Pol Arenas Garcia

Dynamical analysis and modelling of risk assessment in financial networks

Engineering Mathematics and Physics

Final Degree Project

Directed by
Dr. Antonio Garijo Real

UNIVERSITAT ROVIRA i VIRGILI

Tarragona, 2025

Abstract

Abstract (English)

This work presents a discrete analysis of a simplified customer-supplier model, aiming to understand the dynamics of financial distress propagation in interconnected systems. Using a combination of analytical and probabilistic tools, we explore the behavior of the model under various assumptions, including a normal approximation approach to estimate default probabilities. Despite its simplifications, the model captures key mechanisms of systemic risk, such as contagion effects. The results highlight the influence of model parameters on liquidity evolution and default cascades. This study offers an original perspective on a complex problem, based on independent research conducted during the final stage of the degree. It also serves as the culmination of four years of study in Engineering Mathematics and Physics at the Universitat Rovira i Virgili.

Resumen (Español)

Este trabajo presenta un análisis discreto de un modelo simplificado de cliente-proveedor, con el objetivo de comprender la dinámica de la propagación del estrés financiero en sistemas interconectados. Mediante una combinación de herramientas analíticas y probabilísticas, exploramos el comportamiento del modelo bajo diferentes hipótesis, incluyendo una aproximación normal para estimar las probabilidades de impago. A pesar de sus simplificaciones, el modelo capta mecanismos clave del riesgo sistémico, como los efectos de contagio. Los resultados destacan la influencia de los parámetros del modelo en la evolución de la liquidez y en las cascadas de impago. Este estudio ofrece una perspectiva original sobre un problema complejo, fruto de una investigación propia llevada a cabo durante la etapa final del grado. Representa, además, la culminación de cuatro años de estudios en Ingeniería Matemática y Física en la Universitat Rovira i Virgili.

Resum (Català)

Aquest treball presenta una anàlisi discreta d'un model simplificat de client-proveïdor, amb l'objectiu d'entendre la dinàmica de la propagació de l'estrès financer en sistemes interconnectats. Mitjançant una combinació d'eines analítiques i probabilístiques, s'explora el comportament del model sota diferents hipòtesis, incloent-hi una aproximació normal per estimar les probabilitats d'impagament. Tot i les seves simplificacions, el model capta mecanismes clau del risc sistèmic, com ara els efectes de contagi. Els resultats posen de manifest la influència dels paràmetres del model en l'evolució de la liquiditat i les cascades d'impagament. Aquest estudi ofereix una perspectiva original sobre un problema complex, resultat d'una recerca pròpia desenvolupada durant l'etapa final del grau. Representa, a més, la culminació de quatre anys d'estudis en Enginyeria Matemàtica i Física a la Universitat Rovira i Virgili.

Contents

1 Introduction

Since my early years in Spain's high school system, I have always been intrigued by how mathematics can model and explain phenomena we encounter in our daily lives. From predicting how long it will take for a bus to arrive at a specific location, to understanding how light interacts with matter, mathematics seemed to offer answers to a wide array of questions. This interest with the power of mathematics to explain the world around us led me to pursue a degree in engineering mathematics and physics at the Universitat Rovira i Virgili in Tarragona.

Throughout my degree, I have had the opportunity to take advanced courses in real analysis, programming, and probabilistic theory, with the latter capturing my strongest interest. The study of random processes in nature and how to model them from a probabilistic perspective was both new to me and intellectually challenging. This approach not only allowed me to better understand random phenomena, but also opened my eyes to how mathematics can be applied to describe behavior in systems that are seemingly chaotic or uncertain.

In my third year as an undergraduate student, I had the opportunity to join the ALEPHSYS research group through an elective course. This group is primarily focused on the study of complex networks. In this course, I immersed myself in the construction and understanding of contagion models over networks. This fresh perspective on how populations function as networks, and how to apply these models, led me to reflect on how such knowledge could be utilized in real-world scenarios, particularly in economic and financial systems.

It was at this point that I first discovered the field of finance. Finance is one of the most important and foundational sectors in the world, influencing everything from global economies to personal wealth management. The role that finance plays in our daily lives is crucial, ranging from the efficient allocation of resources in an economy to investment decisions that impact the well-being of individuals, companies, and governments. I realized that, just like the other areas I had studied, finance could be understood and analyzed through mathematical models. This realization made me think about how my background in mathematics and engineering could be used to contribute to the financial field.

One of the most critical lessons from recent financial history is the power of interconnectedness in markets. The financial system can be understood as a vast network of institutions and contracts. This network structure means that distress in a single node can propagate and amplify across the entire system, sometimes with devastating effects. A notable example is the 2008 financial crisis, where the failure of a single large insurer, among other factors, acted as a catalyst that triggered one of the deepest economic crises in recent history. The crisis highlighted how systemic risk emerges not just from isolated failures, but from the complex web of dependencies and exposures between financial agents.

Understanding how financial distress propagates through these networks is therefore of paramount importance for risk management and policy design. This insight motivates the study of systemic risk from a quantitative and network-based perspective, seeking to model contagion processes and evaluate the resilience of financial systems. It is precisely this topic that forms

the core of this final degree project.

The work presented here focuses on a simplified model of financial distress propagation, drawing on concepts from quantitative finance, discrete non-linear dynamical systems, and probability theory. Before delving into the model itself, this document will first provide an overview of key financial concepts and revisit mathematical tools necessary for a thorough comprehension of systemic risk and contagion mechanisms.

My goal is to bridge the rigorous mathematical framework I have developed throughout my degree with the complex realities of financial systems, contributing to the understanding of how systemic risk arises and spreads.

2 Some Quantitative Finance Concepts

In the following section we explore different areas of quantitative finance, focusing on systemic risk.

2.1 Asset Pricing and Risk Analysis

One of the most important topics in quantitative finance is asset pricing and risk analysis.

An **asset** is a tradable entity. Examples of assets are: equities (stocks), commodities (oil, gold), and fixed income instruments (bonds, debts).

Asset pricing involves the use of advanced statistical models to predict the value of underlying assets, such as stocks. One of the most significant contributions to this field is the Black–Scholes equation, which provides a theoretical estimate for the price of European-style options, a class of options that can only be exercised at expiration, unlike their American-style counterparts which allow early exercise [1].

An **option** is a financial contract that gives the buyer the right, but not the obligation, to buy or sell an underlying asset at a specific price, on or before a certain date.

Another fundamental area within quantitative finance is risk analysis. This field utilizes advanced probabilistic models and stochastic calculus to evaluate the risk associated with financial transactions. For instance, if we hold a stock at a certain price K, we must determine the optimal price at which to sell it through an options strategy in order to maximize potential profit while minimizing potential losses.

Both topics are closely related and play a crucial role in the broader framework of portfolio management and financial decision-making. The proper understanding and application of asset pricing models and risk analysis techniques are essential for making informed investment decisions and managing financial risk effectively.

2.2 Systemic Risk Assessment

Systemic risk in quantitative finance refers to the risk that the failure of a single financial institution or asset can lead to widespread disruption in the broader financial system. Unlike individual risk, systemic risk arises from the interconnectedness of financial institutions and markets, where problems in one area can quickly propagate to others. A classic reference for systemic risk theory is provided by [2], which offers key insights into the foundations and applications of this field.

The failure of individual financial institutions can lead to the propagation of financial distress. For instance, if a bank goes bankrupt and is unable to repay its debts to other financial institutions, this inability to make payments can trigger a chain reaction. Other firms that were initially unaffected may also find themselves unable to pay their own obligations, resulting in a broader financial crisis. This phenomenon is known as financial distress propagation.

Additionally, assets themselves can contribute to the spread of financial distress. Suppose a bank goes bankrupt and, in an attempt to recover, decides to sell off all of its assets. The massive sale of these assets would likely drive down their market prices, which is known as *fire sales*. Consequently, any firm holding a portion of these assets would incur financial losses. This reduction in some firms' liquidity could lead to financial distress.

The key to understanding systemic risk lies in viewing the market as an interconnected network. In this framework, the nodes of the network represent individual firms, while the connections between them signify various relationships. Depending on the model used to represent the financial network, the edges could represent different factors. For instance, in one model, the connections might reflect the shared exposure to underlying assets between two firms, such as the percentage of a portfolio that they hold in common. In another model, the connections could represent the financial obligations between firms, such as the amounts owed to each other.

Our main objectives as mathematicians and physicists are to identify the most critical nodes within our financial network, those that propagate financial distress the fastest or are central to its spread. Additionally, we aim to understand the overall dynamics of the system in the event of a financial shock, in order to minimize the risk of such an event or mitigate its potential impact.

3 Mathematical Background

In order to proceed with the document, we must first review different concepts utilized on the later study of the model. In the following section, we visit network theory, probability theory and dynamical systems theory.

3.1 Network Theory

In network theory, a network is a mathematical structure consisting of a set of nodes (also called vertices) connected by edges (also called links). It is used to model complex systems where entities (nodes) are interconnected in some way.

A basic reference for complex networks theory is provided by [3], which offers a comprehensive overview of the fundamental concepts and structural properties of complex networks.

Nodes represent entities or individual components of the network, while edges represent the interaction between these.

Examples of networks include social networks, where nodes represent people and edges represent the relationships they have. Our main focus in this document is financial networks, where nodes represent different financial entities, ranging from banks to individual people. Edges within financial networks can represent various things, depending on what we aim to model, as explained earlier. It is often taken to represent the flow of capital or risk between nodes, but many models also consider the share of underlying assets between them.

Networks are typically classified based on their degree distribution. The degree of a single node is the number of edges incident to it. The most well-known and studied synthetic networks are:

- **Random Networks**: These networks exhibit a random degree distribution, meaning each node is connected to a random number of other nodes. However, there exists a mean degree in the network, which is the average number of edges per node. The most famous model for random networks is the *Erdős–Rényi* model, where edges are formed between pairs of nodes with a fixed probability, independently of other edges. This randomness leads to a Poisson degree distribution. Random networks are typically used to model systems where interactions or connections occur randomly or without any specific underlying structure, such as certain communication or social networks. The degree distribution for this model can be expressed as

$$P(k) = \frac{\lambda^k e^{-\lambda}}{k!},$$

  where λ is the average degree.

- **Scale-free Networks**: These networks follow a power-law degree distribution, meaning the number of edges per node scales according to a power law. In other words, there are a few nodes with a very large number of edges, while most nodes have only a few

edges, creating a heavy-tailed distribution. Scale-free networks are often used to model systems where some nodes are highly influential or connected, while most others have relatively few connections. Examples include the internet, citation networks, and many biological systems. A common model for scale-free networks is the *Barabási–Albert* model, which incorporates the principle of preferential attachment, where new nodes are more likely to connect to existing nodes that already have a high degree. The degree distribution for this model can be expressed as

$$P(k) \propto k^{-\gamma},$$

where γ is a positive exponent typically between 2 and 3.

Figure 1 displays the degree distribution of a random network with an average degree of 4 and 100 nodes. Figure 2 shows the degree distribution of a scale-free network generated using the Barabási–Albert model with 100 nodes and $\gamma = 2.5$, plotted on a log-log scale.

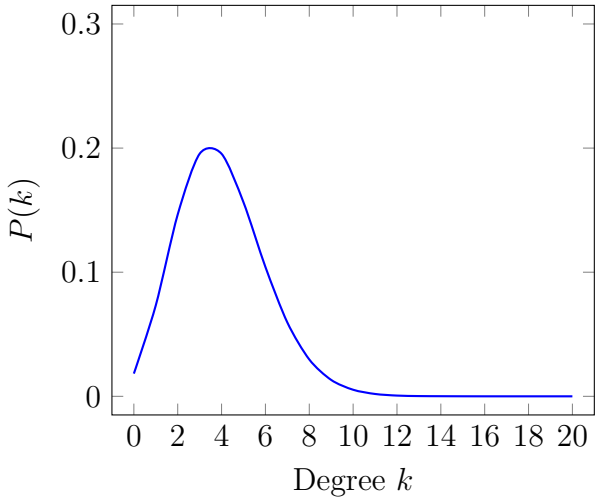

Figure 1: Poisson degree distribution of an Erdős–Rényi (ER) network with average degree $\langle k \rangle = 4$.

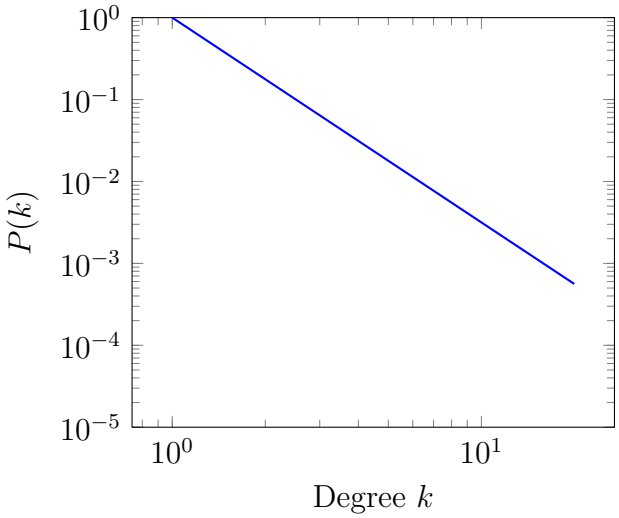

Figure 2: Power-law degree distribution of a Scale-Free (SF) network with exponent $\gamma = 2.5$ (log-log scale).

These topological differences significantly impact the behavior of dynamical processes within the network. For instance, consider a simple contagion model where an *infected* node A can spread its condition to other nodes in the network who are not *infected*. Nodes who are not *infected* are called *susceptible*. Depending on the model parameters and the available connections of node A, we can distinguish between various contagion models. For example, in the Susceptible-Infected (SI) model, infected nodes remain infected forever, whereas the Susceptible-Infected-Susceptible (SIS) model allows infected nodes to return to a susceptible state. Regardless of the model we choose, the dynamics will vary depending on the network's topology.

In a random network, the disease would spread at the same rate across all nodes, irrespective of their positions. However, in a scale-free network, if the disease reaches a highly connected

node (a *hub*), it will spread much faster. Thus, the rate of infection is highly dependent on the network structure.

Systemic risk can be studied via this models. In fact, default propagation has also been studied and modeled as a contagion model following the SIS structure applied to financial networks. See [4] for details.

Network theory is vital, as it provides the different scenarios on which we will apply our model.

3.2 Dynamical Systems Theory

Dynamical systems theory provides valuable insights into the behavior of the model. It allows us to study the fixed points and the orbits of the system.

A good reference for discrete dynamical systems is provided by [5], which presents a thorough overview of the key concept in this area.

3.2.1 Real Discrete Dynamical Systems

We now focus on **real discrete dynamical systems**, that is, systems defined on the real line. Given a function

$$f : \mathbb{R} \to \mathbb{R},$$

we define the discrete dynamical system through iteration:

$$x_0 = \text{initial value}, \quad x_{n+1} = f(x_n) \quad \forall n,$$

This process generates a sequence $\{x_n\}_{n=0}^{\infty}$, and our goal is to understand the nature of this sequence, whether it converges to a fixed point, diverge, exhibit chaotic behavior, or settle into a periodic cycle.

