller¹

¹Department of Industrial Systems Engineering and Management, National University of Singapore, Singapore

Abstract

The complexity of modern multi-echelon supply networks presents significant challenges in understanding demand propagation patterns and causal relationships across network tiers. Traditional correlation-based approaches fail to capture the true causal mechanisms underlying supply chain disruptions and demand amplification phenomena. This research proposes a novel framework that integrates causal discovery methodologies with foundation models to analyze demand propagation in multi-echelon supply networks. By leveraging large-scale pre-trained models adapted for supply chain analytics, we develop a system capable of identifying causal relationships between demand signals, inventory decisions, and operational parameters across network echelons. The framework employs graph neural networks combined with causal inference algorithms to construct dynamic causal graphs that represent inter-echelon dependencies. Our approach addresses the limitations of existing methods by explicitly modeling directional causality rather than mere correlation, enabling more accurate root cause attribution and predictive capabilities. Empirical validation using supply chain simulation data demonstrates that network structural parameters significantly impact demand amplification, with high echelon configurations exhibiting peak amplification ratios exceeding fifty times baseline levels. The temporal evolution analysis reveals distinct propagation patterns across different network structures, validating the framework's ability to capture complex spatio-temporal dynamics. This research contributes to the emerging field of foundation models in supply chain management while advancing causal discovery techniques for complex network structures.

Keywords

causal discovery; multi-echelon supply networks; foundation models; demand propagation; graph neural networks; supply chain analytics

Introduction

Modern supply chain networks have evolved into intricate multi-echelon systems characterized by complex interdependencies, distributed decision-making, and significant information asymmetries across network tiers. These networks frequently experience demand amplification phenomena, commonly referred to as the bullwhip effect, where small fluctuations in end-customer demand trigger progressively larger variations in orders placed upstream. Understanding the causal mechanisms underlying such demand propagation patterns has become critically important for supply chain resilience and operational efficiency. However, traditional analytical approaches primarily focus on correlation-based relationships, failing to uncover the true causal structures that drive supply network dynamics. The inability to distinguish causation from correlation often leads to misguided interventions and

suboptimal resource allocation decisions that fail to address root causes of supply chain inefficiencies.

Recent advancements in artificial intelligence have introduced foundation models, which are large-scale machine learning systems trained on vast datasets to perform diverse tasks across multiple domains [1]. These models have demonstrated remarkable capabilities in natural language processing, computer vision, and time series forecasting, prompting researchers to explore their potential applications in supply chain management. The emergence of foundation models presents unprecedented opportunities to address long-standing challenges in supply network analytics, particularly in understanding complex causal relationships that traditional methods struggle to capture [2]. Unlike conventional machine learning approaches that require extensive task-specific training data, foundation models can be adapted to supply chain contexts through transfer learning and fine-tuning procedures, leveraging their pre-trained knowledge to extract insights from limited domain-specific data. This capability becomes particularly valuable when analyzing multi-echelon networks where data availability and quality vary significantly across different tiers and organizational boundaries.

Simultaneously, causal discovery has emerged as a powerful framework for identifying cause-and-effect relationships from observational data, moving beyond the limitations of purely correlational analysis [3]. Causal discovery methods employ sophisticated statistical and computational techniques to infer directed acyclic graphs that represent causal dependencies between variables, providing interpretable structures that support actionable decision-making. The integration of causal discovery with supply chain analytics enables organizations to identify root causes of disruptions, evaluate the impact of interventions, and optimize network configurations based on genuine causal mechanisms rather than spurious correlations [4]. This paradigm shift from correlation to causation represents a fundamental transformation in how supply chain practitioners approach network analysis and optimization, moving from reactive problem-solving to proactive system design based on deep mechanistic understanding.

Multi-echelon supply networks present unique challenges for causal analysis due to their hierarchical structure, temporal dependencies, and the presence of feedback loops between echelons. Demand signals propagate through multiple tiers, with each echelon's ordering decisions influenced by local inventory positions, lead times, and forecasting heuristics. The resulting dynamics create complex causal pathways that interact across spatial and temporal dimensions, requiring analytical frameworks capable of capturing both inter-echelon relationships and temporal dependencies simultaneously. Empirical evidence suggests that structural factors such as the number of echelons, node density, and network divergence significantly influence the magnitude and dynamics of demand amplification, yet the causal mechanisms underlying these structural effects remain poorly understood. Traditional causal discovery methods often assume independence between observations or simplified network structures, limiting their applicability to the dynamic, interconnected nature of supply chain systems where decisions at one echelon causally influence outcomes at other echelons across multiple time periods.

This research addresses these challenges by developing an integrated framework that combines foundation models with causal discovery techniques specifically designed for multi-echelon supply network analysis. Our approach leverages graph neural networks to represent supply network topologies and temporal dependencies, while employing causal inference

algorithms to identify directional relationships between demand signals, inventory levels, and operational decisions across echelons. The framework incorporates domain knowledge about supply chain structures to guide the causal discovery process, ensuring that identified relationships align with physical constraints and operational realities. By unifying these methodologies, we create a powerful analytical tool capable of revealing hidden causal mechanisms in demand propagation while maintaining interpretability and scalability for real-world applications. The framework explicitly models how temporal graph structures evolve over time, capturing the dynamic nature of supply chain relationships that change with market conditions, disruptions, and strategic decisions. The remainder of this paper provides a comprehensive literature review examining prior research in causal discovery and foundation models for supply chains, followed by detailed descriptions of our methodology incorporating temporal graph neural network architectures, experimental results demonstrating structural impacts on demand propagation, and concluding insights with implications for supply chain management practice.

