

DISS. ETH Nr. 29108

# **Modeling the Dynamics of Distribution Networks**

**A Data-Driven Approach to Supply Chains**

A thesis submitted to attain the degree of  
DOCTOR OF SCIENCES of ETH ZÜRICH  
(Dr. sc. ETH Zürich)

presented by  
AMBRA AMICO

MSc Physics La Sapienza University, Rome, Italy

born on 16.01.1993

citizen of Italy

accepted on the recommendation of  
Prof. Dr. Dr. Frank Schweitzer  
Prof. Dr. Dr. Stefan Thurner  
Prof. Dr. Alexandra Brintrup

2023

## Acknowledgments

Part of this thesis investigates how nodes in a network can face shocks by leveraging each other's resources. During my PhD journey (a significant but beneficial shock in my life), I was fortunate to have the support of several exceptional individuals. They represented my greatest resource.

First and foremost, I want to express my deepest gratitude to my supervisor Frank Schweitzer, who provided me with invaluable guidance. He encouraged me to shape my own research project and do science with clarity, rigour, and intellectual honesty. These values have marked my scientific and personal character.

I would also like to extend my appreciation to my co-examiners, Stefan Thurner and Alexandra Brintrup, for generously sharing their time and expertise in examining my thesis and providing valuable feedback.

Many thanks to all the people at the Chair of Systems Design, enthusiastic colleagues, and supportive friends: Ramona, Christian, George, Sophia, Giuseppe, Christoph, and Laurence. Each of them, for different reasons, was a great source of inspiration and new knowledge. Special thanks go to Giacomo for his invaluable scientific help. His dedication and precision helped me shape my ideas better and enlightened all our discussions, whether deepening matrix transposes or exploring the scientific impact of my research. To Luca, for his time, support, and passion he always guaranteed to my project. To Giona, for the fruitful collaboration and his contagious optimistic attitude.

To the friends from Zurich and abroad, I offer my thanks for the joyful memories we have shared. They enriched my life and provided a welcome respite from thesis writing.

Finally, I owe a deep gratitude to my family. To my parents, Carolina and Mario, for always encouraging my path, fostering my curiosity, and reminding me of my strengths when I needed it most. To my brother, Enrico, and my sister, Selene, for always being by my side. Your presence in my life makes me believe that I can overcome any challenge that appears on my path.

# Contents

<b>Contents</b>	<b>ii</b>
<b>Abstract</b>	<b>vii</b>
<b>Riassunto</b>	<b>viii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Overview . . . . .	1
1.2 Emergence of a systemic perspective . . . . .	3
1.2.1 Questioning the linear perspective. . . . .	3
1.2.2 Supply chains as complex systems. . . . .	4
1.2.3 The need for large-scale empirical analyses . . . . .	6
1.3 Contribution of the thesis . . . . .	7
1.3.1 The opioid distribution system . . . . .	7
1.3.2 Thesis structure . . . . .	7
<b>I Distribution networks: an empirical analysis</b>	<b>11</b>
<b>2 Data and network reconstruction</b>	<b>12</b>
2.1 The ARCOS dataset . . . . .	13

2.2	Reconstructing distribution networks . . . . .	19
2.2.1	From shipping transactions to networks. . . . .	19
2.2.2	Network motifs in directed tree-like structures . . . . .	21
2.3	Local topological changes under product introduction . . . . .	23
2.4	Discussion . . . . .	28
<b>3</b>	<b>Opioid distribution networks: structure and evolution</b>	<b>29</b>
3.1	Introduction. . . . .	30
3.2	The opioid consumption . . . . .	30
3.3	Network characterization . . . . .	34
3.3.1	Network size and density . . . . .	35
3.3.2	Firm heterogeneity . . . . .	37
3.3.3	Firms' positions . . . . .	42
3.4	Discussion . . . . .	44
<b>4</b>	<b>Extracting trajectories of goods from shipping records</b>	<b>47</b>
4.1	Introduction. . . . .	48
4.2	A domain-driven approach . . . . .	48
4.2.1	From shipments to trajectories . . . . .	48
4.2.2	Limitations of temporal network techniques . . . . .	49
4.2.3	Extraction method and boundary specification . . . . .	51
4.3	Descriptive data analysis . . . . .	54
4.4	Ranking nodes according to the empirical flow . . . . .	64
4.5	Discussion . . . . .	68

