

Dynamic Bayesian Networks for Early-Warning Signals of Systemic Risk in FinTech Ecosystems

Chandra Yudhistira Priambodo^{1,*}, Sunyi Praptaningsih²

^{1,2}Magister of Computer Science, Amikom Purwokerto University

ABSTRACT

FinTech ecosystems increasingly exhibit systemic fragility due to tight coupling among platforms, partner banks, payment rails, cloud infrastructure, and user behavior. This study proposes a Dynamic Bayesian Network (DBN) early-warning framework that models systemic risk as a latent temporal regime inferred from multi-channel evidence, including operational reliability (transaction failures, settlement delay), liquidity and funding conditions (buffer drawdowns, spreads), infrastructure incidents, and network exposure proxies. Backtesting under strict time-ordered evaluation demonstrates that the DBN yields decision-grade probabilistic monitoring performance and operationally usable alert behavior. The model achieves ROC-AUC = 0.873 and PR-AUC = 0.417, indicating strong discrimination under rare-event conditions. Probability quality remains stable with a Brier Score = 0.094, supporting threshold-based governance without overconfident miscalibration. In early-warning utility, the system provides an average lead time of 9.6 days before labeled stress episodes, while maintaining a manageable operational burden of 1.7 false alarms per 30 days. Alert behavior is consistent over time, reflected by an Alert Stability Index = 0.931, and stress episode coverage remains high with recall = 0.842. Temporal inspection of the EWS shows sustained elevation during stress windows rather than short-lived spikes, with posterior risk increasing most sharply when operational degradation co-occurs with liquidity/funding pressure—consistent with systemic coupling mechanisms rather than isolated incidents. Overall, the results indicate that DBNs provide a unified, explainable, and deployable mechanism for systemic-risk early warning in FinTech ecosystems, enabling probabilistic escalation policies, evidence-based attribution, and scenario-driven fragility mapping.

Keywords Dynamic Bayesian Networks, Systemic Risk, FinTech Ecosystems, Early-Warning Systems, Probabilistic Graphical Models, Risk Monitoring, Contagion and Dependency Networks, Operational Resilience

INTRODUCTION

FinTech ecosystems have become critical infrastructure for payments, lending, wealth management, and digital asset activity, increasingly operating through tightly coupled platforms, partner banks, payment rails, cloud services, and third-party data providers. This interdependence enables rapid scale but also creates systemic fragility, where localized disruptions—such as settlement delays, liquidity shortfalls, cyber incidents, or abrupt funding shocks—can propagate across institutions and markets faster than traditional supervision and risk controls can respond. The central problem is that systemic risk in FinTech is often nonlinear, cross-channel, and time-varying, making it difficult to detect early using single-indicator monitoring or static scorecards [1], [2].

Early-warning monitoring is therefore a priority, yet it remains challenging because FinTech risk signals are heterogeneous (operational logs, user

Submitted: 5 October 2025
Accepted: 18 December 2025
Published: 12 August 2026

Corresponding author
Chandra Yudhistira Priambodo,
25MA41D024@student.amikom
purwokerto.ac.id

Additional Information and
Declarations can be found on
[page 220](#)

© Copyright
2026 Priambodo and
Praptaningsih

Distributed under
Creative Commons CC-BY 4.0

behavior, funding conditions, market volatility, and network exposures) and often noisy, partially missing, or reported at different frequencies. Conventional approaches frequently rely on threshold rules or supervised classification using past crisis labels; however, crises are rare, labels are ambiguous, and model performance can degrade rapidly under regime shifts. Moreover, the most important question for practitioners is not only “will stress occur,” but “how does stress emerge and spread,” which requires probabilistic reasoning over dependencies rather than isolated feature impacts [3], [4].

Existing systemic-risk analytics in finance provide valuable foundations—network contagion models, stress testing, and econometric early-warning indicators—but are not always well suited to FinTech conditions where platform outages, API dependencies, and rapid user reactions can dominate transmission dynamics. Many data-driven models also treat the system as a black box or assume independent observations, which limits interpretability and operational trust when alerts drive high-stakes decisions. As a result, there is an unresolved methodological gap: a unified approach that models temporal evolution, interdependencies, and uncertainty while remaining explainable and deployable in real time [5], [6].

This research addresses that gap by proposing Dynamic Bayesian Networks (DBNs) as a core modeling framework for early-warning signals of systemic risk in FinTech ecosystems. DBNs represent the ecosystem as a probabilistic graphical model that evolves over time, allowing the latent systemic-risk regime (e.g., normal vs. stressed) to be inferred from multivariate evidence while explicitly encoding dependency structure among operational, liquidity, and market channels. Unlike purely correlational forecasting models, the DBN formalism supports sequential belief updates and scenario reasoning, enabling risk monitoring that is both adaptive and auditable under uncertainty [7], [8].

The primary objective of the study is to design and evaluate a DBN-based early-warning system that (i) integrates multi-source FinTech evidence into a coherent probabilistic state estimate, (ii) produces a continuous risk signal suitable for governance thresholds (warning vs. alarm), and (iii) supports explainability by tracing posterior risk changes to structured evidence patterns. A secondary objective is to establish a backtesting and stress-testing protocol aligned with operational constraints—time-ordered evaluation, rare-event metrics, lead-time analysis, and alert stability—so that model performance reflects deployable utility rather than retrospective optimization [9], [10].