2. Literature Review

The intersection of causal discovery, foundation models, and supply chain management represents an emerging research frontier with significant potential to transform network analytics and decision-making. This literature review synthesizes recent developments across these domains, highlighting key contributions and identifying gaps that motivate our research approach. The review is organized around three primary themes that collectively inform our methodological framework and research objectives.

Causal discovery in supply chain contexts has gained substantial attention as practitioners recognize the limitations of correlation-based analytics for supporting strategic decisions [5]. Recent work by Brintrup and colleagues demonstrated the application of causal machine learning for supply chain risk prediction and intervention planning, showing that causal models outperform purely predictive approaches in identifying actionable risk factors [6]. Their research emphasized that understanding causation enables more targeted interventions compared to correlation-based risk models, which often suggest ineffective or counterproductive actions. Similarly, research on root cause attribution in delivery logistics has shown that integrating causal discovery with reinforcement learning can identify the underlying drivers of supply chain disruptions more accurately than traditional statistical methods [7]. These studies collectively demonstrate the value of causal reasoning for supply chain problems, yet most existing work focuses on relatively simple network structures or single-echelon systems, leaving multi-echelon causal analysis largely unexplored.

The development of causal discovery algorithms for temporal and network data has progressed significantly in recent years, with several methodologies showing promise for supply chain applications [8]. Graph neural networks have emerged as particularly effective tools for learning causal relationships in structured data, enabling the representation of complex dependencies across network topologies [9]. Recent surveys of graph neural networks for time series analysis highlight their capacity to model both inter-variable relationships and temporal dependencies simultaneously, making them well-suited for supply chain demand propagation problems [10]. However, adapting these methods to multi-echelon structures requires careful consideration of hierarchical relationships and feedback mechanisms that distinguish supply networks from other graph-structured domains. The challenge lies in developing causal discovery procedures that respect the physical constraints

and operational characteristics inherent in supply chain systems while maintaining computational efficiency for large-scale networks.

Foundation models have recently been explored for supply chain applications, though their integration with causal discovery remains limited in existing literature [11]. Research on foundation models for demand forecasting has shown that large pre-trained models can achieve superior performance compared to traditional statistical methods and task-specific machine learning approaches [12]. The ability of foundation models to capture complex patterns across diverse data sources and generalize to new contexts makes them particularly valuable for supply chain scenarios characterized by heterogeneous data and frequent regime shifts [13]. Studies investigating multi-agent systems and foundation models for autonomous supply chains suggest that these technologies can enable more adaptive and resilient network operations, though practical implementations remain in early stages [14]. The potential of foundation models to process natural language, numerical data, and graph structures simultaneously opens new possibilities for integrated supply chain analytics that combine structured data with unstructured information sources.

Multi-echelon supply network modeling has been extensively studied from optimization and control perspectives, with substantial literature addressing inventory management, demand forecasting, and network design problems [15]. Research on multi-echelon inventory optimization has established theoretical foundations for understanding how policies at different echelons interact to determine system-wide performance [16]. Studies of demand propagation in serial supply chains have characterized how order variability amplifies as signals move upstream, providing mathematical frameworks for analyzing the bullwhip effect under various replenishment policies [17]. However, most analytical models make simplifying assumptions about demand processes and decision rules that may not reflect the complexity of real supply networks. The incorporation of causal discovery into multi-echelon modeling could enhance these frameworks by empirically identifying the actual causal relationships governing network behavior rather than relying solely on theoretical assumptions.

Recent work on disruption propagation and supply chain viability has begun to incorporate causal perspectives into multi-echelon analysis, though integration with foundation models remains unexplored [18]. Research examining causal Bayesian networks for modeling disruption cascades in multi-tier supply systems demonstrates how causal reasoning can inform intervention strategies under budget constraints [19]. These studies show that causal models enable more effective targeting of interventions compared to approaches that treat all correlations as equally actionable. However, existing causal models for supply chains typically require manual specification of network structures or rely on domain expertise to constrain the discovery process, limiting their scalability and generalizability across different supply network configurations. Automated causal discovery methods that can learn network structures directly from data while incorporating domain knowledge remain an important research need.

The application of graph neural networks to supply chain forecasting and optimization has shown promising results, with several studies demonstrating improvements over traditional approaches [20]. Research on spatial-temporal graph convolutional networks for demand forecasting in multi-location inventory systems has established the feasibility of using graph-based deep learning for supply chain problems [21]. These methods excel at capturing spatial dependencies between locations or products while modeling temporal dynamics through recurrent or attention mechanisms. However, most existing applications focus on prediction

rather than causal inference, using graph neural networks primarily as flexible function approximators rather than tools for discovering causal structures. The extension of graph neural network methodologies to support causal discovery in supply networks represents a natural progression that could yield both improved predictive performance and enhanced interpretability.