More generally, a discrete dynamical system can be defined as

$$f : M \to M,$$

where M is a metric space and f is a continuous function. In our case, however, we restrict the setting to $M = \mathbb{R}$, which suffices to capture the essential behavior of the system we are modeling.

3.2.2 Fixed Points and Periodic Orbits

In dynamical systems, fixed points refer to the points where an arbitrary function $y = f(x)$ reaches a state such that $x = f(x)$, that is, $y = x$. Fixed points can be computed by solving for x in the equation $f(x) = x$.

The function $f(x) = x^2$ intersects the line $y = x$ at the fixed point $x = 1$ and $x = 0$. The diagram in Figure 3 illustrates the fixed points where the function intersects the line.

Periodic orbits exhibit a similar behavior to fixed points but differ in that the function $f(x)$ must be applied multiple times. A periodic orbit occurs when $x = f^n(x)$, where $f^n(x) = f(f^{n-1}(x))$. This describes a periodic orbit of period k, where k is an integer such that $x = f^k(x)$.

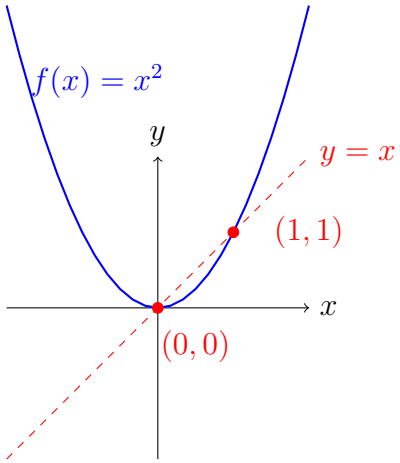

Figure 3: Diagram illustrating the fixed points of the function $f(x) = x^2$ where it intersects the line $y = x$. The red points represent the fixed points.

For the periodic orbit of period $k = 2$, applying $f(x)$ twice will bring us back to the original point. An example of a period-2 orbit is given by the function

$$f(x) = x^2 - 1.$$

Let us consider two points: $x = 0$ and $x = -1$. We observe:

$$f(0) = -1, \quad f(-1) = 0.$$

Applying the function twice:

$$f(f(0)) = f(-1) = 0 = x_0, \quad f(f(-1)) = f(0) = -1 = x_1.$$

This is a valid period-2 orbit: the function cycles between two distinct points. Figure 4 illustrates the periodic orbit where the function intersects the line.

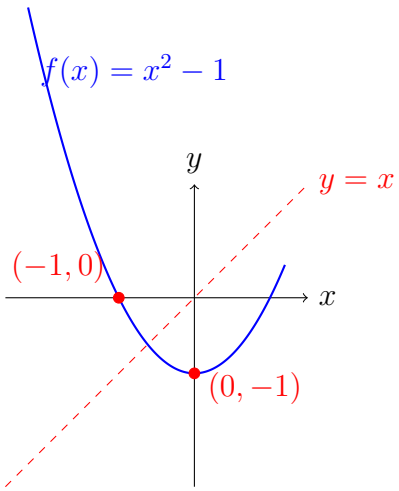

Figure 4: An orbit of period 2 for the function $f(x) = x^2 - 1$. The red points 0 and -1 form a two-step cycle: $f(0) = -1$ and $f(-1) = 0$.

3.2.3 Attractors, Repellers, and Neutral Points

Fixed points and periodic orbits can be classified as either attractors, repellers, or neutral.

1. **Attractor**: If a fixed point or periodic orbit is an attractor, it means that for values of x near the fixed point or periodic orbit, the system will move towards that fixed point or periodic orbit as the iterations of the function $f(x)$ are applied. Mathematically, if $x = x^*$ is an attractor, then for initial points x_0 close to x^*, the sequence $x_{n+1} = f(x_n)$ will converge to x^* as $n \to \infty$. The function pulls values toward the fixed point. This happens when $|f'(x^*)| < 1$.

2. **Repeller**: If a fixed point or orbit is a repeller, it means that for values of x near the fixed point or orbit, the system will move away from that fixed point or orbit as the iterations of the function $f(x)$ are applied. Mathematically, if $x = x^*$ is a repeller, then for initial points x_0 close to x^*, the sequence $x_{n+1} = f(x_n)$ will diverge away from x^*. The function pushes values away from the fixed point. This happens when $|f'(x^*)| > 1$.

3. **Neutral**: If the fixed point or orbit is neither an attractor nor a repeller, we classify it as neutral. In this case, the system does not necessarily move towards or away from the fixed point or orbit. Instead, points remain in a stable oscillation or periodic orbit around the fixed point or orbit. Mathematically, this occurs when $|f'(x^*)| = 1$.

3.2.4 Classification Using Derivative

To determine whether a fixed point or periodic orbit is an attractor, repeller, or neutral, we compute the derivative of $f(x)$ at the fixed point or along the periodic orbit. Specifically, we

evaluate $|(f^k)'(x)|$ either at a fixed point when $k = 1$, or at a point on a period-k orbit when $k > 1$.

- If $|(f^k)'(x)| < 1$, the fixed point or periodic orbit is an **attractor**.
- If $|(f^k)'(x)| > 1$, the fixed point or periodic orbit is a **repeller**.
- If $|(f^k)'(x)| = 1$, the fixed point or periodic orbit is **neutral**.

Dynamic systems theory is fundamental, as we aim to represent our model as a discrete dynamic system. By applying this framework, we can derive valuable insights from the system's behavior.

3.2.5 Iterated Function Systems (IFS)

Iterated Function Systems (IFS) are a class of mathematical models used to generate fractals and describe self-similar structures. A reference for iterated function systems (IFS) is provided by [6]. An IFS consists of a finite set of mappings $f_1, f_2, \ldots, f_n$ on a complete metric space (X, d). Specifically, a map f is called contractive if there exists a constant $c \in [0, 1)$ such that for all $x, y \in X$,

$$d(f(x), f(y)) \leq c \cdot d(x, y).$$

The central result of IFS theory, known as the Banach Fixed-Point Theorem, states that if the system is composed of contractive maps, there exists a unique fixed point x^* such that:

$$f_i(x^*) = x^* \quad \text{for all} \quad i = 1, 2, \ldots, n.$$

This fixed point, or attractor, represents the self-similar structure generated by the system after an infinite number of iterations.

3.3 Probability Theory

Probability theory is fundamental for the study of financial models, as the market is an stochastic process that is constantly subject to small fluctuations and noise, which we refer to as **market volatility**. This volatility encompasses all the uncontrollable changes in the market, such as shifts in investor sentiment, geopolitical events, or sudden economic shocks.

A classic reference for probability theory is provided by [7], which offers a comprehensive introduction to the principles and applications of probability in various contexts.

To understand and model these fluctuations, we need to recall some key concepts from probability theory.

For clearer notation, we first define what a probability space and a probability measure are:

Definition 1. *A **probability space** is a triplet $(\Omega, \mathcal{F}, \mathbb{P})$, where:*

- *Ω is the sample space, the set of all possible outcomes.*

- $\mathcal{F}$ *is a sigma-algebra, a collection of subsets of* Ω *that includes* Ω *itself and is closed under complementation and countable unions.*
- $\mathbb{P}$ *is the probability measure, a function* $\mathbb{P} : \mathcal{F} \to [0,1]$ *that assigns probabilities to events in* $\mathcal{F}$ *such that:*
  - $\mathbb{P}(\Omega) = 1$,
  - $\mathbb{P}(A) \geq 0$ *for all* $A \in \mathcal{F}$,
  - *If* $A_1, A_2, \ldots$ *are mutually exclusive events in* $\mathcal{F}$, *then* $\mathbb{P}\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} \mathbb{P}(A_i)$.

The probability measure $\mathbb{P}$ quantifies the likelihood of various events, ensuring the consistency and mathematical foundation of probability theory.

Definition 2. *A random variable is a function* $X : \Omega \to \mathbb{R}$ *that assigns a real number to each outcome in the sample space* Ω.

Random variables can be discrete or continuous, depending on the framework we are willing to study. In our case, the liquidity of a firm will act as a random variable.

In probability theory, we use the abbreviation *a.s.*, which stands for *almost surely*. This means that if an event $A \in \Omega$ has probability $\mathbb{P}(A) = 1$, then the event occurs with probability 1, or almost certainly.

For example, consider choosing a real number uniformly at random from the interval $[0,1]$. The set of irrational numbers in $[0,1]$ has Lebesgue measure 1, and the set of rational numbers has measure 0. Therefore, the event that the chosen number is irrational occurs almost surely:

$$\mathbb{P}(\text{the number is irrational}) = 1.$$

Even though it is possible to choose a rational number (like $\frac{1}{2}$), the probability of that happening is 0. Hence, the number chosen is almost surely irrational.

The set of rational numbers is countably infinite, which means it can be listed as a sequence $(\ q_1, q_2, q_3, \ldots)$. In the context of Lebesgue measure, any countable set has measure zero. Therefore, the probability of randomly selecting a rational number from the interval $[0,1]$ is zero.

The set of irrational numbers in $[0,1]$ is the complement of the rationals within that interval, $\mathbb{I} = [0,1] \setminus \mathbb{Q}$. Since the total measure (or probability) of $[0,1]$ is 1, we can write:

$$\mathbb{P}(\mathbb{I}) = \mathbb{P}([0,1] \setminus \mathbb{Q}) = \mathbb{P}([0,1]) - \mathbb{P}(\mathbb{Q}) = 1 - 0 = 1.$$

Hence, the probability of choosing an irrational number from $[0, 1]$ is 1 and it occurs almost surely.

3.3.1 Stochastic Processes

The liquidity of a node is not only a random variable, but it is also a stochastic process, as the liquidity evolves over time. That is, the random variable X is a function of time, denoted as $X(t)$ or X_t.

Definition 3. *Formally, a stochastic process is a family of random variables $\{X_t\}_{t\in T}$, where:*

- *T is the index set, typically representing time, which can be discrete ($t = 0, 1, 2, \ldots$) or continuous ($t \in [0, \infty)$).*
- *For each $t \in T$, X_t is a random variable, often representing the state or value of the system at time t.*

Given a stochastic process, if we are able to observe the values of the process from $t = 0$ to $t = n$, we can form a sigma-algebra $\sigma(X_n, X_{n-1}, \ldots, X_1, X_0)$. This sigma-algebra is called a filtration $\mathcal{F}_n$, and it is important because it encapsulates all the information about the random variables up to time n.

Stochastic processes are widely observed in our daily lives, some examples of them are the spread of diseases, stock prices or even weather patterns.

3.3.2 Martingales

Martingales are closely related to stochastic processes. In fact, martingales are stochastic processes that fulfill specific characteristics.

Definition 4. *By definition, $\{X_t\}_{t\in T}$ is a martingale with respect to a filtration $\{\mathcal{F}_t\}_{t\in T}$ if it satisfies the following conditions for all $t \geq 0$:*

- ***Integrability:*** $\mathbb{E}[|X_t|] < \infty$,
- ***Martingale Property:*** $\mathbb{E}[X_{t+1} \mid \mathcal{F}_t] = X_t$ *a.s.*

This means that, given all the past information up to time t, the expected value of the process at the next time step is equal to its current value. In other words, there exists no predictable trend in the process.

One of the most well-known examples of a stochastic process that satisfies the martingale property is the **Brownian motion**. Brownian motion models a process in which the probabilities of moving in either direction are equal, and the magnitude of each step is symmetric. Under these conditions, the expected value of the process at the next time step equals the current value, satisfying the martingale condition:

Suppose a stochastic process $\{X_t\}_{t\in T}$ such that at each time step the process moves by $+1$ or -1 with equal probability:

$$X_{t+1} = \begin{cases} X_t + 1 & \text{with probability } \frac{1}{2} \\ X_t - 1 & \text{with probability } \frac{1}{2} \end{cases} \quad \Rightarrow \quad \mathbb{E}[X_{t+1} \mid \mathcal{F}_t] = \frac{1}{2}(X_t + 1) + \frac{1}{2}(X_t - 1) = X_t.$$

Hence, this process is a martingale.

However, variations of this process can result in upward or downward expected movement. For instance, suppose a process where the step sizes are asymmetric:

$$X_{t+1} = \begin{cases} X_t + 2 & \text{with probability } \frac{1}{2} \\ X_t - 1 & \text{with probability } \frac{1}{2} \end{cases} \quad \Rightarrow \quad \mathbb{E}[X_{t+1} \mid \mathcal{F}_t] = \frac{1}{2}(X_t + 2) + \frac{1}{2}(X_t - 1) = X_t + \frac{1}{2}.$$

Since the conditional expectation is strictly greater than X_t, this process is a **submartingale**, representing a system that tends to increase over time in expectation.

On the other hand, consider a process defined by:

$$X_{t+1} = \begin{cases} X_t + 1 & \text{with probability } \frac{1}{2} \\ X_t - 2 & \text{with probability } \frac{1}{2} \end{cases} \quad \Rightarrow \quad \mathbb{E}[X_{t+1} \mid \mathcal{F}_t] = \frac{1}{2}(X_t + 1) + \frac{1}{2}(X_t - 2) = X_t - \frac{1}{2}.$$

Here, the conditional expectation is less than X_t, indicating a **supermartingale**, which models a system that tends to decrease in expectation over time.

These lead to the formal definitions of related concepts:

- A **submartingale** is a stochastic process satisfying $\mathbb{E}[X_{t+1} \mid \mathcal{F}_t] \geq X_t$. This implies the process tends to increase in expectation over time.
- A **supermartingale** satisfies $\mathbb{E}[X_{t+1} \mid \mathcal{F}_t] \leq X_t$, indicating a tendency to decrease in expectation.
- A **martingale** satisfies the condition $\mathbb{E}[X_{t+1} \mid \mathcal{F}_t] = X_t$, meaning the process is "fair" and exhibits no expected gain or loss.

Thus, martingales, submartingales, and supermartingales provide a flexible and powerful framework for modeling neutral, favorable, and unfavorable processes evolving over time.

3.3.3 Normal Noise

We recall the concept of *normal noise* and its use in modeling financial systems. The expression of the normal distribution is given by:

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}}\, e^{-\frac{(x-\mu)^2}{2\sigma^2}}.$$

Since the market is constantly exposed to volatility, it is natural to expect that the liquidity of any firm is also subject to these fluctuations. To capture how such variations affect liquidity, we often model them using normal noise. There are two main types: **additive noise** and **multiplicative noise**.

Additive Noise. This corresponds to adding a small fluctuation to the process, typically modeled by a random variable $\eta \sim \mathcal{N}(0, \sigma^2)$, a normal distribution with mean zero and variance σ^2. We denote a sample drawn from this distribution as $\eta(0, \sigma)$. If we denote the liquidity of a firm at time t as $L(t)$, the additive noise model is expressed as:

$$L(t+1) = L(t) + \eta(0, \sigma).$$

This can also be interpreted as:

$$L(t+1) = \eta(L(t), \sigma).$$

Since $\mathbb{E}[\eta(0, \sigma)] = 0$, the conditioned expected value of $L(t+1)$ is:

$$\mathbb{E}[L(t+1)|\{L(t), L(t-1), ..., L(0)\}] = L(t).$$

which satisfies the definition of a martingale.

