SupplyNet: Large Language Model-Based Multi-Agent Simulation of Liquidity Risk Contagion in Supply Chain Networks

Anonymous submission

Abstract

Modern supply chains have evolved into digitally enabled value networks that coordinate production, financing, and information flows across globally connected firms. Although this interdependence improves operational efficiency, it also amplifies systemic vulnerability. However, existing simulation models of supply chain risk contagion often rely on mechanistic assumptions: they either assume that risk transmission follows a fixed infection rate or represent supply chain firms as homogeneous nodes that make decisions according to preset rules. These approaches overlook the influence of cognitive factors on supply chain risk propagation in the real world. Large language models (LLMs) can incorporate cognitive factors into supply chain simulations by enabling software agents to interpret information and interact through natural-language context. Accordingly, we propose SupplyNet, a four-phase hybrid LLM/rule framework for simulating liquidity risk contagion in supply chains. LLM calls are used for information interpretation, credit assessment, and payment-term negotiation, while firm accounting, inventory transitions, and liquidation are explicit simulator state transitions. To the best of our knowledge, it is the first LLM-driven simulation framework for liquidity risk contagion in supply chains. We validate the framework on a supply chain topology constructed from real-world data. Building on this foundation, we systematically analyze the propagation pathways of liquidity risk across the supply chain network and evaluate the extent to which various factors mitigate its spread.

1 Introduction

In today's highly globalized economic system, the supply chain has evolved from a linear transmission channel for logistics and information into a complex, digitally enabled global value network (Wiedmer and Griffis 2021; Kano, Tsang, and Yeung 2020; Ricciotti 2020). Against this backdrop, Supply Chain Finance (SCF) has emerged as an innovative model that optimizes capital flows and alleviates financing difficulties of small and medium-sized enterprises (Hofmann 2005; Gelsomino et al. 2016; Wuttke et al. 2013). Liquidity risk refers to the possibility that a firm cannot obtain sufficient available funds in time and therefore cannot meet its obligations when they fall due (Basel Committee on Banking Supervision 2008; Drehmann and Nikolaou 2013). It is not equivalent to any single financial indicator and can be characterized from different perspectives using cash holdings, liquidity ratios, operating cash flows, and liquidity gaps. However, this close financial coupling accelerates the speed and expands the scope of risk transmission (Berloco et al. 2021). The widespread trade credit (such as accounts receivable and prepaid accounts) and mutual guarantee relationships among supply chain enterprises construct a potential risk transmission network: once a core enterprise or key supplier encounters financial distress or defaults, liquidity risk spreads along the transaction chain, forming the ripple effect (Zhu et al. 2019).

In contemporary information-driven markets, changes in market demand, bank lending decisions, and payment and liquidity decisions among firms jointly affect the likelihood of firm default and bankruptcy (Berloco et al. 2021; Zhu et al. 2019). Cognitive factors influence demand forecasts, credit assessments, payment terms, and liquidity allocation, thereby affecting the timing and severity of liquidity shocks as well as the propagation of risk across supply chain networks (Sterman 1989; Li 2019; Yang et al. 2021; Waller, Johnson, and Davis 1999; Bloom and Perry 2001; Strom 2015). The proliferation of social media and digital platforms has accelerated the diffusion of unstructured information, including distorted information and panic-inducing statements (Colicchia et al. 2019; Petratos and Faccia 2023; Kessler et al. 2022). Such information can further amplify the effects of these cognitive factors, producing disproportionate fluctuations in market sentiment and short-term demand (Sadeek and Hanaoka 2023; Dolgui and Ivanov 2020; Schmidt et al. 2020).

Despite the importance of these cognitive factors, existing simulation models remain limited in representing how agents interpret changing information before changing operational or financial decisions. Existing approaches typically assume either a fixed contagion rate borrowed from physics and biology (Bardoscia et al. 2015) or preset decision rules for otherwise homogeneous simulated entities (Ivanov 2017; Taghikhah et al. 2021). Such models cannot represent how heterogeneous decision makers read news, rumors, or policy announcements and adapt demand, payment, or credit decisions.

To address this gap, we introduce SupplyNet, a daily,

This is an anonymized submission for review purposes only. Distribution, citation, or public sharing of this manuscript is strictly prohibited. Copyright and publication details will appear in the final version if accepted.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

four-phase simulation of liquidity-risk contagion. SupplyNet combines (i) synthetic LLM market-participant agents that read forum posts and produce downstream demand, (ii) firm agents that update inventory, trade-credit, and cash states on the supply-chain graph, and (iii) a bank agent that evaluates emergency-loan requests using firm states and market sentiment. LLM outputs are used for information interpretation, credit assessment, and payment-term negotiation; accounting, settlement, inventory updates, liquidation, and bad-debt write-offs follow explicit state-transition rules. This design traces a causal path from external information to demand and credit decisions, firm liquidity pressure, and upstream bad-debt propagation. We validate SupplyNet on a real-world network reconstructed from the CSMAR electronic supply-chain database.

In summary, this paper makes the following contributions:

- We introduce SupplyNet, a four-phase hybrid LLM/rule simulation that connects information interpretation, firm operations, credit decisions, and liquidity-risk contagion.
- We use synthetic market-participant agents to generate information-conditioned demand shocks and trace their propagation through the supply-chain network.
- We systematically quantify different risk-intervention methods in terms of their mitigation effects and implementation costs.

2 Framework Overview

SupplyNet is a daily-step simulator. Each agent is a software entity with a role, state, prompt-conditioned decision rule, and, where applicable, short interaction memory. LLM outputs are converted into bounded actions or coefficients, after which the simulator applies inventory, payment, credit, and accounting rules. The 890 firms are represented by firm agents on the CSMAR graph; market participants are synthetic LLM agents that represent information-consuming market actors; and the bank agent is an LLM-based credit-assessment module with a rule-based fast path. Market demand enters at the most-downstream tier, procurement orders move toward upstream suppliers, and delivered goods move in the opposite direction.