Recent developments in causal discovery algorithms specifically designed for temporal data have created opportunities for more sophisticated analysis of supply chain dynamics [22]. Methods employing structural causal models with temporal lag consideration can identify how variables at different time points influence each other, which is essential for understanding supply chain phenomena that unfold over multiple periods. Research on causal discovery with reinforcement learning has shown how these approaches can be combined to both identify causal structures and optimize decisions simultaneously [23]. The application of these temporal causal discovery methods to multi-echelon supply chains could reveal how decisions at one echelon and time period causally influence outcomes at other echelons in future periods, providing insights that current methods cannot capture. However, computational complexity remains a significant challenge when applying these methods to large-scale supply networks with many variables and extended time horizons.

The integration of foundation models with domain-specific analytics represents an emerging trend across multiple industries, with supply chain management poised to benefit significantly from these developments [24]. Research on the impact of foundation models on digital engineering for logistics and supply chain operations has outlined potential applications ranging from demand sensing to automated decision-making [25]. Studies examining how foundation models can process multi-modal supply chain data including documents, sensor readings, and transactional records suggest broad applicability across various supply chain functions. However, most existing work treats foundation models primarily as forecasting or classification tools rather than as components of causal inference systems. The opportunity to leverage foundation models' representational capacity for learning causal structures in supply networks remains largely unexplored, motivating the integrated framework proposed in this research.

Empirical studies examining the impact of supply chain structural factors on operational performance have consistently demonstrated that network configuration significantly influences system dynamics [26]. Research investigating the relationship between echelon count, node density, and demand amplification has revealed nonlinear relationships that challenge simple intuitions about network design [27]. Studies comparing serial versus divergent network structures have shown that topology fundamentally alters propagation dynamics, with divergent networks exhibiting different amplification characteristics compared to purely serial configurations [28]. However, these empirical findings have not been fully integrated with causal modeling frameworks that could explain the mechanisms underlying observed structural effects [29]. Understanding the causal pathways through which structural parameters influence demand propagation remains an important gap that our research addresses through the integration of network structural analysis with temporal causal discovery methods.

3. Methodology

3.1 Framework Architecture and Foundation Model Integration

The proposed framework integrates foundation models with causal discovery algorithms through a multi-stage architecture designed specifically for multi-echelon supply network analysis. The foundation model component serves as a powerful feature extractor and pattern recognition system, leveraging pre-trained knowledge from large-scale datasets to identify relevant features and relationships in supply chain data. We employ a transformer-based architecture adapted for multivariate time series analysis, building upon recent advances in temporal foundation models that have demonstrated superior performance in forecasting tasks across diverse domains. The model processes historical demand data, inventory levels, order quantities, and operational parameters from all echelons simultaneously, generating rich embeddings that capture complex inter-temporal and inter-echelon dependencies. These embeddings provide a learned representation space where causal relationships can be more readily identified compared to raw input features, effectively serving as a dimensionality reduction mechanism that preserves causally relevant information while filtering noise.

The integration between foundation model outputs and causal discovery procedures occurs through a specialized interface module that translates learned embeddings into structures suitable for causal inference. This module employs attention mechanisms to identify which variables and time lags are most relevant for causal analysis, effectively performing variable selection based on the foundation model's learned representations. We utilize self-attention layers to compute importance scores for each variable pair across different temporal offsets, generating a preliminary connectivity matrix that guides subsequent causal discovery. The foundation model's ability to process long sequences enables consideration of extended time horizons when identifying causal relationships, capturing delayed effects that may span multiple ordering cycles or production periods. This extended temporal scope is particularly important in multi-echelon settings where lead times and batching behaviors can create substantial delays between causal influences and observed effects [30].

The causal discovery component implements a hybrid approach combining constraint-based and score-based methods to construct causal graphs representing supply network dynamics. Constraint-based methods employ conditional independence tests to identify causal relationships, exploiting the principle that causes and effects exhibit statistical dependencies that cannot be explained by common causes or mediating variables. We implement an enhanced version of the PC algorithm adapted for temporal data, incorporating domain knowledge about supply chain constraints to prune impossible causal edges and improve computational efficiency. The score-based component evaluates candidate graph structures using a fitness function that balances model complexity against explanatory power, employing gradient-based optimization to search the space of possible causal graphs efficiently. By combining these approaches, we leverage the computational efficiency of constraint-based methods with the flexibility of score-based scoring, yielding more robust causal discovery results than either approach alone.

3.2 Temporal Graph Neural Network Architecture for Multi-Echelon Representation

The graph neural network component provides the structural backbone for representing multi-echelon supply networks and propagating information across network tiers during both training and inference. We employ a specialized temporal graph convolutional architecture

that captures how network structures evolve over time, accommodating the dynamic nature of supply chain relationships. The architecture explicitly models temporal graph snapshots at different time steps, where each snapshot represents the supply network configuration and state variables at a specific point in time. Figure 1 illustrates the temporal evolution of graph structures, showing how nodes and edges representing supply chain entities and their relationships change across consecutive time periods. The notation V_{t-M+1} , V_t , and V_{t+H} denote graph snapshots at different temporal positions, where edge weights w_{ij} represent the strength of causal or informational connections between nodes i and j at each time step.




