LUISS

Degree Program in Management and Computer Science

Course of Social Network Analysis [MAT/09]

Systemic Risk and the Propagation of Financial Distress: A Network-Based Analysis of Traditional and FinTech Ecosystems

Prof. Xavier Mathieu Raymond Venel

SUPERVISOR

Marta Torella [ID: 284091]

CANDIDATE

Academic Year 2024/2025

Table of Contents

Chapter 1: Introduction

Understanding how financial distress spreads across interconnected institutions has become increasingly important in today's financial landscape. Traditional risk models, which focus primarily on the individual solvency of financial institutions, often fail to account for the complex network of dependencies that define modern financial systems. Moreover, as financial markets continue to evolve, the emergence of new actors such as fintech firms introduces additional complexity, further complicating efforts to assess systemic risk using traditional methods.

This thesis tries to address the challenge of modelling systemic financial risk through a network-based approach, with particular attention to the structural differences between traditional banking institutions and fintech players. Unlike traditional models that treat institutions as isolated entities, a network approach emphasises how the pattern of connections between them shapes the likelihood and extent of financial contagion. Indeed, recent literature suggests that the structure of these connections can influence whether financial shocks are amplified or contained.

To capture these dynamics, this study proposes a novel agent-based simulation framework based on a modified epidemiological model, the SIIS (Susceptible-Infected$_1$-Infected$_2$-Susceptible) contagion model. This approach introduces different levels of infection to mimic the varying degrees of financial distress across three distinct types of financial networks: a Traditional network, a Fintech network, and a Joint network combining both sectors. By comparing how these systems respond to different policies, the research explores how institutional structures and adaptive responses affect the spread of systemic risk.

The thesis is organised as follows: Chapter 2 reviews the theory behind modelling systemic risk in financial networks, comparing traditional and fintech ecosystems. Chapter 3 outlines the methodology for constructing the simulated financial networks, grounded in empirical features observed in real-world interbank markets. Chapter 4 presents the SIIS contagion model, including both its agent-based and mean-field formulations, and describes the simulation scenarios used to evaluate policy responses. Chapter 5 discusses the findings across various contagion scenarios. Finally, Chapter 6 concludes by summarising key insights, acknowledging model limitations, and proposing directions for future research.

Chapter 2: Theoretical Background

This chapter begins by examining systemic risk as a network-driven phenomenon, referencing past literature which proves the important role played by interinstitutional ties in shaping financial stability. It then turns to a comparative approach, focusing on how the structural differences between traditional financial institutions and fintech firms shape their respective vulnerabilities and roles in the propagation of financial distress.

2.1. Modelling Systemic Risk with Social Network Analysis

Systemic risk refers to the potential for the distress or failure of a single financial institution to spread through the entire system, thereby compromising the system's overall financial stability. This risk is not merely defined as the sum of individual weaknesses; rather, it emerges from the complex network of interdependencies between institutions. Most importantly, systemic risk is economically significant because of its ability to interfere with financial intermediation, affect credit markets, and negatively impact the broader economy. In this regard, the structure of financial networks plays a crucial role in determining whether financial distress is amplified or contained.

The first key insight from the literature is that the same network structure can either help stabilise the financial system or make it more vulnerable, depending on the magnitude of shocks. Acemoglu et al. (2015)[1] show that when shocks are relatively small, a more connected network helps spread the risk, making the system more resilient. But when the shocks are large, those same connections can work against the system, allowing distress to quickly spread through many pathways and trigger widespread failures. This structural duality, often described as "robust-yet-fragile", shows the extent to which the configuration of financial networks can influence the scale and impact of a crisis.

Moreover, this dual nature of fragility in financial networks is also evident in the work of Li et al. (2020)[2], which explores how financial distress spreads between fintech firms and traditional institutions under different market conditions. By analysing U.S. financial data through a network-based approach, the study shows that the intensity and direction of risk transmission can vary widely depending on the specific sector and the current state of the market. Notably, the findings reveal that both fintech and traditional firms are not only

[1] Acemoglu, Ozdaglar, and Tahbaz-Salehi, 'Systemic Risk and Stability in Financial Networks'.
[2] Li et al., 'Risk Spillovers between FinTech and Traditional Financial Institutions'.

affected by systemic shocks but also contribute to spreading them in different ways, especially in times of high volatility. These results demonstrate the importance of treating different types of institutions separately in contagion models, and they prove how variations in network structure and market conditions shape the way financial risk moves through the system.

Finally, the work of Bianchi et al. (2023)[3] stresses that financial networks are not just descriptive tools, but they play a key role in assessing how crises unfold. Their review shows that the specific placement of an institution in the network (whether it is a central hub, a connecting bridge, or a more isolated node) can have a major impact on how vulnerable it is to contagion and how much it contributes to spreading distress. Considering this, the way financial institutions are connected matters just as much as their individual financial resilience.

Indeed, social network analysis has proven useful in revealing the network-based drivers of systemic risk. Instead of looking at institutions in isolation, it shifts the focus to the web of relationships that connect them. This approach shows how factors like a firm's position in the network, its level of interconnectedness, and the overall topology of the network can all influence whether a firm is more likely to endure shocks or contribute to spreading them.

2.2. Traditional vs FinTech Financial Ecosystems

Over the past decade, the financial sector has undergone major changes, bringing in new players, new structures, and new types of connections into an already complex system. One of the most notable shifts has been the rise of fintech firms. Once operating on the margins of traditional banking, these companies are now taking on a much more central role in providing essential financial services. As a result, this shift in the industry's structure has significant implications for the way risk propagates through financial networks.

Traditional financial institutions, especially large banks, are often part of centralised, hierarchical networks. Their connections are shaped by regulations, long-standing relationships between banks, and shared financial infrastructure. According to Suprun et al. (2020)[4], these institutions usually maintain strong and stable connections with one another,

[3] Bianchi et al., 'Social Networks Analysis in Accounting and Finance'.
[4] Suprun, Petrishina, and Vasylchuk, 'Competition and Cooperation between Fintech Companies and Traditional Financial Institutions'.

which can help keep the system resilient in normal times. But this stability can be misleading: if instability begins or affects a major player in the network, the impact can spread quickly, revealing the hidden risks that come with having so much connectivity concentrated in a few key nodes.

On the other hand, fintech firms are structurally different. As Siddiqui and Rivera (2022)[5] describe, the fintech ecosystem appears decentralised, fast-moving, and composed of a wide range of players working at the intersection of finance and digital innovation. These firms tend to have more flexible connections and can quickly form or end partnerships. From a network point of view, fintech ecosystems resemble small-world networks, with tighter clusters, shorter paths between nodes, and frequent rewiring of connections. While these features encourage innovation and agility, they can also create tightly interwoven local clusters that allow risk to spread quickly, especially in times of market stress.

Moreover, as fintech becomes more integrated into the core of the financial system, it also creates new paths for risk to spread. Harsono and Suprapti (2024)[6] point out that while fintech firms bring clear benefits (such as better access to services, increased efficiency, and improved customer experience), they also come with potential vulnerabilities. These include growing reliance on technology, gaps in regulation, and operational risks. As a result, the traditional ways of assessing systemic risk are no longer enough.

Ultimately, the structure of the ecosystem in which financial institutions operate has an immediate impact on the system as a whole. Due to their function and central roles, traditional firms often help stabilise the system, until a failure on their part becomes a significant point of collapse. At the same time, fintech firms bring speed and adaptability, but they also introduce some degree of structural volatility that has the potential to escalate minor disruptions into major issues. Thus, systemic risk becomes a shared concern for the entire network as the distinction between the traditional and fintech sectors becomes blurred.

2.3. Toward a Network-Based Analysis of Financial Distress

The previous sections have highlighted two key ideas at the heart of this study. First, systemic financial risk is not just about individual institutions; it is rather a network phenomenon shaped by how financial entities are connected and interact. Second, traditional banks and

[5] Siddiqui and Rivera, 'FinTech and FinTech Ecosystem'.
[6] Harsono and Suprapti, 'The Role of Fintech in Transforming Traditional Financial Services'.

fintech firms operate within distinct financial ecosystems, each with its own structure and way of transmitting risk.

To move from theory to analysis, the next chapters present a network-based simulation framework which reflects the structural features of real-world financial systems. This model serves as the foundation for exploring how financial contagion propagates, and how different configurations of interconnectivity influence the dynamics of systemic risk.

Chapter 3: Network Construction and Data Description

This chapter outlines the construction of financial network models and the data used for simulation. It details the modelling techniques, structural assumptions, and sector-specific variants used to represent Traditional, Fintech, and Joint financial systems.

3.1. Modelling Financial Networks

Building a realistic and reliable financial network is central to understanding how financial distress spreads between interconnected institutions. Based on empirical research on interbank systems and complex network theory, this framework attempts to reflect the structural differences between traditional and fintech sectors, as well as the blended interactions that occur in more integrated financial environments.

3.1.1. Empirical Basis and Structural Characteristics

Empirical studies of interbank networks across multiple national contexts (including Japan, Austria, the United Kingdom, Germany, Hungary, the United States, Italy, and Brazil) demonstrate that financial systems can be accurately characterised as complex networks (Li, He, & Zhuang, 2010)[7]. These systems consistently exhibit three interrelated structural properties: scale-free topology, small-world connectivity, and preferential attachment with embedded randomness.

Scale-Free Structure

In a scale-free network, the distribution of links across nodes follows a power law, where a small number of nodes (or institutions) maintain a disproportionately high number of connections, while most nodes have only a few. These highly connected nodes, known as hubs, are crucial for holding the system together. In financial networks, hubs often represent major players like central banks or other systemically important institutions. Because of their central roles, these entities can act as stabilisers in normal times, but they also pose serious risks: if one of them experiences distress, the effects can quickly spread through the network. This kind of hierarchical structure mirrors what we see in the real world, where institutions

[7] Li, He, and Zhuang, 'A Network Model of the Interbank Market'.

with more resources, stronger reputations, or regulatory importance tend to attract more connections and influence[8].

Small-World Property

Even though financial networks can include thousands of institutions, they often show what's known as the small-world property, meaning that most institutions are only a few steps away from one another, and many are grouped into tight clusters. This structure reflects the types of connected relationships present in actual banking groups or industry alliances. In practice, it allows information and liquidity to move quickly across the system, which helps with coordination and day-to-day operations. However, it also implies that issues can spread just as quickly, leaving little time to respond once a crisis begins[9].

Randomness and Preferential Attachment

Two mechanisms control the creation of links in financial networks: preferential attachment, which makes it statistically more likely for new nodes to connect to nodes that are already well-connected, and a degree of randomness, which represents random interactions or opportunistic trading. However, this behaviour is not strictly arbitrary: in order to obtain stability and access to resources, smaller or newer institutions frequently look to form alliances with more established players. Over time, this pattern reinforces the dominance of key players and creates a "hub-and-spoke" structure in the network. By concentrating risk around a small number of crucial nodes, this structure increases systemic fragility, even though it might render the network more resilient under normal conditions[10].

These three structural features (scale-free structure, small-world property, and preferential attachment) collectively define the topology of the financial systems modelled in this study. They form the empirical and theoretical basis for building simulated networks that can replicate the dynamics of actual financial contagion, as explored in the following sections.

3.1.2. Modelling with Barabási-Albert and Watts-Strogatz Networks

To realistically simulate the topology of financial systems, this study employs two well-established network generation models: the Barabási-Albert (BA) model and the Watts-

[8] Barabási and Albert, 'Emergence of Scaling in Random Networks.'
[9] Duncan J. Watts and Strogatz Steven H. Strogatz, 'Collective Dynamics of Small-World Networks'
[10] Albert and Barabási, 'Statistical Mechanics of Complex Networks'.

Strogatz (WS) model. Each provides a complementary set of structural properties that together help replicate the observed complexity of real-world financial networks.

Barabási-Albert model

The Barabási-Albert (BA) model is a network generation mechanism that captures two fundamental properties observed in real-world systems: growth and preferential attachment. In this model, a network evolves over time by continuously adding new nodes, each of which forms links preferentially to existing nodes that already have a higher degree of connectivity. This "rich-get-richer" dynamic leads to the emergence of a scale-free topology, where the degree distribution follows a power law. That is, most nodes have few connections, while a small number of hubs accumulate disproportionately many links. In the context of financial systems, these hubs often correspond to systemically important institutions, such as large commercial banks and major payment platforms. Their central role makes them both stabilisers under normal conditions and critical risk factors under distress. Using the BA model to construct the network backbone ensures that the simulation faithfully reproduces this asymmetry in systemic influence, that is, the unequal distribution of connections within the network.

Watts-Strogatz model

The Watts-Strogatz (WS) model provides a generative mechanism for constructing networks that simultaneously exhibit high clustering and short average path lengths, two defining features of so-called small-world networks. Starting from a regular lattice where each node is connected to its nearest neighbours, the model randomly rewires a fraction of the edges with a certain probability. This rewiring process maintains the local connectivity patterns typical of regular graphs while introducing shortcuts that drastically reduce the typical distance between nodes. In the context of financial systems, small-world properties are particularly relevant for modelling flexible and interconnected ecosystems, such as those emerging in fintech sectors. High clustering reflects the tendency of firms to form dense ties, while short paths between institutions facilitate rapid information diffusion and potential contagion. By incorporating the WS model into the network construction process, the simulation framework captures the coexistence of local resilience and global vulnerability.