<b>II</b>	<b>Formation, evolution, and growth</b>	<b>70</b>
<b>5</b>	<b>Modeling network formation and evolution</b>	<b>71</b>
5.1	Introduction . . . . .	72
5.2	Model . . . . .	73
5.2.1	Two driving mechanisms . . . . .	73
5.2.2	Deriving the interaction rules . . . . .	75
5.2.3	Synthetic case . . . . .	79
5.3	Model validation . . . . .	81
5.3.1	Optimal parameters' estimation . . . . .	81
5.3.2	Statistical validation . . . . .	85
5.4	Discussion . . . . .	88
<b>6</b>	<b>Firms' growth and their determinants</b>	<b>90</b>
6.1	Introduction . . . . .	91
6.2	Growth rates and their functional form . . . . .	91
6.3	Assessment of the determining factors . . . . .	95
6.3.1	An econometric approach . . . . .	95
6.3.2	Results from the OLS regression . . . . .	97
6.4	Discussion . . . . .	99
<b>7</b>	<b>Modeling network growth: topology and goods flow</b>	<b>101</b>
7.1	Introduction . . . . .	102
7.2	Empirical evidences . . . . .	102
7.3	Model . . . . .	105
7.3.1	Demand and supply fluctuations . . . . .	105
7.3.2	Parameter calibration: an analytic solution . . . . .	107

7.4	Model validation . . . . .	110
7.4.1	Role of the upstream constraint. . . . .	110
7.4.2	Volatility of growth rates . . . . .	112
7.5	Discussion . . . . .	115
<b>III</b>	<b>Cascade dynamics and network responses</b>	<b>117</b>
<b>8</b>	<b>Upstream and downstream shock propagation</b>	<b>118</b>
8.1	Introduction. . . . .	119
8.2	Modeling cascades in distribution systems. . . . .	120
8.2.1	Related works . . . . .	120
8.2.2	Model: underlying principles. . . . .	121
8.2.3	Model: system dynamics . . . . .	123
8.2.4	Model: system response to shocks . . . . .	125
8.2.5	Measuring the cascade impact . . . . .	127
8.3	Results . . . . .	129
8.3.1	Illustrative example for system dynamics . . . . .	129
8.3.2	Firms' dependencies under cascade effects . . . . .	130
8.3.3	Indirect loss estimation. . . . .	135
8.4	Discussion . . . . .	136
<b>9</b>	<b>Flexibility: a pillar of supply chain resilience</b>	<b>139</b>
9.1	Introduction. . . . .	140
9.2	Modeling flexibility through higher-order networks. . . . .	141
9.2.1	Capturing the empirical flow . . . . .	141
9.2.2	Higher-order representation . . . . .	143
9.2.3	Formalizing flexibility . . . . .	145
9.2.4	An illustrative example . . . . .	147

9.2.5	System dynamics . . . . .	149
9.2.6	Measuring the shock impact . . . . .	150
9.3	Stress-testing the opioid distribution system . . . . .	152
9.3.1	Model initialization . . . . .	152
9.3.2	Tracing the efficient frontier . . . . .	154
9.4	Discussion . . . . .	158
<b>10</b>	<b>Summary and Conclusions</b>	<b>162</b>
10.1	Overview . . . . .	162
10.2	Summary in perspective . . . . .	163
10.3	Scientific contribution . . . . .	167
10.4	Outlook . . . . .	169
	<b>Bibliography</b>	<b>173</b>
	<b>Appendix</b>	<b>187</b>
<b>A</b>	<b>Supplementary Material to Chapter 2</b>	<b>188</b>
<b>B</b>	<b>Supplementary Material to Chapter 3</b>	<b>191</b>
<b>C</b>	<b>Supplementary Material to Chapter 9</b>	<b>193</b>
	<b>List of Figures</b>	<b>195</b>
	<b>List of Tables</b>	<b>203</b>

## **Abstract**

This thesis examines the formation, growth, and resilience of large-scale distribution systems. We investigate the interactions among manufacturers, distributors, and consumers, and show how these interactions shape the growth and resilience of these systems.