The following terminology is used throughout the framework. Cash denotes immediately available funds. Accounts receivable (AR) are amounts owed to a firm by customers for delivered goods, whereas accounts payable (AP) are amounts owed by the firm to its suppliers. Trade credit is the resulting delivery-before-settlement arrangement. An inventory position is on-hand inventory minus any outstanding backlog. A liquidity shortfall occurs when available funds cannot cover obligations and operating expenses due at the relevant date; receivables that are not yet due are not treated as immediately available cash. A bad debt is an AR written off after the customer defaults. Liquidation is the simulator's operational failure state: the firm stops normal trading and attempts to repay creditors with remaining cash, inventory salvage value, and subsequently collected receivables.

Each simulated day applies the following ordered state transitions (Figure 1): market information creates downstream demand; demand creates orders and trade credit; due obligations trigger collection, settlement, and emergency-credit decisions; and a failed settlement followed by rejected emergency credit creates liquidation and upstream bad debt. Read the diagram as this one-day loop. Blue, green, orange, and red arrows correspond to market, supply-chain, financial, and liquidation flows, respectively.

1. **Market phase.** At the start of day t, each synthetic market-participant agent reads the top-K forum posts, its role-conditioned persona, product affinity, and recent interaction history. An LLM returns a willingness coefficient; multiplying this coefficient by baseline demand gives the agent's individual demand. The simulator sums these demands and sends the aggregate to the most-downstream firms. Agents may also generate role-conditioned posts and choose to upvote, downvote, or pass. The forum content evolves endogenously through these agents' posting and voting decisions.
2. **Supply-chain phase.** Downstream firms process market demand, while firms at higher tiers process procurement demand from downstream firms. Each firm applies its (s, S) inventory policy. If its inventory position is at or below s, it selects up to three active upstream suppliers using transaction history and supplier power labels, allocates the order, and records the transaction before negotiating its settlement delay. A delivery creates an AR for the supplier and an equal AP for the buyer.
3. **Financial phase.** The simulator first collects receivables and settles payables, loan repayments, and operating expenses that fall due. If the firm's estimated near-term funding gap remains positive, it submits an emergency-loan request. The bank agent observes the request, the firm's financial state, the current industry default rate, and forum-derived market sentiment, then approves or rejects the request using the hybrid credit mechanism described in Section 3. An approved loan increases cash; a rejected request leaves the firm exposed to the liquidation phase.
4. **Liquidation phase.** A firm that still cannot meet a due obligation after a rejected loan enters liquidation and stops normal production, procurement, and trading. Its cash and the salvage value of remaining inventory form a liquidation pool, which pays bank claims first and supplier claims next. Any supplier receivable that the failed firm cannot repay is written off as bad debt. If that loss causes another failed settlement and rejected loan request, the same process repeats upstream.

3 Methodology

The framework overview defines the daily information-to-cash pathway. This section specifies the state variables and transitions for each module in the same order: market demand and forum updates, firm operations and trade credit, bank lending, and liquidation with bad-debt propagation.

3.1 Market Participant Agents

Each market participant is modeled as a synthetic LLM-driven software agent with an individual persona, interaction

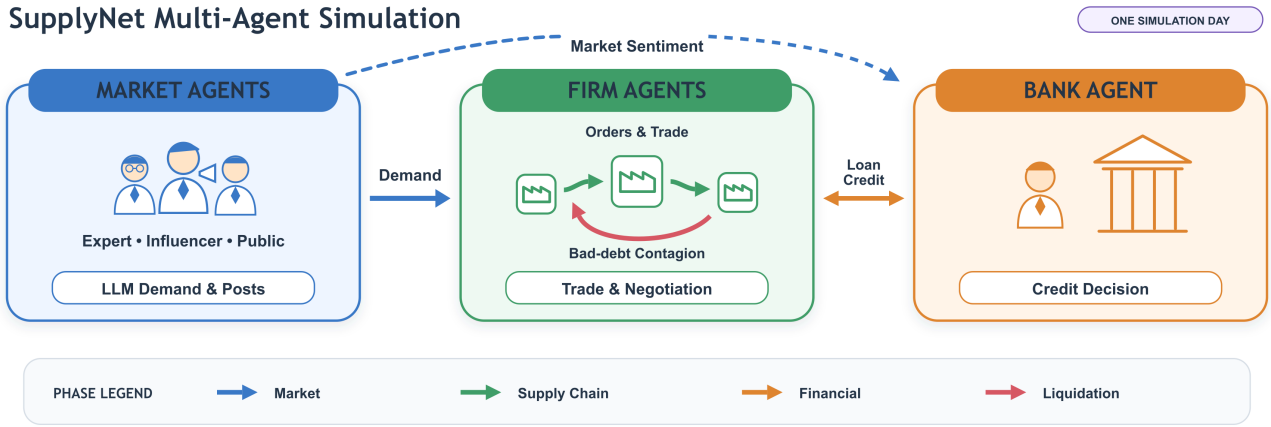


Figure 1: One simulation day in SupplyNet. Synthetic LLM market agents convert forum posts into downstream demand and post interactions; firm agents execute orders, inventory updates, and trade-credit negotiation; liquidation state transitions propagate unrecovered receivables upstream; and the bank agent evaluates emergency credit. Blue, green, orange, and red arrows denote market, supply-chain, financial, and liquidation flows, respectively. The dashed blue arrow carries market-sentiment feedback to the bank, and the red arrow carries bad-debt exposure from a failed firm to its upstream suppliers.

history, and product affinity (Park et al. 2023; Li et al. 2024). On day t, participant m reads the top-K forum posts ranked by heat and combines their content with its own attributes as input to the LLM. The LLM returns a willingness coefficient ω_m^t, which determines the participant's individual demand:

$$D_m^t = d_{\text{base}}\omega_m^t, \tag{1}$$

$$\omega_m^t = \mathcal{L}\left(\psi_m, \Omega_t^K, M_m^t, a_m\right), \tag{2}$$

where $\mathcal{L}(\cdot)$ denotes the LLM decision function, ψ_m is the participant's persona description, Ω_t^K is the set of top-K posts on day t, M_m^t is the participant's interaction history, and $a_m \in [0,1]$ is the participant's product-affinity score, which is converted into a textual description before being included in the LLM prompt.