The novelty of this work lies in framing systemic risk in FinTech as a latent temporal regime inferred from interdependent evidence streams, rather than as a static classification of isolated snapshots. Methodologically, the contribution is the integration of (a) dependency-aware probabilistic modeling (DBN structure), (b) sequential inference for real-time monitoring, and (c) operationally grounded evaluation emphasizing calibration, false-alarm burden, and lead time. Practically, the framework is designed to be extensible: nodes can be added for new infrastructures (e.g., additional rails, custodians, cloud regions) and the model can be stress-tested under targeted shocks to identify fragility amplifiers such as concentration and liquidity coupling.

Literature Review

Systemic risk research has traditionally been grounded in banking and capital

markets, where crises emerge through leverage cycles, liquidity spirals, correlated exposures, and confidence shocks [11], [12]. Early strands emphasize macro-financial linkages and procyclicality, documenting how common shocks and balance-sheet constraints can synchronize distress across institutions and markets [13], [14]. This literature establishes the core premise that systemic risk is not reducible to idiosyncratic failures; instead, it is an emergent property of interconnection and feedback loops that can amplify relatively small perturbations into widespread disruption [15].

A second major line of work formalizes contagion through financial networks, modeling interbank liabilities, counterparty exposures, and clearing-and-settlement dependencies. Network-based studies show that topology (e.g., concentration, core-periphery structure) and shock allocation mechanisms can strongly determine whether losses remain localized or cascade. These findings are directly relevant to FinTech ecosystems, where dependency graphs often extend beyond balance sheets to include payment rails, shared cloud infrastructure, API providers, and platform-mediated user flows. The network literature motivates representing interdependence explicitly rather than assuming independence among observed signals [16], [17].

Early-warning systems (EWS) in finance commonly rely on econometric indicators—credit-to-GDP gaps, asset-price misalignments, volatility shifts, and liquidity measures—often estimated using logistic models, probit specifications, or panel regressions. While these approaches can be effective for macroprudential surveillance, they are less aligned with FinTech settings where risk can materialize through operational disruption, cyber/fraud shocks, or platform-specific run dynamics that evolve at higher frequency and through different transmission channels. Moreover, indicator-based EWS often struggle when labels are sparse or when structural breaks occur, which is typical in rapidly evolving FinTech markets [18], [19].

Recent machine learning research has expanded EWS toolkits with tree ensembles, deep learning, and anomaly detection to capture nonlinearities and high-dimensional patterns. These methods frequently improve predictive performance but can be difficult to govern operationally because their outputs are not inherently probabilistic, not naturally sequential, and often less interpretable under shifting regimes. In addition, “black-box” models may conflate correlation with causal propagation, leading to alerts that are difficult to justify, hard to debug, and prone to instability when feature distributions drift. This tension between predictive power and decision-grade interpretability is central to systemic risk monitoring [20], [21].

Probabilistic graphical models provide a complementary foundation by encoding dependencies through conditional independencies, supporting reasoning under uncertainty and partial observability. Bayesian Networks have been applied to risk assessment, fault diagnosis, and financial decision support because they produce interpretable probability statements and can combine heterogeneous evidence sources. However, static Bayesian Networks are insufficient when risk evolves over time and when past conditions shape future vulnerability. This motivates temporal extensions—Hidden Markov Models, state-space models, and, most directly, Dynamic Bayesian Networks—which can represent evolving regimes and time-indexed dependencies [22], [23].

Within temporal probabilistic modeling, DBNs are particularly suitable for systemic-risk early warning because they unify (i) regime persistence via transition structure, (ii) multi-channel evidence integration via observation dependencies, and (iii) sequential updating via filtering. The key methodological advantage is that the model's risk belief is updated as evidence accumulates, which aligns with real-time monitoring requirements. Furthermore, DBNs support scenario reasoning by perturbing evidence nodes or transition dynamics and quantifying posterior shifts, enabling a controlled assessment of fragility amplifiers such as liquidity coupling, settlement bottlenecks, and concentration risk [24], [25].

Evaluation practices in early-warning research increasingly emphasize that high headline discrimination (e.g., AUC) is insufficient for deployment. For rare-event settings, calibration and operational stability matter because alerts must be credible and consistent under uncertainty. Metrics such as Brier score, precision-recall analysis, lead time, and false-alarm rate over time better reflect the governance trade-offs faced by regulators and platform operators. This literature implies that systemic-risk monitoring should be assessed as a decision-support system, not solely as a predictive classifier [26], [27].

Finally, FinTech-specific scholarship highlights additional risk channels that are underrepresented in traditional systemic-risk models: API and cloud concentration, cyber-fraud adversarial dynamics, instantaneous user response to rumors, and platform-driven liquidity transformation that can mimic banking-like maturity mismatch. These characteristics reinforce the need for a model class that can (a) integrate operational and behavioral telemetry with financial indicators, (b) remain robust under incomplete data, and (c) provide explainable pathways from evidence to systemic risk escalation. DBNs offer a principled mechanism to meet these requirements while remaining extensible as new infrastructures and interaction patterns emerge [28], [29].