Our study begins with an empirical analysis, where we reconstruct the complete distribution networks of opioids in the United States using data from nearly half a billion shipping records. We then examine the main topological properties of these networks and analyze their stability over a nine-year period. Surprisingly, we find that despite the increasing demand for opioids, the main topological properties of the distribution networks remain stable.

To investigate how distribution systems form and evolve, we develop an evolutionary network growth model that simulates strategic link formation between firms. Testing the model against the empirical data, we show that two mechanisms are essential for the emergence of the observed networks: centralization and multi-sourcing. While centralization enhances efficiency, multi-sourcing fosters local resilience to shocks. Next, we discuss firm growth dynamics and examine how previous economic theories can be applied to the supply chain domain.

Finally, we analyze system resilience to possible disruptions. We model the propagation of supply shocks at the firm-level and discuss various system responses to mitigate them. Our focus is on the role of supply substitution as a quick strategy that we show can effectively reduce the shock impact. Our research offers a valuable tool for managers and policymakers, enabling them to devise effective mitigation strategies that can be implemented after disruptions occur.

Through a rigorous approach that combines both empirical analysis and data-driven modeling, we are able to unveil the underlying mechanisms that govern these systems. Our results contribute to both network science and supply chain management. In our attempt to bridge the gap between the two fields, we provide new methodologies based on high-resolution data to study the dynamics of large-scale distribution networks.



## Riassunto

Questa tesi esamina la formazione, la crescita e la resilienza dei sistemi di distribuzione su larga scala. Investigando le interazioni tra produttori, distributori e consumatori, analizziamo come tali interazioni plasmino la crescita e la resilienza di questi sistemi.

Iniziamo la tesi con un'analisi empirica sul sistema di distribuzione di farmaci contenuti oppioidi, negli Stati Uniti. Ricostruiamo questo sistema basandoci su circa 500,000,000 records di spedizioni, e ne esaminiamo le principali proprietà topologiche in un periodo di nove anni. Scopriamo che nonostante l'aumento della domanda di oppioidi, le principali proprietà topologiche di questi sistemi rimangono stabili.

Per capire come tali sistemi si formano ed evolvono, sviluppiamo un modello analitico che simula la formazione strategica di alleanze tra le imprese nei sistemi di distribuzione. Dimostriamo che esistono due meccanismi principali alla base della formazione dei sistemi osservati: il meccanismo di centralizzazione e quello di multi-sourcing. Il primo migliora l'efficienza di tali sistemi, mentre il secondo la loro resilienza a possibili shocks. Discutiamo poi la crescita delle imprese e l'applicazione di alcune teorie economiche al settore delle reti di approvvigionamento.

Infine, analizziamo la resilienza di questi sistemi. Sviluppiamo un modello analitico che descrive come carenze di beni possono propagarsi da un'azienda ad un'altra, e analizziamo possibili strategie di risposta a tale evento. In particolare, esploriamo il ruolo di sostituzione di beni come risposta rapida a tale shock e dimostriamo che tale risposta è in grado di mitigare l'impatto dello shock. Forniamo quindi un nuovo strumento per aiutare manager e politici a definire strategie di mitigazione di possibili carenze di beni.

Grazie ad un approccio rigoroso che combina analisi empirica e modellizzazione basata sui dati, siamo in grado di delineare i meccanismi che governano le dinamiche principali di tali sistemi. I nostri risultati contribuiscono sia alla scienza delle reti che alla gestione della supply chain. Nel nostro tentativo di colmare il divario tra i due campi, forniamo nuove metodologie basate su dati ad alta risoluzione per studiare i sistemi di distribuzione su larga scala.

# Chapter 1

## Introduction

“The essence of strategy is choosing what not to do.”