Hence, to build a network that mirrors both the scale-free and small-world features found in real-world financial systems, this thesis uses a hybrid modelling strategy. The core structure

is created using the Barabási-Albert model, which helps capture the formation of hubs and variation in connectivity across institutions. Then, to introduce more clustering and reduce the average distance between nodes (without disrupting the overall power law structure), some edges are randomly rewired based on the Watts-Strogatz model. This layered approach captures the dual nature of today's financial systems: the dominance of a few key institutions and the evolving connections between smaller players (especially where traditional banks and fintech platforms intersect through contractual ties). Together, the BA and WS models offer a realistic framework for simulating different network types (Traditional, Fintech, and Joint) and for testing how each responds to financial contagion in the scenarios presented and analysed in Chapters 4 and 5.

3.1.3. Node and Edge Definitions, Weights, and Directionality

In this simulation framework, financial institutions are modelled as nodes within a directed, weighted graph, where edges represent financial exposures between them and the edge weights quantify the magnitude of these exposures.

In particular, each node in the network represents a financial entity, either a traditional institution or a fintech firm, depending on the simulation scenario. For the purposes of the simulation, each node is assigned a unique identifier, a list of directly connected institutions to represent its financial relationships and a counter tracking consecutive periods in financial distress.

As for the edges, since they represent the channels through which financial distress propagates, their weights are determined according to a simple risk-sharing assumption: each outgoing edge distributes risk equally among the node's neighbours. Formally, the weight w_{ij} of an edge from node i to node j is defined as:

$$w_{ij} = \frac{1}{\text{out-degree}(i)}$$

This approach bases weights solely on the number of outgoing edges, rather than the total degree, reflecting the intuition that a node spreads its obligations across all its outgoing links. In financial networks, an outgoing edge from node i to node j signifies a financial exposure of j to i (for instance, a loan or a liquidity provision that j depends on). If node i were to default, the loss would propagate outward to all connected institutions according to the distributed weights. Importantly, this structure captures the directionality of exposures: a directed edge

from node i to node j indicates that cash flows from i to j, implying that j bears potential losses if i defaults. This design mirrors real-world financial contagion dynamics, where shocks travel downstream through credit or liquidity. Assigning weights based only on outgoing edges emphasises that institutions transmit risk outward, while incoming edges reflect the risks they absorb from others.

Using degree-based weighting offers a straightforward way to model how institutions share risk, but it doesn't take into account differences in their size. Since larger organisations are more likely to have an impact on a larger range of counterparties, taking into account the originating institution's size (such as its total assets or capital buffer) could increase realism. However, including this kind of detail would make the model more complex and would require reliable data that may not always be available or easy to compare. Still, the current approach effectively shows how institutions spread risk through their connections while also capturing how they absorb risk from others. In this way, the network doesn't just show who's connected to whom, but it also reflects the direction in which financial risk moves and how it can lead to systemic contagion.

3.2. Data Generation and Parameter Definition

In this thesis, three distinct types of financial networks are constructed and analysed: a Traditional network, a Fintech network, and a Joint network. To ensure comparability, the number of nodes across all three configurations is maintained constant. However, two key structural parameters are varied to reflect the differing patterns of connectivity and risk-sharing dynamics empirically observed. In particular, we adjust the number of connections that new nodes establish upon entry and the probability of edge rewiring for each network.

Number of nodes

The network size in the simulations is fixed at N = 1,000 nodes across all configurations. This choice is driven by precedents in the financial contagion literature, seeking a balance between realistic structure and manageable complexity. Gai and Kapadia (2010)[11] explicitly use networks of 1,000 institutions to represent national banking systems and study systemic fragility under different shock scenarios. Similarly, Amerongen et al. (2019)[12] generate a

[11] Gai and Kapadia, 'Contagion In Financial Networks'.
[12] Amerongen et al., 'Agent-Based Models for Assessing the Risk of Default Propagation in Interconnected Sectorial Financial Networks'.

synthetic financial client-supplier network with 10^3 nodes based on transaction-level data from 140,000 Spanish firms, later using subsets of this size for running simulations. Fixing N = 1,000 allows us to realistically capture sectoral heterogeneity while ensuring convergence of simulations across scenarios.

Number of entry connections

The number of entry connections *m* that a new node establishes upon joining the network influences the system's resilience, clustering, and vulnerability to contagion. In this thesis, different values of this parameter for each network type are assigned following empirical studies on real-world interbank networks. According to data compiled by the European Systemic Risk Board (2017)[13], which surveyed 13 national financial systems, the average node degree typically ranges between 3 and 10, with median degrees around 3 to 8, reflecting sparse but considerable interconnectedness across institutions. These findings provide a baseline for constructing networks that are neither overly dense nor unrealistically disconnected. For the Traditional banking network, the number of entry connections *m* is set in order to achieve an average degree $<k> = 8$, aligning with both ESRB's observations and prior modelling work that highlighted the hierarchical structure of traditional interbank networks (Suprun et al., 2020; Gai and Kapadia, 2010)[14,15]. Conversely, for the Fintech network, a lower entry degree *m* is assigned in order to achieve an average degree $<k> = 4$, reflecting the more decentralised character of fintech ecosystems, as described by Siddiqui and Rivera (2022)[16]. Finally, the Joint network, integrating both types of institutions, achieves an intermediate average degree $<k> = 7$ to mimic hybrid sector structures.

Rewiring probability

The rewiring probability (which regulates the rewiring mechanism during the Watts–Strogatz phase) governs the extent to which the network deviates from a regular structure by randomly reassigning connections between nodes and introducing shortcuts that significantly reduce the average path length while maintaining a high level of clustering. As a result, the network acquires small-world properties, characterised by a combination of local cohesion and global

[13] European Systemic Risk Board, "*How Does Risk Flow in the Credit Default Swap Market?*"
[14] Suprun, Petrishina, and Vasylchuk, 'Competition and Cooperation between Fintech Companies and Traditional Financial Institutions'.
[15] Gai and Kapadia, 'Contagion In Financial Networks'.
[16] Siddiqui and Rivera, 'FinTech and FinTech Ecosystem'.

reachability. In this thesis, different rewiring probabilities are assigned across the three network types to reflect their distinct structural dynamics. For the Traditional banking network, a low rewiring probability is used, consistent with the relatively rigid and hierarchical structures observed in conventional interbank systems, while the Fintech network is assigned a higher rewiring probability, capturing the decentralised and dynamic connections of fintech ecosystems. Finally, for the Joint network, an intermediate rewiring probability is selected to represent a blended topology, where traditional stability and fintech dynamism coexist.

3.3. Network Variants and Sectoral Representations

Having established the general framework and the data generation process (including the parameters used to construct the networks), we now turn to the implementation of the different network variants. This section outlines how each network is generated in practice. Additionally, a set of structural metrics is introduced to assess the networks, highlighting how their properties realistically capture the key differences between traditional and fintech financial ecosystems.

3.3.1. Traditional vs FinTech Financial Network

As already discussed in the previous paragraphs of this chapter, the Traditional and Fintech financial networks are constructed following a structured procedure aimed at replicating the key topological features observed in real-world financial systems. Following the already presented method in Section 3.1.2., the network generation process is based on the Barabási-Albert (BA) and Watts-Strogatz (WS) models, and it is formalised through the pseudo-code outlined below for clarity and reproducibility.

Step 1: Generate Base Graph (Scale-Free Structure)
CREATE scale-free graph using Barabási–Albert model with (*num_nodes, num_connections_entry_banks*)
CONVERT graph to directed graph (*DiGraph*)

Step 2: Assign Weights and Attributes to Nodes
FOR each node in graph:
    GET list of neighbouring nodes
    SET num_connections = number of neighbours

    IF num_connections > 0:
        SET weight = 1 / *num_connections*
    ELSE:

SET weight = 0

ADD attributes to node
ADD weighted directed edges to graph

Step 3: Rewire Edges for Clustering (Watts-Strogatz Style)
FOR each node in graph:
GET list of outgoing edges (successors)

IF node has at least 2 neighbours:
FOR each pair of neighbours:
WITH probability *rewire_prob*:
IF edge does not exist between neighbours:
ADD directed edge between them with weight = 1 / *num_neighbours*

Therefore, while the basic construction methodology is shared by both networks, the intrinsic topological characteristics differ significantly between the traditional financial system and the fintech ecosystem thanks to the use of distinct parameter settings (as discussed in Section 3.2.).

These differences are examined through a range of structural metrics and centrality measures, which are presented below, with their corresponding values reported in Table 1. Thus, the goal is to assess whether the networks generated for this study effectively replicate the characteristics observed in real-world financial systems.

Distance Metrics

The *average shortest path length* (defined as the average number of steps along the shortest paths between all pairs of nodes in a network) in the Traditional banking network is lower, meaning that institutions can generally reach each other through fewer intermediaries. This reflects the presence of central hubs in traditional systems, enabling rapid transmission of financial flows or contagion across the network. Conversely, the Fintech network displays a higher average shortest path length, indicating a more fragmented structure where institutions may require more steps to connect, consistent with a decentralised topology.
The *diameter* (measured as the maximum shortest path between any two nodes) follows the same pattern: it is smaller in the Traditional network and larger in the Fintech network. This confirms that traditional systems, despite being less locally clustered, achieve faster global connectivity through centralised hubs, while fintech systems exhibit slower cross-network reachability.

Structural Properties

The Fintech network exhibits a higher *average clustering coefficient* (measured as the tendency of a node's neighbours to also be connected to each other) compared to the Traditional network. This indicates that institutions in the fintech sector tend to form densely interconnected groups, a feature driven by the collaborative nature of fintech firms, which often establish flexible partnerships. In contrast, the Traditional banking network displays a lower clustering coefficient, reflecting its more centralised and hierarchical organization around a few dominant hubs.
Similarly*, transitivity* (defined as the likelihood that two institutions connected to a common third party are themselves directly connected) is higher in the Fintech network. This outcome aligns with the fintech sector's decentralised configuration, where firms frequently engage in mutual partnerships, fostering strong local interconnectivity. The Traditional network, on the other hand, shows lower transitivity, consistent with the "hub-and-spoke" system where smaller institutions are typically connected through major banks rather than directly to each other.

Centrality Measures

The Traditional network achieves a higher *maximum closeness centrality* (which measures how easily a node can reach all other nodes in the network based on the inverse of the average shortest path length from that node to all others), suggesting that in the traditional banking system, key institutions are able to reach other nodes more efficiently across the network, thanks to the centralised hub structure that reduces the number of intermediaries needed. On the other hand, the Fintech network, despite its decentralised design, shows lower closeness centrality, indicating that, on average, institutions are relatively less efficient in reaching the broader system, likely due to the fragmented nature of fintech interconnections.

Thus, while the fintech sector maintains a more decentralised and locally clustered structure, the traditional banking sector facilitates faster global connectivity through central hubs. Overall, the Fintech network exhibits features of a modular system, characterised by high clustering, high transitivity, longer average path lengths, a larger diameter, and a more decentralised distribution of centrality. This architecture fosters innovation, flexibility, and local resilience, but at the cost of slower global reachability and increased vulnerability to fragmentation between clusters. In contrast, the Traditional banking network presents

characteristics typical of a scale-free architecture, with lower clustering and transitivity, shorter path lengths, and a smaller diameter. Its organisation around centralised hubs enables more efficient system-wide communication and faster global contagion spread, but also concentrates systemic risk within a few dominant institutions.

Accordingly, the metrics obtained from the generated networks are consistent with the structural patterns observed in the real-world financial systems discussed in Sections 3.1.1. and 3.2.