Multiplicative Noise. In this case, the random fluctuation is proportional to the current value of the process. Again, let $\eta \sim \mathcal{N}(0, \sigma^2)$, then:

$$L(t+1) = L(t)(1 + \eta(0, \sigma)).$$

This can also be expressed as:

$$L(t+1) = \eta(L(t), \sigma L(t)).$$

The conditioned expected value, if $L(t)$ is independent of η, is:

$$\mathbb{E}[L(t+1)|\{L(t), L(t-1), ..., L(0)\}] = \mathbb{E}[L(t)(1 + \eta(0, \sigma))|\{L(t), L(t-1), ..., L(0)\}] = L(t).$$

so this model also satisfies the martingale condition. However, the variance behaves differently. While the additive noise has a constant variance σ, the multiplicative noise has a variance proportional to the current liquidity $\sigma \cdot L(t)$.

The choice between additive and multiplicative noise has significant consequences for the dynamics of the model. For instance, when $0 < L(t) < 1$, multiplicative noise leads to a distribution for $L(t+1)$ that is narrower and more peaked, while additive noise results in a distribution with constant spread centered around $L(t)$. Thus, depending on the desired characteristics of the model, one must choose the appropriate noise structure accordingly.

3.3.4 Z-Score Standardization

The Z-Score is a statistical measure defined as

$$Z = \frac{X - \mu}{\sigma_{std}},$$

where X is the value of the variable, μ is the mean, and σ_{std} is the standard deviation of the distribution. It represents how many standard deviations X is away from the mean. Using the cumulative distribution function (CDF) of the standard normal distribution, the Z-Score allows us to calculate probabilities. For instance, a Z-Score of 0 corresponds to the mean, and the CDF at this point is 0.5, indicating a 50% probability that the variable is less than or equal to the mean.

4 Modeling Financial Distress Propagation

Financial distress propagation is one of the core challenges in the study of systemic risk. It can be mathematically modeled in various ways, depending on the desired level of accuracy and complexity. These approaches range from detailed models that account for the specific assets of individual firms to more abstract mean-field models applied across the entire network.

Models also differ in terms of what aspects of the network they aim to capture. The interpretation of nodes and connections may vary, although it is common for nodes to represent firms or assets, and edges to reflect relationships such as inter-firm debt, portfolio overlap, or asset dependencies. Different models based on different interpretations of nodes and connections are given in [8] [9].

It is important to emphasize that there is no universally best model, appropriateness depends on the goals of the analysis and the required level of detail.

In our case, we adopt a **customer–supplier model** implemented on a synthetic, heterogeneous network of **N** nodes. The distribution of these nodes will be defined later in order to explore different scenarios.

4.1 Modeling Financial Distress Propagation: The Model

The model presented in this work focuses directly on the liquidity of the various firms within the network. The liquidity of a firm represents the amount of money an entity has available to carry out transactions. We assume that a firm's liquidity is a time-dependent variable that is not directly observable or easily quantifiable. It is influenced by two primary factors: market volatility and the flow of funds within the collection and payment network. Volatility can directly impact a firm's ability to make payments, depending on its exposure to market fluctuations.

We adopt the model introduced in [10], which is designed to capture how the default of a single node can propagate to its neighbors and, potentially, to the entire network.

Thus, we model the evolution of the liquidity of node i at time t, denoted by $L_i(t)$, as a discrete stochastic equation:

$$L_i(t+1) = L_i(t)(1+\eta(0,\sigma)) + \sum_{j \in N_i} w_{ji} P_j^1(t) H(L_j(t) - w_{ji}) - \sum_{j \in N_i} w_{ij} P_j^2(t) H(L_i(t) - w_{ij}). \quad (1)$$

However, we introduce a slight modification to the original framework. Specifically, we consider each Bernoulli distribution P to be independent across nodes, redefining it as P^i for each firm i in the network. Also, η follows a normal distribution with mean 0 and standard deviation σ, accounting for the market volatility. The money that firm i owes to firm j is expressed by the weight w_{ij}. $P^k(t)$ on $k \in \{1, 2\}$ follows a Bernoulli distribution with

parameter p, which accounts for the network's activation parameter. One can think of p as the market's operation parameter, since when the market is open, money exchange will occur, given that:

$$P^k(t) = \begin{cases} 1 & \text{if } P^k < p. \\ 0 & \text{if } P^k \geq p. \end{cases}$$

The Heaviside function $H(x)$ is defined by:

$$H(x) = \begin{cases} 0 & \text{if } x < 0. \\ 1 & \text{if } x \geq 0. \end{cases}$$

When applied at the point $x - a$ for $a \geq 0$, the Heaviside ensures a payment is only made if there is enough liquidity. Figure 5 shows the Heaviside step function.

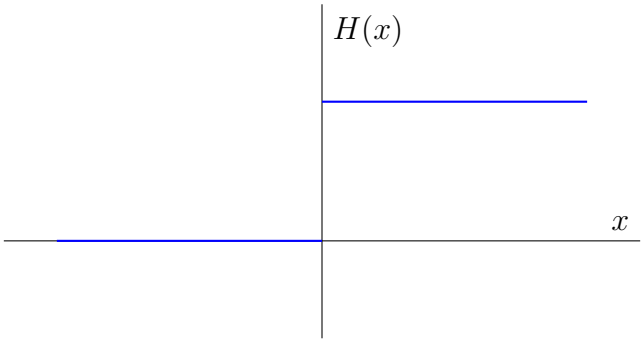

Figure 5: The Heaviside step function $H(x)$.

This prevents firms from making payments that would bankrupt them.

4.1.1 The Model : Examples

To clarify the proposed model, let us manually perform a one-step simulation over a network with only two nodes. Consider the graph illustrated in Figure 6:

Let $L_i(0) = 8$ for $i \in \{1, 2\}$, meaning each firm starts off with a liquidity value of 8. Starting with node A, we take Equation (1) and substitute:

$$L_A(1) = L_A(0)\left(1 + \eta(0, \sigma)\right) + w_{BA}P_B^1(t)H\left(L_B(t) - w_{BA}\right) - w_{AB}P_B^2(t)H\left(L_A(t) - w_{AB}\right).$$

Substituting with the given values:

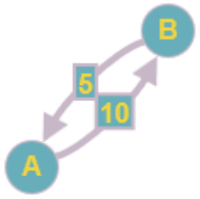

Figure 6: Diagram illustrating a 2-node graph with directed connections.

$$L_A(1) = 8\,(1 + \eta(0, \sigma)) + 5P_B^1(1)H(8-5) - 10P_B^2(1)H(8-10).$$

Recalling the definition of the Heaviside function, only the term that sums survives:

$$L_A(1) = 8\,(1 + \eta(0, \sigma)) + 5P_B^1(1).$$

Given the randomness of the normal and Bernoulli distributions that appear in the final expression, without a proper simulation, we can only provide an approximate range for $L_A(1)$:

First we fix $\sigma = 0.1$ for medium volatility of the market. The term $8(1 + \eta(0, 0.1))$ can range between $[7.2, 8.8]$, we now add the term $5P_B^1$. Since $P_B^1 = \{0, 1\}$:

$$L_A(1) \in [7.2, 13.8].$$

The same process can now be performed for node B. This process will be repeated for as many time steps as we choose to simulate. It is important to note that L_B is still operating at $t = 0$. Therefore, the new value of L_A that we have just computed is not yet available for B to access.

One potential issue in the problem setup is illustrated using a star graph with 5 nodes, as shown in Figure 7, $G = (V, E)$, where a central node is connected to all other nodes. Let the initial liquidity of this central node, A, be $L_A(0) = 10$. We assume that node A owes a quantity of 3 to each of the other nodes, that is, $w_{Ai} = 3\ \forall i \in V - A$.

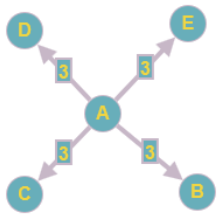

Figure 7: Diagram illustrating a 5-node star graph with directed connections.

In a one-step simulation with parameter $p = 1$, meaning the network is fully activated, we observe that $L_A(1) < 0$. This is because the Heaviside function evaluates the liquidity of node A at time t, while updating $L_A(t+1)$. The function does not account for changes in liquidity at each step but only considers the previous time step.

The Heaviside function continually evaluates $H(10-3)$ instead of accounting for the successive updates, such as $H(10) \xrightarrow{w_{AB}} H(10-3) \xrightarrow{w_{AC}} H(7-3) \xrightarrow{w_{AD}} H(4-3) \xrightarrow{w_{AE}} H(1-3)$. Where the arrow with w_{ij} stands for the weight update or transformation applied to the argument of the Heaviside function, which iteratively changes the value inside the function, reflecting the evolution of the system's state over time. As a result, node A will always appear to have sufficient liquidity at time t, as $L_A(t) > w_{Ai}\ \forall i \in V - A$, failing to reflect the actual change in liquidity at time $t+1$.

At first glance, this might appear to be a fatal flaw in the model. However, this issue is explainable, given the large number of nodes we are working with in a real network. Despite node A going bankrupt at time $t+1$, it is likely that this node will be saved by its neighbors. If not, its bankruptcy will trigger the propagation of financial distress, which is the phenomenon we aim to model in this study. Hence, the model seems to be correct.

4.1.2 The Model: Simulations

To gain a clearer understanding of the model's behavior, we simulate the dynamical process on two distinct network topologies. The first is a synthetic random network, and the second is a scale-free network generated using the Barabási–Albert model.

All simulations are performed using the following parameters configuration: $p = 0.5$ and $\sigma = 0.1$. The initial liquidity of each node is randomly assigned within the interval $[1000, 10000]$. The weights of the connections between nodes are also randomly initialized, uniformly distributed in the interval $[200, 5000]$.

The process is simulated 1000 times over 1000 time steps. For each time step, we compute the mean liquidity value across all simulations for every node. These mean values are then used to plot the temporal evolution of node liquidity, allowing us to observe how liquidity levels evolve and potentially stabilize or diverge over time under different network structures.

For the random network, we set each node to have a probability of 0.25 of forming a directed

connection with any other node in the network. For the scale-free network, generated using the Barabási–Albert model, the average number of connections per node is set to 3. The random network is a directed graph, while the Barabási–Albert network is undirected.

The following Figure 8 displays the evolution of $L_i(t)$ for 20 nodes interconnected via a random network.

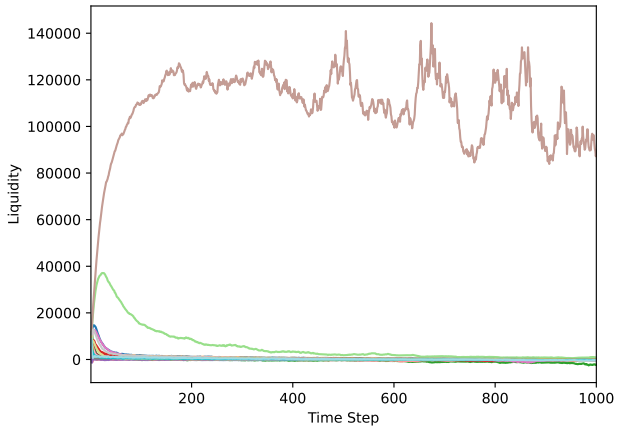

Figure 8: Evolution of $L_i(t)$ for 20 nodes interconnected via a random network topology. Each node is represented by a distinct color.

Figure 8 shows how the dynamics of a single node diverges completely, fluctuating over time. While the remaining 19 nodes converge to what appears to be a stable point around $t = 75$. The probability of default for this specific case is $\mathbb{P}(L_i(1000) < 0) = 0.4401$. This probability has been computed by summing over all the nodes that ended with negative liquidity and dividing it by the total number of nodes, for each simulation.

Figure 9 displays the evolution of $L_i(t)$ for 20 nodes interconnected via a Barabási–Albert network.

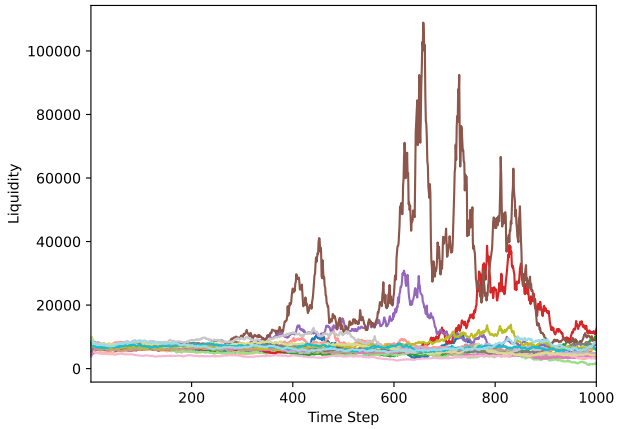

Figure 9: Evolution of $L_i(t)$ for 20 nodes interconnected via a Barabási-Albert network topology. Each node is represented by a distinct color.

Figure 9 illustrates the slow fluctuations of liquidity in the Barabási–Albert structure, mainly due to the small number of nodes with many links. The distinct topology of the Barabási–Albert network causes liquidity fluctuations to appear later than in the random network. At time $t = 400$, few nodes begin to fluctuate while the others remain stable. The probability of default for this specific case, calculated using the same procedure as previously described, is $\mathbb{P}(L_i(1000) < 0) = 0.387$, which is significantly lower than the one computed for the random network.

Understanding why the dynamics of the process differ across network topologies primarily depends on the structure of the networks themselves. However, analyzing the complete model proves to be too complex for deriving general insights. Therefore, to better understand the overall behavior, we reduce the problem to a simplified version of the dynamics.

5 The 2-Node Problem

In this section, we study the 2-Node Problem, which is a simplified approximation of the entire network reduced to just two nodes. To our knowledge this study is entirely novel, meaning it has not been conducted before and constitutes original research. Despite having thoroughly reviewed every step of the process, the exploratory nature of the results means they may be subject to errors introduced during the research process.

For calculus simplicity, we focus on a 2-node network, considering node A and node B. We will also consider node B to be the market node. The market node accounts for the entire network besides node A. Thus, given a graph G:

$$G = (V, E) : A \in V, B = V - A.$$

In fact, despite the name of this problem being the 2-node problem, we are only interested in the study of node A. Since node B works as the market node, we will assume that the market never goes into bankruptcy.

In this scenario, our model Equation (1) gets modified to:

$$L_i(t+1) = L_i(t)(1 + \eta(0, \sigma)) + w_{ji}P_j^1(t) - w_{ij}P_j^2(t). \tag{2}$$

Notice that both Heaviside functions have disappeared; since we only have two nodes, the Heaviside function would make our firms never go into bankruptcy, which is not what we aim to study. Also, the summations have been simplified, since we only have 2 nodes. For simplicity, let us define $w_{AB} = b$ and $w_{BA} = c$. Also, since we are interested only in node A, $L_i = L_A$. Figure 10 displays the 2-node problem graph.