Methodology

Research Design and System Boundary Definition

This study adopts a quantitative, model-driven early-warning design in which systemic risk is treated as a latent regime (e.g., normal, stressed, crisis) that evolves over time and is partially observed through FinTech ecosystem signals (payments, lending, market microstructure, liquidity proxies, network exposures). The core methodological choice is a Dynamic Bayesian Network (DBN) because it represents temporal dependence, conditional independence, and probabilistic uncertainty in a single coherent framework, enabling both prediction and explanation of risk drivers.

The system boundary is operationalized by specifying the ecosystem entities (FinTech firms, banks/partners, payment rails, crypto/market venues, major merchants, infrastructure providers) and the interaction channels (transaction flows, credit relationships, liquidity channels, shared technology dependencies, common investor/funding shocks). These entities are represented as nodes in a time-indexed graph, while edges encode hypothesized causal or dependency pathways (e.g., funding stress → liquidity shortfall → transaction failures → reputational run dynamics). Formally, we define a latent systemic-risk state vector X_t and observed evidence E_t at time t . The DBN assumes a first-order Markov property:

$$P(X_t | X_{0:t-1}) = P(X_t | X_{t-1}) \quad (1)$$

This equation states that (conditional on the immediately previous risk state) the future state does not directly depend on older states. Methodologically, this is a pragmatic balance: it preserves temporal memory through state recursion while keeping inference tractable and data requirements reasonable. If empirical diagnostics show longer memory, the model can be extended with higher-order lags or “memory” nodes, but the baseline design starts with first-order dynamics to maintain identifiability. To ensure the methodology is reproducible and auditable, the boundary definition is recorded in a structured inventory that maps each node to its data source, update frequency, and measurement quality.

Data Collection, Preprocessing, and Feature Engineering

Data are assembled as synchronized time series and event streams representing ecosystem health. Typical inputs include: (i) transaction throughput and failure rates, (ii) liquidity and settlement delay proxies, (iii) credit performance (delinquencies, roll rates), (iv) funding conditions (spreads, withdrawal rates, deposit flows), (v) market risk signals (returns, volatility, order imbalance), and (vi) network-level exposures (counterparty concentration, dependency centrality). Data ingestion standardizes timestamps, handles missingness explicitly (missing-as-information when appropriate), and aligns all variables to a common temporal resolution (e.g., daily or hourly).

Preprocessing emphasizes comparability across heterogeneous signals. Numerical streams are winsorized or robustly scaled to reduce extreme-noise dominance; categorical events (e.g., outages, regulatory actions) are encoded into indicator variables with meaningful persistence windows (e.g., event decays). Where the DBN uses discrete nodes, continuous variables are discretized using quantile bins or risk-relevant thresholds (e.g., “normal / elevated / extreme”), with binning rules documented to avoid leakage and maintain interpretability. A canonical transformation for market-like series is the log return:

$$r_t = \ln \left(\frac{p_t}{p_{t-1}} \right) \quad (2)$$

This converts price levels into additive changes and reduces nonstationarity. In the systemic-risk setting, returns (and their higher moments) are used as stress amplifiers: persistent negative returns or volatility clustering often co-move with funding pressure and liquidity impairment. The methodological point is not to “predict prices,” but to treat market stress as probabilistic evidence influencing latent systemic regimes. Feature engineering then aggregates micro-signals into interpretable macro-indicators (e.g., rolling volatility, tail-risk counts, settlement delay percentiles, concentration metrics).

Figure 1 documents the transformation chain from raw operational/market streams into DBN-compatible evidence. It clarifies where alignment, missingness handling, and discretization occur, which is essential because DBN inference quality is highly sensitive to consistent time indexing and stable measurement definitions. In systemic-risk monitoring, “preprocessing” is not merely technical hygiene; it is governance. The pipeline makes explicit where

design decisions can introduce false alarms (e.g., aggressive binning, over-smoothing) or miss early signals (e.g., overly coarse aggregation). This diagram therefore supports reproducibility and auditability.




Figure 1 Feature Pipeline from Raw Streams to DBN Evidence

DBN Structure Specification and Parameter Learning

The DBN is defined as a two-slice temporal Bayesian network with (a) intra-slice dependencies capturing contemporaneous relationships among variables within t , and (b) inter-slice dependencies capturing temporal transitions from $t - 1$ to t . Structure specification is conducted using a hybrid approach: domain constraints (to enforce plausible directionality and prevent nonsensical edges) plus data-driven selection (to discover dependencies not obvious a priori). This mitigates overfitting while preserving causal interpretability. Let $Z_t = \{Z_t^{(1)}, \dots, Z_t^{(n)}\}$ be the set of DBN variables at time t , including the latent systemic state S_t and observed indicators. The joint distribution factorizes as:

$$P(Z_{1:T}) = \prod_{t=1}^T \prod_{i=1}^n P(Z_t^{(i)} \mid \text{Pa}(Z_t^{(i)})) \quad (3)$$

This factorization is the methodological engine of the DBN: it decomposes a high-dimensional joint distribution into local conditional probability tables (CPTs) or conditional densities. Practically, it allows the model to learn how specific signals (e.g., “settlement delay = high”) change the probability of being in a stressed regime, while respecting conditional independencies.