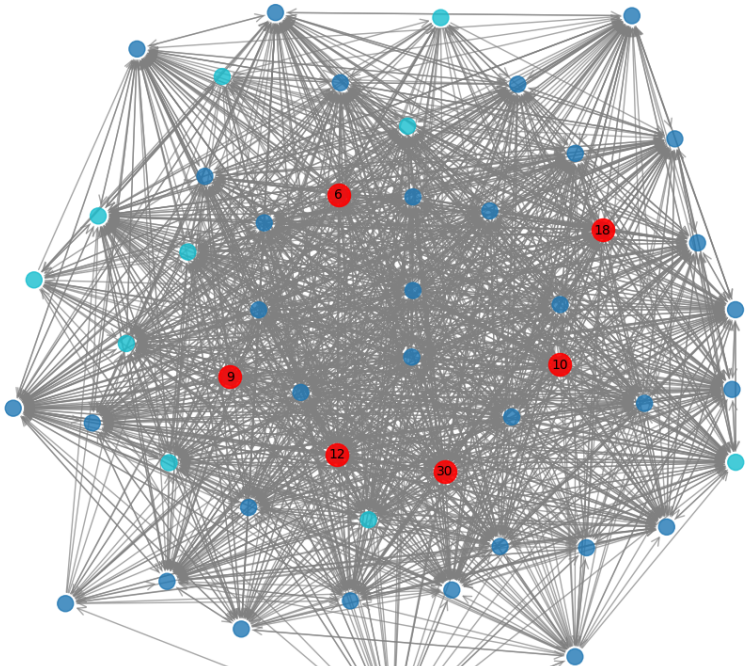

Figure 1 Sampled Traditional Financial Network with Highlighted Hubs and Clusters

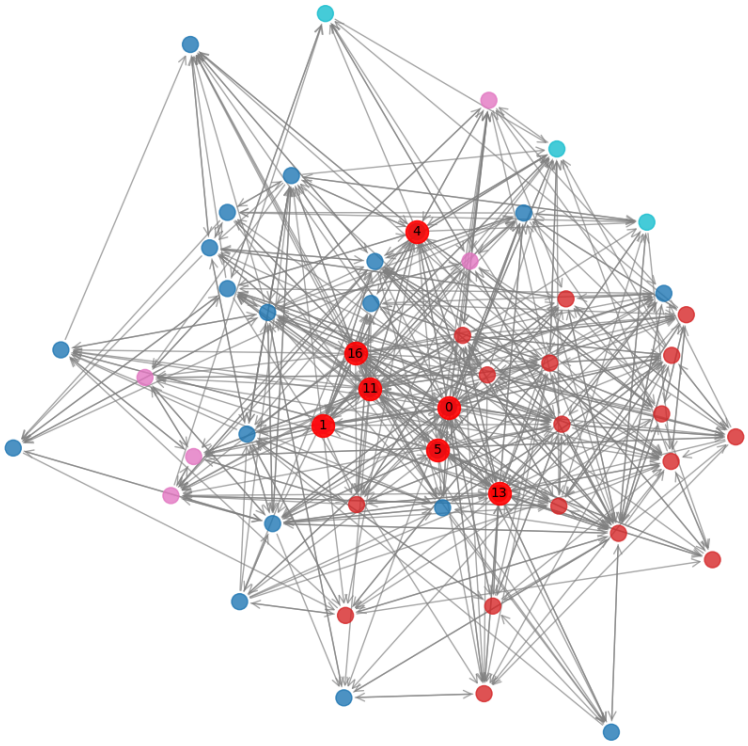

Figure 2 Sampled Fintech Financial Network with Highlighted Hubs and Clusters

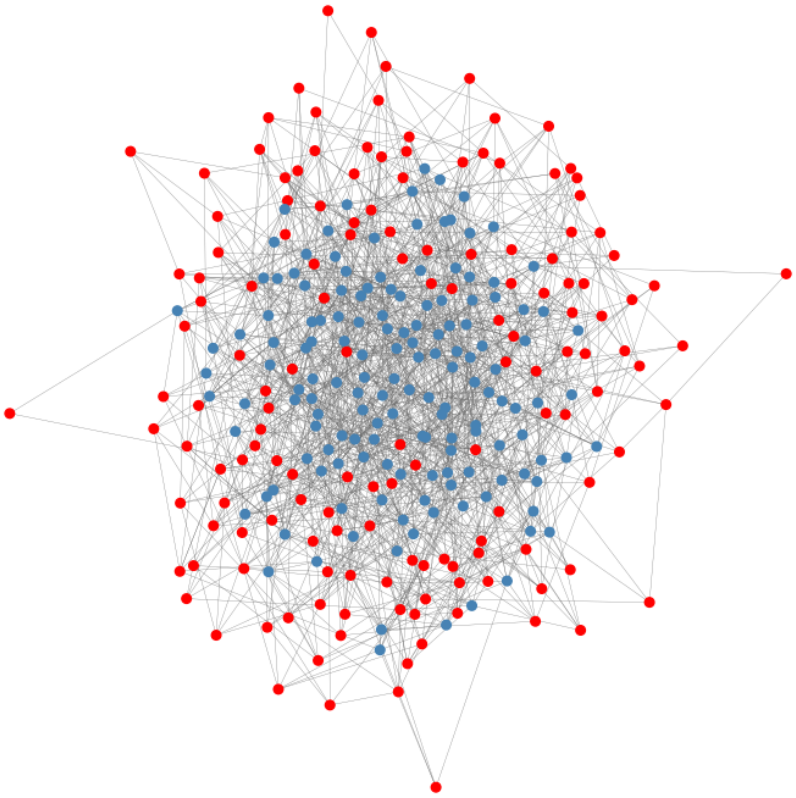

Figure 3 Sampled Joint Financial Network

3.3.2. Joint Financial Network

The Joint financial network is designed to capture the interconnected structure of modern financial systems, where traditional institutions and fintech firms coexist and interact within a common environment. The construction methodology follows the general procedure outlined previously, combining the Barabási-Albert (BA) growth process with the Watts-Strogatz (WS) rewiring procedure, as discussed in Section 3.1.2. and in Section 3.3.1.

However, to the pseudo-code presented in Section 3.3.1, a fourth step is added to explicitly distinguish between traditional and fintech entities inside the Joint network:

Step 4: Assign Fintech (F) and Traditional (T) Labels
*SET n_fintech = fintech_fraction * num_nodes*

SORT nodes by degree in ascending order
SELECT first *n_fintech* nodes as fintech firms
FOR each node in selected fintech firms:
    SET node['type'] = 'F'
FOR each remaining node:
    SET node['type'] = 'T'

Thus, following the generation of the base graph and the assignment of weights and node attributes (Steps 1 to 3), nodes are partitioned into traditional and fintech institutions based on their degree centrality. A fraction of nodes, corresponding to 40% of the total, is designated as fintech firms (represented with the colour red in Figure 3) by assigning this role more often to nodes with lower degrees. This proportion reflects the increasing presence of fintech institutions in the global financial landscape, which, while growing rapidly, still represent a minority compared to traditional players. Moreover, this approach mirrors the empirical observation that fintech firms typically maintain peripheral positions compared to established banks, which often occupy highly central positions in the network.

Structurally, the Joint network exhibits intermediate topological features between those observed in the pure Traditional and Fintech systems, as shown by the values reported in Table 1. While the presence of highly connected traditional hubs guarantees wider reachability, the inclusion of decentralised fintech nodes promotes localised clustering and flexibility. As a result, the Joint network displays a balance between long-range connectivity

and local cohesiveness, combining the structural stability of traditional financial systems with the modular adaptability of fintech ecosystems.

Centrality measures within the Joint network similarly reflect this hybrid nature. Fintech nodes tend to occupy peripheral positions, while traditional institutions dominate in terms of connectivity. However, the overall distribution of centrality is more balanced compared to the pure traditional network, indicating that systemic influence is more evenly distributed, although key hubs continue to play a significant role in maintaining the cohesion of the system.

Metric	**Traditional Network**	**Fintech Network**	**Joint Network**
Average Clustering Coefficient	0.048	0.059	0.053
Transitivity	0.022	0.029	0.023
Average Shortest Path Length	4.051	7.757	5.041
Diameter	7	17	9
Maximum Closeness Centrality	0.002	0.0008	0.001

Table 1 Structural and centrality metrics for Traditional, Fintech, and Joint financial networks.

Chapter 4: Methodological Framework

This chapter presents the methodological framework used to model financial contagion through the SIIS approach. It introduces both micro-level (agent-based) and macro-level (mean-field) contagion dynamics, alongside the simulation design, calibration, and scenarios evaluated. The aim is to enable a comprehensive assessment of how systemic risk spreads.

4.1. The SIIS contagion model

Having established a network formation process that mimics the structural properties of real-world financial systems, this section introduces the SIIS contagion model. It adapts epidemiological principles to simulate how financial distress spreads through both micro-level and macro-level contagion dynamics.

4.1.1. Overview of Classical Epidemiological Models and SIIS

Epidemiological modelling is a mathematical approach traditionally used to study how diseases spread within populations. These models, such as the Susceptible-Infected-Susceptible (SIS) and the Susceptible-Infected-Recovered (SIR) frameworks, simplify individuals into compartments based on "health" states and simulate transitions between them over time. In recent years, this approach has been applied to financial systems to understand systemic risk, where "infection" corresponds to financial distress spreading through a network of institutions. Just as diseases spread through social contact, financial contagion propagates via interbank exposures. Both models have been instrumental in assessing systemic risk, allowing researchers to simulate how initial defaults can propagate through financial networks.

SIS contagion model

The Susceptible-Infected-Susceptible (SIS) model is a classical compartmental framework in epidemiology designed to capture scenarios in which individuals, after being infected and recovering, return to a state of susceptibility rather than acquiring immunity. This cyclical structure makes it particularly relevant for modelling processes characterised by recurrent exposure and vulnerability. Formally, agents transition between two states: susceptible (S), meaning healthy but exposed to potential infection, and infected (I), meaning currently infected and capable of spreading the disease. Over time, infected agents recover and become susceptible again, and the cycle continues, or it stops when an eventual steady state is

reached. In financial systems, the SIS model has been adapted to study contagion dynamics where institutions oscillate between solvency (S) and financial distress (I). This modelling approach has been effectively applied in the study by Van Amerongen, Mir Mora, and Sánchez de la Blanca Contreras (2019)[17], where an agent-based SIS framework was developed to simulate the spread of default risk across several financial networks, offering insights into how distress propagates through these systems.

SIR contagion model

The Susceptible-Infected-Recovered (SIR) model is another foundational framework in epidemiological modelling, designed to represent processes in which individuals, once infected and recovered, gain immunity and do not re-enter the susceptible state. The model consists of three states: susceptible (S), infected (I), and recovered (R). Susceptible individuals may contract the infection and transition to the infected state, from which they eventually recover, and enter the recovered state permanently. This one-way transition reflects contagion scenarios where agents are removed from further propagation cycles after infection. In financial systems, the SIR model has been adapted to study contagion in contexts where financial distress leads to the permanent removal (R) of institutions from the market (such as bankruptcies or long-term regulatory interventions), or where the recovered (R) state may symbolise a bank resolution or merger that makes the institution no longer vulnerable to further contagion. For instance, Gai and Kapadia (2010)[18] introduced a network-based contagion model that mirrors SIR progression to show how small shocks can lead to widespread defaults depending on network density and connectivity.

SIIS contagion model

For the purpose of this thesis, a modified epidemiological framework is implemented, the Susceptible-Infected$_1$-Infected$_2$-Susceptible (SIIS) model, to better capture the heterogeneity of financial distress across complex financial networks. While classical SIS and SIR models have proven effective in modelling contagion, they assume uniform severity of infection. However, real-world financial systems often experience multiple stages of distress, such as liquidity pressure followed by insolvency risk, especially in interconnected ecosystems

[17] Amerongen et al., 'Agent-Based Models for Assessing the Risk of Default Propagation in Interconnected Sectorial Financial Networks'.
[18] Gai and Kapadia, 'Contagion In Financial Networks'.

involving both traditional banks and fintech entities. The SIIS model extends the SIS structure by introducing two distinct states of infection, allowing for the simulation of escalating risk propagation. That is, at each iteration, a node can be found in one of four possible states: Susceptible (S), First-degree Infected (I_1), Second-degree Infected (I_2), or Removed (R), with transitions governed by probabilistic rules based on exposure and institutional resilience. This allows for a more granular representation of contagion dynamics, where institutions may move from mild distress (e.g., short-term funding strain) to severe distress (e.g., insolvency) before potentially recovering. The choice of SIIS over SIIR is motivated by the need to capture recurrent exposure and recovery cycles, reflecting the empirical observation that institutions may return to stability and re-engage in the financial network after distress, rather than gaining permanent immunity. In contrast, nodes in the Removed (R) state represent institutions that have exited the system and are, from a network perspective, no longer active participants, thus physically removed from the graph and excluded from further contagion dynamics.

4.1.2. Agent-Based Modelling of SIIS: Micro-Level Contagion Dynamics

From a micro-level perspective, the SIIS model is formulated as a stochastic agent-based model (ABM). A stochastic ABM simulates the behaviour of a system by modelling the interactions of individual agents, each following probabilistic rules. Indeed, the stochasticity reflects real-world uncertainty (such as unexpected defaults) by introducing random variation in each agent's state transitions at every time step. Unlike deterministic models, where the system evolves in a fixed and predictable way, stochastic ABMs allow different simulation runs to yield different outcomes, even under identical initial conditions.

In the case of the SIIS model, transitions between the different states depend not only on the institution's own characteristics and current state, but also on its network neighbours and their level of distress. At every time step, a random value is generated (from a uniform distribution between 0 and 1) for each possible transition (infection, escalation, recovery, or removal). This random value is then compared to the corresponding computed transition probability, which is derived from factors such as exposure intensity (edge weights), infection rate, and institutional persistence. If the random draw falls below the threshold, the transition occurs.

To operationalise these transitions, the model uses a set of core parameters that regulate the dynamics of contagion and recovery:

- Λ (lambda): the external entry rate, representing the inflow of new healthy institutions into the financial system over time.
- β (beta): the infection rate, determining the likelihood that a susceptible node becomes infected based on its exposure to distressed neighbours.
- α (alpha): the escalation probability, which governs the transition from first-degree (I_1) to second-degree (I_2) infection under persistent distress.
- μ (mu): the recovery probability, controlling how likely a node is to return to the susceptible state from either I_1 or I_2.
- ν (nu): the failure probability, governing the chance that a severely distressed node (I_2) exits the system.
- τ (tau): the sector-specific persistence threshold, representing the number of consecutive distressed periods a node can withstand before forced removal once it enters the severely distressed state I_2.