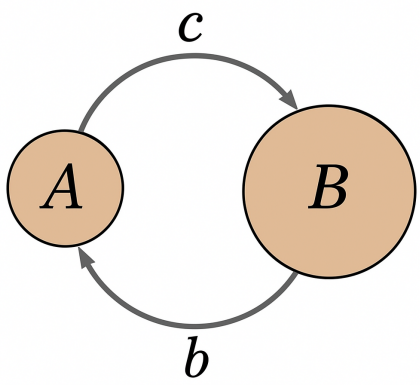

Figure 10: Diagram illustrating 2-Node problem graph. Where B is the market node, A is the node we are studying, and b, c are the weights between them

This new formulation is particularly interesting, as it allows us to express the model in a much simpler way, while also enabling us to discretize it into four different cases.

$$L_A(t+1) = \begin{cases} L_A(t)(1+\eta(0,\sigma)) & \text{if } P^1, P^2 > p. \\ L_A(t)(1+\eta(0,\sigma)) + b & \text{if } P^2 \geq p, P^1 < p. \\ L_A(t)(1+\eta(0,\sigma)) - c & \text{if } P^2 < p, P^1 \geq p. \\ L_A(t)(1+\eta(0,\sigma)) + b - c & \text{if } P^1, P^2 < p. \end{cases} \tag{3}$$

Our new simplified model is very convenient since it only depends on 4 parameters:

- σ: Represents the volatility and ranges from $(0, \infty)$.
- p: Represents the network's activation parameter and ranges from $(0, 1)$.
- b: Represents the amount of liquidity node A owes to market node B, and ranges from $(0, \infty)$.
- c: Represents the amount of liquidity market node B owes to node A, and ranges from $(0, \infty)$.

5.1 2-Node Problem: Simulations

In this section, we aim to study the parameters that most influence the results of the simulations. To do so, we simulate the 2-node problem using two different groups of parameters. The first group includes b and c, while the second group includes the remaining parameters, σ and p.

In the first set of simulations, we fix the second group of parameters and explore how changing the relationship between the parameters in the first group affects the system. In the subsequent simulations, we fix an equality relation between the parameters in the first group in order to focus on the variation of the parameters in the second group.

To ensure reliable results, we first define the initial liquidity of node A for every simulation as $L_A(0) = 0.1$.

All simulations were computed using Python programs. The main code of the programs are available in [11].

Based on theory, the most interesting simulations will involve the following parameter configurations, as they represent the lower, middle, and upper ranges of our parameters:

Activation Parameter p	**Volatility** σ	**Liquidity Relation** b **vs** c
$p = 0.5$	$\sigma = 0.05$	$b = c$
$p \to 1$	$\sigma = 0.1$	$b < c$
$p \to 0$	$\sigma = 0.3$	$b > c$

Table 1: Parameter combinations used in the simplified model.

Firstly, we simulate the liquidity relation b vs c. For this, we fix $\sigma = 0.1$ and $p = 0.5$, as they represent the mid-range values for the other two parameters.

For each simulation we evaluate the system 1000 times, we conserve the final value given and we repeat this process 1000 times more. We then plot the density of the final values of L_A. The case $b = c$ accounts for $b = 0.1$, $c = 0.1$. The seconds case, $b > c$ accounts for $b = 0.15$, $c = 0.1$. For the final case $b < c$, the values taken are $b = 0.1$, $c = 0.15$:

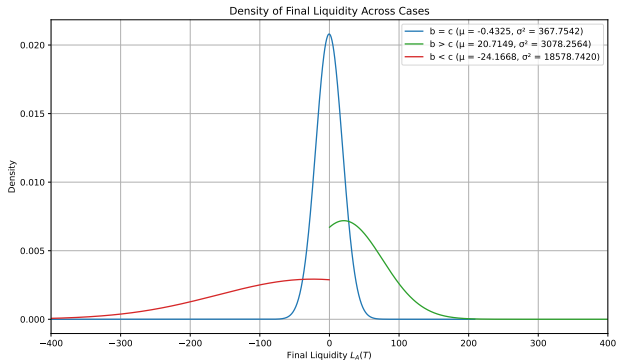

Figure 11: Diagram illustrating the comparison of the final values of L_A under different configurations of b and c. Blue accounts for the case $b = c$, red for the case $b < c$, and green for the case $b > c$.

The results shown in Figure 11 are intuitive and align with expectations. When firm A owes more to the market node B than B owes to A ($b < c$), the density of L_A peaks at approximately -22.905 and exhibits a heavy tail extending toward the negative domain. Conversely, when B owes more to A than A owes to B ($b > c$), the density peaks around 24.027, with its tail extending into the positive domain.

The most interesting scenario arises when A owes exactly the same amount to B as B owes to A ($b = c$). In this symmetric case, the density is expected to peak at the initial value $L_A(0)$, with minimal dispersion. However, this is not observed in our simulations due to the multiplicative stochastic nature of the process. Accurately estimating the mean in such a system requires a very large number of simulations, which may exceed practical computational limits.

A practical alternative is to compute the median, which represents the central value across all simulations. For multiplicative stochastic processes, the median often provides a more

robust and representative measure of central tendency than the mean. Moreover, increasing or decreasing the number of time steps primarily affects the mean and variance: the mean evolves as dictated by the process dynamics, while the variance grows rapidly—an expected consequence of the system's multiplicative structure.

We can then compute use the expected value of $L_A(t+1)$ to estimate the mean of L_A for each case:

$$\mathbb{E}(L_A(t+1)) = \mathbb{E}(L_A(t)) + p(b-c) = \mathbb{E}(L_A(t-1)) + 2p(b-c) = \cdots = L_A(0) + tp(b-c).$$

Thus, we have the following cases:

- If $b > c$, then $E(L_A(t+1)) \to \infty$ as $t \to \infty$.
- If $c > b$, then $E(L_A(t+1)) \to -\infty$ as $t \to \infty$.
- If $b = c$, then $E(L_A(t+1)) = L_A(0)$ for all t.

Although we have not yet explored the other variations of the parameters, we will focus on the case $b = c$, as it provides us with a balanced distribution of values on both the negative and positive sides. This allows for more interesting results to analyze. In fact, varying the parameters b and c will likely only affect the movement of the mean of L_A over time.

For the following simulations, we vary the market volatility parameter σ. We analyze three types of markets: a low-volatility market with $\sigma = 0.05$, a medium-volatility market with $\sigma = 0.1$, and a high-volatility market with $\sigma = 0.3$. Additionally, we vary the network activation parameter p, considering the values $p = 0.05$, $p = 0.5$, and $p = 0.95$.

We run 100000 simulations for 200 time steps. As previously explained, a large number of simulations is required to gain precise insights into the system's dynamics. The number of time steps, however, does not significantly affect the qualitative behavior of the process. Therefore, we decide to perform a high number of simulations with a reduced number of time steps compared to earlier experiments, in order to achieve faster and more computationally efficient results. We also assume $b = c$, since—as observed earlier—this condition allows liquidity values to be distributed on both the negative and positive sides. The y-axes represent the density of the final step of the liquidity, while the x-axes represent the value for the final step of the liquidity.

The plots for $p = 0.05$ and different volatility σ are shown in Figure 12:

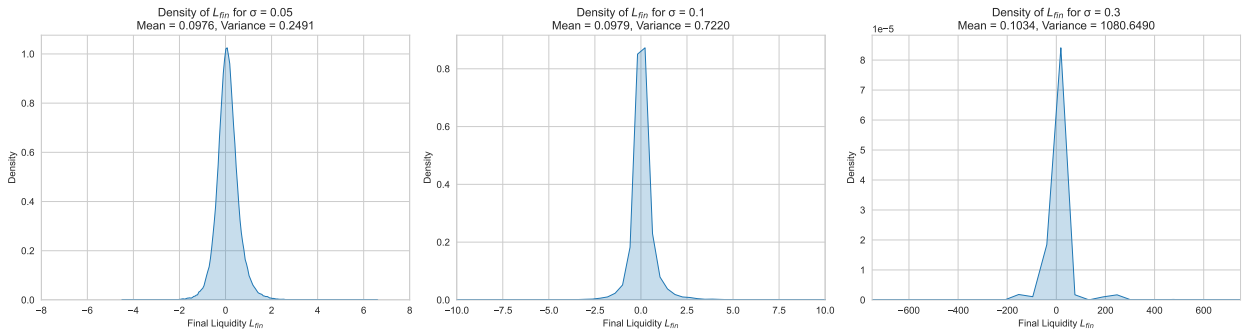

Figure 12: Comparison of the density of final values of L_A for different values of σ, with $p = 0.05$.

The plots for $p = 0.5$ and different volatility σ are shown in Figure 13:

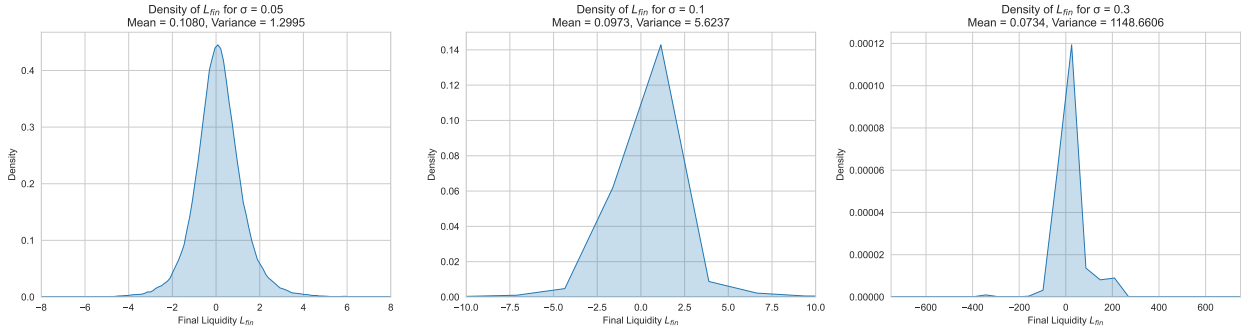

Figure 13: Comparison of the density of final values of L_A for different values of σ, with $p = 0.5$.

The plots for $p = 0.95$ and different volatility σ are shown in Figure 14:

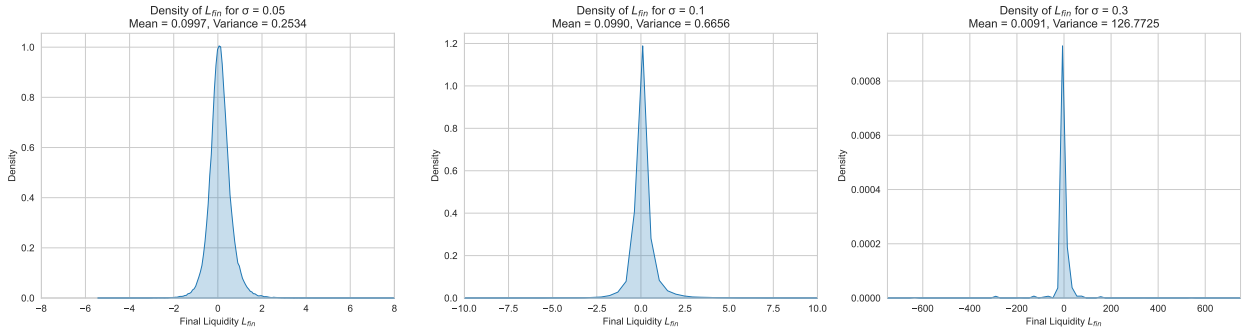

Figure 14: Comparison of the final values of L_A for different values of σ, with $p = 0.95$.

We first observe that the mean remains stable for low and medium market volatility for all values of the activation parameter p. However, for high volatility $\sigma = 0.3$, the mean shifts significantly. The explanation for this behavior can be found in basic stochastic process

theory. Given the multiplicative noise nature of the system, high volatility leads to larger random fluctuations at each time step, which accumulate over the course of the simulation. As a result, the system becomes more sensitive to these fluctuations, causing the final value of L_A to deviate substantially from its expected value.

We also observe that the variance increases as σ increases, with the variance reaching its maximum for $p = 0.5$ and decreasing for other values of p. This phenomenon is explained further below in section 5.2.2.

An important point to mention is that, if there were no shifts introduced in our simulations (no $+b$ or $-c$ additions), increasing σ would cause the density of the final values to become narrower, and the variance would be smaller. This behavior arises from the multiplicative nature of the process, which is discussed in more detail later. However, given that we do have these shifts, a higher volatility results in a higher variance for the process.

Finally, we will not delve deeply into the case for high values of σ, as this would require a large number of simulations to accurately capture the behavior of the system. High volatility leads to increased randomness in the process, which in turn results in greater variability in the outcomes. As the volatility increases, the final values L_A become more widely distributed, and the system becomes less predictable. To obtain meaningful statistical conclusions, a significantly higher number of simulations would be required to account for the broader spread of possible outcomes. Moreover, the computational cost of simulating such high volatility scenarios is substantially higher, making it impractical for our current analysis. Thus, for the sake of clarity and efficiency, we focus on the lower and medium volatility cases, where the system behaves more predictably and the results are more easily interpretable.

5.2 2-Node Problem: Dynamical Study

In the following section, we study the 2-node problem from a dynamical systems perspective. That is, we apply tools from the theory of discrete dynamical systems to identify the model's critical points. To proceed, we first simplify the model.

We begin by removing the randomness introduced by the term $(1 + \eta(0, \sigma))$, which appears in Equation (2) and follows a normal distribution. However, to preserve the structure of the model without introducing stochastic behavior, we do not eliminate the term entirely. Instead, we assume it to be constant and denote it by a. As a result, Equation (3) becomes:

$$L_A(t+1) = \begin{cases} aL_A(t) & \text{if } P^1, P^2 > p. \\ aL_A(t) + b & \text{if } P^2 \geq p, P^1 < p. \\ aL_A(t) - c & \text{if } P^2 < p, P^1 \geq p. \\ aL_A(t) + b - c & \text{if } P^1, P^2 < p. \end{cases} \tag{4}$$

5.2.1 Dynamical Study: Fixed Points

Although we have removed the normal noise term, the model still contains randomness, as $L_A(t+1)$ depends on which of the four cases occurs. To shed some light on the behavior of

the system, we now analyze the regions of the parameter space that determine which case is selected based on p.

- If $p \to 1$, then $P^1, P^2 < p$, *a.s.* Hence, $L_A(t+1) = aL_A(t) + b - c$, *a.s.*
- If $p \to 0$, then $P^1, P^2 > p$, *a.s.* Hence, $L_A(t+1) = aL_A(t)$, *a.s.*
- If $0 < p < 1$, the evolution of $L_A(t+1)$ cannot be determined analytically and must be studied through simulation.

Given this partition into three regimes, we can apply basic dynamical systems theory to compute the fixed points of the first and second cases by solving $L_A(t+1) = L_A(t)$.