Parameters are learned via maximum likelihood or Bayesian estimation. For discrete CPTs, a common Bayesian approach uses a Dirichlet prior to stabilize estimates under sparse regimes:

$$\theta_{ijk} = P(Z^{(i)} = k \mid \text{Pa}(Z^{(i)}) = j), \quad \theta_{ij.} \sim \text{Dirichlet}(\alpha_{ij.}) \quad (4)$$

Here, α controls smoothing: larger α reduces variance but can underreact to emerging patterns; smaller α is more responsive but noisier. This trade-off is critical for early-warning objectives, where sensitivity to nascent stress must be balanced against false alarms.

Figure 2 depicts the DBN as a two-slice temporal structure: nodes at $t - 1$ influence nodes at t through transition edges, and the latent regime $S(t)$ influences multiple observed variables within the current slice. This makes the modelling logic explicit: systemic risk is not directly observed but manifests through correlated disruptions and funding/liquidity dynamics.




Figure 2 Two-Slice Dynamic Bayesian Network Structure

In early-warning contexts, the two-slice design is critical because it separates persistence (state transitions $S(t - 1) \rightarrow S(t)$) from contemporaneous evidence coupling (e.g., $S(t) \rightarrow \text{Failures}(t)$). This structure improves interpretability: one can distinguish “risk remains elevated due to inertia” from “risk rises sharply due to new adverse evidence.” Figure 3 demonstrates the operational output of the DBN: a probability-valued early-warning signal over time. The two horizontal lines represent governance thresholds (warning vs. alarm), and shaded intervals represent stress episodes used for backtesting (or discovered episodes in monitoring).




Figure 3 Early-Warning Signal with Thresholds and Stress Episodes

From a methodological standpoint, this visualization supports threshold selection and stability assessment. It reveals whether the EWS increases early enough (lead time), whether it remains elevated during stress (persistence), and whether it produces excessive oscillations outside stress (false alarms). These properties cannot be judged from aggregate metrics alone; time-series inspection is essential for deployable early-warning systems.

Inference and Early-Warning Signal Construction

Given a learned DBN, early-warning requires online inference: updating beliefs about systemic risk as new evidence arrives. The primary quantity of interest is the filtered posterior of the latent systemic state S_t given all observations up to time t . Filtering supports real-time monitoring, while prediction supports lead-time warnings. The filtering recursion can be expressed as:

$$P(S_t \mid e_{1:t}) \propto P(e_t \mid S_t) \sum_{S_{t-1}} P(S_t \mid S_{t-1}) P(S_{t-1} \mid e_{1:t-1})$$

(5)

This equation formalizes how the warning signal is constructed: yesterday’s belief $P(S_{t-1} \mid e_{1:t-1})$ is propagated through the transition model $P(S_t \mid S_{t-1})$, then corrected by the likelihood of today’s evidence $P(e_t \mid S_t)$. Methodologically, it ensures the warning signal is probabilistic, sequential, and evidence-consistent, rather than a static score. The early-warning indicator is defined as a probability (or a function of it), for example:

$$EWS_t = P(S_t = \text{“stressed”} \mid e_{1:t})$$

(6)

This yields an interpretable monitoring variable in $[0,1][0,1][0,1]$. It can be extended to expected-loss form by weighting regimes by severity costs (e.g., operational disruption cost, liquidity shortfall cost). Thresholds and alert rules are selected to prioritize lead time while controlling false positives, and can be adapted by risk appetite (regulator vs. platform operator).

Table 1 provides a menu of operational alerting policies that can be layered on top of DBN-derived probabilities. Each rule is defined in an implementation-ready form, which is necessary to ensure that backtesting and production alerting match exactly. The table also makes governance trade-offs explicit. In systemic-risk contexts, the “best” threshold is not purely statistical; it depends on who bears the cost of false positives (operational disruption) versus false negatives (missed crises). By documenting rule strengths and limitations, the methodology supports defensible deployment decisions.

Table 1 Alert Rule Variants: Static Threshold, Percentile-Based, and Cost-Sensitive				
Rule Type	Definition	Parameters	Strength	Limitation
Static Threshold	Alarm if $EWS_t \geq \tau$	$\tau \in [0, 1]$	Simple, transparent, easy to govern.	May underperform under regime drift or seasonality.
Two-Level Threshold	Warning if $EWS_t \geq \tau_w$, Alarm if $EWS_t \geq \tau_a$	$\tau_w < \tau_a$	Supports staged response playbooks.	Requires operational maturity to act on warnings.
Percentile-Based	Alarm if $EWS_t \geq Q_p(EWS \text{ history})$	p (e.g., 95th)	Adapts to baseline shifts in probability levels.	Can mask broad risk increases if history already elevated.
Cost-Sensitive	Choose τ minimizing $C(\tau) = c_{FP}FP(\tau) + c_{FN}FN(\tau)$	c_{FP}, c_{FN}	Aligns thresholds with explicit risk appetite.	Needs reliable event labels and stable cost assumptions.
Lead-Time Optimized	Select τ maximizing average lead time subject to FP cap	FP cap, horizon H	Directly optimizes early-warning usefulness.	May reduce calibration quality if over-optimized.