When $\omega_m^t > 1$, individual demand exceeds the baseline; when $\omega_m^t < 1$, demand falls below the baseline; and when $\omega_m^t = 1$, demand remains unchanged. Individual demand is aggregated across all market participants to obtain the market demand on day t:

$$D_{\text{market}}^t = \sum_m D_m^t. \tag{3}$$

Forum posts are ranked by heat. For post p, the heat score combines community feedback with time decay (Yang et al. 2024):

$$h(p) = \frac{\log_{10}\left(\max(v_p^+ - v_p^-, 0) + 1\right)}{(T_p + 1)^\gamma}, \tag{4}$$

where v_p^+ and v_p^- denote the numbers of upvotes and downvotes, respectively, T_p is the normalized age of the post, and γ is the heat-decay exponent. Each synthetic agent observes only the top-K posts and chooses to upvote, downvote, or pass according to the post content and its alignment with its individual attributes.

Following prior work (Park et al. 2023; Yang et al. 2024; Cho, Guntuku, and Ungar 2025), the synthetic agents are Expert, Influencer, or Public agents, which respectively generate evidence-based, emotive, and experience-based posts. The supplementary material reports the exact role counts, posting probabilities, prompts, and examples.

3.2 Firm Agents

Each supply-chain firm is represented by an independent rule-based state-transition agent. Firm agents occupy different tiers of the supply chain and maintain their cash, inventory, backlog, average cost, selling price, accounts receivable, accounts payable, outstanding loans, and transaction history. In each simulation day, firms at the most downstream end of the supply chain receive the market demand generated by the synthetic market agents, whereas firms at other tiers receive procurement demand from downstream firms. When inventory is insufficient, a firm fulfills only the quantity supported by its available inventory. The simulator applies these accounting and inventory transitions explicitly; the LLM is used only for specified interactions such as payment-term negotiation.

Each firm also incurs daily operating expenses consisting of inventory holding costs and base operating costs. The operating expense of firm i on day t is

$$E_{i,\text{op}}^t = \kappa_{\text{inv}} I_i^t + e_{\text{base}} z_i, \tag{5}$$

where κ_{inv} is the daily holding cost per unit of inventory, I_i^t is the inventory of firm i on day t, e_{base} is the base operating cost, and z_i is the firm-size factor. The operating expense is deducted directly from the firm's cash balance.

Firms employ an (s, S) inventory management policy. The reorder point and target inventory level of firm i are

$$s_i^t = \bar{d}_{i,W_d}^t \beta_i, \tag{6}$$

$$S_i^t = s_i^t + \bar{d}_{i,W_d}^t c, \tag{7}$$

where $\bar{d}^t_{i,W_d}$ is the average daily demand received by firm i over the preceding W_d days, β_i is its risk-buffer period, and c is the replenishment cycle. When the inventory position is no higher than the reorder point, the firm initiates procurement and replenishes its inventory to the target level.

After determining the procurement quantity, a firm selects suppliers from the active upstream firms. During network initialization, each firm is assigned a size-based power label $g_i \in \{\mathrm{S}, \mathrm{M}, \mathrm{L}\}$: small (S), medium (M), or large (L). We initialize these labels using the firm-size variable *BusinessOutput* from the CSMAR database and set quantile thresholds at 0.40 and 0.75. Firms below the 0.40 quantile are classified as S, firms between the 0.40 and 0.75 quantiles are classified as M, and firms at or above the 0.75 quantile are classified as L. Firms with missing or non-positive *BusinessOutput* values are classified as S.

This initialization allows power labels to reflect firms' relative sizes in the empirical data. The labels represent structural power asymmetry among supply chain firms: larger firms typically have stronger market positions and may obtain more favorable conditions in supplier selection and payment-term negotiation (Waller, Johnson, and Davis 1999; Bloom and Perry 2001; Strom 2015; Wuttke et al. 2013). The 0.40 and 0.75 cutoffs are transparent quantile-partition choices for three size groups; sensitivity analysis is left for future work.

The power labels are further mapped to supplier-selection weights, with $\rho(\mathrm{S}) = 1$, $\rho(\mathrm{M}) = 2$, and $\rho(\mathrm{L}) = 5$. These S/M/L labels are simulation labels derived from empirical relative size, rather than legal or regulatory ratings. The mapping is a heuristic relative weight of 1:2:5, while transaction history supplies an additional reinforcement term; sensitivity analysis is left for future work. For buyer i and candidate supplier j, the supplier-selection weight is

$$w^t_{ij} = \left(N^t_{ij} + 1\right) \rho(g_j), \tag{8}$$

where N^t_{ij} is the number of previous transactions between firm i and supplier j by day t. The probability of selecting supplier j is

$$p^{t,\mathrm{sel}}_{ij} = \frac{w^t_{ij}}{\sum_{\ell \in \mathcal{N}^+_i} w^t_{i\ell}}, \tag{9}$$

where $\mathcal{N}^+_i$ is the set of active upstream suppliers available to firm i. The firm selects at most three suppliers using weighted sampling without replacement and allocates the procurement quantity in proportion to their relative weights. This mechanism allows firms' empirical sizes and historical relationships to jointly shape order allocation across the supply chain network.