Case 1: $L_A(t+1) = aL_A(t) + b - c$.

Solving for the fixed point:

$$L_A(t) = aL_A(t) + b - c \quad \Rightarrow \quad L_A^1(t) = \frac{b-c}{1-a}.$$

Case 2: $L_A(t+1) = aL_A(t)$.

Solving for the fixed point:

$$L_A(t) = aL_A(t) \quad \Rightarrow \quad L_A^2(t) = 0.$$

Thus, the system for this configuration has two fixed points:

$$(1)\ L_A^1 = \frac{b-c}{1-a}, \quad (2)\ L_A^2 = 0.$$

To determine the nature of these fixed points (attractors, repellers, or indifferent points), we compute the derivative of each case with respect to $L_A(t)$.

For both cases:

$$\frac{df(L_A(t))}{dL_A(t)} = a.$$

Therefore:

- If $|a| < 1$, the fixed point is an attractor.
- If $|a| > 1$, the fixed point is a repeller.

In the marginal case where $|a| = 1$, we consider only the scenario $a = 1$, since $a = -1$ would correspond to an unrealistically large volatility. Consequently, the dynamics simplify to:

$$L_A(t+1) = L_A(t) + P^1 b - P^2 c.$$

For $0 < p < 1$, and given a sufficiently large number of simulations, we approximate:

$$L_A(t) = L_A(0) + tp(b-c),$$

which corresponds to the expected value of $L_A(t)$, as discussed earlier. If $|a| < 1$, the point is an atractor, if $|a| > 1$, the point is a repeller, and for $|a| = 1$ the point diverges and gives no information. However, if $|a| = 1$ that means $L_A(t+1) = L_A(t) + P^1 b - P^2 c$, and for $0 < p < 1$ and enough simulations: $L_A(t+1) = L_A(t) = L_A(0) + tp(b-c)$, which is simply the expected value of $L_A(t+1)$, as seen earlier.

Figure 15 show the phase diagrams of the fixed point L_A^1 as a function of parameter a for different values of b and c.

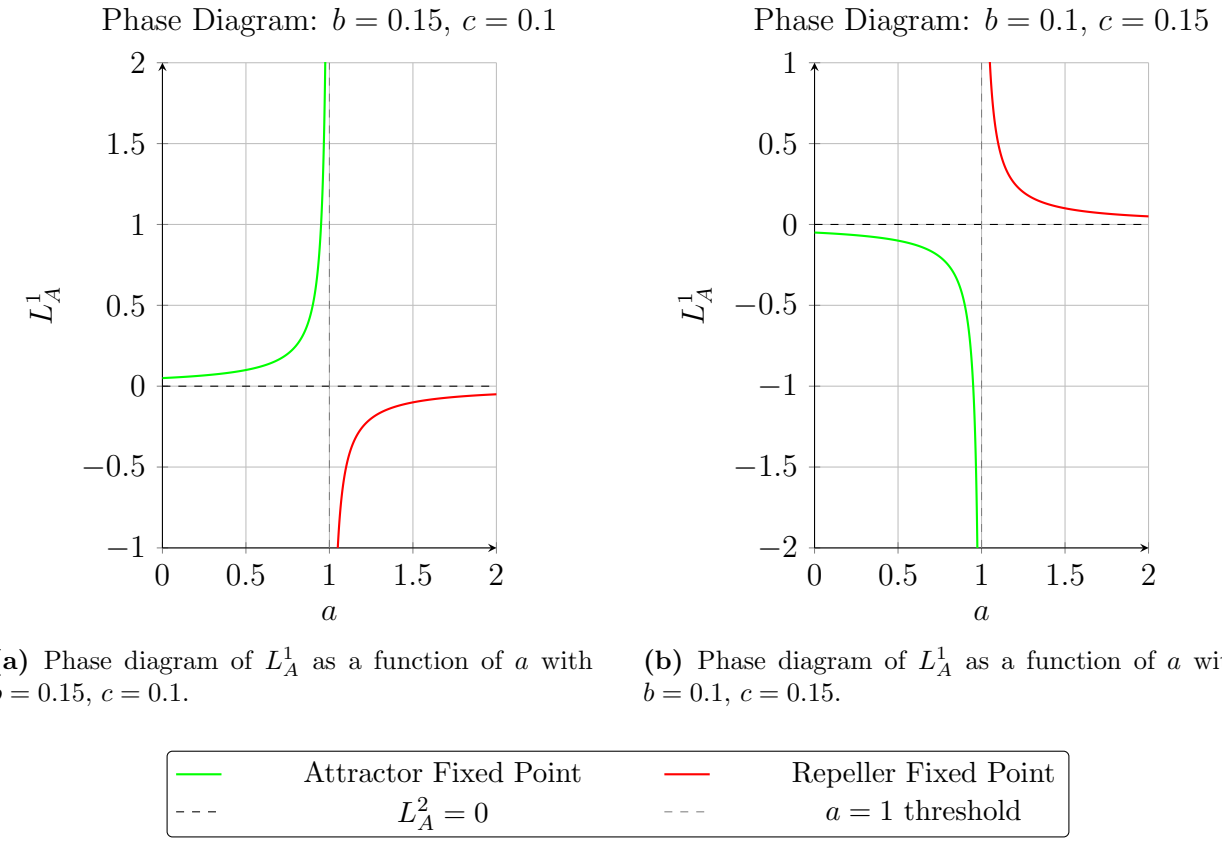

(a) Phase diagram of L_A^1 as a function of a with $b = 0.15$, $c = 0.1$.

(b) Phase diagram of L_A^1 as a function of a with $b = 0.1$, $c = 0.15$.

Figure 15: Phase diagrams of the fixed point L_A^1 as a function of parameter a for two different configurations of b and c.

The main difference between both plots is the difference in the sign of L_A^1. The difference arises from the values of the parameters b and c:

- In the case where $b = 0.15$ and $c = 0.1$ (Figure 15a), the fixed point $L_A^1 = \frac{b-c}{1-a}$ is **positive** for $a < 1$, since $b - c = 0.05$, resulting in $L_A^1 > 0$.

- In the case where $b = 0.1$, $c = 0.15$ (Figure 15b), the fixed point $L_A^1 = \frac{b-c}{1-a}$ is **negative** for $a < 1$, since $b - c = -0.05$, resulting in $L_A^1 < 0$.

Thus, the green curve is above 0 in the first plot and below 0 in the second plot due to the sign of $b - c$. The same reasoning explains the difference of sign for the unstable fixed points.

For the third case, when $p \in (0, 1)$, there exists a fixed point as well. However, due to the inherent randomness of the model, this fixed point is never exactly reached; instead, we oscillate around it. To find this fixed point, we can apply the same technique as before, but instead of using the expression for $L_A(t + 1)$, we compute its expected value, denoted as $\mathbb{E}(L_A(t + 1))$. That is, we assume:

$$\mathbb{E}(L_A(t + 1)) = \mathbb{E}(L_A(t)).$$

Solving for the fixed point:

$$\mathbb{E}(L(t + 1)) = a\mathbb{E}(L_A(t)) + p(b - c) = \mathbb{E}(L_A(t)) \quad \Rightarrow \quad L_A^3(t) = \frac{p(b - c)}{1 - a}.$$

Thus, the general fixed point for this case is:

$$L_A^3 = \frac{p(b - c)}{1 - a}.$$

Notice that this result is also consistent with the cases studied previously. Specifically:

- If $p \to 1$, we recover the fixed point from case (1):

$$L_A^1 = \frac{b - c}{1 - a}.$$

- If $p \to 0$, the fixed point vanishes, which is the fixed point from case (2):

$$L_A^2 = 0.$$

This shows that as p varies between 0 and 1, the fixed point smoothly transitions between these two extremes.

Figure 16 shows how the value of $L_A(T)$ (for $T = 1000$) oscillates around the fixed point. This simulation was performed with the parameters $b = 0.15$, $c = 0.1$, $p = 0.5$, and $a = 0.9$. Thus, $L_A^3 = 0.25$.

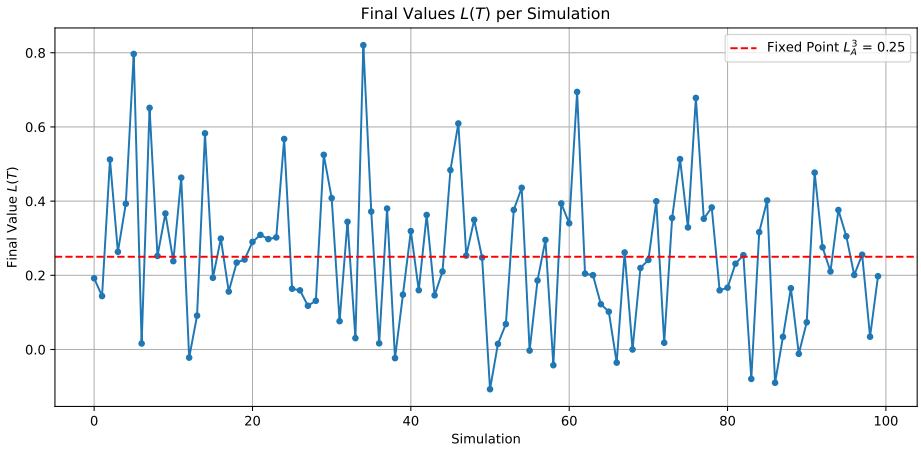

Figure 16: Diagram illustrating the oscillation around the fixed point for 100 final values of $L_A(T)$ (for $T = 1000$). Blue points are the final value of L_A for simulation i, the red line accounts for the fixed point $y = L_A^3 = 0.25$. The parameters are fixed as $b = 0.15$, $c = 0.1$, $p = 0.5$, and $a = 0.9$.

5.2.2 Dynamical Study: Contractive Map Approximation

Consider the case where $p \approx 0$. In order to reintroduce some of the randomness inherent in the normal distribution, we consider two different values for the constant a. Thus, *a.s.*:

$$L_A(t+1) = \begin{cases} a_1 L_A(t) & : a_1 = 1 + \epsilon_1. \\ a_2 L_A(t) & : a_2 = 1 - \epsilon_2. \end{cases}$$

where $0 < \epsilon_1, \epsilon_2 < 1$.

Suppose we iterate the first case n times. Then:

$$f(L_A(0)) = (1+\epsilon_1)L_A(0) \quad \Rightarrow \quad f^n(L_A(0)) = (1+\epsilon_1)^n L_A(0).$$

Since $(1+\epsilon_1) > 1$, the multiplicative factor satisfies $(1+\epsilon_1)^n \xrightarrow{n\to\infty} \infty$. Therefore, $L_A(t) \xrightarrow{t\to\infty} \infty$.

Conversely, if we iterate the second case n times:

$$f(L_A(0)) = (1-\epsilon_2)L_A(0) \quad \Rightarrow \quad f^n(L_A(0)) = (1-\epsilon_2)^n L_A(0).$$

Since $(1-\epsilon_2) < 1$, the multiplicative factor satisfies $(1-\epsilon_2)^n \xrightarrow{n\to\infty} 0$. Hence, $L_A(t) \xrightarrow{t\to\infty} 0$.

Now, suppose we alternate randomly between the two cases with equal probability. Over n iterations, we can expect to have applied each case approximately $n/2$ times:

$$\begin{aligned} f^n(L_A(0)) &= (1+\epsilon_1)^{\frac{n}{2}}(1-\epsilon_2)^{\frac{n}{2}} L_A(0) = [(1+\epsilon_1)(1-\epsilon_2)]^{\frac{n}{2}} L_A(0) \\ &= (1+\epsilon_1-\epsilon_2-\epsilon_1\epsilon_2)^{\frac{n}{2}} L_A(0). \end{aligned}$$

We now distinguish between three different regimes based on the relative sizes of ϵ_1 and ϵ_2:

- **Case 1:** $(1 + \epsilon_1 + \epsilon_2 - \epsilon_1\epsilon_2) > 1$
  In this case, the term $(1+\epsilon_1-\epsilon_2-\epsilon_1\epsilon_2) > 1$. Therefore, the multiplicative factor grows:

  $$(1 + \epsilon_1 - \epsilon_2 - \epsilon_1\epsilon_2)^{\frac{n}{2}} \xrightarrow{n\to\infty} \infty$$

  As a result, $L_A(t) \xrightarrow{t\to\infty} \infty$. The system behaves **expansively**.

- **Case 2:** $(1 + \epsilon_1 - \epsilon_2 - \epsilon_1\epsilon_2) < 1$
  Here, the term $(1 + \epsilon_1 - \epsilon_2 - \epsilon_1\epsilon_2) < 1$. Thus, the multiplicative factor decays:

  $$(1 + \epsilon_1 - \epsilon_2 - \epsilon_1\epsilon_2)^{\frac{n}{2}} \xrightarrow{n\to\infty} 0.$$

  Consequently, $L_A(t) \xrightarrow{t\to\infty} 0$. The system is **contractive**.

- **Case 3:** $(1 + \epsilon_1 - \epsilon_2 - \epsilon_1\epsilon_2) = 1$
  In this case, the term $(1 + \epsilon_1 - \epsilon_2 - \epsilon_1\epsilon_2) <= 1$, which provides no information. The system is indifferent.

- **Case 4:** $\epsilon_1 = \epsilon_2 = \epsilon$
  Substituting gives:

  $$(1 + \epsilon)(1 - \epsilon) = 1 - \epsilon^2 < 1,$$

  so the multiplicative factor becomes:

  $$(1 - \epsilon^2)^{\frac{n}{2}} \xrightarrow{n\to\infty} 0.$$

  Even in the balanced case, the decay dominates due to the ϵ^2 term. Therefore, $L_A(t) \xrightarrow{t\to\infty} 0$, and the system is still **contractive**.

This last case is both surprising and useful. Since we are working with a normal distribution centered at zero, every value is sampled with the same probability as its negative counterpart. Therefore, the case where $\epsilon_1 = \epsilon_2$ naturally arises over enough iterations.