Validation, Backtesting, and Stress-Testing Protocol

Validation is performed using temporal backtesting with strict time ordering (train on past, evaluate on future) to prevent leakage. Outcomes are defined as stress events (e.g., sustained transaction failure spikes, liquidity drawdowns, major outage cascades, funding run indicators) using pre-registered rules.

Performance is evaluated on both classification quality and early-warning utility, emphasizing lead time and stability rather than only point accuracy. A core probabilistic calibration metric is the Brier score:

$$BS = \frac{1}{T} \sum_{t=1}^T (EWS_t - y_t)^2 \quad (7)$$

where $y_t \in \{0,1\}$ indicates whether a stress event occurs in the defined horizon. This metric directly penalizes overconfident wrong warnings and rewards well-calibrated probabilities—essential for decision-making under uncertainty. Complementary metrics include ROC-AUC/PR-AUC (for ranking), expected lead time, and false-alarm rate per unit time. Stress testing is conducted via scenario perturbations to key evidence nodes or transition dynamics (e.g., funding shock, settlement latency spike, counterparty default), examining how $P(S_t | e_{1:t})$ responds. A compact measure of scenario impact is the KL divergence between baseline and stressed posteriors:

$$D_{KL}(P || Q) = \sum_s P(s) \ln \frac{P(s)}{Q(s)} \quad (8)$$

Large divergence indicates that the scenario materially changes inferred systemic risk. This is used to rank vulnerabilities, identify “risk accelerators,” and test whether the DBN produces sensible responses under controlled shocks.

Result and Discussion

Experimental Setup and Backtesting Protocol

The DBN early-warning system was evaluated using a strictly time-ordered backtesting design to preserve causal validity and prevent information leakage. All evidence streams (operations, liquidity, funding, infrastructure incidents, and market stress indicators) were aligned to a single time index, then split into training and testing segments such that the testing window occurs entirely after the training window. The prediction objective was defined as a forward-looking early-warning task: the model continuously updates a posterior risk score and triggers alerts when risk exceeds a policy threshold.

Stress events were labeled using operationally grounded criteria (e.g., sustained settlement delay elevation, transaction failure spikes, liquidity drawdowns, or combined service degradation). This event definition is critical: the DBN is not evaluated on “perfect crisis dates,” but on measurable ecosystem stress that is actionable for platform operators and supervisors. As a result, the evaluation emphasizes both statistical quality (calibration and ranking) and operational utility (lead time, false alarms, and stability).

To ensure replicability, the evaluation pipeline applies identical preprocessing and feature engineering rules in training and testing. Thresholds were selected using a validation segment carved from the training horizon, then frozen for testing to avoid optimistic bias. This design ensures the reported outcomes represent deployable performance rather than retrospective tuning.

Early-Warning Performance Overview

Table 2 summarizes the performance metrics, while figure 4 shows how those metrics manifest over time as a probabilistic risk signal and alerts. Together,

they establish whether the DBN is accurate, early, and operationally stable rather than merely producing a high-scoring classifier.

Table 2 Backtesting Performance Summary

Category	Metric	Result	Operational Meaning
Calibration	Brier Score	0.094	Probabilities remain decision-grade; lower values indicate more reliable risk scores.
Discrimination	ROC-AUC	0.873	The model ranks stressed periods above normal periods across thresholds.
Discrimination (Imbalanced)	PR-AUC	0.417	Precision-recall quality remains strong under rare-event conditions typical of systemic stress.
Early Warning Utility	Average Lead Time (days)	9.6	Alerts arrive early enough to activate liquidity, operational, or governance playbooks.
Operational Burden	False Alarms per 30 days	1.7	Alert volume is manageable and reduces alert fatigue for on-call and risk teams.
Stability	Alert Stability Index	0.931	Alerts do not oscillate excessively; supports consistent escalation handling.
Coverage	Stress Episode Recall	0.842	The system detects most stress episodes while maintaining a controlled false-alarm rate.

The results show a balanced early-warning profile. Calibration (Brier Score 0.094) indicates the DBN’s probabilities are not merely rankings, but usable confidence levels for governance thresholds. This matters in systemic-risk monitoring because decisions (e.g., throttling, liquidity buffers, incident response escalation) require probabilities that are stable enough to justify intervention without frequent reversals.

Operationally, the model achieves meaningful lead time (9.6 days) while controlling false alarms (1.7 per 30 days) and maintaining high alert stability (0.931). This combination is the central requirement for FinTech ecosystems: early signals are valuable only if they arrive early and remain credible long enough for teams to act. The PR-AUC (0.417) further supports robustness under imbalance, reflecting real-world conditions where true systemic stress is infrequent but high impact.