After a transaction is completed, the buyer updates its inventory and weighted average cost using the quantity and purchase price of the newly received goods. When firm i purchases quantity q at unit price p_{in}, its updated average cost is

$$\bar{c}_{i,\mathrm{new}} = \frac{I_{i,\mathrm{old}} \bar{c}_{i,\mathrm{old}} + q p_{\mathrm{in}}}{I_{i,\mathrm{old}} + q}, \tag{10}$$

where $I_{i,\mathrm{old}}$ and $\bar{c}_{i,\mathrm{old}}$ are the inventory quantity and average cost before delivery. The firm then updates its selling price using a fixed markup rate:

$$p_{i,\mathrm{out}} = \bar{c}_{i,\mathrm{new}}(1 + \mu), \tag{11}$$

where μ is the markup rate. The supplier records an account receivable for the delivered goods, while the buyer records an account payable with the same amount and due date. Physical flows between firms therefore generate corresponding financial claims and obligations.

Payment delay negotiation. Once a transaction is established, buyer i and supplier j negotiate the payment term through an LLM. The negotiation prompt includes the transaction amount, the number of previous transactions, the power labels of both firms, the buyer's cash status, the supplier's inventory status, and the negotiation persona of each party. A negotiation persona combines the firm's aggressiveness level with either a Cash-King or Relationship-First strategy. The negotiated payment delay is

$$\delta^t_{ij} = \mathcal{L}_{\mathrm{neg}}\left(V^t_{ij}, N^t_{ij}, g_i, g_j, C^t_i, I^t_j, \pi_i, \pi_j\right), \tag{12}$$

subject to

$$1 \le \delta^t_{ij} \le \delta_{\max}, \tag{13}$$

where V^t_{ij} is the transaction amount, C^t_i is the buyer's cash status, π_i and π_j are their negotiation personas, and $\delta_{\max}$ is the longest payment term permitted by the system. Representative negotiation inputs and recorded outcomes are provided in the supplementary material.

3.3 Bank Agents

When the estimated near-term funding gap is positive, the firm submits a loan application to the bank agent. The gap is computed from payables due within the next 30 days and projected operating costs, net of current cash and receivables due within that horizon; an immediate operating-cost shortfall can also trigger an emergency request. The bank agent is the LLM-based credit-assessment module, supplemented by a rule-based fast path. For small, low-risk loans whose requested amount is below the threshold θ_{small} and whose debt-to-assets ratio is below the safe threshold ϕ_{safe}, the bank directly approves the application through this fast path, bypassing the LLM and heuristic score. For other loan applications, the bank agent considers the requested amount, debt-to-assets condition, current industry default rate, and visible forum posts to decide whether to approve or reject the application. Bank loans are generally assigned a repayment term of 120 days and an annualized interest rate of 5%. If the loan application is rejected and the firm remains unable to cover the corresponding funding shortfall, it proceeds to the subsequent liquidation stage.

To compare LLM-based credit decisions with traditional rule-based models, the framework also includes a heuristic bank based on a risk score. The risk score of firm i on day t is

$$\mathrm{RiskScore}^t_i = \frac{L^t_i}{A^t_i}\left(1 + \frac{\mathrm{DR}^t_{\mathrm{ind}}}{100} - 2.0\sigma^t\right), \tag{14}$$

Here L^t_i is the firm's outstanding accounts payable plus the scheduled total repayments of its loans, and A^t_i is current

cash plus inventory valued at average cost plus outstanding accounts receivable. If $A_i^t \leq 0$, the implementation sets $L_i^t/A_i^t = 1$; otherwise it uses the raw ratio, which is not clipped at one. The current industry default rate is $\mathrm{DR}_{\mathrm{ind}}^t$, and σ^t is the heat-weighted mean sentiment of the visible top-K forum posts, clipped to $[-1, 1]$. The bank approves the loan application when the risk score is below the approval threshold τ_{bank} and rejects it otherwise. The coefficient 2.0 is a fixed configuration constant controlling this heuristic bank's sensitivity to the sentiment signal; sensitivity analysis is left for future work.

3.4 Bankruptcy Liquidation and Bad-Debt Contagion

When a firm is unable to meet its due obligations and its loan application is rejected, it enters bankruptcy liquidation, the simulator's operational failure state, and ceases its normal production, procurement, and trading activities. The firm's cash and the discounted value of its remaining inventory jointly form the initial liquidation pool. For firm i entering liquidation on day t, the liquidation pool is

$$\Lambda_i^t = C_i^t + \eta_{\mathrm{salv}} I_i^t \bar{c}_i^t, \tag{15}$$

where $\bar{c}_i^t$ is the average inventory cost and η_{salv} is the inventory salvage rate. After entering liquidation, the firm repays only the portion of its obligations that can be covered by its current funds, and any accounts receivable collected subsequently are added to the liquidation pool to continue debt repayment. The liquidation pool is distributed according to a fixed priority order: bank loans and associated interest are repaid first, followed by payments to upstream suppliers in the order in which the corresponding accounts payable fall due. When the firm no longer holds accounts receivable awaiting collection, the liquidation process ends and the firm formally exits the supply chain network.

Once a firm enters bankruptcy liquidation, its upstream suppliers immediately write off their accounts receivable against the failed firm as bad debts. These write-offs reduce the cash inflows expected by upstream suppliers and may intensify their liquidity pressure. If an affected supplier consequently becomes unable to meet its own due obligations and likewise fails to secure emergency financing, it also enters bankruptcy liquidation. A liquidity crisis originating from a single firm can therefore propagate upstream through interfirm accounts receivable and accounts payable, producing successive bad-debt contagion and cascading bankruptcies (Berloco et al. 2021; Ivanov 2017).

4 Experiments

Our experiments are organized around three research questions (RQs). **RQ1:** Can the proposed framework reproduce established supply-chain phenomena on a supply-chain network constructed from the CSMAR database? **RQ2:** How does liquidity risk propagate through the supply-chain network, from an initial information or demand shock to firm failure and upstream bad-debt losses? **RQ3:** Which intervention mechanisms can mitigate liquidity-risk propagation, and how do their effectiveness and implementation costs differ?