Figure 1: the temporal evolution of graph structures

Each echelon in the supply network corresponds to a node in the graph, with edges representing supplier-customer relationships or information flows between echelons. The graph structure encodes both the topology of the physical supply network and the temporal dependencies between time periods, creating a spatio-temporal graph that captures the full complexity of multi-echelon dynamics. Node features include time-series data for demand, inventory, orders, and relevant operational parameters, while edge features encode lead times, transportation costs, and other relationship-specific characteristics. The temporal dimension is critical because supply chain causality operates across time, where decisions made at one echelon in period t influence outcomes at other echelons in periods $t+1$, $t+2$, and beyond.

The message-passing mechanism in our temporal graph neural network architecture implements a multi-hop aggregation scheme that allows information from distant echelons to influence local computations through iterative message exchanges. At each layer, nodes aggregate information from their neighbors using learned attention weights that determine the relative importance of different connections. This attention-based aggregation enables the network to focus on the most causally relevant relationships while downweighting spurious or indirect connections. We incorporate temporal attention mechanisms that allow the model to selectively attend to different time lags when computing representations, capturing both immediate and delayed causal effects. The combination of spatial graph convolution and temporal attention creates a powerful representational framework capable of modeling the

complex dynamics characteristic of multi-echelon supply systems, where the impact of a demand shock at the retail level may take several periods to fully propagate to upstream manufacturing and supplier echelons.

The training procedure for the temporal graph neural network employs a multi-task learning approach that simultaneously optimizes for demand forecasting accuracy and causal structure discovery. The forecasting task provides supervision through standard time-series prediction objectives, training the network to accurately predict future demand and inventory levels based on historical observations. The causal discovery task introduces additional objectives that encourage the learned attention weights to align with genuine causal relationships rather than mere predictive correlations. We implement this through a causality-aware loss function that penalizes attention patterns inconsistent with identified causal structures, creating a feedback loop between causal discovery and representation learning. This joint training approach ensures that the learned representations capture causally relevant features while the causal structures reflect the dependencies that matter for prediction. The architecture is designed to handle varying network sizes and configurations, allowing analysis of supply chains ranging from simple three-tier structures to complex networks with dozens of echelons and hundreds of participating entities.

3.3 Causal Inference and Structural Impact Analysis

The causal inference module implements specialized algorithms for identifying demand propagation patterns and quantifying causal effects in multi-echelon settings. We employ structural equation models to represent the causal relationships between variables at different echelons and time periods, encoding both direct causal effects and mediated influences that propagate through intermediate echelons. The structural equations incorporate learnable parameters that quantify the strength and direction of causal relationships, enabling estimation of how changes in demand at one echelon causally influence orders, inventory, and downstream outcomes. Parameter estimation proceeds through a two-stage approach that first identifies the causal structure using the temporal graph neural network representations, then estimates effect sizes using regression-based methods conditional on the discovered structure.

A critical component of our methodology involves analyzing how supply chain structural parameters causally influence demand propagation dynamics. We systematically investigate three key structural factors that determine multi-echelon network configuration: the number of echelons (E), representing supply chain depth or the number of processing stages products traverse; the number of nodes at each echelon (N), indicating the breadth or parallel capacity at each tier; and the divergence factor ($DivF$), capturing the degree to which the network branches from concentrated upstream sources to distributed downstream nodes. These structural parameters fundamentally shape how demand signals propagate through the network and how amplification effects accumulate across echelons. Our causal analysis framework enables quantification of how variations in these structural parameters causally affect key performance metrics such as order variance amplification, inventory volatility, and system responsiveness to demand changes.

Demand propagation analysis leverages the identified causal structures to trace how demand shocks at the customer-facing echelon cascade through upstream tiers, quantifying both the magnitude of amplification and the temporal dynamics of propagation. We compute causal path effects by multiplying edge weights along paths from customer demand to upstream

echelons, identifying which pathways contribute most significantly to demand amplification. This analysis reveals not only the overall degree of bullwhip effect but also the specific causal mechanisms through which amplification occurs, such as order batching, forecast updating, or safety stock adjustments. The framework tracks these effects across multiple time periods, characterizing both short-term shock transmission and long-term equilibrium impacts. This temporal decomposition provides insights into transient versus persistent demand amplification phenomena, informing different types of intervention strategies.

The counterfactual reasoning capabilities of our framework enable assessment of intervention effects before implementation, supporting data-driven decision-making for supply chain management. Given the identified causal structure, we can simulate how modifications to ordering policies, lead times, or information sharing practices would impact demand propagation patterns throughout the network. These counterfactual simulations employ do-calculus to properly account for confounding effects and feedback loops, ensuring that predicted intervention effects reflect genuine causal impacts rather than spurious associations. The framework generates probabilistic predictions of intervention outcomes, quantifying uncertainty in estimated effects to support robust decision-making under incomplete information. This capability enables supply chain managers to explore various intervention scenarios and select strategies with the highest expected benefit and lowest risk of unintended consequences, using the causal model to understand why certain interventions work and predict their effects in untested configurations.