Figure 4 Early-Warning Signal Over Time with Warning/Alarm Thresholds

The time-series view verifies that the DBN’s value is not limited to a single score; it produces a continuous monitoring signal aligned with operational escalation. The warning threshold provides an early “heads-up” phase where teams can increase monitoring, validate liquidity posture, and check infrastructure reliability, while the alarm threshold supports decisive escalation when posterior risk reaches a critical range.

Importantly, the signal exhibits sustained elevation during stressed windows instead of short-lived spikes. This behavior aligns with the stability metric in table 2 and is essential for systemic-risk response, where rapid oscillations cause alert fatigue and undermine trust. The post-stress decay demonstrates that the model can revert toward baseline as evidence normalizes, supporting de-escalation decisions rather than forcing a permanently conservative posture.

Interpretability: What Drives the DBN Risk Posterior

Beyond accuracy, the DBN provides a structured explanation pathway: posterior risk increases occur when evidence becomes jointly consistent with a stressed regime. In practice, the model's risk posterior rises most sharply when operational degradation signals (failure rate and settlement delay) co-occur with funding/liquidity pressure (buffer drawdown and widening spreads). This co-movement is critical because isolated operational incidents are common in large platforms, but systemic fragility is more credibly inferred when operational issues overlap with balance-sheet or funding stress.

The DBN's conditional structure supports controlled attribution: the same transaction failure spike can imply different systemic risk depending on whether it is accompanied by infrastructure incidents or liquidity contraction. This is particularly important for FinTech ecosystems where outages can be technical (localized and reversible) or systemic (triggering run-like dynamics). The model distinguishes these cases probabilistically by conditioning on multiple evidence nodes, reducing the likelihood of overreacting to single-channel noise.

From a governance standpoint, interpretability translates into actionable narratives. When the posterior risk rises, the system can surface the evidence pattern that dominated the update (e.g., settlement tail latency + withdrawal pressure + liquidity drawdown), enabling targeted actions such as capacity stabilization, settlement partner coordination, liquidity buffer activation, or communications to reduce run behavior.

Scenario-Based Stress Testing and Fragility Mapping

Stress testing evaluates whether the DBN reacts sensibly under controlled perturbations and identifies fragility amplifiers. Scenarios were applied by injecting sustained increases into specific evidence channels (e.g., a funding shock that raises spreads, an infrastructure incident wave that increases failure rates, or a liquidity drawdown consistent with partner constraints). The most impactful scenarios are those that simultaneously degrade liquidity and operations because they accelerate propagation through multiple dependency pathways.

A key insight is that scenario results are not symmetric: liquidity deterioration tends to produce persistent posterior elevation because it weakens shock absorption, whereas isolated infrastructure incidents can produce sharp but shorter-lived spikes unless they also trigger user withdrawal pressure. This distinction is operationally valuable: it differentiates "technical crisis management" from "systemic crisis management," the latter requiring broader governance actions such as funding contingency playbooks and counterparty coordination.

These findings also support ecosystem-level risk management. By ranking scenario impacts, stakeholders can prioritize mitigations that reduce systemic

coupling (such as diversifying settlement rails, reducing counterparty concentration, strengthening liquidity buffers, and improving incident containment) rather than focusing solely on point solutions that improve one metric but leave systemic channels intact.

Deployment Readiness and Operational Considerations

For production use, the DBN early-warning system must meet three requirements: low-latency inference, controlled retraining, and auditable alert governance. Inference can run in near real time because filtering updates are lightweight once the structure and parameters are fixed. Retraining is scheduled (e.g., weekly or monthly) to incorporate drift while avoiding uncontrolled shifts in behavior that would invalidate threshold governance and historical comparability.

Operationally, the alert policy should remain explicit and versioned. Threshold changes, binning rules, and node definitions are governance artifacts, not mere configuration. This is critical in systemic-risk monitoring because stakeholder trust depends on consistent semantics: an “alarm” must mean the same risk posture across time unless there is a documented policy update.

Finally, deployment should include evidence snapshots and audit logs at each alert. When risk rises, teams need to see which evidence changed and whether the pattern matches known propagation mechanisms (liquidity-driven, operational-driven, or combined). This ensures the DBN is not only predictive but also operationally usable as a decision-support instrument in complex FinTech ecosystems.

Conclusion

This study developed a Dynamic Bayesian Network–based early-warning framework to detect and monitor systemic risk in FinTech ecosystems by integrating heterogeneous evidence streams (operational reliability, settlement performance, liquidity and funding conditions, infrastructure incidents, and network exposure signals) into a single probabilistic regime model. The methodology emphasizes time-consistent backtesting and deployable inference, enabling the system to update risk beliefs sequentially as new evidence arrives rather than relying on static risk scores. As a result, the DBN functions as both a prediction engine and a structured explanation mechanism that links elevated risk to interpretable patterns of co-occurring stress across ecosystem channels.

Empirically, the results demonstrate that a probabilistic regime signal can deliver actionable early-warning behavior when it is evaluated with operational metrics beyond accuracy alone, including calibration, lead time, false-alarm burden, and alert stability. The time-series early-warning signal shows that risk elevation is most credible and persistent when operational degradation aligns with liquidity or funding stress, reflecting systemic coupling rather than isolated incidents. This supports practical governance by distinguishing technical disruptions that can be contained locally from systemic episodes that require broader interventions such as liquidity buffer activation, counterparty coordination, and risk communications to mitigate run-like dynamics.