4.1 Setup

We use the China Listed Companies Green Supply Chains Research Database in CSMAR. The database aggregates information on major customers, suppliers, and related transaction and management data disclosed in annual reports, environmental, social, and governance reports, and corporate social-responsibility reports of Chinese listed companies, and can be used to construct firm-level supply-chain networks. We primarily use its 2024 Listed Company Customer/Supplier Information table (*GSC_CustomerSupplierInfo*) to identify interfirm transaction relationships from disclosed major customers and suppliers, business-relation types, transaction amounts, and shares of annual business. We also use the operating output, fund-utilization rate, and inventory-utilization rate from the Listed Company Green Supply Chain Management Index table (*GSC_ManagementIndex*) to support firm-size classification and the initialization of firms' financial states. We convert the customer–supplier disclosures into a directed network following the direction of physical flows, apply a Disparity Filter to retain locally salient transaction relationships, and map the network into a downstream-to-upstream multitier supply-chain topology after processing its strongly connected components. The resulting network contains 890 firms and 482 directed links across four tiers. Below, we use d to denote simulation days; for example, $d105$ denotes simulation day 105. Firm identifiers follow the format Tk_j, where Tk denotes the k-th supply-chain tier, ordered from downstream to upstream, and j denotes the within-tier firm index. Detailed construction rules, firm initialization, and the network visualization are provided in the supplementary material.

We conduct a 135-day simulation on this supply-chain topology. The model includes 80 market-participant agents; detailed experimental parameters are provided in the supplementary material. We set this population as a controlled compromise between behavioral coverage and simulation tractability. This size provides multiple independent instances of each role and sustains a continuously updated top-K forum for aggregate demand generation; the role composition and posting probabilities are reported in the supplementary material. Days 0–99 constitute the warm-up period. Beginning on d100, negative-news posts are injected into the forum to create an external information shock. Representative negative-news posts are also provided in the supplementary material.

4.2 Main Experimental Results

During the simulation, four classic supply-chain phenomena emerge naturally: the Matthew effect, bullwhip effect, payment-term bullying, and cascade failure. Their emergence successfully demonstrates the effectiveness of our method. Definitions of these phenomena are provided in the supplementary material. We analyze each phenomenon below.

Matthew effect. Figure 2a compares the orders received by the four T3 firms in two equal 30-day windows. During d0–d29, orders are relatively evenly distributed across all four firms. During d71–d100, T3_01 and T3_02 each receive ap-

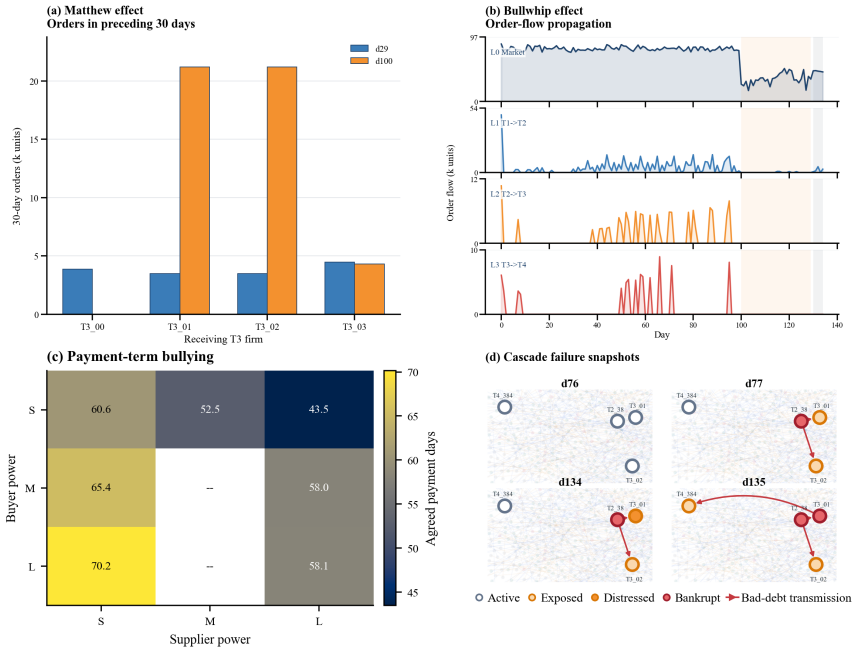


Figure 2: Framework validation on the CSMAR-derived network: (a) order concentration, (b) upstream order amplification, (c) payment terms by buyer–supplier power, and (d) a representative bankruptcy cascade.

proximately 45% of the total orders in that window. This shift from relatively dispersed to persistently concentrated order allocation shows that transaction opportunities increasingly accumulate among a small subset of firms, squeezing the operating space of the remaining firms and reproducing the Matthew effect.

Bullwhip effect. The order-to-demand coefficient of variation (CoV) ratio compares the variability of each tier's orders with that of market demand, and a value above one indicates amplification. The full-window ratios increase from 5.81 for T1–T2 orders to 7.61 for T2–T3 orders and 10.56 for T3–T4 orders, all substantially above one. As shown in Figure 2b, the experiment successfully reproduces the bullwhip effect: batch replenishment and local inventory decisions amplify market-demand variability as orders propagate upstream.

Payment-term bullying. As shown in Figure 2c, large buyers negotiating with small suppliers obtain an average payment term of 70.2 days, whereas small buyers negotiating with large suppliers obtain 43.5 days. Small-to-small transactions average 60.6 days. This pattern is consistent with empirical evidence that weaker suppliers grant important customers more favorable payment terms (Fabbri and Klapper 2016). Because of limitations in the network topology, some firm-pair combinations do not transact in the experiment. In addition, our experiments show that foundation-model choice affects cooperative and compromising tendencies in negotiation. Gemini exhibits the strongest tendency to compromise, whereas Claude is the least willing to compromise and the most likely to produce an impasse. The complete experimental setup and model-pairing results are reported in the supplementary material.