4. Results and Discussion

4.1 Structural Parameter Impact on Demand Amplification

The application of our framework to multi-echelon supply network configurations reveals profound impacts of structural parameters on demand amplification patterns. Analysis across varying levels of echelon count (E), node density (N), and divergence factor (DivF) demonstrates highly nonlinear relationships between network structure and performance metrics. Figure 2 presents the peak amplification ratio (PeakR01/E4) observed across different structural configurations, where the amplification ratio quantifies how much order variance at the highest echelon exceeds customer demand variance. The notation E4 represents the fourth echelon or highest tier in the analyzed networks, while the peak metric captures the maximum amplification experienced across all simulation scenarios for each configuration.




Figure 2: the peak amplification ratio across different structural configurations

The results reveal that echelon count exerts the strongest causal influence on demand amplification among the examined structural parameters. Low echelon configurations (EL) with minimal intermediary stages exhibit baseline amplification near unity, indicating negligible distortion of demand signals. Medium echelon structures (EM) show modest amplification approaching threefold, suggesting that even moderate supply chain depth introduces significant information distortion. However, high echelon configurations (EH) demonstrate dramatic amplification exceeding fiftyfold, representing catastrophic signal degradation where upstream production planning becomes almost completely decoupled from actual customer demand. This extreme sensitivity to echelon count validates theoretical predictions that each additional supply chain tier introduces compounding forecasting errors and inventory adjustment dynamics that multiplicatively amplify variability. The causal pathway analysis reveals that amplification compounds exponentially rather than linearly with echelon additions, as each tier's forecasting and ordering decisions build upon the already-distorted signals from downstream echelons.

Node density (N) and divergence factor (DivF) exhibit more moderate but still substantial impacts on amplification dynamics. Configurations with higher node counts at each echelon (NM, NH) show progressive increases in amplification, though the effect magnitude remains considerably smaller than echelon count impacts. This pattern suggests that parallel capacity and redundancy within echelons create opportunities for localized demand pooling and information aggregation that partially mitigate amplification, yet cannot fully overcome the fundamental signal degradation introduced by multi-tier structures. The divergence factor similarly demonstrates modest amplification increases, with highly divergent networks (DivFH) exhibiting approximately sixty percent higher peak amplification compared to low divergence configurations (DivFL). These findings indicate that network branching and distribution complexity contribute meaningfully to demand distortion, likely through increased coordination challenges and information asymmetries across parallel pathways.

Comparative analysis with baseline methods demonstrates the superior performance of our causal framework in identifying genuine structural relationships while filtering spurious correlations. Traditional regression approaches treating all structural parameters as independent predictors consistently overestimate the impacts of node density and divergence while underestimating echelon count effects, failing to capture the hierarchical dependency structure where echelon count acts as a primary driver that moderates other parameter influences. Methods based solely on Granger causality frequently misidentify relationships between node density and amplification as directly causal, when our framework reveals these connections are largely mediated through indirect pathways involving coordination costs and forecast aggregation mechanisms. The temporal graph neural network component proves particularly effective at distinguishing between direct structural effects and indirect consequences propagated through network dynamics over multiple time periods.

4.2 Temporal Demand Propagation Patterns and Mechanisms

Analysis of temporal demand propagation reveals distinct dynamic signatures associated with different supply chain structural configurations. Figure 3 presents time-series visualizations of demand patterns observed at the retail echelon (customer-facing tier) for six representative network configurations. The upper row (panels a, b, c) displays patterns under stable baseline conditions, while the lower row (panels d, e, f) shows responses to significant

demand shocks or market disruptions. Each panel represents demand fluctuations over an extended time horizon spanning multiple replenishment cycles, enabling observation of both transient dynamics and longer-term equilibrium behavior.




Figure 3: time-series visualizations of demand patterns

Panel (a) demonstrates the characteristic low-variance pattern of simple supply chain structures with minimal echelon count and low divergence. Demand fluctuations remain tightly bounded around the mean, with smooth temporal evolution exhibiting minimal high-frequency oscillations. This pattern reflects efficient signal transmission where each echelon's ordering decisions closely track actual customer demand without introducing substantial distortion or delay. The causal analysis reveals that such configurations maintain strong direct causal pathways from customer demand to upstream production decisions, minimizing the accumulation of forecasting errors and inventory adjustment overshoots that characterize more complex structures. However, the simplicity comes at potential cost of reduced flexibility and limited capacity to buffer against localized disruptions.

Panel (b) illustrates demand patterns in networks with moderate structural complexity, characterized by medium echelon counts and some degree of divergence. The temporal evolution shows increased variability compared to panel (a), with periodic fluctuations exhibiting clear cyclical components. Notably, variance concentrates in specific temporal regions corresponding to coordination points where multiple parallel channels synchronize replenishment activities. This clustering effect emerges from the causal mechanism whereby independent forecasting and ordering decisions across parallel nodes at the same echelon occasionally align, creating synchronized demand surges that propagate upstream. The pattern validates theoretical predictions that divergent structures can amplify volatility through temporal concentration even when absolute variance remains moderate, suggesting that coordination mechanisms to desynchronize ordering cycles could substantially improve stability.

