The Network Foundations of Credit Counterparty Risk

Theory and Evidence

Jonathan Brogaard and Belinda (Chen) Chen

Discussant: Zhuo Chen, PBC School of Finance, Tsinghua University

SKBI / UBS Quant Investment Forum: AI and Finance

June, 16 2026

Motivation and Overview

- **Research question and mechanism**
  - Do ***production-network linkages shape credit counterparty risk*** and help explain firms' CDS spreads?
  - Supplier shocks lower cash flows, raise equity volatility, reduce distance-to-default, and thus widen CDS spreads.
- **Theoretical foundation and ML estimation**
  - Theory: default intensity and the CDS spread load on an n×n production-network spillover matrix A = [a_g(i),g(j)] — generalizing two-industry (mapped to firms), tail-only models to cascading, asymmetric, ordinary-shock contagion.
  - ML estimation: a deliberately parsimonious GNN of which nodes ***carry firm characteristics and macro variables*** (112 node features) and of which directed edges carry the ***spillovers*** a_ij — uniquely using the full network topology; built for interpretability and benchmarked v.s. a node-only CNN (edges removed).
- **Main findings**
  - Incorporating network-edge information markedly raises explanatory power for firm credit spreads: GNN2 RMSE 0.888 < 1.338 for the node-only CNN2 — a 21.8% network-attributable spread change and 0.56 incremental R^2.
  - The network's role is state-dependent — largest during production-network disruptions: 2008 financial crisis, the 2009-10 post-crisis reconfiguration, 2018 trade tensions, and 2020 COVID-19 period.
  - It is concentrated in intermediate, logistics-heavy, low-substitutability industries (electrical equipment, shipbuilding/railroad, defense), and is stronger for investment-grade (larger) firms.
- **Contribution**
  - First comprehensive structural framework for how the ***inter-industry (firm) production network shapes credit risk*** — a microfoundation for counterparty risk.
  - Introduces graph neural networks to credit-risk research as a tool for ***pricing the full network topology*** — though in the current draft that topology is estimated at the industry (sector) level and mapped to firms within each sector.

Comment 1: Prediction or pricing?

- **Prediction may not be pricing**
  - GNNs forecast log CDS spreads better than node-only and non-network benchmarks.
  - But calling network risk "priced" requires an ***asset-pricing framework/interpretation*** beyond lower prediction error.
  - CDS spreads are not pure default risk — they also embed liquidity, dealer/market-maker frictions, etc.; the largest GNN gains (e.g., 2008 financial crisis or COVID-19 period) coincide with both liquidity and risk premia spiking (*fundamental, liquidity, or both?*).
- **Bring liquidity and dealers into the model**
  - Theory: the spread prices only default (distance-to-default) — extend it with a ***liquidity / non-default premium*** so network shocks can also transmit through liquidity.
  - ML: add CDS liquidity proxies (bid-ask, depth, number of contributing dealers, etc.) and market-wide liquidity as additional features — does the network's incremental power survive?
  - Dealers (if data allowed): a name's CDS is quoted by a varying, overlapping set of dealers — control for dealer identity / shared-dealer exposure (the paper sets dealer cores aside) to absorb common market-maker shocks.
- **Pin down pricing due to supplier shocks due to production network**
  - Alternative dependent variable: remove the impact of liquidity from CDS spreads and isolate the component driven by fundamentals.
  - Use predetermined network exposure to predict CDS changes / excess returns around supply-chain shocks.
  - Separate a genuine ***production-network default*** channel from ***liquidity- or dealer-driven*** comovement in CDS spreads.

Comment 2: Industry-level network topology, or industry information?

- **Core concern: *network topology* may partly capture *industry information***
  - The edge matrix A is an industry-to-industry network mapped to firms: $a_{ij,t} = A_{g(i),g(j),t}$.
  - Firms in the same industry share the ***same edge profile*** to other industries, so the GNN receives an industry fingerprint.
  - This is still a network, but it also embeds industry identity. Example: semiconductor firms may have very different CDS dynamics and edge profiles from retailers; if the CNN benchmark is not given industry dummies or embeddings, part of the GNN–CNN gap may reflect learning "semiconductor vs. retail," not only supply-chain propagation.
- **Incremental value to isolate**
  - Current comparison: firm characteristics only v.s. firm characteristics plus industry-network edges.
  - Cleaner comparison: industry-aware node model v.s. the same information plus ***directed network topology***.
  - If the GNN still wins, evidence is stronger that ***directional industry linkages*** matter beyond industry membership.
- **Suggested tests**
  - Industry-aware node baseline: control for NAICS-2 industry dummies/embeddings, inbound/outbound exposure, industry size, and industry-by-month effects.
  - Placebo edges: shuffled industries, transposed A, or dense equal-weight edges, with architecture fixed.
  - Validate top upstream edges against real supplier-customer links and concrete shocks such as tariffs or recent semiconductor shortages.

Comment 3: Additional comments

- **Interpretability and validation**
  - For highlighted episodes, connect the top edges to tariffs, defense contracts, or semiconductor shortages.
  - Add the April 2025 tariff war as a pure out-of-sample test — a fresh production-network shock (the sample ends in 2020) mirroring 2018 — to see whether the network effect persists.
- **Benchmarks and specification**
  - Add industry size (e.g., sales or total market cap) as a node feature and test whether larger industries show higher NSC — separating a size/granularity channel from network position.
  - Clarify Table 6: several alternatives beat the main baseline — how can we think of these alternatives?
- **Exposition and reporting**
  - A brief timeline of the estimation steps — rank normalization, rolling-VAR edges, macro variables, validation, and t+1 prediction — would help readers follow the design.
  - Since the target is log spreads, it would also be helpful to report level-spread RMSE/MAE and basis-point magnitudes for economic interpretation.