The following part of this section formalises the agent-based implementation of the SIIS contagion process by detailing the specific transition dynamics that govern the behaviour of each financial institution (node) in the network. At each time step, every agent updates its state based on a set of probabilistic rules presented below, according to the stochastic process mentioned above. Thus, the logic is executed at the node level and encoded directly into the agent-based simulation engine, enabling the model to replicate non-linear contagion pathways under stochastic uncertainty.

The general transition dynamics for both the Traditional and Fintech sectors are described as follows:

1. *Susceptible (S):* a node in the S state represents a financially healthy institution that is not currently in distress, but it is exposed to risk from its distressed neighbours.

    - *Transition to First-Degree Infected (I_1):* if the institution has connections to infected nodes (I_1 or I_2), it may transition to financial distress with a probability dependent on the infection rate β and on the financial exposure (represented by the edge weights) to distressed neighbours.

The probability of infection for a healthy node due to the influence of its neighbours is computed iteratively by initialising $q_i = 1$ (which represents the probability of avoiding infection) and updating it for each infected neighbour based on the infection rate β and the edge weight between the two nodes. The formula $q_i(t) = \Pi^{N}_{j=1} (1-\beta rji)$ progressively lowers q_i as more infected neighbours apply their influence, where *rji* represents the weight of the edge between node *i* and node *j*. The final infection probability is given by $1 - q_i$, meaning that a node is more likely to be infected if it has multiple distressed neighbours or strong financial ties to them.

- *Remains in S:* if no transmission occurs, the node stays in S.

```
IF node is Susceptible (S):
    SET q_i = 1  "" probability of NOT getting infected ""

    FOR each neighbour in neighbours:
        IF neighbour is infected (I1 or I2):
            IF an edge exists between neighbour and current node:
                RETRIEVE edge weight
                UPDATE q_i using infection probability formula:
                 q_i *= (1 - beta * edge_weight)

    COMPUTE infection probability: infection_prob = (1 - q_i)

    GENERATE random value rand_val between 0 and 1
    IF rand_val < infection_prob:
        SET node's state to I1
    ELSE:
            node remains S
```

2. *First-Degree Infected (I_1):* nodes in I_1 have entered a state of financial distress, representing institutions facing liquidity issues, delayed payments, or other financial instability.

    - *Transition to Second-Degree Infected (I_2):* if distress persists, the institution escalates to I_2 with probability α, reflecting a worsening financial condition.
    - *Recovery to Susceptible (S):* the institution may recover and return to S with probability μ, indicating that it has regained stability.

- *Remains in I_1:* if neither escalation nor recovery occurs, the node stays in I_1 while accumulating defaulted consecutive turns, increasing its risk of being removed.

ELIF node is First-Degree Infected (I1):
    INCREMENT defaulted consecutive turns
    GENERATE random value *rand_val* between 0 and 1

    IF rand_val < escalation rate (α):
        SET node's state to I2
        INCREMENT defaulted consecutive turns
    ELIF rand_val < (α + recovery rate μ):
        SET node's state to S
        RESET defaulted consecutive turns
    ELSE:
        node remains in I1

3. *Second-Degree Infected (I_2):* institutions in I_2 face increased financial instability and a greater likelihood of failure.

- *Transition to Removed (R):* if distress persists, the institution fails with probability ν and is permanently removed from the network. An institution could also transition to R from I_2 if the number of consecutive turns it has been in I_2 exceeds the removal threshold τ.
- *Recovery to Susceptible (S):* the institution may recover and return to S with probability μ, indicating that it has regained stability.
- *Remains in I_2:* if neither failure nor recovery occurs, the node stays in I_2 while accumulating defaulted consecutive turns, increasing its risk of being removed.

ELIF node is Second-Degree Infected (I2):
    INCREMENT defaulted consecutive turns
    GENERATE random value *rand_val* between 0 and 1

    IF rand_val < removal probability (ν):
        ADD node to removal list
    ELIF rand_val < (ν + recovery rate μ):
        SET node's state to S
        RESET defaulted consecutive turns
    ELIF defaulted consecutive turns exceed removal threshold:
        ADD node to removal list
    ELSE:

node remains in I2

4. *Removed (R)*: institutions in R have permanently exited due to failure, bankruptcy, or regulatory intervention. Once a node enters R, it is removed from the network, and its edges might be rewired to maintain connectivity among the remaining institutions, according to the selected reconnection policy presented in section 4.2.3.

FOR each node in *nodes_to_remove*:
REMOVE node from graph
REWIRE remaining neighbours to maintain network connectivity

Finally, to further reflect the intrinsic differences between traditional and fintech institutions within the Joint financial network simulation, an additional modification to the standard SIIS contagion mechanism is introduced. Specifically, while susceptible traditional institutions exposed to financial distress transition first into a mildly distressed state ($Infected_1$), fintech firms in the Joint setting escalate immediately to the severely distressed state ($Infected_2$) after infection. This sector-specific escalation rule is applied only in the Joint network to capture the hybrid dynamics of an integrated financial system. Note that in the distinct Traditional and Fintech network simulations, all institutions initially transition to the $Infected_1$ stage regardless of their type. Formally, under the Joint network specification:

GENERATE random value *rand_val* between 0 and 1
IF rand_val < *infection_prob*:
node becomes infected
IF the node is Fintech (`F`):
SET node's state to I2
ELSE:
SET node's state to I1
ELSE:
node remains S

4.1.3. Mean-Field Approximation of SIIS: Macro-Level Contagion Dynamics

While the agent-based model captures contagion dynamics through local interactions and stochastic transitions at the node level, it is often analytically useful to approximate the aggregate behaviour of the system using a deterministic framework. This is achieved through a mean-field approximation, a technique that replaces the individual stochastic processes with average quantities representing the behaviour of a large population. In the context of the SIIS model, the mean-field approach translates the micro-level contagion mechanics into a system of three ordinary differential equations (ODEs), describing the time evolution of the proportions of nodes in each state, that is, Susceptible (S), First-degree Infected (I_1), and Second-degree Infected (I_2). For simplicity, the total population is assumed to remain constant over time, such that $N = S + I_1 + I_2$, with removed nodes (R) permanently excluded from the system.

$$\begin{cases} \frac{dS}{dt} = \Lambda - \beta\langle k\rangle \frac{(I_1 + I_2)S}{N} + \mu(I_1 + I_2) \\ \frac{dI_1}{dt} = \beta\langle k\rangle \frac{(I_1 + I_2)S}{N} - \mu I_1 - \alpha I_1 \qquad (1) \\ \frac{dI_2}{dt} = \alpha I_1 - \mu I_2 - \nu I_2 \qquad (2) \end{cases} \qquad (3)$$

The first equation models the overall change in the population of financially healthy institutions. It balances new entries into the system, losses due to contagion from distressed neighbours, and gains from institutions recovering from either mild or severe distress:

- $+ \Lambda$ represents the external inflow of new solvent institutions (e.g., market entrants).
- $- \beta\langle k\rangle \frac{(I_1+I_2)S}{N}$ is the expected infection rate: susceptible nodes become infected through interactions with both I_1 and I_2 neighbours.
- $+ \mu(I_1+I_2)$ accounts for recovery: infected institutions that regain stability re-enter the susceptible pool.

The second equation tracks the total share of institutions in early-stage distress. It reflects how many new cases arise due to contagion, and how many agents leave this state either by recovering or deteriorating further:

- $+ \beta\langle k\rangle \frac{(I_1+I_2)S}{N}$ models new infections entering the I_1 state due to exposure from infected neighbours.
- $-\mu I_1$ represents recovery from mild distress back to the susceptible state.
- $-\alpha I_1$ reflects escalation: institutions whose condition worsens move from I_1 to I_2.

The third equation measures the evolution of the severely distressed institutions in the network. It accounts for inflows from escalating I_1 nodes, and outflows due to either recovery or failure:

- $+\alpha I_1$ models the inflow into I_2 from escalation of I_1 institutions.
- $-\mu I_2$ represents recovery from severe distress.
- $-\nu I_2$ captures permanent removal of highly distressed institutions.

Note that in the mean-field formulation, the persistence threshold τ used in the agent-based model is indirectly reflected in the choice of the removal rate parameter ν. The continuous system does not explicitly track time spent in a given state, but rather aggregates expected outflows based on average duration.

Steady-State Analysis

Understanding the steady-state behaviour of the SIIS contagion model is essential for assessing the long-term resilience or vulnerability of financial networks under distress. The steady state refers to the equilibrium condition where the proportions of susceptible, mildly distressed, and severely distressed institutions stabilise over time, and no further significant changes occur.

In this section, the equilibrium values of the system are derived to identify the conditions under which financial distress stabilises. By setting the time derivatives of the state variables equal to zero, we determine the steady-state values of *susceptible institutions* S^**, mildly distressed institutions* I_1^*, and *severely distressed institutions* I_2^*.

1. The first step in the steady-state analysis focuses on solving for the proportion of severely distressed institutions, I_2, under equilibrium conditions. At steady state, by definition, the rate of change of each state variable becomes zero. Rearranging the third equation from the system allows us to solve explicitly for I_2 in terms of I_1:

$$\frac{dI_2}{dt} = \alpha I_1 - \mu I_2 - \nu I_2 = 0$$

$$\alpha I_1 = (\mu + \nu) I_2$$

$$I_2 = \frac{\alpha I_1}{\mu + \nu}$$

2. In the second step, we solve for the steady-state proportion of mildly distressed institutions, I_1, ensuring a non-trivial equilibrium condition where $I_1 > 0$. Then, we solve for S in order to obtain the first steady-state condition for **S***:

$$\frac{dI_1}{dt} = \beta\langle k\rangle \frac{(I_1 + I_2)S}{N} - \mu I_1 - \alpha I_1 = 0$$

$$\beta\langle k\rangle \frac{\left(I_1 + \frac{\alpha I_1}{\mu + \nu}\right) S}{N} - \mu I_1 - \alpha I_1 = 0$$

$$I_1 \left(\beta\langle k\rangle \frac{\left(1 + \frac{\alpha}{\mu + \nu}\right) S}{N} - \mu - \alpha \right) = 0$$

$$\beta\langle k\rangle \frac{\left(1 + \frac{\alpha}{\mu + \nu}\right) S}{N} = \mu + \alpha$$

$$\boldsymbol{S^* = \frac{N(\mu + \alpha)}{\beta\langle k\rangle \left(1 + \frac{\alpha}{\mu + \nu}\right)}}$$

3. In the third step, we determine the steady-state value of $\mathbf{I_1^*}$ by substituting the previously derived expressions for **S*** and I_2 into the steady-state equation for S, and solving for I_1 explicitly:

$$\Lambda - \beta\langle k\rangle \frac{\left(I_1 + \frac{\alpha I_1}{\mu + \nu}\right) S^*}{N} + \mu \left(I_1 + \frac{\alpha I_1}{\mu + \nu}\right) = 0$$

$$I_1\left(\beta\langle k\rangle\frac{\left(1+\frac{\alpha}{\mu+\nu}\right)S^*}{N}-\mu\left(1+\frac{\alpha}{\mu+\nu}\right)\right)=\Lambda$$

$$I_1\left(\beta\langle k\rangle\frac{\left(1+\frac{\alpha}{\mu+\nu}\right)}{N}\cdot\frac{N(\mu+\alpha)}{\beta\langle k\rangle\left(1+\frac{\alpha}{\mu+\nu}\right)}-\mu\left(1+\frac{\alpha}{\mu+\nu}\right)\right)=\Lambda$$

$$I_1\left(\mu+\alpha-\mu\left(1+\frac{\alpha}{\mu+\nu}\right)\right)=\Lambda$$

$$I_1\left(\alpha-\frac{\mu\alpha}{\mu+\nu}\right)=\Lambda$$

$$I_1\alpha\frac{\nu}{\mu+\nu}=\Lambda$$

$$\boldsymbol{I_1^*=\frac{\Lambda(\mu+\nu)}{\alpha\nu}}$$

4. In this step, we compute the steady-state value of $\mathbf{I_2^*}$ by substituting the previously obtained expression for $\mathbf{I_1^*}$ into the relationship between I_1 and I_2 derived in Step 1:

$$I_2^*=\frac{\alpha}{\mu+\nu}\times\frac{\Lambda(\mu+\nu)}{\alpha\nu}$$

$$\boldsymbol{I_2^*=\frac{\Lambda}{\nu}}$$

5. In the final step, we substitute the expressions for $\mathbf{I_1^*}$ and $\mathbf{I_2^*}$ back into the steady-state condition for S to solve for the explicit equilibrium value $\mathbf{S^*}$, since $N^* = S^* + I_1^* + I_2^*$:

$$S^*=\frac{\left(S^*+\frac{\Lambda(\mu+\nu)}{\alpha\nu}+\frac{\Lambda}{\nu}\right)(\mu+\alpha)}{\beta\langle k\rangle\left(1+\frac{\alpha}{\mu+\nu}\right)}$$

$$S^*=\frac{\left(S^*+\frac{\Lambda}{\nu}\cdot\frac{\mu+\nu+\alpha}{\alpha}\right)(\mu+\alpha)}{\beta\langle k\rangle\left(1+\frac{\alpha}{\mu+\nu}\right)}$$

$$S^*\cdot\beta\langle k\rangle\left(1+\frac{\alpha}{\mu+\nu}\right)=\left(S^*+\frac{\Lambda}{\nu}\cdot\frac{\mu+\nu+\alpha}{\alpha}\right)(\mu+\alpha)$$

$$S^* \cdot \beta\langle k\rangle \cdot \frac{\mu+\nu+\alpha}{\mu+\nu} = S^*(\mu+\alpha) + \frac{\Lambda}{\nu} \cdot \frac{(\mu+\nu+\alpha)(\mu+\alpha)}{\alpha}$$

$$S^* \left[\beta\langle k\rangle \cdot \frac{\mu+\nu+\alpha}{\mu+\nu} - (\mu+\alpha)\right] = \frac{\Lambda}{\nu} \cdot \frac{(\mu+\nu+\alpha)(\mu+\alpha)}{\alpha}$$

$$S^* = \frac{\frac{\Lambda}{\nu} \cdot \frac{(\mu+\nu+\alpha)(\mu+\alpha)}{\alpha}}{\beta\langle k\rangle \cdot \frac{\mu+\nu+\alpha}{\mu+\nu} - (\mu+\alpha)}$$

$$\mathbf{S^* = \frac{\Lambda}{\nu} \cdot \frac{(\mu+\nu+\alpha)(\mu+\alpha)}{\alpha\left[\beta\langle k\rangle \cdot \frac{\mu+\nu+\alpha}{\mu+\nu} - (\mu+\alpha)\right]}}$$

As demonstrated by the computations above, each steady-state variable is determined by a specific combination of model parameters: the infection rate β, escalation probability α, recovery rate μ, failure rate ν, network connectivity ⟨k⟩, and external inflow Λ.