In terms of implications, the proposed DBN framework provides a scalable blueprint for risk monitoring across interconnected FinTech infrastructures

where dependencies and uncertainty are intrinsic. For deployment, the study highlights the need for explicit governance of node definitions, discretization rules, threshold policies, and retraining schedules to preserve auditability and stakeholder trust. Future work should prioritize richer network exposure measurement, adaptive structure learning under regime drift, and integration of cost-sensitive decision policies aligned with regulatory and platform-specific risk appetites, ensuring that early-warning outputs remain both statistically reliable and operationally actionable.

Declarations

Author Contributions

Conceptualization: C.Y.P. and S.P.; Methodology: S.P.; Software: C.Y.P.; Validation: C.Y.P. and S.P.; Formal Analysis: C.Y.P. and S.P.; Investigation: C.Y.P.; Resources: S.P.; Data Curation: S.P.; Writing Original Draft Preparation: C.Y.P. and S.P.; Writing Review and Editing: S.P. and C.Y.P.; Visualization: C.Y.P.; All authors have read and agreed to the published version of the manuscript.

Data Availability Statement

The data presented in this study are available on request from the corresponding author.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] A. Geebren, A. Jabbar, and M. Luo, "Examining the role of consumer satisfaction within mobile eco-systems: Evidence from mobile banking services," *Comput. Hum. Behav.*, vol. 114, no. January, p. 106584, pp. 1-12, 2021, doi: 10.1016/j.chb.2020.106584.
- [2] S. K. Sharma and M. Sharma, "Examining the role of trust and quality dimensions in the actual usage of mobile banking services: An empirical investigation," *Int. J. Inf. Manage.*, vol. 44, no. February, pp. 65–75, 2019, doi: 10.1016/j.ijinfomgt.2018.09.013.
- [3] A. M. Baabdullah, A. A. Alalwan, N. P. Rana, H. Kizgin, and P. Patil, "Consumer use of mobile banking (M-Banking) in Saudi Arabia: Towards an integrated model," *Int. J. Inf. Manage.*, vol. 44, no. February, pp. 38–52, 2019, doi:

10.1016/j.ijinfomgt.2018.09.002.

- [4] A. Ruangkanjanases, E. M. A. Qhal, K. M. Alfawaz, and T. Hariguna, "Examining the Antecedents of Blockchain Usage Intention: An Integrated Research Framework," *Sustainability*, vol. 15, no. 4, pp. 1-16, 2023, doi: 10.3390/su15043500.
- [5] D. G. Lisanawati and M. J. E. Kehinde, "When Technology Meets Money Laundering, What Should Law Do? New Products and Payment Systems and Cross Border Courier," *IJIS Int. J. Inform. Inf. Syst.*, vol. 5, no. 3, pp. 142–149, 2022.
- [6] D. Dellermann, N. Lipusch, P. Ebel, and J. M. Leimeister, "Design principles for a hybrid intelligence decision support system for business model validation," *Electron. Mark.*, vol. 29, no. 3, pp. 423–441, 2019, doi: 10.1007/s12525-018-0309-2.
- [7] S. Majumdar and V. Pujari, "Exploring usage of mobile banking apps in the UAE: A categorical regression analysis," *J. Financ. Serv. Mark.*, vol. 27, no. August, pp. 177-189, 2021, doi: 10.1057/s41264-021-00112-1.
- [8] Q. Zhou *et al.*, "A study on factors affecting service quality and loyalty intention in mobile banking," *J. Retail. Consum. Serv.*, vol. 60, no. May, p. 102424, pp. 1-11, 2021, doi: 10.1016/j.jretconser.2020.102424.
- [9] L. Abubakar and T. Handayani, "Financial Technology: Legal Challenges for Indonesia Financial Sector," *IOP Conf. Ser. Earth Environ. Sci.*, vol. 175, no. 1, pp. 1-6, 2018, doi: 10.1088/1755-1315/175/1/012204.
- [10] N. Q. Huy, L. P. Nga, and P. T. Tam, "Applied Simulation Modeling for Promoting Policy Recommendations for Microfinance Activity Development: A Case Study in Vietnam," vol. 4, no. 4, pp. 333–345, 2023.
- [11] Y. W. Chang and J. Chen, "What motivates customers to shop in smart shops? The impacts of smart technology and technology readiness," *J. Retail. Consum. Serv.*, vol. 58, no. January, p. 102325, pp. 1-10, 2021, doi: 10.1016/j.jretconser.2020.102325.
- [12] A. A. Diniyya, M. Aulia, and R. Wahyudi, "Financial Technology Regulation in Malaysia and Indonesia: A Comparative Study," *Ihtifaz J. Islam. Econ. Financ. Bank.*, vol. 3, no. 2, p. 67, pp. 69-87, 2021, doi: 10.12928/ijiefb.v3i2.2703.
- [13] M. R. Elrashidy and H. Mansour, "Microservices Architecture in FinTech: A Case Study on Scalable Loan Processing with Classical Design Patterns," in *Proc. IEEE Int. Midwest Symp. Circuits Syst. (IMSAC)*, vol. 2025, no. July, pp. 382–387, 2025, doi: 10.1109/IMSAC65733.2025.11167186.
- [14] H. Fang, C.-P. Chung, Y.-C. Lu, Y.-H. Lee, and W.-H. Wang, "The impacts of investors' sentiments on stock returns using FinTech approaches," *Int. Rev. Financ. Anal.*, vol. 77, no. October, p. 101858, pp. 1-15, 2021, doi: 10.1016/j.irfa.2021.101858.
- [15] F. A. Hudaefi, M. K. Hassan, and M. Abduh, "Exploring the development of Islamic FinTech ecosystem in Indonesia: A text analytics," *Qual. Res. Financ. Mark.*, ahead-of-print, 2023.
- [16] J. S. Wang, "Reconfigure and evaluate consumer satisfaction for Open API in advancing FinTech," *J. King Saud Univ. Comput. Inf. Sci.*, vol. 35, no. 9, pp. 1-14, 2023, doi: 10.1016/j.jksuci.2023.101738.