*M.E. Porter*

### 1.1 Overview

As humans rely on each other’s capabilities to achieve goals and deliver projects, so do firms. Within a supply chain, firms depend on each other’s capabilities to produce and deliver goods to end consumers (Brintrup and Ledwoch, 2018). This is accomplished through the well-known outsourcing practice, in which production and distribution responsibilities are shared among firms (Schoenherr, 2010).

Over the past two decades, outsourcing has emerged as a popular business paradigm for several companies in response to the growing market competition. Firms gained various benefits by collaborating as part of a supply chain rather than in isolation. They could reduce costs through lower inventory buffers and economizing the supply base while maintaining high profits (Altıparmak *et al.*, 2009; Mizgier *et al.*, 2013). In addition, outsourcing the distribution process

has facilitated efficient deliveries and shorter lead times, ultimately improving customer service performance (Martínez-Olvera *et al.*, 2015).

The emergence of highly interconnected supply chains has resulted in far-reaching transformations, impacting not only business operations but also national economies, international trade, and society. Most supply chains are not confined to the border of a single country but span the globe, with various stages of production and distribution taking place in different regions. This global context has significantly altered international trade, promoting economic growth in low-income countries (Gereffi, 2014). As Gereffi (2019) noted the global dimension of trade is not a new phenomenon, but the speed and scale of these interactions are.

Besides firms and countries, the good functioning of a supply chain directly affects a third actor: the end consumer. People's lives rely on several supply chains to access goods and services. Disruptions to these systems can adversely impact people's well-being and, in some cases, even threaten their lives (Schueller *et al.*, 2022). Examples of critical supply chains include those providing essential goods such as food and medicines. The recent Covid-19 pandemic is a major example. Production shutdowns due to quarantine policies generated significant shortages of food and medical equipment, thus harming millions of lives. Before the Covid-19 pandemic, various incidents such as natural disasters, for instance, the earthquake that hit Japan in 2011, and labour issues, such as the fire at Toyota's supplier in 1997, have affected supply chains and, in turn, entire populations.

All these events have raised awareness of the fragility of today's supply chains (Diem *et al.*, 2022; Reisch *et al.*, 2022). They look unprepared to address modern challenges, including the ongoing climate crisis and future pandemics. As Professor J. Byrnes stated, "many supply chains are perfectly suited to the needs that the business had 20 years ago". However, whether they suit today's needs is still unclear. Many questions have been raised. How do supply chain structures evolve? To what extent can they bear or amplify local shocks? What policies can be implemented to enhance their resilience to disruptions? This thesis aims to explore possible answers to these questions.

## 1.2 Emergence of a systemic perspective

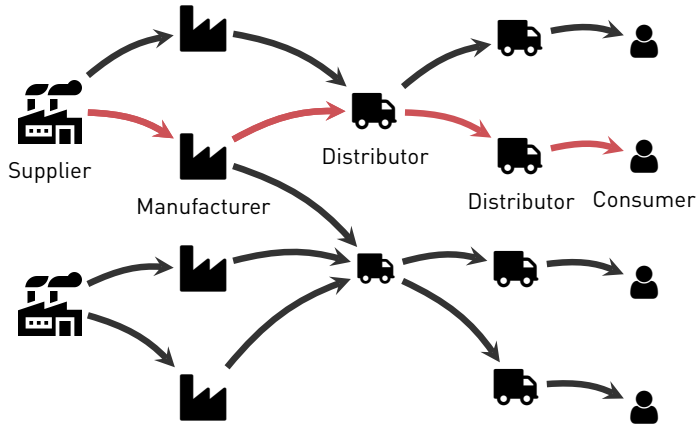
### 1.2.1 Questioning the linear perspective

Most prior supply chain research has assumed a linear structure, where a “chain” connects a focal manufacturer to its upstream suppliers and downstream distributors. However, modern supply chains are far more complex. Firms often establish connections with multiple suppliers and distributors, who in turn are connected to other firms. This results in the emergence of large-scale inter-firm networks, overlooked by this linear perspective. A schematic representation of a supply network is shown in Fig. 1.1. The linear perspective approach considers only the single “chain” of firms connected by the red links in the representation, thus overlooking the whole network.