Moreover, we compute the probability of having a contractive case versus an expansive one. To achieve this, we take:

$$(1 + \epsilon_1 - \epsilon_2 - \epsilon_1\epsilon_2) < 1$$
$$\epsilon_1 - \epsilon_2 - \epsilon_1\epsilon_2 < 0$$
$$\epsilon_1 < \frac{\epsilon_2}{1 - \epsilon_2.}$$

Since $0 < \epsilon_1 < 1$ and both ϵ_1, ϵ_2 are independent, we can compute the probability of being in the contractive regime as:

$$\mathbb{P}^{(0,1)}_{\text{contractive}} = \int_0^1 \int_0^{\min\left(1, \frac{\epsilon_2}{1-\epsilon_2}\right)} d\epsilon_1\, d\epsilon_2.$$

We split the integral:

$$\mathbb{P}^{(0,1)}_{\text{contractive}} = \int_0^{1/2} \frac{\epsilon_2}{1 - \epsilon_2}\, d\epsilon_2 + \int_{1/2}^1 1\, d\epsilon_2.$$

Evaluating:

$$\int_0^{1/2} \frac{\epsilon_2}{1-\epsilon_2}\, d\epsilon_2 = \ln 2 - \frac{1}{2}, \qquad \int_{1/2}^{1} 1\, d\epsilon_2 = \frac{1}{2}.$$

Hence, the total probability of contraction is:

$$\mathbb{P}^{(0,1)}_{\text{contractive}} = \ln 2 \approx 0.6931.$$

Notice that if we assume ϵ_1, ϵ_2 have a broader range, specifically $\epsilon_1, \epsilon_2 \in (0, 2)$, the integral limits change accordingly, which leads to a different result. In this case, the contractive condition:

$$\epsilon_1(1-\epsilon_2) < \epsilon_2,$$

defines a different region over the domain $(0, 2) \times (0, 2)$. The probability of being in the contractive regime would now be:

$$\mathbb{P}^{(0,2)}_{\text{contractive}} = \frac{1}{Area} \int_0^2 \int_0^{\min\left(2, \frac{\epsilon_2}{1-\epsilon_2}\right)} d\epsilon_1\, d\epsilon_2 = \frac{1}{Area}(\ln 3 + 2) \approx \frac{3.098}{4} = 0.7747.$$

Therefore, the probability of being contracted for the domain $(0, 2) \times (0, 2)$ is:

$$\mathbb{P}^{(0,2)}_{\text{contractive}} \approx 0.7747.$$

Which is bigger than the one computed previously for epsilon with a narrower range.

In general, if we integrate over $\epsilon_1, \epsilon_2 \in (0, N)$ where $N \in \mathbb{R}$. We obtain the probability of being in a contractive map is given by following expression:

$$\mathbb{P}^{(0,N)}_{(\text{contractive})} = \frac{N^2 - N + ln(N+1)}{N^2}. \tag{5}$$

Thus, if the range of epsilon is bigger, so is the probability of being a contractive map. This can be seen in Figure 17, where the area under the curve represents the probability of an expansive map, and the area on top of the under the curve represents the probability of a contractive map.

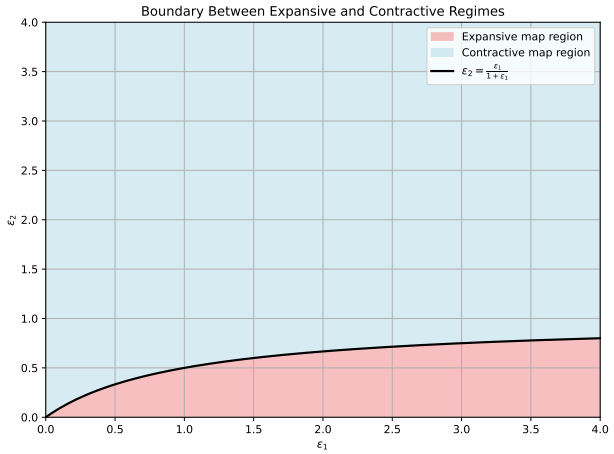

Figure 17: Plot of the function $\epsilon_2 = \frac{\epsilon_1}{1+\epsilon_1}$, illustrating the boundary between contractive (blue) and expansive (red) regimes in the model, where the area under the curve represents the probability of an expansive map and the area outside the curve represents the probability of a contractive map.

This result is of great interest as it explains why, in the presence of multiplicative normal noise in the model, increasing the normal deviation,what we refer to as a more volatile market,causes the liquidity density to become narrower.

Moreover, we are interested in studying when is the map expansive. The probability of being in a contractive map is given by Equation (5). We now investigate when this probability is less than 0.5, whether there exists a value of σ such that the map becomes expansive:

$$\mathbb{P}^{(0,N)}(\text{contractive}) = \frac{N^2 - N + \ln(N+1)}{N^2} < 0.5.$$

There is no need to solve this inequality analytically; instead, we can visualize it through a plot, where the x-axis represents $N = \sigma$ and the y-axis shows the probability given by the formula.

From Figure 18, we observe that there is no value of σ for which the map is not contractive. The only values of σ that would make the inequality hold are strictly $\sigma \leq 0$, which are not meaningful in this context.

If the map is always contractive, we can refer to the general fixed point result we studied previously (L_A^3).

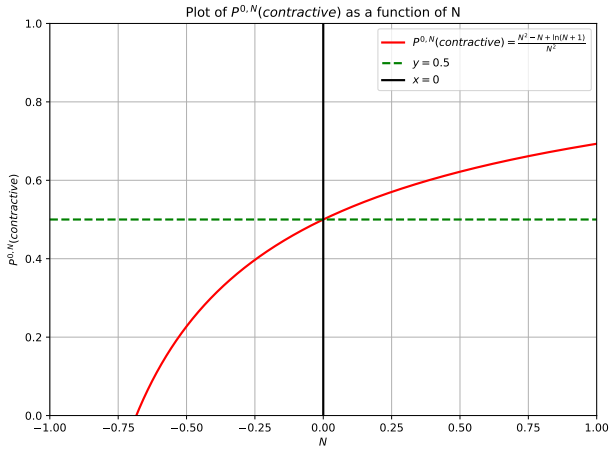

Figure 18: The red line is a visualization of the contractive map probability given by Equation (5) as a function of $N = \sigma$. The green line is the function $y = 0.5$. The plot confirms that the probability remains above 0.5 for all meaningful values of σ, indicating that the map is always contractive.

Let us first assume that $\sigma = 0.1$. This implies that approximately 95.45% of the data lies within the interval $\pm 2\sigma$. Therefore, we set $\epsilon = \pm 0.2$.

If we consider the third case discussed previously, where $\epsilon_1 = \epsilon_2 = \epsilon$, the multiplicative factor becomes $(1-\epsilon^2)^{\frac{n}{2}}$. For the extreme case $\epsilon = 2\sigma$, we can compute a representative contraction factor as

$$(1 - 0.2^2)^{\frac{n}{2}} = 0.96^{\frac{n}{2}}.$$

This corresponds to the decay rate previously defined as

$$a = 0.96^{\frac{n}{2}}.$$

Due to the randomness inherent in the process, it is not possible to define a single stable fixed point; instead, the process tends toward a different fixed point at each iteration. However, we can establish approximate bounds within which these fixed points lie. These bounds correspond to the limiting cases $a = 0.96^{\frac{n}{2}}$ and $a = 1$. The order of these factors may vary depending on the parameters b, c, n, p. Since n represents the number of time steps, it can equivalently be denoted as t.

Thus, the fixed point will approximately range in the interval:

$$L_A^* \in \left[p\frac{b-c}{1 - 0.96^{\frac{t}{2}}},\ L_A(0) + tp(b-c) \right].$$

Now consider the full model with a single coefficient a, defined as $a = 1 + \epsilon_i$, where $\epsilon_i \in (-0.75\sigma, 0.75\sigma)$. This choice captures 75% of the market volatility range for each ϵ_i. At each iteration, one of four affine functions is selected at random and applied:

$$f^n(L_A) \in \left\{ af^{n-1}(L_A),\ af^{n-1}(L_A) + b,\ af^{n-1}(L_A) - c,\ af^{n-1}(L_A) + b - c \right\}.$$

We note that we can use a single coefficient a to represent all four outcomes. Strictly speaking, we should denote them as a_i for $i \in \{1, 2, 3, 4\}$ to preserve the model's structure. However, since these coefficients are independent and identically distributed normal random variables with the same mean μ and variance σ, and since we will be iterating over them a large number of times, we can simplify the analysis.

Suppose we apply the functions k_1, k_2, k_3, k_4 times respectively, such that $k_1+k_2+k_3+k_4 = n$. Then, the expected value of the product is:

$$\mathbb{E}(a_1^{k_1} a_2^{k_2} a_3^{k_3} a_4^{k_4}) = \mu^n$$

which is exactly the same as the expected value for a single coefficient raised to the n-th power:

$$\mathbb{E}(a^n) = \mu^n.$$

Therefore, we are justified in replacing the set $\{a_1, a_2, a_3, a_4\}$ with a single representative coefficient a.

Moreover, we will define $f_i(x) = ax + d_i$, where $d_1 = 0$, $d_2 = b$, $d_3 = -c$, and $d_4 = b - c$. Thus, each iteration corresponds to the form:

$$L_t = aL_{t-1} + d_t,$$

where d_t is a randomly chosen from the set $\{d_1, d_2, d_3, d_4\}$ for each step.

After n iterations, the composition of these functions results in:

$$f^n(L_A(0)) = a^n L_A(0) + \sum_{k=1}^{n} a^{n-k} d_k.$$

This expression represents the accumulated effect of both multiplicative scaling and additive shifts introduced by the chosen affine transformations.

Therefore, the state at time n becomes independent of the initial condition, as the product of the a-terms decays to zero over time. Moreover, note that only the final iterations contribute to the final result, since all the previous terms will be multiplied by an a^{x_i} term that tends to zero rapidly.

Theoretically, we should then observe the following behavior: when $p \approx 1$, only the term $b - c$ is likely to survive, as this event is the most frequent and thus highly likely to appear within the final iterations. Conversely, when $p \approx 0.5$, the terms $+b$ or $-c$ are more likely to survive one at a time. Note that the absolute value of $b - c$ is always lower than the absolute value of either b or c.

This reasoning extends the probability framework introduced earlier. Given that $L(t+1) = \eta(L(t), \sigma L(t))$, we can draw the following conclusions:

- When $p \approx 1$, we have $L_A(t) \approx b - c$. In this case, the width of the distribution of $L_A(t)$ depends directly on the difference $b - c$.
- When $p \approx 0.5$, the expression becomes $L_A(t) = b$ or $L_A(t) = -c$, which has greater absolute magnitude than the previous case, leading to a wider distribution.
- When $p \approx 0$, the expression tends to $L_A(t) = 0$. Leading to the narrowest distribution

However, a key limitation remains in our analysis: we have focused solely on the final iteration of the process, implicitly assuming that the influence of prior steps can be neglected. This assumption is problematic. Because the multiplicative noise is drawn from a normal distribution centered at zero, each perturbation has an equal probability of being contractive or expansive.

Formally, let us define $a = 1 + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \sigma^2)$. Then, in approximately half of the realizations, we have $a > 1$, leading to local expansion rather than contraction. While the expected value of the system may trend toward zero due to the predominance of contractive effects over time, this ignores the cumulative impact of rare but persistent expansive episodes. For example, in 50,000 independent simulations, we expect roughly 48 instances in which $a > 1$ occurs for 10 consecutive steps, since $0.5^{10} \approx 0.00097$.

Such sequences, though statistically infrequent, can cause significant divergence in individual trajectories. As market volatility increases, these rare events become more consequential, amplifying the tails of the distribution and contributing disproportionately to the overall variance.

For greater market volatility, represented by σ, the density of the final liquidity becomes wider compared to lower volatility levels. Additionally, the density becomes narrower for smaller values of the network activation parameter p. Theoretically, the distribution reaches its maximum spread at $p = 0.5$. However, in practice, we observe that the final variance of the system continues to grow as p increases beyond 0.5.

This discrepancy arises due to the nonlinear and multiplicative nature of the model dynamics. While the variance of the additive perturbations introduced by the network is indeed proportional to $p(1 - p)$, which is maximized at $p = 0.5$, the system's response to these perturbations is far from linear. When p increases, the perturbations (positive and negative liquidity shocks) are applied more frequently, and although their individual variance decreases past $p = 0.5$, their cumulative impact grows.

More importantly, the multiplicative stochastic process amplifies any early deviations in liquidity over time. For high values of p, the system experiences near-constant drift due to almost continuous shocks, which shifts the base level of liquidity. This, in turn, modifies the scale of the multiplicative noise, making the system increasingly sensitive to fluctuations. As a result, even small differences in early stages can lead to large dispersion in the final states, causing the overall variance to increase with p, contrary to what would be expected in a purely additive, linear system.

Therefore, the interaction between frequent perturbations and multiplicative amplification explains the increasing variance for $p > 0.5$, despite the local variance of the shocks themselves decreasing in that regime.

5.3 2-Node Problem: Probabilistic Approximation

The following section aims to provide an analytical formula for computing the probability of default, along with other relevant results. We approximate our model using a probabilistic framework, and we seek to compute the mean and variance of the final liquidity distribution for this approximation. Therefore, the analysis presented below relies on basic principles of probability theory.

As shown in Equation (3), the model can be expressed in terms of four distinct outcomes, which depend on the values of P^1 and P^2. In this approximation, our goal is to express the outcome as a single expression. To achieve this, we define the following cases:

$$L_i(t+1) = \begin{cases} L_i(t)\big(1+\eta(0,\sigma)\big) & \text{if } P^1, P^2 > p \quad \text{(Case A)}. \\ L_i(t)\big(1+\eta(0,\sigma)\big) + b & \text{if } P^2 \geq p,\ P^1 < p \quad \text{(Case B)}. \\ L_i(t)\big(1+\eta(0,\sigma)\big) - c & \text{if } P^2 < p,\ P^1 \geq p \quad \text{(Case C)}. \\ L_i(t)\big(1+\eta(0,\sigma)\big) + b - c & \text{if } P^1, P^2 < p \quad \text{(Case D)}. \end{cases} \tag{6}$$

Thus, Case A can be seen as a random variable L_A, Case B as L_B, and similarly, Case C and Case D are represented as random variables L_C and L_D, respectively.

Since $L_i(t+1)$ is also a random variable, we can express it as a linear combination of the four cases:

$$L_i(t+1) = \alpha L_A(t) + \beta L_B(t) + \gamma L_C(t) + \delta L_D(t).$$

This means that the random variable $L_i(t+1)$ is the pound sum of the four different outcomes, each of which is a random variable. In fact, not only are these random variables, but they also represent stochastic processes, as they depend on $L_i(t-1)$.

Our next goal is to determine the values of the parameters α, β, γ, and δ, as well as the distributions followed by L_A, L_B, L_C, and L_D.

First, we need to determine the distribution of L_A. Recall that L_A is defined by Case A, where $L_A(t+1) = L_i(t)(1+\eta(0,\sigma))$. This implies $L_A(t+1) \sim N(L_i(t), L_i(t)\sigma)$. Similarly, applying this process to the other three cases, we obtain the following distributions:

$$L_A(t+1) \sim N(L_i(t), |L_i(t)|\sigma).$$
$$L_B(t+1) \sim N(L_i(t)+b, |L_i(t)+b|\sigma).$$
$$L_C(t+1) \sim N(L_i(t)-c, |L_i(t)-c|\sigma).$$
$$L_D(t+1) \sim N(L_i(t)+b-c, |L_i(t)+b-c|\sigma).$$

Here, the '|' represent the absolute values, which are necessary since we can have $L_i(t) < 0$ or simply $L_i(t) + b - c < 0$, and a negative variance would make no sense.