The share of severely distressed institutions at equilibrium (I_2*) is driven by the ratio of the inflow rate Λ to the failure rate ν, indicating that systems with high external inflows or slow failure dynamics are more likely to accumulate vulnerable nodes. Similarly, I_1* depends on both escalation (α), and recovery plus failure parameters (μ, ν), reflecting the tension between deterioration and resilience among mildly distressed institutions. Finally, the equilibrium proportion of healthy institutions S* is jointly influenced by the infection pressure (captured by β⟨k⟩), the internal recovery capacity of the system (μ), and escalation dynamics (α, ν); this points out how stability depends not only on the network's structural connectivity, but also on the balance between internal resilience and the pressure applied by incoming shocks.

These steady-state relationships offer a quantitative benchmark for interpreting simulation outcomes. In Chapter 5, they will help assess whether the micro-level dynamics align with the macro-level expected results derived from the mean-field model.

4.2. Simulation Design

After defining how contagion spreads through the network, this section outlines the simulation setup used to implement the agent-based SIIS model. It details the parameter calibration strategy and the policy interventions embedded within the agent-based environment.

4.2.1. Parameter Calibration

To simulate the contagion dynamics accurately across the Traditional, Fintech, and Joint financial networks, the key parameters of the SIIS model were calibrated to reflect the empirical characteristics observed in each sector. In particular, constant values were set for infection probability (β) and recovery rate (μ) across all simulations, while the values for escalation probability (α) and removal probability (ν) were varied to capture sector-specific resilience and vulnerability patterns.

The infection probability β and recovery rate μ are kept constant across all network scenarios to reflect the assumption that general economic conditions are common to all institutions in the simulated environment, consistent with the modelling choices proposed by Cheng and Zhao (2019)[19] in their study. This choice ensures that differences in contagion outcomes are due to network structure and specific sectoral fragilities.

In contrast, the escalation probability α and removal probability ν are varied between traditional and fintech institutions to reflect their distinct financial architectures. Higher escalation rates (α) are assigned to fintech firms to account for their generally lower buffers against shocks and their faster deterioration once distress begins. Lower escalation rates are chosen for traditional banks, which benefit from deeper capital reserves and historically slower paths toward failure. This distinction mirrors findings by Li, Tan, and Huang (2023)[20], which show that fintech institutions move more quickly from stability to distress and are more fragile when facing contagion shocks than traditional financial firms. Similarly, removal probabilities (ν) are set higher for fintech firms due to weaker systemic safeguards, while lower values are assigned to traditional institutions, reflecting their stronger chances of recovery or support from regulators.

Finally, both the external inflow rate of new healthy institutions (Λ) and the removal threshold (τ) are also differentiated across network types. The external inflow rate is set higher for the fintech network to capture its dynamic nature, where continuous innovation drives a more frequent arrival of new institutions compared to the more mature traditional banking sector. Conversely, the removal threshold τ is set lower for fintech firms to reflect

[19] Cheng and Zhao, 'Modeling, analysis and mitigation of contagion in financial systems'.
[20] Li, Tan, and Huang, ‘Research on Risk Contagion Mechanism of Big Fintech Based on the SIRS Model’.

their higher vulnerability and limited capacity to resist prolonged distress, whereas traditional banks are allowed longer persistence under stress, consistent with their stronger financial buffers.

4.2.2. Reconnection Policies

Following the removal of a node due to persistent distress or failure, its former neighbours are allowed to reconnect to maintain network continuity. Inspired by the framework proposed by Van Amerongen, Mir Mora, and Sánchez de la Blanca Contreras (2019)[21], three different reconnection policies are presented to model alternative network adaptation processes.

- *None*: the node's neighbours do not rewire their disconnected edges after the node's removal, leading to progressive fragmentation of the network.
- *Random*: the node's neighbours select a new partner uniformly at random from all available nodes in the network, regardless of the health state of the target node (that is, both healthy and distressed nodes can be selected).
- *Risk-Aware:* the node's neighbours reconnect only to nodes that are currently healthy (i.e., nodes in the susceptible state 'S'), avoiding connections to already distressed or defaulted institutions.

For the purpose of this thesis, both the None and Risk-Aware reconnection policies are employed in the simulation for each network configuration. The random reconnection policy is intentionally excluded from the main simulations, as it introduces excessive noise and unrealistic dynamics by allowing distressed institutions to reconnect randomly with both healthy and unhealthy nodes. Such behaviour would not align with the rational risk-averse strategies typically observed in financial systems, where entities aim to minimise exposure to already vulnerable counterparties during crises. Focusing on None and Risk-Aware reconnection thus ensures a more realistic and policy-relevant exploration of systemic risk evolution.

[21] Amerongen et al., 'Agent-Based Models for Assessing the Risk of Default Propagation in Interconnected Sectorial Financial Networks'.

4.2.3. Simulation Scenarios

To study how network structure and reconnection rules affect contagion, we run simulations combining different policies with different network setups. The goal is to compare how the system behaves depending on whether agents reconnect in certain ways and whether a precautionary saving rule is used or not.

In all simulations, initial shocks are assigned at random to a subset of nodes. The structural parameters of the SIIS model (e.g., infection rate, escalation rate, removal threshold) are customised for each network configuration (Traditional, Fintech, Joint) as described in section 4.2.1., but kept constant across the three policy scenarios to ensure comparability. Each simulation is run until a steady state is reached, defined as the point at which changes in the states of nodes fall below a tolerance threshold of 0.1%, indicating convergence of the system's dynamics.

To ensure reproducibility, the full set of parameters used in the simulations is presented in Table 2. This tabular overview serves as a reference point for interpreting simulation outcomes and for evaluating the relative impact of policy design across different financial network topologies. These parameters were chosen after preliminary simulation testing, as they produced the clearest differentiation in systemic outcomes across network types and policy regimes, while preserving internal consistency with the model's behavioural assumptions.

	Λ	β	α	μ	ν	τ
Traditional	2	0.7	0.3	0.03	0.07	15
Fintech	4	0.7	0.4	0.03	0.1	10
Joint	3	0.7	0.35	0.03	0.08	13

Table 2 Simulations Parameters

Each network type is tested under the following three scenarios:

- Scenario 1) *No Reconnection, No Saving Mechanism*: agents that lose connections due to the removal of distressed neighbours do not attempt to rewire. The network is allowed to fragment as contagion progresses, and no saving mechanism is implemented to mitigate shocks. This scenario serves as a baseline to evaluate the unmitigated spread of systemic risk in the absence of adaptive behaviour.

- Scenario 2) *Risk-Aware Reconnection without Saving Mechanism*: agents that lose links only reconnect to susceptible nodes. However, there is no external intervention for systemically important institutions, making the system reliant solely on agent-level adaptation.
- Scenario 3) *Risk-Aware Reconnection with Saving Mechanism (Too-Big-to-Fail)*: in addition to risk-aware rewiring, the simulation activates a too-big-to-fail (TBTF) mechanism, which means that nodes identified as systemically important (e.g., highly connected hubs) are protected from removal through external intervention. Upon default, these nodes are automatically transitioned back to a recovered state, simulating emergency support aimed at preserving systemic stability.

Combining the three policy scenarios with the three network configurations (Traditional, Fintech, and Joint) results in 9 simulation runs in total. These runs allow us to isolate both the effect of network structure and the efficacy of policy designs under conditions of systemic vulnerability.

Chapter 5: Findings and Discussion

This chapter presents and interprets the key simulation outcomes of the contagion dynamics modelled in Chapter 4. The discussion focuses on how systemic risk develops and how severe it becomes across different network types and policy setups, as well as on how well the mean-field model matches the outcomes of the agent-based simulations.

5.1. Failure Cascades Analysis

In the context of financial contagion modelling, *failure cascades* describe the chain reaction of defaults that can be triggered after one or more institutions in a network fail. Because financial entities are connected through relationships like credit or trading, the collapse of a single node can spread distress across the system. As each institution fails, it puts additional pressure on its neighbours, potentially pushing them into failure too, amplifying the initial shock and threatening the stability of the entire network.

In this thesis, failure cascades are used as a key measure of systemic fragility. They help assess systemic risk by showing how far contagion spreads in simulations. For the purpose of this study, a failure cascade is evaluated using its *failure density*, defined as the proportion of institutions that fail at each time step relative to the total number of institutions present before failure occurs. Specifically, it is computed as the number of nodes removed in a given iteration divided by the sum of active nodes and nodes marked for removal at that step. Unlike cumulative counts, failure density provides a normalised measure that adjusts for the system's size at each time step. Hence, this metric captures the intensity of systemic distress over time, indicating the extent of propagation in each round of contagion.

This dynamic view helps compare the resilience of different financial network structures and evaluate how well mitigation tools, like reconnection strategies or too-big-to-fail protections, can help contain systemic risk.

Scenario 1: No Reconnection, No Saving Mechanism

Scenario 1 represents the baseline contagion environment, where no reconnection policies or internal saving mechanisms (such as the too-big-to-fail intervention) are in place. The observed failure cascades in this setting reflect the system's raw vulnerability to shock propagation. The Traditional network displays a slower but more persistent failure trajectory,

with the failure density peaking around iteration 13 and unsteadily declining. This pattern reflects a hub-and-spoke architecture in which risk initially concentrates within central nodes but eventually propagates broadly, amplifying systemic risk as central institutions fail. In contrast, the Fintech network exhibits a slightly earlier spike in failures (around iteration 7), followed by a more rapid attenuation. This suggests a more fragile early-stage response due to its decentralised structure, yet its modular design limits the contagion's reach, restricting systemic risk to localised clusters. The Joint network reveals an intermediate behaviour: although it reflects the decentralisation of Fintech and the connectivity of Traditional banks, it experiences both an early rise in failure density and a sustained level of systemic distress. This combination indicates a mix of vulnerabilities: rapid shock spread from fintech nodes and wide contagion through traditional network hubs. Overall, the results demonstrate that systemic risk spreads more extensively in centralised networks, while modular systems localise but do not fully neutralise shock propagation.