Cascade failure. The simulation successfully reproduces cascade failure. Unpaid obligations generated by downstream firm bankruptcies impose bad-debt losses on upstream suppliers, causing liquidity risk to propagate along the supply chain. Figure 2d presents a representative cascade pattern. The next subsection analyzes the liquidity-risk propagation process in detail.

Panel (a) shows the four-node local path T1_49 → T2_38 → T3_01 → T4_384. Panels (b) and (c) track the three financially exposed firms T2_38, T3_01, and T4_384 over the same d81–d134 window. The branches share the same pre-d100 state; after d99, [S] applies the negative-news shock with the baseline bank, whereas [C] disables exogenous shocks and uses the specified loan-approval counterfactual. Panel (b) shows when the three firms' 30-day funding gaps become positive and whether their loan requests are accepted; panel (c) shows how those events change their operating status.

4.3 Liquidity-Risk Propagation

We select a two-stage cascade from the simulation results and compare it with a matched no-shock counterfactual. Panel (a) of Figure 3 shows the local topology T1_49, T2_38, T3_01, and T4_384 from downstream to upstream. The first red edge records the loss transferred from the failure of T2_38 to T3_01; the second records the loss transferred from the failure of T3_01 to T4_384. Panel (b) reports the corresponding funding gaps and loan outcomes, and panel (c) records the operating-status changes.

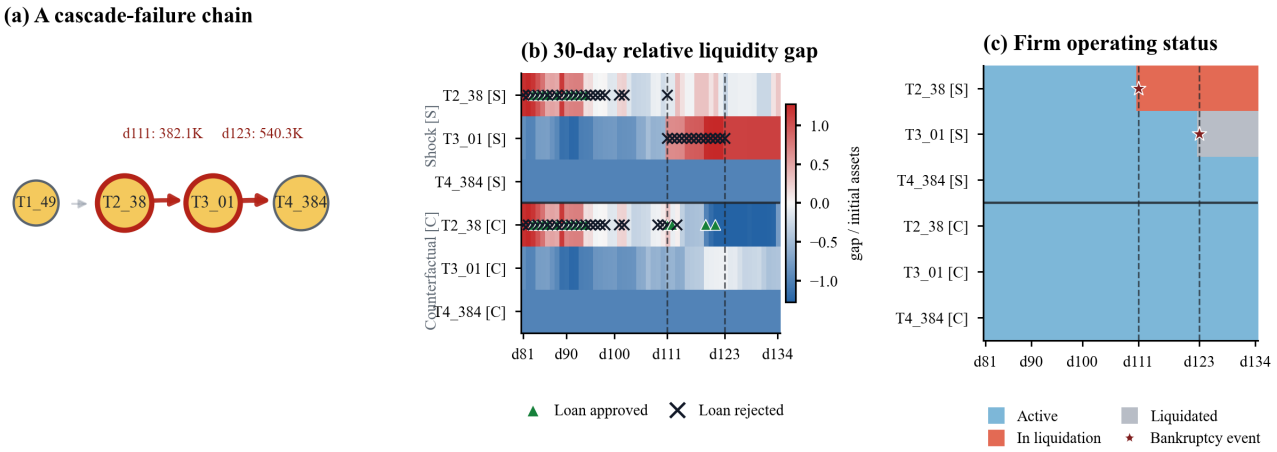


Figure 3: Matched shock/control view of a two-stage liquidity-risk cascade. [S] denotes the run with negative-news and declining demand, and [C] denotes the no-shock loan-approval counterfactual. (a) Directed local path from downstream to upstream; red arrows show bad-debt transmission and labels give the transmission day and nominal amount. (b) The 30-day relative liquidity gap for T2_38, T3_01, and T4_384 from d81 to d134; red cells indicate positive funding shortfalls, blue cells indicate negative gaps, crosses mark rejected loan applications, triangles mark approved applications, and dashed lines mark d111 and d123. (c) Operating status over time for the same three firms; blue is active, orange is in liquidation, gray is liquidated, and stars mark bankruptcy events.

Here, we use the relative liquidity gap over the next 30 days to observe liquidity risk:

$$\mathrm{RLG}_i^t = \frac{\mathrm{AP}_i^{t:t+30} + E_{i,\mathrm{op}}^{t:t+30} - \left(C_i^t + \mathrm{AR}_i^{t:t+30}\right)}{A_i^0}, \quad (16)$$

where $\mathrm{AR}_i^{t:t+30}$ and $\mathrm{AP}_i^{t:t+30}$ are the receivables and payables due within the next 30 days, respectively, $E_{i,\mathrm{op}}^{t:t+30}$ is projected operating expenditure, and A_i^0 denotes initial assets. A positive value indicates a funding shortfall, and a larger relative liquidity gap indicates greater liquidity risk.

T2_38 depends on emergency loans for operating costs in the lead-up to d111. On d111 in [S], its 13.703K emergency loan application is rejected because of payable settlement. Panel (c) records the resulting bankruptcy event and transition out of active operation, while panel (b) shows the associated funding gap and rejected request. The same day, the unpaid 382.116K payable to T3_01 creates the first red arrow in panel (a). At that time, T2_38 has 382.116K in payables to T3_01, 624.082K in receivables from downstream firms, and 352.955K in bank-loan obligations. Although the receivables exceed the payables, they are not immediately available to cover the settlement shortfall.