According to Brintrup and Ledwoch (2018), there are two reasons why this linear perspective is still vastly used. The first reason is the lack of data. Many analyses are restricted to the data companies provide, if they provide. Firms tend to keep their partnership information private; therefore, the data available for analysis are usually limited in scale. This is because they fear that disclosing such information may weaken their bargaining power and benefit their competitors.

The second reason is what we can call an “emotional attachment” to established techniques. In the engineering supply chain domain, there are well-established tools for optimizing goods flow concerning some firm-level variables, e.g., costs or delivery times. For computational reasons, these models hardly deal with large-scale analysis. Further, they are not designed to incorporate network effects.

We argue that complexity science can provide the appropriate tools to move beyond the oversimplified linear perspective in favour of a more realistic and large-scale one.



**Figure 1.1:** Schematic representation of a supply network. A single supply chain, consisting of one supplier, one manufacturer and two distributors, is highlighted by the red links in the figure.

### 1.2.2 Supply chains as complex systems

In this thesis, we use the tools from complex networks to study supply chains. Following the pioneer work of Choi *et al.* (2001), we view supply chains as complex systems: they emerge spontaneously from the interactions of a large set of entities and are not controlled by a single one. A complex network, then, is a specific formal representation of a complex system.

Scholars have used this representation to successfully study the structure and the dynamics of several economic systems, e.g., financial systems (Battiston *et al.*, 2012b), global trades (Burkholz and Schweitzer, 2019), and inter-firms collaboration networks (Tomasello *et al.*, 2014). We argue that also supply chains can be universally abstracted to networks, where links represent supply relations and nodes represent firms (i.e., suppliers, manufacturers, distributors) and consumers. We identify four key factors that allow us to qualify such systems as complex systems. These are listed below. Note that throughout the thesis, we mainly use the term “supply networks” to refer to supply chains.

**Self-organized systems.** Supply networks are not designed by a single company but emerge spontaneously. Firms can select their direct suppliers and distributors, but they have limited control over indirectly related firms, i.e., those that are connected to their suppliers and distributors. These indirect connections are still very important for firms because they can influence their performance and risk exposure.

**Heterogeneity of economic actors.** The economic actors that are part of a supply network are very diverse, as are the tasks they perform. We identify four economic actors: suppliers, producers, distributors and consumers. The inherent diversity of these actors fosters the formation of supply links. Manufacturers connect to suppliers because they need raw materials for their production inputs; distributors connect to producers because they need finished goods for delivery; and consumers connect to distributors to access goods and services. Then, these connections enable supply networks to function.

**Product heterogeneity.** Several products are supplied using interconnected and even overlapping supply chains. Depending on the given product, the same firm may contribute differently to the overall production and distribution process. Let's consider the example presented by Brintrup and Ledwoch (2018): the production chain of a car. A given supplier produces both fabrics for car seats and assembled seats. This supplier sells the seats to one manufacturer and only the fabrics to a second manufacturer, which has not outsourced seat production. Any disruptions at this supplier may affect the two manufacturers differently. A similar situation may occur on the distribution side of a supply chain. A given distributor may sell cars directly to consumers and, for other vehicles, outsource this to retail distributors. Product heterogeneity adds further complexity, leading to more complex structures than simple chains.

**High dynamism.** Supply networks are intrinsically dynamic (Wycisk *et al.*, 2008). Their structure constantly changes as firms adjust their business operations. New firms may enter the market while others exit. Supply links may also form or terminate, and some firms may introduce or withdraw

products (Wycisk *et al.*, 2008). Moreover, supply networks may display non-linear behavior (Choi *et al.*, 2001). Local shocks may propagate through the supply links even with amplification mechanisms (Carvalho *et al.*, 2021). Dynamic models from complexity science can serve as a foundation for analyzing such phenomena.