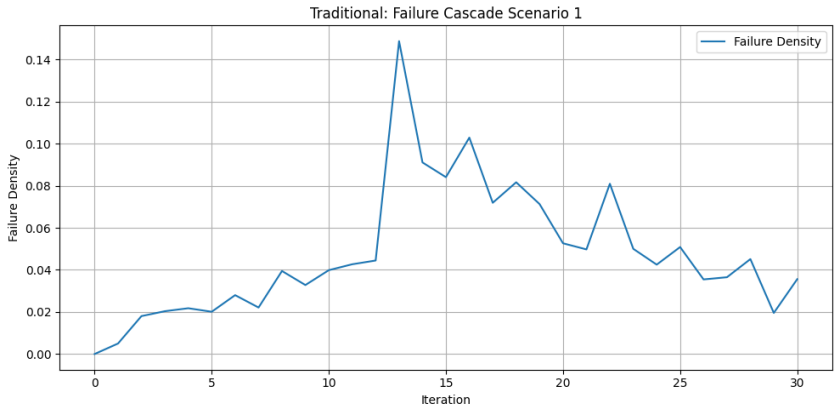

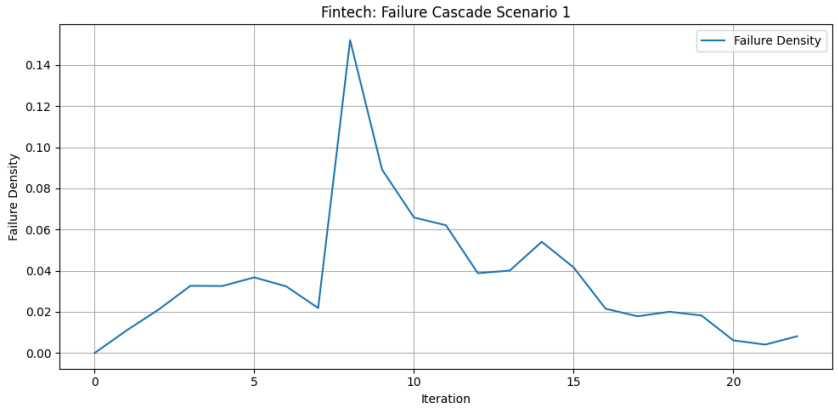

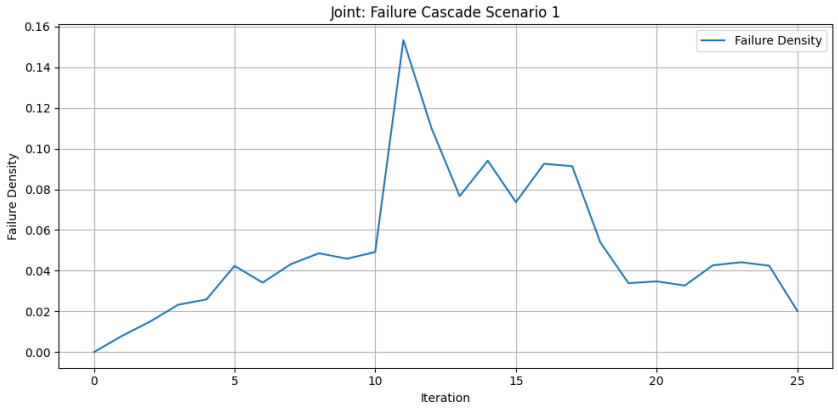

Scenario 2: Risk-Aware Reconnection without Saving Mechanism

In Scenario 2, the introduction of a risk-aware reconnection mechanism (where agents reconnect to healthy neighbours) leads to a slight improvement in systemic resilience compared to the uncontrolled contagion observed in Scenario 1. This adaptive behaviour helps reduce failure cascades across all three network types. In the Traditional network, failure density still peaks around iteration 13, but the decline is slightly more sustained, showing that even centralised systems can benefit from dynamic reconnection by reducing the persistence of contagion. The Fintech network shows some earlier and sharper failure spikes compared to Scenario 1, but in Scenario 2 the system recovers more quickly and ends with a lower failure density. This faster recovery shows how modular networks with adaptive rewiring can contain shocks and limit wider contagion. In the Joint network, the interaction between traditional and fintech nodes initially leads to both central vulnerability and local sensitivity, resulting in a higher failure peak in Scenario 2. However, the decline that follows is steadier than in Scenario 1, indicating that the risk-aware reconnection helps contain the cascade more effectively over time. This suggests improved shock absorption and a reduction in long-term systemic risk. Overall, these findings show that risk-aware reconnection helps reduce systemic risk by reshaping the network, improving resilience even in mixed financial systems.

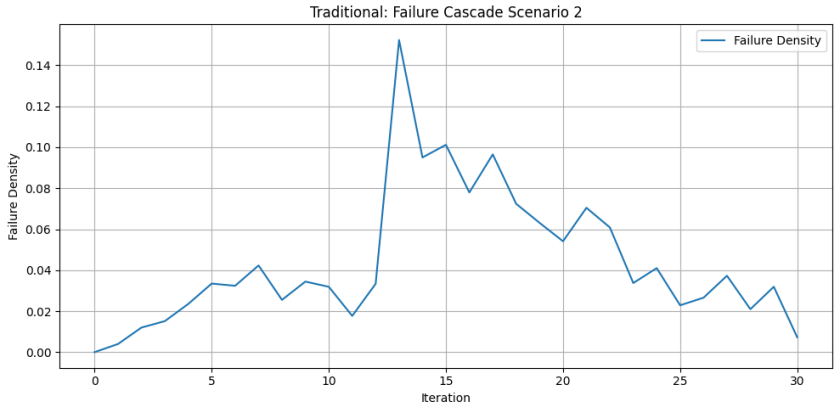

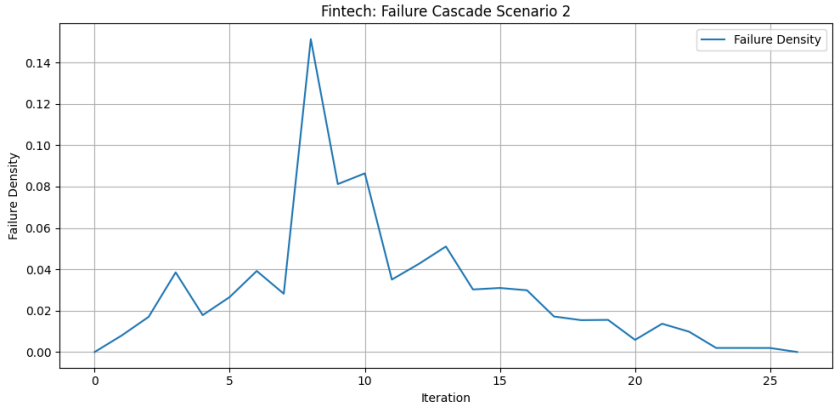

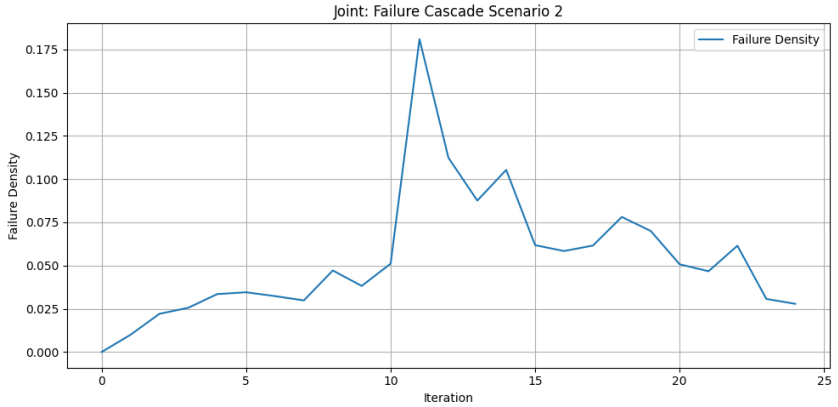

Scenario 3: Risk-Aware Reconnection with Saving Mechanism (Too-Big-to-Fail):

In this scenario, the system uses both targeted reconnection and a saving rule that protects key nodes, like major hubs, from failing. This setup reflects real-world emergency support for critical institutions. Notably, Scenario 3 differs from the previous ones in terms of timing: equilibrium is reached significantly later across all network configurations. While in Scenarios 1 and 2 the system typically stabilises within 20-30 iterations, Scenario 3 sees persistent but declining failure activity that extends up to iteration 180 in the Traditional

network. This slower convergence can be attributed to the introduction of the saving mechanism, which prevents the immediate removal of key nodes. While this intervention limits sudden collapses, it also allows mild levels of distress to propagate at a slower pace. As a result, the system avoids sharp cascades but takes longer to fully stabilise. In the Traditional network, the failure peak in Scenario 3 is slightly lower than in the earlier scenarios, showing that the saving mechanism helps soften the initial shock. By protecting key hubs from failing, the system avoids a sudden rise in default, but this also slows down recovery. The Fintech network experiences a similarly sharp initial failure peak (around 0.13 at iteration 8), but it recovers even more quickly than the Traditional network, with failure density returning to near zero by iteration 30. This fast recovery shows the strength of the saving mechanism in decentralised systems, where local clusters can stabilise more effectively when key nodes are protected. The Joint network shows the most significant improvement. Unlike Scenarios 1 and 2 (where failure peaks exceeded 0.15 and 0.17), Scenario 3 exhibits a smaller peak (just above 0.12 around iteration 12), while the system reaches near-zero failure density by iteration 60. This outcome highlights again the effectiveness of combining saving mechanisms and adaptive reconnection.

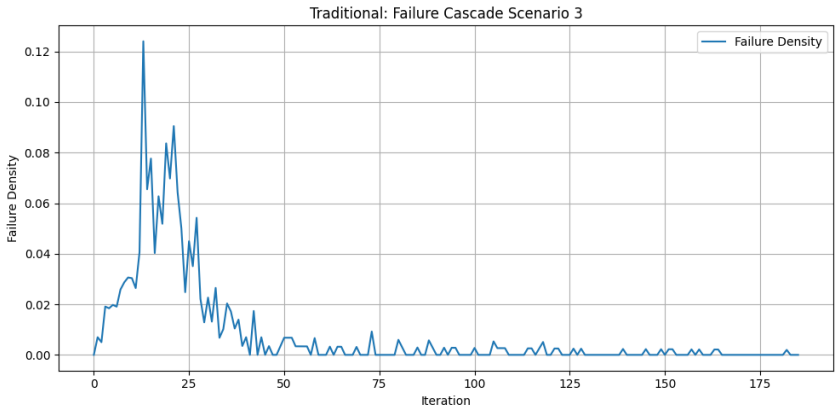

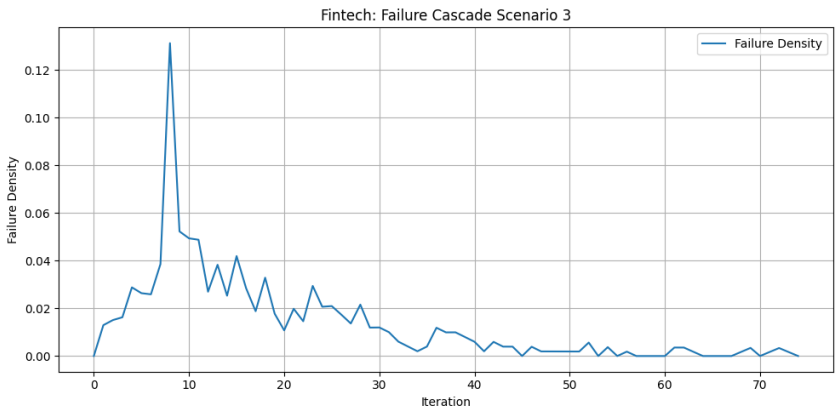

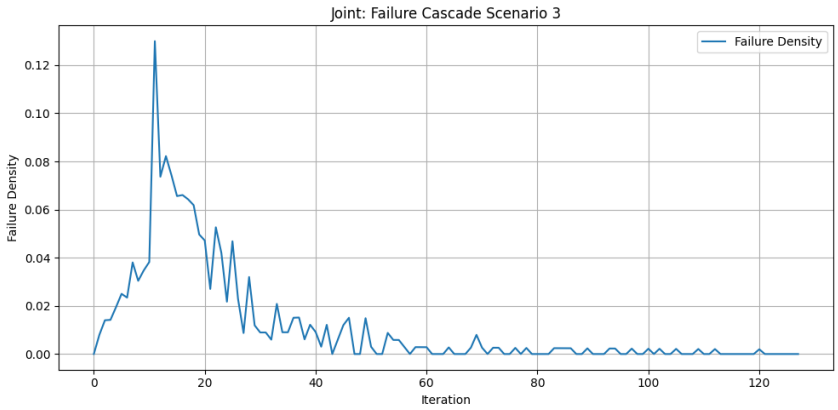

Comparing the three scenarios shows how policy mechanisms progressively improve systemic stability and reduce failure cascades. In Scenario 1, with no reconnection or saving rules, contagion spreads freely, leading to higher and prolonged failure rates, especially in the Traditional and Joint networks. This highlights how fragile centralised or mixed structures can be when left unregulated. In Scenario 2, adding risk-aware reconnection helps slow down the spread of failures, particularly in the Fintech network, which benefits from its more decentralised structure. Still, reconnection alone doesn't fully prevent systemic risk, especially in networks with key hubs. Scenario 3 provides the most effective containment: failure peaks are lower, and a more sustained stabilisation is consistent across all network types. Even the Joint network, previously the most vulnerable, stabilises well when critical hubs are protected. However, this improved resilience comes with a longer time to reach equilibrium. Failures decline more slowly, but this reflects a healthier adjustment process, where the network absorbs shocks gradually rather than facing concentrated distress early on. Thus, by combining reconnection and hub protection, the system stays flexible and responsive for longer before fully stabilising.

5.2. State Evolution Over Time

While the previous section looked at how failure spreads over time, it's also important to see how the system's internal composition changes as it moves toward stability. By tracking the shares of susceptible nodes (S), mildly distressed (I_1), and severely distressed (I_2) nodes, we can better understand how shocks move through the network and eventually fade.