The bankruptcy of T2_38 causes T3_01 to write off 382.116K on d111, forming the first liquidity-risk transmission. The loss sharply increases the relative liquidity gap of T3_01 in panel (b). It submits regular loan applications on each day from d111 to d122, but all applications are rejected. On d123, T3_01's 40.268K emergency loan application is rejected, and panel (c) marks the second bankruptcy event. The firm holds 21.494K in pre-liquidation cash against 540.310K in payables to T4_384, with no receivables or bank-loan obligations, and becomes bankrupt. This second failure creates the second red arrow in panel (a) and causes T4_384 to write off 540.310K. T4_384 immediately recovers 79.806K from the liquidation funds, leaving 460.504K unrecovered, and remains active because its liquidity buffer is sufficient to absorb the loss.

The specified no-shock loan-approval counterfactual produces different results. It disables the post-d99 negative-news shock and approves the positive loan applications submitted by T2_38, including the payable-settlement request on d111. T3_01 therefore avoids the bad-debt-induced liquidity shortfall; consequently, neither of the two bad-debt transmissions occurs. The three financially exposed firms remain active through the end of the control simulation on d134, while T1_49 appears only as the downstream predecessor in panel (a).

These results explain the transmission process of liquidity risk. After negative information emerges, market demand declines sharply; restricted credit access further converts firms' operating funding shortfalls into bankruptcy. The unpaid trade credit of bankrupt firms subsequently imposes bad-debt losses on upstream suppliers, ultimately forming the two-stage cascade T2_38 $\rightarrow$ T3_01 $\rightarrow$ T4_384.

4.4 Liquidity-Risk Mitigation

Drawing on the three key factors identified in the preceding liquidity-risk propagation pathway, we propose three corresponding mitigation measures. **Opinion control** mitigates information shocks, a **temporary credit guarantee** provides an additional financing channel, and **bad-debt insurance** absorbs realized losses. We evaluate these measures from two perspectives: mitigation effectiveness and implementation cost.

Opinion control. Opinion control manages market conditions by moderating public discussion and restricting posted

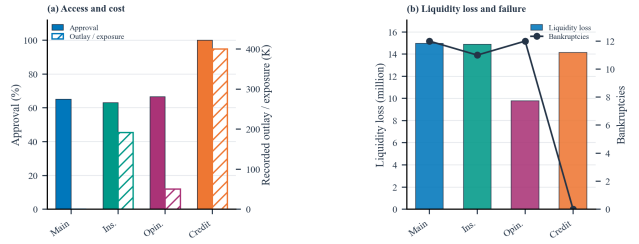


Figure 4: Mitigation effects and implementation costs of different liquidity-risk interventions. (a) Loan approval rates and policy costs under the baseline, bad-debt-insurance, opinion-control, and credit-guarantee scenarios. (b) Liquidity loss and post-shock bankruptcies under the same scenarios.

content. In the experiment, half of the most popular panic-inducing and negative-sentiment posts are removed from the forum each day, while the injection of exogenous crisis information is reduced. Because this intervention has a relatively low implementation cost, we assign it a small fixed cost.

Credit guarantee. A temporary credit guarantee expands financing access for high-risk firms and temporarily alleviates their liquidity risk. In the experiment, policy announcements increase the bank's willingness to approve applications that would otherwise be classified as high risk. At the same time, financial institutions assume additional risk exposure.

Bad-debt insurance. Bad-debt insurance compensates realized losses after a downstream default, with the aim of interrupting liquidity-risk propagation. In the experiment, a hypothetical insurer reimburses part of the uncollected receivables. This mechanism does not directly prevent the initial bankruptcy, but it reduces the harm caused by liquidity risk as it propagates through the supply chain. Its implementation cost equals the realized claim payments.

The final experimental results are shown in Figure 4. Compared with the main scenario, all three interventions reduce liquidity losses to varying degrees, but their effects on post-shock bankruptcies differ substantially. We measure liquidity loss as the decline in firms' liquidity buffers, where the liquidity buffer is defined as

$$\mathrm{LB}_i^t = C_i^t + \mathrm{AR}_i^{t:t+30} - \mathrm{AP}_i^{t:t+30} - E_{i,\mathrm{op}}^{t:t+30}. \qquad (17)$$

The temporary credit guarantee is the most effective intervention for preventing bankruptcies, reducing the post-shock bankruptcy count from 12 to zero. By contrast, bad-debt insurance reduces the count only from 12 to 11, indicating a limited mitigation effect, while opinion control does not reduce the number of bankruptcies. However, both the temporary credit guarantee and bad-debt insurance have only limited effects on liquidity risk. Although opinion control does not prevent bankruptcies, it substantially reduces liquidity losses. By curbing the spread of panic-inducing information, opinion control mitigates the decline in market demand, which in turn leads the bank to assess market conditions more favorably and increases the loan approval rate.

These results further reveal a clear trade-off among intervention timing, mitigation effectiveness, and implementation cost: the earlier emerging problems are identified and addressed, the more likely an intervention is to reduce the cost of containing liquidity risk and improve mitigation effectiveness.

5 Conclusion

This paper introduced SupplyNet, an LLM-based multi-agent simulation framework that connects cognitive market and bank agents with the physical and financial operations of a supply chain. On a network reconstructed from CSMAR data, classic supply-chain phenomena, including Matthew concentration, bullwhip amplification, payment-term bullying, and cascading insolvency, emerge naturally within the framework. The traced pathway shows that cognition-driven shocks contract market demand and tighten credit, thereby accelerating cash depletion, bankruptcy, and upstream bad-debt propagation. The intervention experiments further show that different policies have distinct effects on reducing liquidity losses and preventing firm bankruptcy, revealing trade-offs among intervention timing, mitigation effectiveness, and implementation cost. The results also indicate that the earlier emerging risks are identified and addressed, the more likely interventions are to improve mitigation effectiveness at lower cost.

References

Bardoscia, M.; Battiston, S.; Caccioli, F.; and Caldarelli, G. 2015. DebtRank: A microscopic foundation for shock propagation. *PLOS ONE*, 10(6): e0130406.

Basel Committee on Banking Supervision. 2008. Principles for Sound Liquidity Risk Management and Supervision. Technical report, Bank for International Settlements, Basel, Switzerland.