To determine the values of the coefficients, we first impose the condition that these are normalized, meaning that:

$$\alpha + \beta + \gamma + \delta = 1.$$

Since our aim is to consolidate the outcome into a single type, rather than dealing with four different cases, and because this outcome is closely related to the values of both P^1 and P^2, it is expected that the normal coefficients will have some relationship with p. This expectation is indeed correct, and we can assign values to these coefficients following the reasoning outlined earlier. Which is also explained here.

If $p \approx 0$, the network is effectively deactivated, and it is most likely that only Case A will occur. On the other hand, if $p \approx 1$, Case D becomes the most likely outcome. For values where $0 < p < 1$, the other two cases become more probable. Based on this reasoning, we conclude that $\beta = \gamma$, since both have the same likelihood of occurring. We then assign the following values, which correspond to the are of the rectangle represented by Case A, B, C, and D as a function of p:

$$\alpha = (1-p)^2, \quad \beta = \gamma = (1-p)p, \quad \delta = p^2.$$

These coefficients are also normalized.

Now, in order to ensure our approximation fits the real model as closely as possible, we compute the expected values of both the approximation and the real model, and then equate them:

For the real model:

$$\mathbb{E}[L_i(t+1)] = \mathbb{E}[L_i(t)] + p(b-c).$$

For the approximated model:

$$\mathbb{E}[L_i(t+1)] = \alpha\mathbb{E}[L_i(t)] + \beta\left(\mathbb{E}[L_i(t)] + b\right) + \gamma\left(\mathbb{E}[L_i(t)] - c\right) + \delta\left(\mathbb{E}[L_i(t)] + b - c\right).$$

Substituting the values of α, β, γ, and δ:

$$\mathbb{E}[L_i(t+1)] = (1-p)^2\mathbb{E}[L_i(t)] + (1-p)p\left(\mathbb{E}[L_i(t)] + b\right) + (1-p)p\left(\mathbb{E}[L_i(t)] - c\right) + p^2\left(\mathbb{E}[L_i(t)] + b - c\right).$$

Distribute the terms:

$$\mathbb{E}[L_i(t+1)] = (1-p)^2\mathbb{E}[L_i(t)] + (1-p)p\mathbb{E}[L_i(t)] + (1-p)pb + (1-p)p\mathbb{E}[L_i(t)] - (1-p)pc + p^2\mathbb{E}[L_i(t)] + p^2b - p^2c.$$

Now group the terms with $\mathbb{E}[L_i(t)]$ and the constant terms:

$$\mathbb{E}[L_i(t+1)] = \left((1-p)^2 + 2(1-p)p + p^2\right)\mathbb{E}[L_i(t)] + \left((1-p)p + p^2\right)b - \left((1-p)p + p^2\right)c.$$

Simplify the coefficients:

$$(1-p)^2 + 2(1-p)p + p^2 = 1.$$
$$(1-p)p + p^2 = p.$$

Thus, the final simplified expression is:

$$\mathbb{E}[L_i(t+1)] = \mathbb{E}[L_i(t)] + p(b-c).$$

Which is the same as the original model. Thus, we have chosen correct coefficients and distributions for the random variables.

If we now reconstruct our model equation:

$$\begin{aligned} L_i(t+1) = {} & (1-p)^2\, N(L_i(t),\, \sigma|L_i(t)|) \\ & + (1-p)p\left[N(L_i(t)+b,\, |(L_i(t)+b)|\sigma) + N(L_i(t)-c,\, |(L_i(t)-c)|\sigma)\right] \\ & + p^2\, N(L_i(t)+b-c,\, |(L_i(t)+b-c)|\sigma). \end{aligned} \tag{7}$$

There is no analytical way to solve this equation directly; therefore, we must once again rely on stochastic simulations to obtain any meaningful results.

5.3.1 Probabilistic Approximation: Normal Distribution Approach

Since we are working with the sum of normal distributions, we make the assumption that the final liquidity also follows a normal distribution for small values of σ. This implies we can utilize the normal CDF to obtain the probability that a node goes into bankruptcy, that is, $\mathbb{P}(L(t+1) < 0)$. To further explore this hypothesis, we have simulated our simplified model

10,000 times and empirically computed the proportion of nodes that, at the final time step, were in default. By dividing this number by the total number of simulations, we obtain an empirical probability of default we can use to test our hypothesis.

Moreover, we aim to compare this value to the theoretical probability computed via the Z-score of our gaussian. Since we do not have an analytical formula for the standard deviation, we have also computed the variance of the final liquidity values in the simplified model. We then use the theoretical mean and the empirical standard deviation to compute the probability of default. Table 2 below compares the empirical and theoretical probabilities. The parameter configuration used for these simulations is: $L_A(0) = 0.1,\ b = 0.1,\ c = 0.12,\ t = 200,\ \sigma = 0.1,\ p = k \cdot 0.05$ for $k = 0, \ldots, 14$.

p	Z-Score Probability of Default	Empirical Probability of Default	Relative Error
0.00	0.343	0.0000	N/A
0.05	0.548	0.5866	0.0659
0.10	0.600	0.6670	0.1005
0.15	0.628	0.7208	0.1286
0.20	0.662	0.8280	0.2005
0.25	0.692	0.9104	0.2396
0.30	0.702	0.9273	0.2427
0.35	0.720	0.9485	0.2407
0.40	0.736	0.9644	0.2367
0.45	0.740	0.9879	0.2510
0.50	0.754	0.9937	0.2410
0.55	0.768	0.9949	0.2281
0.60	0.778	0.9934	0.2166
0.65	0.788	0.9938	0.2061
0.70	0.778	0.9927	0.2164

Table 2: Z-score approximation versus empirical probability of default as a function of p, with relative error.

This table clearly shows that the assumption of normality is not fully accurate, with a relative error or about 20% between our probabilities. However, it suggests the possibility of a skewed normal distribution. Thus, our next focus will be to determine its skewness parameter.

To compute the skewness parameter, we must first compute the true variance of the model. The model does have a theoretical variance, however the multiplicative nature of the process prevents our simulations from converging to the theoretical value without an infinite number of simulations. Therefore, we need a more robust and empirical method of estimating it. Nonetheless, we begin by analyzing the theoretical variance.

The theoretical variance of our model is strictly

$$\mathrm{Var}(L(t+1)) = \mathbb{E}[L(t+1)^2] - (\mathbb{E}[L(t+1)])^2.$$

Expanding the squared terms and computing the expected values, as previously derived, leads to the following recurrence:

$$\mathrm{Var}(L(t+1)) = (1+\sigma^2)\mathrm{Var}(L(t)) + \sigma^2(\mathbb{E}[L(t)])^2 + p(1-p)(b^2+c^2).$$

We solve this recurrence analytically using an external solver (DeepSeek), applying the characteristic polynomial method. Using the initial condition $\mathrm{Var}(L(0)) = 0$, the solution is:

$$\mathrm{Var}(L(t)) = \sigma^2 \sum_{k=0}^{t-1} (1+\sigma^2)^{t-1-k} \left(L(0) + kp(b-c)\right)^2 + p(1-p)(b^2+c^2)\frac{(1+\sigma^2)^t - 1}{\sigma^2}.$$

Since we are in a low-σ regime, we approximate the variance by taking the limit as $\sigma \to 0$ and perform a first-order Taylor expansion on $(1+\sigma^2)^{t-1-k}$ and $(1+\sigma^2)^t$:

$$\begin{aligned}
\mathrm{Var}(L(t)) &= \sigma^2 \sum_{k=0}^{t-1} (1+\sigma^2)^{t-1-k} (L(0) + kp(b-c))^2 + p(1-p)(b^2+c^2)\frac{(1+\sigma^2)^t - 1}{\sigma^2} \\
&\approx \sigma^2 \sum_{k=0}^{t-1} \left[1 + (t-1-k)\sigma^2\right] (L(0) + kp\delta)^2 + p(1-p)(b^2+c^2)\frac{t\sigma^2}{\sigma^2} \\
&= \sigma^2 \sum_{k=0}^{t-1} (L(0) + kp\delta)^2 + p(1-p)(b^2+c^2)t \\
&= \sigma^2 \left[tL(0)^2 + L(0)p\delta t(t-1) + p^2\delta^2 \frac{t(t-1)(2t-1)}{6} \right] + p(1-p)(b^2+c^2)t,
\end{aligned}$$

where $\delta = b - c$. The final approximation is:

$$\mathrm{Var}(L(t)) \approx \underbrace{p(1-p)(b^2+c^2)t}_{\text{independent term}} + \underbrace{\sigma^2 \left[tL(0)^2 + L(0)p(b-c)t(t-1) + \frac{p^2(b-c)^2 t(t-1)(2t-1)}{6} \right]}_{\text{dependence in } \sigma^2}.$$

Notice that we can determine which distribution dominates by comparing $p(1-p)(b^2 + c^2)$ and $L(0)^2\sigma^2 t$, since the other terms multiplied by σ^2 are small enough to be neglected for small σ and t. To validate this approximation, we compute both the exact and the approximated variance for different values of σ and evaluate the relative error. We observe excellent agreement for $\sigma \leq 0.1$, with Error $< 0.1\%$. The results are presented in Table 3.

σ	Exact Variance	Approximated Variance	Error (%)
0.0100	0.00052134	0.00052134	0.0000
0.0500	0.00130435	0.00130417	0.0138
0.1000	0.00260869	0.00260634	0.0901

Table 3: Comparison between exact and approximated variance.

Therefore, in the limit $\sigma \to 0$, the variance can be accurately approximated by a constant term plus a linear term in σ^2. We now proceed with a robust numerical approach: we simulate

the actual variance of the model for different values of σ and p, while keeping the parameters b, c, t, and $L(0)$ fixed. The goal of these simulations is to find a good linear fit to the form $\mathrm{Var}(L(t+1)) = A + B\sigma^2$. If this approximation holds, we can use the analytical expression of the variance to compute the skewness of our distribution.

The following figures display the evolution of variance with respect to σ for $\sigma \in [0.001, 0.1]$ in steps of 0.005. We have performed 50000 simulations for each value of σ, since this number of simulations ensures an error of approximately 1% in the variance.

Figure 19 shows the evolution of the variance as a function of σ for $p = 0$. As expected, the linear fit performs poorly, since $p = 0$ implies that only the multiplicative component is active in the system. As discussed in the probability refresher, when only multiplicative noise is present, the variance tends to zero. However, to observe this effect in our simulations, a significantly higher number of iterations is required for the data to converge toward zero. Since our approximated variance is not exactly 0, the error between our simulated variance and the true variance is given by:

$$\text{Error} = \frac{|\text{Approx Var} - 0|}{|0|} \to \infty,$$

which further explains why the linear fit is not correct for this case. This result is not ideal for analytical comparison, but since $p = 0$ does not yield any meaningful or realistic behavior in the context of our model, it holds limited relevance to our study.

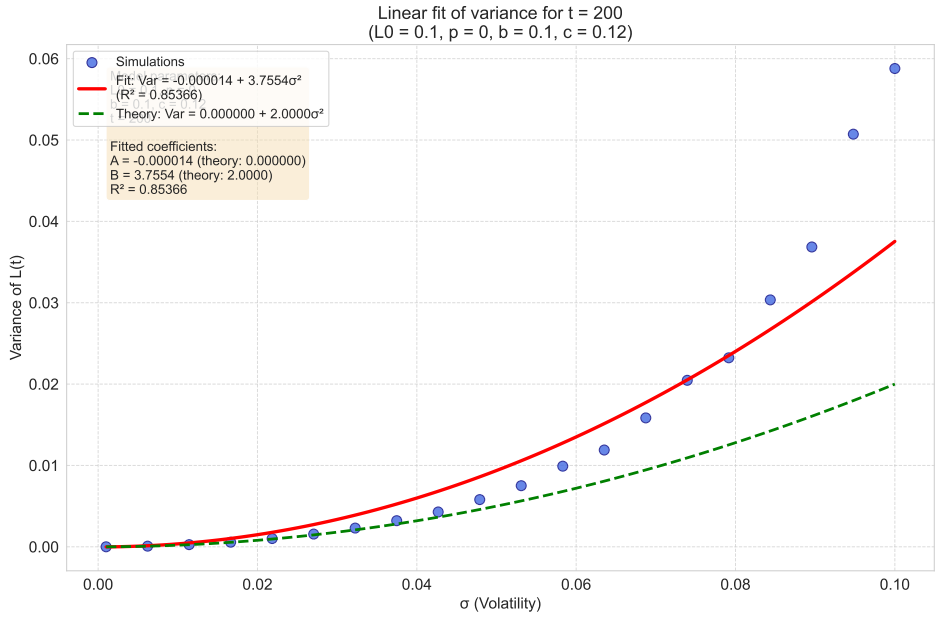

Figure 19: Variance of $L_A(t+1)$ versus σ in the real model. The red line represents the linear fit, the green line shows the theoretical variance, and the blue dots correspond to the simulated variance values. The fitted parameters are $A = 3.7554$, $B = -0.000014$, with an $R^2 = 0.8536$.

Figure 20 shows the evolution of the variance as a function of σ for $p = 0.5$. This linear fit performs impressively well, with an $R^2 = 0.972$. The case $p = 0.5$ provides the best fit

among all possible values of p, which can be attributed to the dominance of the Bernoulli noise in the system. Given our parameter configuration, the ratio between the Bernoulli and Gaussian contributions to the variance reaches its maximum at $p = 0.5$. This ratio is:

$$\text{Ratio} = \frac{\text{Bernoulli}}{\text{Gaussian}} = \frac{0.5^2(b^2 + c^2)}{L(0)^2\sigma^2 t}.$$

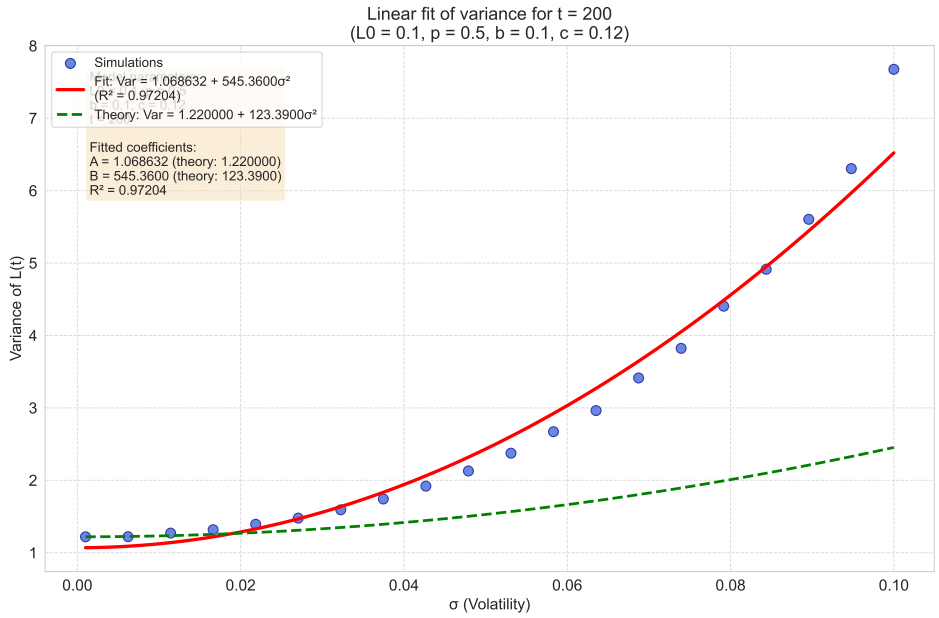

Figure 20: Variance of $L_A(t+1)$ versus σ in the real model. The red line represents the linear fit, the green line shows the theoretical variance, and the blue dots correspond to the simulated variance values. The fitted parameters are $A = 545.36$, $B = 1.0686$, with an $R^2 = 0.974$.