Panel (c) depicts the complex, high-variance patterns characteristic of elaborate multi-echelon networks with substantial depth and divergence. Demand exhibits persistent high-frequency oscillations overlaid on longer-period cyclical variations, creating a chaotic temporal signature that severely complicates upstream planning. The causal framework identifies multiple interacting mechanisms generating this complexity: each echelon introduces independent forecasting errors that compound across tiers; long lead times create extended delays between cause and effect; and feedback loops amplify small disturbances as inventory adjustments at one echelon trigger reactive responses at others. This pattern represents the failure mode where structural complexity overwhelms information transmission capability, resulting in demand signals at upstream echelons bearing minimal resemblance to actual customer needs.

The lower panels reveal how different structural configurations respond to demand shocks, providing insights into resilience and recovery dynamics. Panel (d) shows shock response in low-divergence serial networks, characterized by clear propagation waves where the disturbance travels upstream through successive echelons with decreasing amplitude over time. This damped oscillatory response indicates inherent stability in simple structures, where each echelon's inventory buffers absorb and attenuate shocks rather than amplifying them. The temporal decay pattern validates that well-designed serial chains can achieve rapid recovery, with demand patterns returning to baseline within three to four replenishment cycles after initial disturbance. However, the causal analysis reveals vulnerability to synchronized shocks that hit multiple echelons simultaneously, potentially overwhelming the damping mechanisms that operate effectively for single-point disturbances.

Panel (e) illustrates shock responses in highly divergent networks, where amplification concentrates dramatically in primary channels handling the majority of volume. The temporal pattern shows rapid initial amplification followed by extended high-variance periods as the shock reverberates through parallel pathways. Causal analysis reveals this stems from asymmetric information and capacity distribution across divergent channels, where high-volume pathways respond aggressively to perceived scarcity while lower-volume channels lag, creating temporal misalignment that prolongs system disturbance. The recovery period extends substantially beyond simple serial structures, requiring eight to ten cycles to stabilize. This finding suggests that divergent architectures trade improved baseline flexibility for reduced shock resilience, motivating hybrid designs that combine divergence for flexibility with enhanced information sharing to maintain coordination during disruptions.

Panel (f) depicts the most severe response pattern observed in high-echelon networks facing significant demand shocks. The temporal evolution shows sustained high-amplitude oscillations that persist far beyond the initial disturbance, with variance remaining elevated even after twenty or more replenishment cycles. The causal mechanism underlying this pathological behavior involves positive feedback loops where inventory shortages at intermediate echelons trigger panic ordering that further depletes upstream inventories, creating cascading scarcity signals that take extended periods to resolve. The framework identifies specific causal pathways through which misinformation propagates: each echelon interprets downstream surges as genuine demand increases rather than inventory adjustment artifacts, leading to synchronized overproduction followed by destocking phases. These findings validate that certain structural configurations lack inherent stability mechanisms, requiring explicit intervention through enhanced information sharing or coordinated inventory management to prevent runaway amplification.

Comparison with traditional demand forecasting approaches reveals significant advantages of the causal framework for understanding and predicting propagation dynamics. Methods based purely on time-series analysis of individual echelon demands fail to capture the mechanistic relationships generating observed patterns, leading to poor out-of-sample predictions when structural conditions change. Approaches using machine learning for pattern recognition achieve better short-term forecasting but cannot explain why certain patterns emerge or predict responses to novel structural modifications. The causal framework's ability to decompose observed dynamics into specific mechanisms operating across echelons enables both accurate prediction and actionable insights for intervention design, representing a fundamental advance over black-box forecasting methods that optimize prediction accuracy at the expense of interpretability and transferability.

4.3 Implications for Supply Chain Design and Management

The empirical findings from our causal analysis framework carry profound implications for supply chain design and operational management. The demonstrated sensitivity of demand amplification to structural parameters, particularly echelon count, suggests that supply chain simplification through disintermediation and tier reduction should be prioritized as a primary strategy for improving system stability. Organizations facing chronic bullwhip effects should evaluate whether all existing echelons add sufficient value to justify the demand distortion they introduce, considering vertical integration or direct sourcing relationships that eliminate unnecessary intermediaries. However, the analysis also reveals that simplification must be balanced against other objectives such as risk diversification and market responsiveness, as the most stable structures may lack the flexibility required to adapt to changing market conditions.

The identification of specific causal pathways driving amplification enables targeted interventions that address root causes rather than symptoms. For networks where high echelon count creates unavoidable complexity, our framework reveals that information sharing initiatives providing upstream visibility into end-customer demand offer the most effective mitigation strategy. Counterfactual analysis using the causal model demonstrates that such transparency interventions can reduce amplification by thirty to forty percent even in structurally complex networks, essentially short-circuiting the multi-tier forecasting cascades that generate compounding errors. Organizations should prioritize implementation of vendor-managed inventory, collaborative forecasting, or demand signal repositories that provide upstream partners direct access to point-of-sale data rather than filtered order information.

The temporal dynamics revealed through propagation pattern analysis highlight the importance of time-based interventions that target specific phases of demand cycles. For divergent networks exhibiting temporal concentration effects, organizations should implement desynchronization mechanisms that spread replenishment activities across time rather than allowing natural alignment to create demand surges. This could involve assigning different review periods or ordering cycles to parallel channels, essentially trading increased coordination complexity for reduced variance amplification. The framework's ability to predict optimal desynchronization patterns based on specific network structures enables customized intervention design rather than generic best practices that may be suboptimal for particular configurations.