Scenario 1: No Reconnection, No Saving Mechanism

Without any policy measures, in Scenario 1, the Traditional and Joint networks show a steep drop in healthy nodes (S) and a rapid rise in severely distressed nodes (I_2), indicating a fast and intense spread of contagion. The Joint network worsens the quickest, with I_2 overtaking S early on, highlighting its structural fragility due to both central hubs and dense clustering. The Traditional network follows a similar pattern, with slightly slower dynamics. By contrast, the Fintech network shows a more gradual decline in S and a lower peak in I_2. While distress still spreads, the process appears less severe overall. The number of susceptible nodes begins to stabilise earlier, suggesting that the Fintech network's more decentralised structure helps slow contagion and limit its reach. Overall, the absence of intervention in Scenario 1 leads to high infection levels across all networks. While decentralisation offers some delay, structural design alone is not enough to prevent widespread systemic distress.

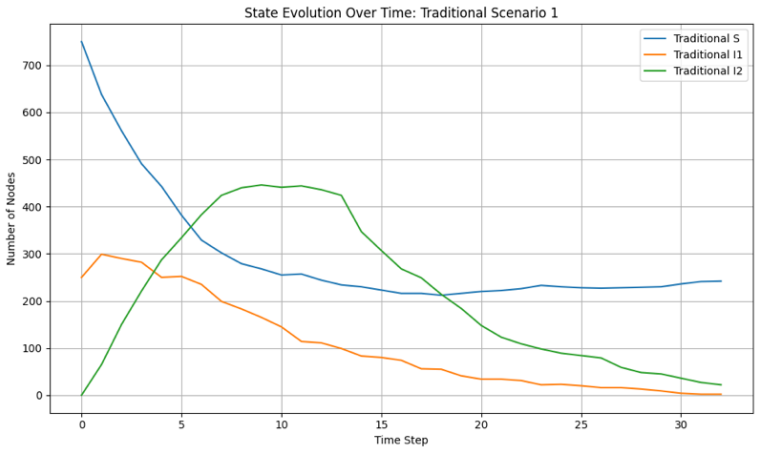

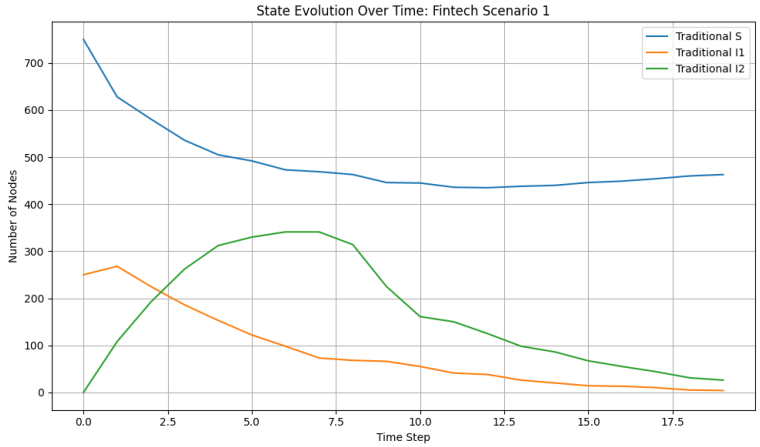

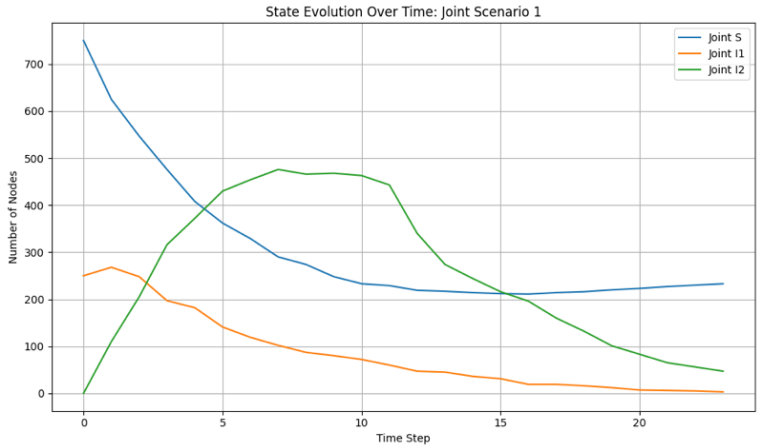

Scenario 2: Risk-Aware Reconnection without Saving Mechanism

In Scenario 2, the introduction of risk-aware reconnection leads to a moderate improvement in system behaviour, though the differences with Scenario 1 are less evident than expected. Across the networks, the decline in susceptible nodes (S) remains steep, and the rise of infected nodes is still substantial. In the Traditional network, the S population still drops rapidly, and I_2 and I_1 both reach similar peaks as in Scenario 1, although they begin to decline slightly earlier. This suggests only a modest containment effect from reconnection. The Joint network shows similar dynamics: I_2 again surpasses both S and I_1, but it falls more steadily. These results indicate that while reconnection may help delay or smooth the peak, it does not drastically alter the trajectory of contagion. In contrast, the Fintech network benefits more visibly from reconnection. While S still declines, it visibly rises again when converging to equilibrium, and the I_2 curve grows more slowly than in Scenario 1. This supports the idea that decentralised networks are better suited to benefit from adaptive behaviours like link

rewiring, which help contain contagion locally and reduce systemic exposure. Overall, Scenario 2 shows that risk-aware reconnection offers some benefit, especially in networks with decentralised structures, even if in more centralised or hybrid systems, its impact is limited.

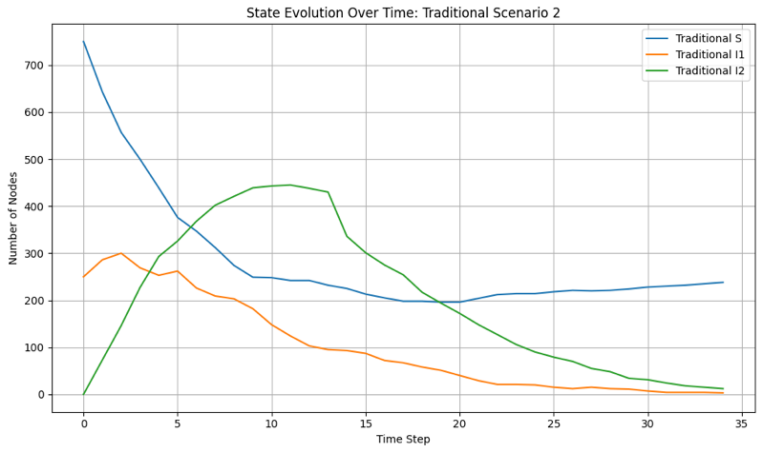

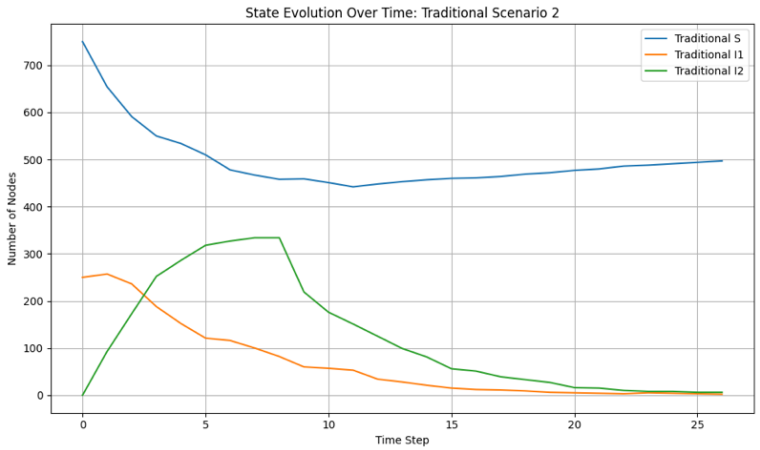

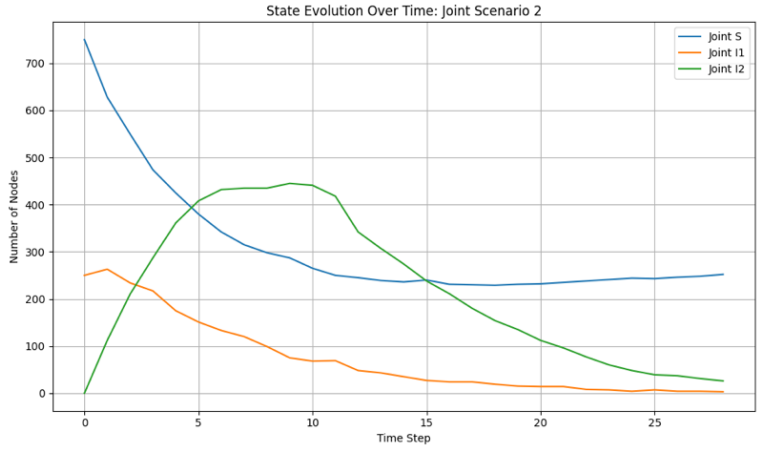

Scenario 3: Risk-Aware Reconnection with Saving Mechanism (Too-Big-to-Fail):

In Scenario 3, the combination of risk-aware reconnection and a saving mechanism leads to a different system trajectory. All networks show an initial spike in contagion (reflected in the rise of I_1 and I_2 values), especially in the Traditional and Joint configurations. However, this time, the contagion spread is more effectively contained. In the Traditional network, I_1 and I_2 rise quickly but then fall steadily, while the number of healthy nodes (S) begins to recover. Unlike previous scenarios, this recovery is strong and sustained, showing that saved nodes can reintegrate and support system stability over time. The Fintech network also experiences an early contagion spike, but the recovery in S begins even earlier and is smoother, likely thanks to its decentralised structure. The Joint network shows behaviour between the two: the early spike in distress resembles the Traditional case, but recovery starts sooner. This suggests that hybrid networks benefit significantly from having both protection for key nodes and dynamic reconnection paths. By the end of the simulation, 314 traditional and 218 fintech institutions remain active, meaning around 41% of the fintech nodes survive, closely matching their original proportion in the system. This outcome suggests that the combined use of reconnection and targeted protection not only limits systemic damage but also preserves the original composition of the network, highlighting the efficiency of these intervention strategies in maintaining long-term resilience.

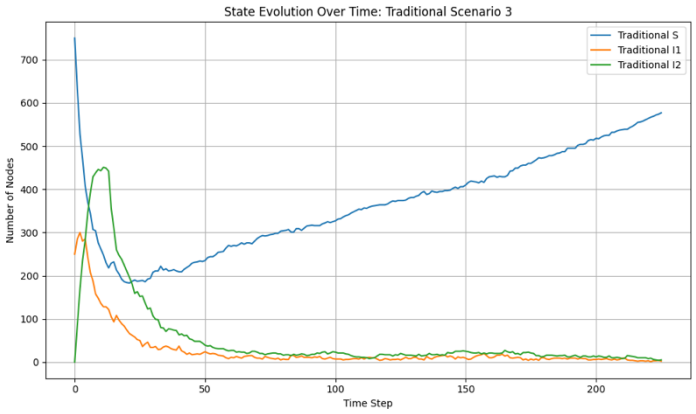

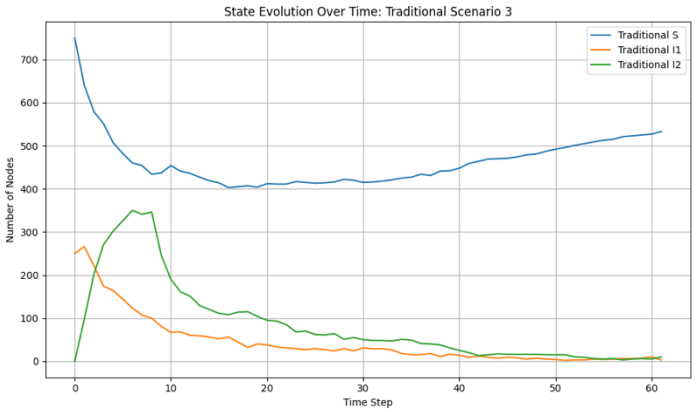

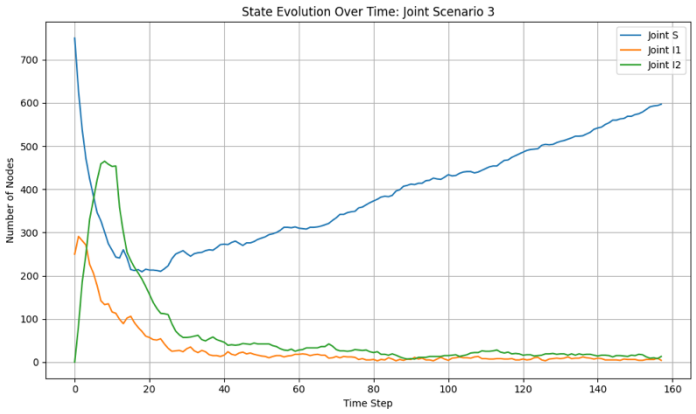

5.3. Steady-State Benchmarking

To check how well the simulation model performs, this section compares the steady states obtained from the simulations once equilibrium was reached to the values predicted by the mean-field approximation introduced in Section 4.1.3. The focus is on Scenario 1, which has no reconnection or saving mechanisms. This scenario is chosen because it matches the assumptions of the mean-field model best: a fixed network structure and no adaptive behaviours. These conditions reduce complexity, making it easier to compare the simulated and analytical results.