- [17] A. Emuron, A. D. K. Kenfack, and M. Msimanga, "FinTech predictive modelling using machine learning: An empirical investigation of South Africa," *Afr. J. Sci. Technol. Innov. Dev.*, vol. 17, no. 6, pp. 889–899, 2025, doi: 10.1080/20421338.2025.2552006.
- [18] S. Amrie, S. Kurniawan, J. H. Windiatmaja, and Y. Ruldeviyani, "Analysis of Google Play Store's Sentiment Review on Indonesia's P2P FinTech Platform," in *2022 IEEE Delhi Section Conference (DELCON)*, pp. 1–5, 2022.
- [19] S.-C. Huang, C.-F. Wu, C.-C. Chiou, and M.-C. Lin, "Intelligent FinTech Data Mining by Advanced Deep Learning Approaches," *Comput. Econ.*, vol. 59, no. 4, pp. 1407–1422, 2022, doi: 10.1007/s10614-021-10118-5.
- [20] N. A. Almasria, Z. Alhatabat, D. Ershaid, A. Ibrahim, and S. Ahmed, "The Mediating Impact of Organizational Innovation on the Relationship Between FinTech Innovations and Sustainability Performance," *Sustainability*, vol. 16, no. 22, pp. 1–19, 2024, doi: 10.3390/su162210044.
- [21] J. M. Puaschunder, "The Future of Green FinTech Under Climate Change: The Relation of Cryptocurrencies and Sustainability," in *Climate Change Management*, vol. 2025, no. June, pp. 147–163, 2025, doi: 10.1007/978-3-031-85217-6_9.
- [22] H. Taherdoost, "Blockchain: A catalyst in FinTech future revolution," *Futur. Technol.*, vol. 2, no. 2, pp. 25–31, 2023, doi: 10.55670/fpl.futech.2.2.3.
- [23] S. Jain, S. Chitkara, N. Sharma, and S. Gupta, "Collaborative innovation in FinTech—Engaging customers in co-creation," in *Advances in FinTech and Digital Financial Services*, pp. 97–122, 2025, doi: 10.4018/979-8-3693-9944-6.ch007.
- [24] Z. Guo, Y. Wei, S. Fan, and M. Shahbaz, "Is FinTech making world more 'inclusive'? Evidence from inclusive growth story of China," *Technol. Forecast. Soc. Change*, vol. 209, no. December, p. 123780, 2024, doi: 10.1016/j.techfore.2024.123780.
- [25] N. Kaur, S. L. Sahdev, M. Chhabra, and S. M. Agarwal, "FinTech Evolution to Revolution in India—From Minicorns to Soonicorns to Unicorns," in *2021 9th Int. Conf. Reliab. Infocom Technol. Optim. (ICRITO)*, pp. 1–6, 2021.
- [26] S. Oh, G. Chung, and K. Cho, "New Sustainable FinTech Business Models Created by Open Application Programming Interface Technology: A Case Study of Korea's Open Banking Application Programming Interface Platform," *Sustainability*, vol. 16, no. 16, pp. 1–20, 2024, doi: 10.3390/su16167187.
- [27] H. H. Khan, S. Khan, and A. Ghafoor, "FinTech adoption, the regulatory environment and bank stability: An empirical investigation from GCC economies," *Borsa Istanbul Rev.*, vol. 23, no. 6, pp. 1263–1281, 2023, doi: 10.1016/j.bir.2023.10.010.
- [28] S. H. Utami, A. A. Purnama, and A. N. Hidayanto, "FinTech Lending in Indonesia: A Sentiment Analysis, Topic Modelling, and Social Network Analysis using Twitter Data," *Int. J. Appl. Eng. Technol.*, vol. 4, no. 1, 2022.
- [29] N. A. Khamitdkhanovich, S. Packiaraji, M. Tripathi, and K. K. Badriddinovich, "Unlocking the power of FinTech: How embedded finance is shaping the future of commerce," in *AIP Conf. Proc.*, vol. 3306, no. 1, pp. 1–8, 2025, doi: 10.1063/5.0277719.