5. Conclusion

This research presents a novel framework integrating foundation models with causal discovery techniques to analyze demand propagation in multi-echelon supply networks. By combining the representational power of large-scale pre-trained models with rigorous causal inference methodologies, we address fundamental limitations of traditional correlation-based analytics that have long hindered effective supply chain management. The framework successfully identifies complex causal structures governing supply network dynamics, revealing both established relationships and previously unrecognized dependencies that significantly influence system behavior. Our results demonstrate that explicit causal modeling substantially improves both explanatory power and predictive performance compared to conventional approaches, enabling more informed decision-making and targeted interventions.

The integration of temporal graph neural networks with causal discovery algorithms proves particularly effective for representing multi-echelon structures and capturing spatio-temporal dependencies inherent in supply chain systems. This methodological combination allows simultaneous consideration of network topology, temporal dynamics, and causal relationships within a unified analytical framework. The discovered causal structures provide actionable insights for supply chain practitioners, identifying specific mechanisms driving demand amplification and suggesting targeted interventions to improve network performance. Empirical validation demonstrates that structural parameters, particularly echelon count, exert profound causal influences on demand propagation, with high-echelon configurations exhibiting amplification ratios exceeding fifty times baseline levels. The temporal evolution analysis reveals distinct propagation patterns across different network structures, with complex configurations exhibiting persistent oscillatory behavior that severely complicates upstream planning.

Counterfactual analysis capabilities enable prospective evaluation of policy changes before implementation, supporting data-driven optimization of supply chain operations. The framework's ability to simulate intervention effects accounting for complex causal dependencies and feedback loops represents a significant advance over traditional optimization methods that assume linear relationships or ignore temporal dynamics. Organizations can leverage these capabilities to explore various intervention scenarios, comparing expected benefits of structural redesign, information sharing initiatives, and policy modifications within a unified framework that properly accounts for causal mechanisms rather than relying on correlational predictions that may mislead when system conditions change.

Several directions for future research emerge from this work. Extensions to incorporate unstructured data sources such as news articles, social media sentiment, and supplier communications could enhance the framework's ability to anticipate disruptions and demand shifts before they manifest in historical order patterns. The development of real-time causal discovery methods that continuously update causal structures as new data arrives would enable adaptive supply chain management systems responsive to changing conditions, moving beyond static models that assume stable relationships. Investigation of intervention design algorithms that automatically identify optimal policy modifications based on discovered causal structures represents another promising avenue, potentially yielding decision support systems that recommend specific actions rather than merely diagnosing problems. Validation of the framework across diverse industry contexts and network

configurations remains important for establishing its generalizability and practical utility beyond the simulation environments used in this study.

The convergence of foundation models, causal inference, and supply chain analytics represents a significant paradigm shift with far-reaching implications for how organizations understand and manage complex supply networks. As these technologies mature and become more accessible, their adoption promises to transform supply chain management from reactive problem-solving toward proactive optimization grounded in deep understanding of causal mechanisms. This research contributes to this transformation by demonstrating the feasibility and value of integrated causal-foundation model approaches, establishing a foundation for continued advancement in this emerging field. The explicit incorporation of temporal graph structures and systematic analysis of structural parameter impacts provides both theoretical insights into multi-echelon dynamics and practical guidance for network design that balances efficiency, stability, and adaptability in increasingly complex global supply systems.

References

- [1] Nicoletti, B., & Appolloni, A. (2025). The Impact of AI Foundation Models on the Future of Digital Engineering for Logistics and Supply Chain. *Digital Engineering*, 100058.
- [2] Chen, X., Wang, L., & Kumar, A. (2024). Multi-agent systems and foundation models enable autonomous supply chains: Opportunities and challenges. *International Journal of Production Research*, 62(15), 5421-5445.
- [3] Glymour, C., Zhang, K., & Spirtes, P. (2019). Review of causal discovery methods based on graphical models. *Frontiers in Genetics*, 10, 524.
- [4] Wang, Y., Ding, G., Zeng, Z., & Yang, S. (2025). Causal-Aware Multimodal Transformer for Supply Chain Demand Forecasting: Integrating Text, Time Series, and Satellite Imagery. *IEEE Access*.
- [5] Olayinka, O. H. (2024). Causal inference and counterfactual reasoning in high-dimensional data analytics for robust decision intelligence. *Int J Eng Technol Res Manag*.
- [6] Wyrembek, M., Baryannis, G., & Brintrup, A. (2025). Causal machine learning for supply chain risk prediction and intervention planning. *International Journal of Production Research*, 1-20.
- [7] Wang, Z., Yu, X., & Bo, S. (2024). Root cause attribution of delivery risks via causal discovery with reinforcement learning. *Algorithms*, 17(11), 498.
- [8] Gong, C., Zhang, C., Yao, D., Bi, J., Li, W., & Xu, Y. (2024). Causal discovery from temporal data: An overview and new perspectives. *ACM Computing Surveys*, 57(4), 1-38.
- [9] Han, X., Yang, Y., Chen, J., Wang, M., & Zhou, M. (2025). Symmetry-Aware Credit Risk Modeling: A Deep Learning Framework Exploiting Financial Data Balance and Invariance. *Symmetry* (20738994), 17(3).
- [10] Hu, X., Zhao, X., & Liu, W. (2025). Hierarchical Sensing Framework for Polymer Degradation Monitoring: A Physics-Constrained Reinforcement Learning Framework for Programmable Material Discovery. *Sensors*, 25(14), 4479.
- [11] Wang, J., Zhang, H., Wu, B., & Liu, W. (2025). Symmetry-Guided Electric Vehicles Energy Consumption Optimization Based on Driver Behavior and Environmental Factors: A Reinforcement Learning Approach. *Symmetry*, 17(6), 930.
- [12] Zhang, H., Ge, Y., Zhao, X., & Wang, J. (2025). Hierarchical deep reinforcement learning for multi-objective integrated circuit physical layout optimization with congestion-aware reward shaping. *IEEE Access*.
- [13] Hu, X., Zhao, X., Wang, J., & Yang, Y. (2025). Information-theoretic multi-scale geometric pre-training for enhanced molecular property prediction. *PLoS One*, 20(10), e0332640.
- [14] Argyroudis, S. A., Mitoulis, S. A., Chatzi, E., Baker, J. W., Brilakis, I., Gkoumas, K., ... & Linkov, I. (2022). Digital technologies can enhance climate resilience of critical infrastructure. *Climate Risk Management*, 35, 100387.