Berloco, C.; De Francisci Morales, G.; Frassineti, D.; Greco, G.; Kumarasinghe, H.; Lamieri, M.; Massaro, E.; Miola, A.; and Yang, S. 2021. Predicting corporate credit risk: Network contagion via trade credit. *PLOS ONE*, 16(4): e0250115.

Bloom, P. N.; and Perry, V. G. 2001. Retailer power and supplier welfare: The case of Wal-Mart. *Journal of Retailing*, 77(3): 379–396.

Cho, Y.-M.; Guntuku, S. C.; and Ungar, L. 2025. Herd Behavior: Investigating Peer Influence in LLM-based Multi-Agent Systems. arXiv:2505.21588.

Colicchia, C.; Creazza, A.; Noe, C.; and Strozzi, F. 2019. Information sharing in supply chains: a review of risks and opportunities using the systematic literature network analysis (SLNA). *Supply Chain Management: An International Journal*.

Dolgui, A.; and Ivanov, D. 2020. Exploring supply chain structural dynamics: New disruptive technologies and disruption risks. *International Journal of Production Economics*, 229: 107886 – 107886.

Drehmann, M.; and Nikolaou, K. 2013. Funding Liquidity Risk: Definition and Measurement. *Journal of Banking & Finance*, 37(7): 2173–2182.

Fabbri, D.; and Klapper, L. F. 2016. Bargaining Power and Trade Credit. *Journal of Corporate Finance*, 41: 66–80.

Gelsomino, L. M.; Mangiaracina, R.; Perego, A.; and Tumino, A. 2016. Supply chain finance: a literature review. *International Journal of Physical Distribution & Logistics Management*, 46(4).

Hofmann, E. 2005. *Supply Chain Finance: some conceptual insights*, 203–214. Deutscher Universitätsverlag. ISBN 9783322821652.

Ivanov, D. 2017. Simulation-based ripple effect modelling in the supply chain. *International Journal of Production Research*, 55(7): 2083–2101.

Kano, L.; Tsang, E. W.; and Yeung, H. W.-c. 2020. Global value chains: A review of the multi-disciplinary literature. *Journal of International Business Studies*, 51(4): 577–622.

Kessler, M.; Arlinghaus, J. C.; Rosca, E.; and Zimmermann, M. 2022. Curse or Blessing? Exploring risk factors of digital technologies in industrial operations. *International Journal of Production Economics*, 243: 108323.

Li, M. 2019. Overconfident distribution channels. *Production and Operations Management*, 28(6): 1347–1365.

Li, N.; Gao, C.; Li, M.; Li, Y.; and Liao, Q. 2024. EconAgent: Large Language Model-Empowered Agents for Simulating Macroeconomic Activities. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*, 15523–15536. Association for Computational Linguistics.

Park, J. S.; O'Brien, J.; Cai, C. J.; Morris, M. R.; Liang, P.; and Bernstein, M. S. 2023. Generative Agents: Interactive Simulacra of Human Behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23)*, 1–22. ACM.

Petratos, P.; and Faccia, A. 2023. Fake news, misinformation, disinformation and supply chain risks and disruptions: risk management and resilience using blockchain. *Annals of Operations Research*, 1 – 28.

Ricciotti, F. 2020. From value chain to value network: a systematic literature review. *Management Review Quarterly*, 70(2): 191–212.

Sadeek, S. N.; and Hanaoka, S. 2023. Assessment of text-generated supply chain risks considering news and social media during disruptive events. *Social Network Analysis and Mining*, 13: 1–18.

Schmidt, C. G.; Wuttke, D. A.; Ball, G. P.; and Heese, H. S. 2020. Does social media elevate supply chain importance? An empirical examination of supply chain glitches, Twitter reactions, and stock market returns. *Journal of Operations Management*, 66: 646–669.

Sterman, J. D. 1989. Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment. *Management Science*, 35(3): 321–339.

Strom, S. 2015. Big companies pay later, squeezing their suppliers. *The New York Times*, 6.

Taghikhah, F.; Voinov, A.; Shukla, N.; Filatova, T.; and Anufriev, M. 2021. Integrated modeling of extended agro-food supply chains: A systems approach. *European Journal of Operational Research*, 288(3): 852–868.

Waller, M.; Johnson, M. E.; and Davis, T. 1999. Vendor-managed inventory in the retail supply chain. *Journal of Business Logistics*, 20(1): 183.

Wiedmer, R.; and Griffis, S. E. 2021. Structural characteristics of complex supply chain networks. *Journal of Business Logistics*, 42(2): 264–290.

Wuttke, D. A.; Blome, C.; Foerstl, K.; and Henke, M. 2013. Managing the Innovation Adoption of Supply Chain Finance: Empirical Evidence From Six European Case Studies. *Journal of Business Logistics*, 34(2): 148–166.

Yang, Y.; Lin, J.; Liu, G.; and Zhou, L. 2021. The behavioural causes of bullwhip effect in supply chains: A systematic literature review. *International Journal of Production Economics*, 236: 108120.

Yang, Z.; Zhang, Z.; Zheng, Z.; Jiang, Y.; Gan, Z.; Wang, Z.; Ling, Z.; Chen, J.; Ma, M.; Dong, B.; Gupta, P.; Hu, S.; Yin, Z.; Li, G.; Jia, X.; Wang, L.; Ghanem, B.; Lu, H.; Lu, C.; Ouyang, W.; Qiao, Y.; Torr, P.; and Shao, J. 2024. OASIS: Open Agent Social Interaction Simulations with One Million Agents. arXiv:2411.11581.

Zhu, Y.; Zhou, L.; Xie, C.; Wang, G.-J.; and Nguyen, T. V. 2019. Forecasting SMEs' credit risk in supply chain finance with an enhanced hybrid ensemble machine learning approach. *International Journal of Production Economics*, 211: 22–33.