Traditional Network

Using the analytical expressions presented in Section 4.1.3 and the parameter values listed in Section 4.2.3, we compute the theoretical steady-state distribution for the Traditional network. From the simulation in Scenario 1, we obtain an average degree of $\langle k \rangle = 0.63$ and a

total system size of N = 246 at equilibrium. Substituting these values into the steady-state formulas:

$$S^* = \frac{N(\mu + \alpha)}{\beta\langle k \rangle \left(1 + \frac{\alpha}{\mu + \nu}\right)} = \frac{246(0.03 + 0.3)}{0.7 \cdot 0.63 \left(1 + \frac{0.3}{0.03 + 0.07}\right)} = \frac{81.18}{0.441 \cdot 4} = \frac{81.18}{1.764} \approx 46$$

$$I_1^* = \frac{\Lambda\ (\mu + \nu)}{\alpha\ \nu} = \frac{2\ (0.03 + 0.07)}{0.3 \cdot 0.07} = \frac{0.2}{0.021} \approx 9$$

$$I_2^* = \frac{\Lambda}{\nu} = \frac{2}{0.07} \approx 28$$

Comparing the theoretical predictions to the steady-state results from Scenario 1 shows a noticeable difference in the number of susceptible nodes. In one of the closest simulation runs, the system converged at S = 205, I_1 = 12, and I_2 = 29. While the computed values for I_1 (9) and I_2 (28) closely match the simulated results, the mean-field approximation underestimates the number of susceptible nodes. This difference is likely due to the Traditional network's centralised structure, which accelerates early contagion but also causes it to fade quickly, allowing more nodes to avoid long-term distress. In contrast, the mean-field model assumes uniform behaviour, which can't fully account for this early saturation effect. Still, the strong alignment in the infected groups suggests that the model captures the core contagion dynamics fairly well, supporting its use as a reasonable baseline.

Fintech Network

From the simulation of the Fintech network under Scenario 1, we extract an average degree of ⟨k⟩ = 1.2 and a total population size of N = 487 at equilibrium. Substituting these values and the parameters presented in Section 4.2.3 into the steady-state formulas derived from the maen-field approximation we obtain:

$$S^* = \frac{N\,(\mu + \alpha)}{\beta\,\langle k \rangle \left(1 + \frac{\alpha}{\mu + \nu}\right)} = \frac{487\,(0.03 + 0.4)}{0.7 \cdot 1.2 \left(1 + \frac{0.4}{0.03 + 0.1}\right)} = \frac{487 \cdot 0.43}{0.84 \cdot (1 + 3.08)}$$

$$= \frac{209.41}{0.84 \cdot 4.08} = \frac{209.41}{3.4272} \approx 61$$

$$I_1^* = \frac{\Lambda(\mu + \nu)}{\alpha\nu} = \frac{4(0.03 + 0.1)}{0.4 \cdot 0.1} = \frac{4 \cdot 0.13}{0.04} = \frac{0.52}{0.04} = 13$$

$$I_2^* = \frac{\Lambda}{\nu} = \frac{4}{0.1} = 40$$

When comparing the computed steady state for the Fintech network to the closest result from repeated simulations under Scenario 1, which are S = 432, I_1 = 7, and I_2 = 48, we observe again partial alignment. The infected groups (I_1 = 7 and I_2 = 48) are fairly close to the derived values (13 and 40, respectively), but the number of healthy nodes (S = 432) is much higher than the mean-field model's estimate of 61. This gap might come from the Fintech network's sparse and decentralised structure, which limits how easily contagion can spread. Because the mean-field model assumes that every node has an equal chance of interacting with others, it overestimates infection levels in networks that are more fragmented and locally clustered.

Joint Network

From the simulation of the Joint network under Scenario 1, the average degree is approximately ⟨k⟩ = 0.97, and the total population size at equilibrium is N = 323. Applying the steady-state equations from Section 4.1.3 and the parameter values reported in Section 4.2.3 we compute:

$$S^* = \frac{N(\mu + \alpha)}{\beta\langle k \rangle \left(1 + \frac{\alpha}{\mu + \nu}\right)} = \frac{323(0.03 + 0.35)}{0.7 \cdot 0.97 \left(1 + \frac{0.35}{0.03 + 0.08}\right)} = \frac{323 \cdot 0.38}{0.679 \cdot (1 + 3.18)} = \frac{122.74}{2.84} \approx 43$$

$$I_1^* = \frac{\Lambda(\mu + \nu)}{\alpha\nu} = \frac{3(0.03 + 0.08)}{0.35 \cdot 0.08} = \frac{3 \cdot 0.11}{0.028} = \frac{0.33}{0.028} \approx 12$$

$$I_2^* = \frac{\Lambda}{\nu} = \frac{3}{0.08} = 38$$

When we compare the computed steady states to the closest simulation results for the Joint network (S = 273, I_1 = 10, and I_2 = 40) we again see a good match for the infected segments. However, the number of susceptible nodes is still much higher in the simulation than the value obtained from the mean-field approximation. A possible reason lies in the network's hybrid structure, which combines centralised hubs with more loosely connected nodes. These peripheral areas act as a buffer, slowing down contagion and helping more nodes avoid infection. As a result, the simulation shows more healthy agents than the homogeneous mean-field model can account for, highlighting how network heterogeneity can moderate the spread of systemic risk.

Looking across the steady states of the Traditional, Fintech, and Joint networks, it appears clear that a system's structure plays a central role in how vulnerable it is to long-term

distress. The Traditional network, with its lower number of susceptible nodes at the equilibrium reached through the simulation, shows how centralised systems allow contagion to spread widely; indeed, tight connections mean that once distress starts, it travels fast. On the other hand, the Fintech network retains a much larger group of nodes unaffected, thanks to its decentralised and loosely connected design. This structure acts like a natural buffer, slowing the spread and reducing overall risk, though it might also mean less coordination or mutual support between agents. The Joint network falls somewhere in between. Its final infection levels, once the simulation reaches equilibrium, are close to the Fintech case, but it ends with more removed nodes and fewer remaining susceptible ones. This suggests that while it benefits from decentralisation, it still carries vulnerabilities from its more centralised components.

Across all three networks, the mean-field model, even with its simplified assumptions, captures the system's behaviour with fair accuracy. It closely matches the number of infected nodes and reflects the overall patterns seen in the simulations. While it tends to underestimate how many nodes finish as susceptible once the equilibrium is reached, especially in decentralised networks, this gap is likely due to structural details like clustering and non-uniform connectivity, which are not captured by the homogeneous mixing assumption of the mean-field model. This assumption treats all nodes as equally likely to interact, failing to consider the uneven interaction patterns that shape real network dynamics. Still, the strong alignment in infection levels shows that the simulation model is a reliable tool for understanding how different network structures shape long-term systemic risk.

Chapter 6: Conclusion

This chapter summarises the key findings of the thesis and discusses their implications for understanding financial contagion. It also considers the strengths and limitations of the model and outlines possible directions for future improvements.

6.1. Summary of Key Findings

This thesis introduced and applied a new agent-based SIIS contagion model to simulate how systemic risk spreads through three types of financial networks: Traditional, Fintech, and Joint. Using a set of different simulation scenarios, the model examined how network structure and policy responses affect both the severity and control of financial contagion. The analysis focused on three key areas: failure cascades, state evolution over time, and comparison between simulation outcomes and steady-state values derived from the mean-field model.

The failure cascade analysis (Section 5.1.) showed that centralised networks, like the Traditional configuration, are more exposed to deep and prolonged contagion. In contrast, the Fintech network, despite being vulnerable to early failures, tends to localise shocks due to its decentralised structure. The Joint network combined vulnerabilities from both types, experiencing wide contagion and early instability. Policy measures introduced in Scenario 2 (risk-aware reconnection) and Scenario 3 (reconnection with a saving mechanism) strengthened system resilience, with Scenario 3 proving most effective, especially in stabilising the Joint network by protecting key hubs.

The state evolution analysis (Section 5.2.) supported these results. In Scenario 1, all networks showed clear signs of systemic stress, while Scenarios 2 and 3 enabled higher recovery and more effective stabilisation. The Fintech network consistently retained more healthy nodes, and Scenario 3 stood out for lowering the contagion peak and enabling a more gradual transition to equilibrium, thanks to the combined effect of adaptive rewiring and targeted node protection.

In Section 5.3., the model's steady-state results from simulations were compared with mean-field predictions. While the mean-field approach underestimated the number of healthy nodes, especially in decentralised networks, it closely matched the infection levels, validating the core model assumptions. Network structure emerged as a key driver of outcomes: the

Traditional network ended with more infections, the Fintech network preserved more healthy nodes, and the Joint network reflected a mixed pattern.

Overall, these findings highlight the critical role of network topology in shaping systemic risk and demonstrate how targeted, adaptive interventions can meaningfully reduce contagion in complex financial systems.

6.2. Limitations

While the simulation model provides useful insights into how systemic risk spreads, it's important to recognise its limitations. First, the mean-field approximation assumes that all nodes interact equally, ignoring real-world structural features like clustering or uneven connectivity. This leads it to underestimate the share of susceptible nodes, especially in sparse or decentralised networks, making it less accurate in those settings.

Second, the model treats all financial institutions as identical agents that follow fixed rules. It doesn't capture differences in size, risk tolerance, regulation, or how institutions might change their behaviour in response to stress. It also assumes a constant inflow of new nodes (through the parameter Λ), which simplifies growth but doesn't reflect how institutions might enter or exit the market under real economic pressures.

Finally, the network structures (Traditional, Fintech, and Joint) are static throughout the simulations. In reality, connections between institutions often shift as they respond to shocks, regulations, or market conditions. These simplifications were necessary to keep the model manageable, but they also point to opportunities for future improvements, especially in making the model more dynamic and behaviourally realistic.

6.3. Directions for Future Research

Based on the results presented in this thesis, several extensions could enhance the analysis of systemic risk in financial networks. One important step would be to introduce agent heterogeneity, allowing institutions to differ in size, risk exposure, and how they respond to contagion. This would make it easier to study the effects of targeted interventions and how shocks impact different parts of the system.

Another valuable improvement would be to make network formation dynamic, so agents can form or break connections based on market conditions or learning over time. This would

reflect real-world behaviour more accurately and test how policy tools hold up as networks evolve.

Finally, future research could use real-world data from interbank networks or fintech ecosystems to calibrate the model. This would allow for more realistic testing and evaluation of policy measures under actual historical shock scenarios.

Bibliography

Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. ‘Systemic Risk and Stability in Financial Networks’. *American Economic Review* 105, no. 2 (1 February 2015): 564–608.

Albert, Reka, and Albert-Laszlo Barabàsi. ‘Statistical Mechanics of Complex Networks’. *Reviews of Modern Physics* 74, no. 1 (30 January 2002): 47–97.

Amerongen, Philippe VAN, Ramon MIR Mora, Blanca Contreras, and Jordi Nin. ‘Agent-Based Models for Assessing the Risk of Default Propagation in Interconnected Sectorial Financial Networks’.

Barabàsi, Albert-Laszlo, and Reka Albert. ‘Albert, R.: Emergence of Scaling in Random Networks. Science 286, 509-512’. *Science (New York, N.Y.)* 286 (1 November 1999): 509–12.

Bianchi, Pietro A., Monika Causholli, Miguel Minutti-Meza, and Valbona Sulcaj. ‘Social Networks Analysis in Accounting and Finance’. *Contemporary Accounting Research* 40, no. 1 (2023): 577–623.

Cheng, Xian, and Haichuan Zhao. ‘Modeling, Analysis and Mitigation of Contagion in Financial Systems’. *Economic Modelling* 76 (January 2019): 281–92.

European Systemic Risk Board. *How Does Risk Flow in the Credit Default Swap Market?.No 33 / December 2016* . LU: Publications Office, 2016.

Harsono, Iwan, and Ida Suprapti. ‘The Role of Fintech in Transforming Traditional Financial Services’. *Accounting Studies and Tax Journal (COUNT)* 1 (22 January 2024): 81–91.

Li, Jianping, Jingyu Li, Xiaoqian Zhu, Yinhong Yao, and Barbara Casu. ‘Risk Spillovers between FinTech and Traditional Financial Institutions: Evidence from the U.S.’ *International Review of Financial Analysis* 71 (October 2020): 101544.

Li, Shouwei, Jianmin He, and Yaming Zhuang. ‘A Network Model of the Interbank Market’. *Physica A: Statistical Mechanics and Its Applications* 389, no. 24 (December 2010): 5587–93.

Li, Yutong, Zhongming Tan, and Chenyu Huang. ‘Research on Risk Contagion Mechanism of Big Fintech Based on the SIRS Model’. *PLOS ONE* 18, no. 9 (8 September 2023): e0291230.

Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of ‘small-world’ networks. Nature, 393(6684), 440–442.

Siddiqui, Zakia, and Claudio Andres Rivera. ‘FinTech and FinTech Ecosystem: A Review of Literature’. *Risk Governance and Control: Financial Markets and Institutions* 12, no. 1 (2022): 63–73.

Suprun, Anatoly, Tetiana Petrishina, and Iryna Vasylchuk. ‘Competition and Cooperation between Fintech Companies and Traditional Financial Institutions’. *E3S Web of Conferences* 166 (1 January 2020): 13028.