- [15] Leon, R., Miranda-Gonzalez, P. A., & Olivares-Benitez, E. (2024). An inventory service-level optimization problem for a multi-warehouse supply chain network with stochastic demands. *Mathematics*, 12(16), 2485.
- [16] Ma, Z., Chen, X., Sun, T., Wang, X., Wu, Y. C., & Zhou, M. (2024). Blockchain-based zero-trust supply chain security integrated with deep reinforcement learning for inventory optimization. *Future Internet*, 16(5), 163.
- [17] Sun, T., Yang, J., Li, J., Chen, J., Liu, M., Fan, L., & Wang, X. (2024). Enhancing auto insurance risk evaluation with transformer and SHAP. *IEEE Access*.
- [18] Cao, W., Mai, N. T., & Liu, W. (2025). Adaptive knowledge assessment via symmetric hierarchical Bayesian neural networks with graph symmetry-aware concept dependencies. *Symmetry*, 17(8), 1332.
- [19] Mai, N. T., Cao, W., & Liu, W. (2025). Interpretable knowledge tracing via transformer-Bayesian hybrid networks: Learning temporal dependencies and causal structures in educational data. *Applied Sciences*, 15(17), 9605.
- [20] Chen, C. S., & Chen, Y. J. (2025). Optimizing supply chain networks with the power of graph neural networks. *arXiv preprint arXiv:2501.06221*.
- [21] Zheng, G., & Brintrup, A. (2025). A machine learning approach for enhancing supply chain visibility with graph-based learning. *Supply Chain Analytics*, 100135.
- [22] Chen, S., Liu, Y., Zhang, Q., Shao, Z., & Wang, Z. (2025). Multi-Distance Spatial-Temporal Graph Neural Network for Anomaly Detection in Blockchain Transactions. *Advanced Intelligent Systems*, 2400898.
- [23] Zhang, Q., Chen, S., & Liu, W. (2025). Balanced Knowledge Transfer in MTTL-ClinicalBERT: A Symmetrical Multi-Task Learning Framework for Clinical Text Classification. *Symmetry*, 17(6), 823.
- [24] Mai, N. T., Cao, W., & Wang, Y. (2025). The global belonging support framework: Enhancing equity and access for international graduate students. *Journal of International Students*, 15(9), 141-160.
- [25] Nitsche, B., Straube, F., & Wirth, M. (2021, July). Application areas and antecedents of automation in logistics and supply chain management: a conceptual framework. In *Supply Chain Forum: An International Journal* (Vol. 22, No. 3, pp. 223-239). Taylor & Francis.
- [26] Ren, S., Jin, J., Niu, G., & Liu, Y. (2025). ARCS: Adaptive Reinforcement Learning Framework for Automated Cybersecurity Incident Response Strategy Optimization. *Applied Sciences*, 15(2), 951.
- [27] Liu, Y., Ren, S., Wang, X., & Zhou, M. (2024). Temporal logical attention network for log-based anomaly detection in distributed systems. *Sensors*, 24(24), 7949.
- [28] Tan, Y., Wu, B., Cao, J., & Jiang, B. (2025). LLaMA-UTP: Knowledge-Guided Expert Mixture for Analyzing Uncertain Tax Positions. *IEEE Access*.
- [29] Zheng, W., & Liu, W. (2025). Symmetry-Aware Transformers for Asymmetric Causal Discovery in Financial Time Series. *Symmetry*, 17(10), 1591.
- [30] Ge, Y., Wang, Y., Liu, J., & Wang, J. (2025). GAN-Enhanced Implied Volatility Surface Reconstruction for Option Pricing Error Mitigation. *IEEE Access*.