Figure 21 shows the evolution of the variance as a function of σ for $p = 1$. Once again, the linear fit performs well, with an $R^2 = 0.955$, supporting the validity of our approximation.

Now that we have derived an analytical expression for the variance as a function of σ for different values of p, we proceed to compute the skewness of our distribution and apply the CDF once again.

The skewness is defined as:

$$\text{Skew}(X) = \frac{\mathbb{E}[(X - \mu)^3]}{\text{Var}(X)^{3/2}}.$$

Since here $X = L(t)$ and its exact skewness is likely not analytically tractable, we assume an asymptotic approximation by considering separately the contributions from the multiplicative and additive components.

First, the multiplicative process, which resembles a log-normal distribution, contributes a skewness of $\frac{3\sigma}{\sqrt{t}}$.

Next, we consider the additive part, which consists of a sum of t i.i.d. random variables of the form $B_k \cdot b - C_k \cdot c$, where $B_k, C_k \sim \text{Bernoulli}(p)$ are independent. The skewness of such

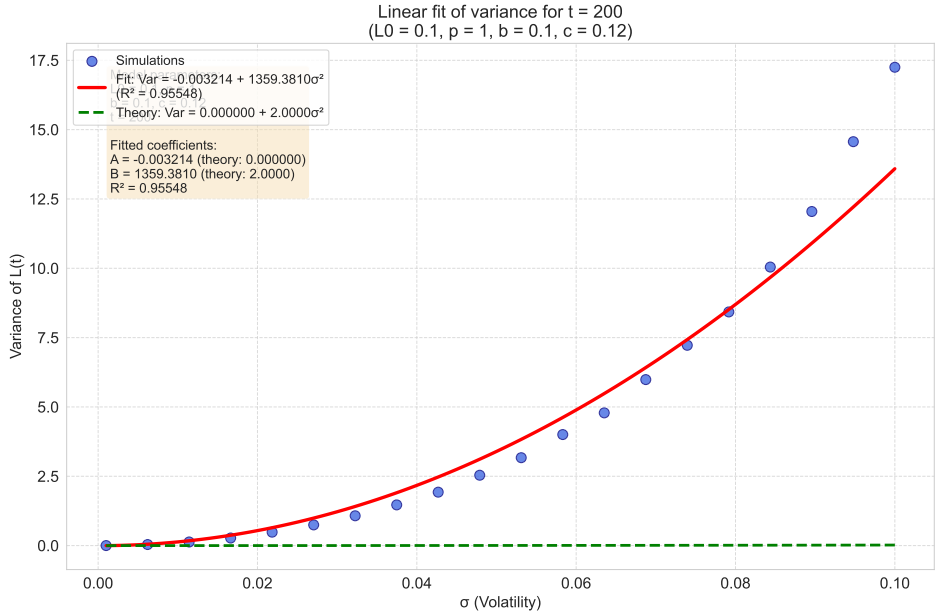

Figure 21: Variance of $L_A(t+1)$ versus σ in the real model. The red line represents the linear fit, the green line shows the theoretical variance, and the blue dots correspond to the simulated variance values. The fitted parameters are $A = 1359.38$, $B = -0.003214$, with an $R^2 = 0.955$.

a sum scales with $\frac{1}{\sqrt{t}}$, and the third central moment of the individual terms can be computed analytically. The resulting skewness of the additive component is:

$$\text{Skew}_{\text{add}}(L(t)) = \frac{(1-2p)(b^3+c^3)}{[p(1-p)(b^2+c^2)]^{3/2} \cdot \sqrt{t}}.$$

Combining both contributions, the total skewness is approximated by:

$$\text{Skew}(L(t)) \approx \frac{3\sigma}{\sqrt{t}} + \frac{(1-2p)(b^3+c^3)}{[p(1-p)(b^2+c^2)]^{3/2} \cdot \sqrt{t}}.$$

Now that we have obtained an analytical expression for the skewness, let us assume the distribution is approximately normal for small values of σ. Moreover, small values of t will make the distribution more normal-like, since fewer multiplicative noise terms will have been accumulated.

Finally, we compute the probability of default for a single node using the CDF of the skew-normal distribution. Specifically, we evaluate the CDF numerically using the Python function `skewnorm.cdf` to obtain the probability of default. 4 displays the results for different values of p and σ, comparing the probability of default using the fitted and theoretical parameters A and B of the approximated variance.

If we compare the results from Table 4 with those in Table 2, we observe that incorporating

p	σ	$P(L(t) < 0)$ Fitted Parameters	$P(L(t) < 0)$ Theoretical Parameters
0	0.010	0.000000	0.000000
0	0.033	0.055845	0.014788
0	0.055	0.173911	0.099283
0	0.077	0.252689	0.180780
0	0.100	0.302883	0.239750
0.5	0.010	0.963497	0.956510
0.5	0.033	0.930769	0.948982
0.5	0.055	0.875421	0.933871
0.5	0.077	0.819007	0.912571
0.5	0.100	0.771553	0.887416
1	0.010	1.000000	1.000000
1	0.033	0.999440	1.000000
1	0.055	0.972822	1.000000
1	0.077	0.913896	1.000000
1	0.100	0.854950	1.000000

Table 4: Comparison of $P(L(t) < 0)$ using fitted and theoretical parameters for different values of p and σ.

skewness shifts the estimated probabilities in the expected direction. Specifically, it reduces the probability of default for $p = 0$ and increases it for $p = 0.5$.

Furthermore, the theoretical parameters give probabilities that are generally more accurate than those obtained from the fitted parameters. This suggests that the actual distribution of $L(t)$ may not follow a skewed-normal distribution exactly, but rather a distribution with heavier tails, such as a skewed Student's t-distribution. This interpretation is consistent with the observation that the fitting performed well for $p > 0$, where the tail behavior plays a more significant role.

While skewness correction improves the model, a more flexible distribution with heavier tails might be needed for a more precise representation of the probability of default.

5.3.2 Probabilistic Approximation: Propagation of Default

In this section, we aim to demonstrate how, under the normal approximation of liquidity, the default of a single node can propagate to other nodes in the network.

Although the distribution does not perfectly adhere to normality, largely due to the process's multiplicative nature, the normal approximation remains a useful tool, particularly when σ is very small and t is sufficiently large, which is often the case in practical scenarios.

When $\sigma = 0$, the general process becomes additive rather than multiplicative:

$$L(t+1) = L(t) + P^1 b - P^2 c,$$

which can be rewritten as:

$$L(t) = L(0) + \sum_{k=1}^{t} (P_k^1 b - P_k^2 c).$$

Since P^1 and P^2 are independent at each time step, and the expression has finite mean and variance, the Central Limit Theorem applies. This theorem states that the sum of many independent and identically distributed random variables with finite mean and variance converges to a normal distribution.

Under this assumption, our analysis remains valid.

We begin by computing the Z-score for zero, where the mean is given by Equation (7):

$$\mu = L_A(0) + tp(b - c),$$

and the standard deviation is treated as an unknown parameter, σ_{std}.

Thus, the probability of default for a single node is approximated as:

$$\mathbb{P}(L_A(t+1) < 0) = \Phi\left(\frac{0 - (L_A(0) + tp(b-c))}{\sigma_{\text{std}}}\right).$$

We now investigate how the default of one node affects the probability that neighboring nodes also default, which leads us to consider conditional probabilities.

From Bayes' Theorem:

$$\mathbb{P}(B \mid A) = \frac{\mathbb{P}(A \mid B) \cdot \mathbb{P}(B)}{\mathbb{P}(A)},$$

where $A = \{L_A(t+1) < 0\}$ and $B = \{L_B(t+1) < 0\}$.

We are particularly interested in:

$$\mathbb{P}(L_B(t+1) < 0 \mid L_A(t+1) < 0).$$

Assuming normality and that the default of node B results in $b^* < b$ for its neighbors (since it can no longer make payments), we write:

$$\mathbb{P}(L_A(t+1) < 0 \mid L_B(t+1) < 0) = \Phi\left(\frac{0 - (L_A(0) + tp(b^* - c))}{\sigma_{\text{std}}}\right).$$

We now compare this conditional probability to the unconditional probability $\mathbb{P}(L_B(t+1) < 0)$ to assess the effect of a neighboring default:

$$\frac{\Phi\left(\frac{-(L_A(0)+tp(b^*-c))}{\sigma_{\text{std}}}\right) \cdot \Phi\left(\frac{-(L_A(0)+tp(d-e))}{\sigma_{\text{std}}}\right)}{\Phi\left(\frac{-(L_A(0)+tp(b-c))}{\sigma_{\text{std}}}\right)} \quad > \quad \Phi\left(\frac{-(L_A(0) + tp(d-e))}{\sigma_{\text{std}}}\right).$$

Canceling the common term on both sides yields:

$$\frac{\Phi\left(\frac{-(L_A(0)+tp(b^*-c))}{\sigma_{\text{std}}}\right)}{\Phi\left(\frac{-(L_A(0)+tp(b-c))}{\sigma_{\text{std}}}\right)} \quad > \quad 1.$$

This holds if:

$$\Phi\left(\frac{-(L_A(0)+tp(b^*-c))}{\sigma_{\text{std}}}\right) > \Phi\left(\frac{-(L_A(0)+tp(b-c))}{\sigma_{\text{std}}}\right)$$

Because Φ is monotonically increasing, this simplifies to:

$$-(L_A(0)+tp(b^*-c)) > -(L_A(0)+tp(b-c)) \Rightarrow tp(b^*-c) < tp(b-c) \Rightarrow b^* < b.$$

Since it is assumed that $b^* < b$, the inequality holds. This implies that the probability of default for a node increases when a neighboring node is already in default. In other words, defaults propagate through direct connections.

This effect can extend beyond immediate neighbors. For example, the default of node A may increase the probability of default for its neighbor B, which in turn may raise the risk for B's neighbor C, and so on. Such propagation can be expressed using nested conditional probabilities:

$$\mathbb{P}(L_C(t+1) < 0 \mid L_B(t+1) < 0 \mid L_A(t+1) < 0).$$

Applying the chain rule:

$$\mathbb{P}(L_C(t+1) < 0 \mid L_B(t+1) < 0 \cap L_A(t+1) < 0) = \frac{\mathbb{P}(L_C(t+1) < 0 \cap L_B(t+1) < 0 \cap L_A(t+1) < 0)}{\mathbb{P}(L_B(t+1) < 0 \cap L_A(t+1) < 0)}.$$

This analysis shows that the likelihood of default cascades through the network in a manner shaped by its topology and the interdependencies between nodes. Not only do immediate neighbors experience increased default risk, but also nodes that are several steps away, highlighting the systemic nature of default propagation in interconnected networks.

6 Conclusions

The primary objective of this work was to provide an analytical understanding of the dynamics underlying the proposed model. Although a complete analytical treatment of the full model proved intractable, we conducted an in-depth exploration of a simplified approximation.

We demonstrated that the simplified model behaves as a contractive map, analyzed the existence and stability of a dynamic fixed point, and provided an approximate range for its values. Furthermore, we introduced a probabilistic approximation grounded in basic martingale principles, modeling the system's behavior through a normal distribution. In doing so, we investigated the relationship between the true dynamics and their approximated mean and variance.

Finally, we showed that, within a regime of small σ, where the normal distribution approximation holds reasonably well, the probability of default exhibits contagion effects. This illustrates how financial distress can propagate through a network, highlighting the systemic risk posed by interdependencies between agents.

These findings offer valuable insights into the qualitative behavior of the system and lay the groundwork for further research into more complex networked financial systems.

6.1 Ethical Principles and Social Responsability

After a thorough analysis of the project, it is concluded that none of the four indicators associated with equality, environment, social responsibility, and ethics, have direct implications in the work carried out. The project's predominantly technical and methodological nature does not involve elements that could lead to gender inequalities, environmental impacts, significant social consequences, or ethical or deontological conflicts in its development or application.

Personal Note

Personally, I am glad, and proud, to have directed my Bachelor's thesis toward a research-oriented experience. As I am not planning to continue along an academic path, I leave this work as a contribution to future research. It represents the culmination of four years of study in Engineering Mathematics and Physics at the Universitat Rovira i Virgili.

7 Future Work

This work provides a foundation upon which future research can build. Several potential directions for extending this research are outlined below.

First, a more detailed examination of tail events—particularly the probability of sustained expansive sequences despite the system's overall contractive behavior—could be highly valuable, especially in risk-sensitive contexts.

Second, different distributions could be used to explore how well the real distribution behaves, making it possible to compute the exact probability of default of a single node.

Lastly, this study focuses exclusively on the two-node case, which constitutes a significant simplification of the full model. Extending the contractive map approximation to the entire network could yield a deeper understanding of the global system dynamics and how interactions propagate across nodes.

References

[1] Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.

[2] Olivier de Bandt and Philipp Hartmann. Systemic risk: A survey. *European Central Bank Working Paper Series*, (35), 2000.

[3] M. E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45(2):167–256, 2003.

[4] Alex Arenas, Adrià Barja, Alejandro Martínez, Pablo Fleurquin, Jordi Nin, José J. Ramasco, and Elena Tomás. Assessing the risk of default propagation in interconnected sectoral financial networks. *EPJ Data Science*, 8(1):35, 2019.

[5] Sara Fernandes, Carlos Ramos, Gyan Bahadur Thapa, and Luís Mário Lopes. Discrete dynamical systems: A brief survey. *Journal of the Institute of Engineering*, 2018.

[6] Arkadiusz Jadczyk. On quantum iterated function systems. *Central European Journal of Physics*, 2(3), 2004.

[7] Dimitri P. Bertsekas and John N. Tsitsiklis. *Introduction to Probability*. Athena Scientific, 2008.

[8] Antonio Cabrales, Piero Gottardi, and Fernando Vega-Redondo. Risk-sharing and contagion in networks. *Review of Financial Studies*, 30(9):3086–3127, 2017.

[9] Stefano Battiston and Guido Caldarelli. Systemic risk in financial networks. *Journal of Financial Management, Markets and Institutions*, 7(2):155–176, 2013.

[10] Jordi Nin, Bernat Salbanya Rovira, Pablo Fleurquin, Elena Tomás, Alex Arenas, and José J. Ramasco. Modeling financial distress propagation on customer–supplier networks. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 31(5):053119, 2021.

[11] Pol Arenas. Tfg-codes: Code repository for the bachelor's thesis, 2025. You can access the repository here: https://github.com/PolArenasG/TFG-Codes.