### 1.2.3 The need for large-scale empirical analyses

Over the past two decades, several works have been published to support the view of supply chains as complex systems and to encourage the use of network-based tools (Brintrup *et al.*, 2015; Hearnshaw and Wilson, 2013; Pathak *et al.*, 2007; Perera *et al.*, 2017a; Wycisk *et al.*, 2008). Although the importance of this emergent perspective has been vastly acknowledged, few works have adopted it in practice. Complexity science is still in its infancy in the supply chain domain compared to other fields, e.g., social science, political science, biology, or ecology.

The progress of complexity science in the field of supply chains has been hindered by a lack of comprehensive data. Currently, there is a lack of not only large-scale data, but also fine-grained data capturing, for instance, temporal interactions between firms. As a result of limited data availability, previous research has been constrained to case study analyses of a small number of firms or static network approaches (Brintrup *et al.*, 2015; Dong *et al.*, 2020; Kim *et al.*, 2011; Perera *et al.*, 2017a; Potter and Wilhelm, 2020).

To deepen our understanding of how these networks form, function on a daily basis, and respond to shocks, more empirical insights and dynamic models tested on large-scale data are needed. In response to this need, this thesis contributes to the field by providing novel empirical analysis on a large-scale distribution system, as detailed in the following Section.

## **1.3 Contribution of the thesis**

### **1.3.1 The opioid distribution system**

This thesis examines a prominent example of a large-scale distribution system. We study, for the first time, the US distribution system of opioid drugs, which serves over 200,000 pharmacies, hospitals, and practitioners nationwide and connects over 1,500 firms, i.e., opioid manufacturers and distributors. To our knowledge, this is the largest distribution system analyzed to date. Based on approximately 500,000,000 shipping records, we empirically reconstruct and characterize the distribution system of various opioids over a nine-year period.

Unlike most previous research, our study does not focus on the supply chain of a single company. Instead, we examine the interconnected supply chains of numerous firms within a single, very wide nation. Through this comprehensive and systemic approach, our research makes a unique empirical contribution to the field of supply chain management.

In addition to empirical results, this thesis presents a theoretical framework for the analysis and modeling of large-scale distribution systems. The extensive available data allow for comprehensive and rigorous testing of our models. It is worth noting that our goal is to use the opioid distribution system as a test bed for our analysis. As we do not want to limit our contribution to the pharmaceutical industry, we strive to formulate general hypotheses that go beyond the narrow opioid market.

### **1.3.2 Thesis structure**

The thesis comprises 10 Chapters. The current Chapter introduces the thesis by explaining the underlying motivations and providing an overview of the theoretical background. Then, the thesis is structured in three Parts according to three different research lines: (i) empirical analysis, (ii) dynamic models to explain network formation and growth, (iii) and dynamic models to assess network resilience.



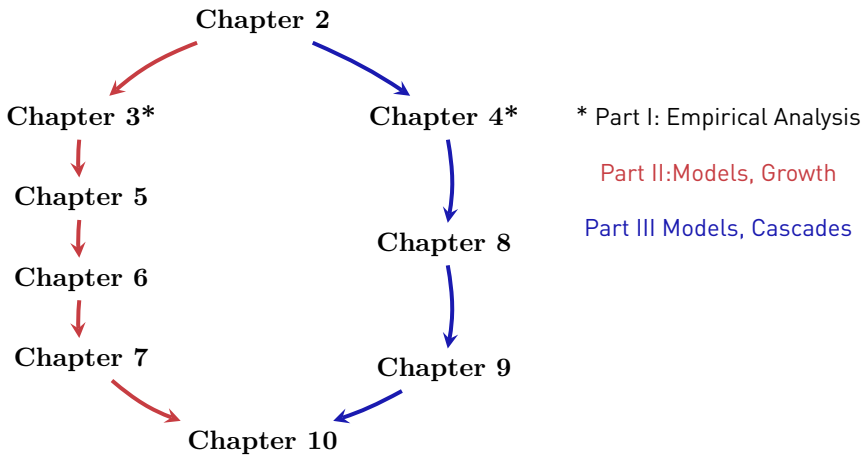
**Part I: large-scale empirical analysis.** The first Part of the thesis provides an in-depth characterization of the system under study. We first describe the data used to reconstruct the opioid distribution networks in Chapter 2. In Chapter 3, we present a topological characterization of those networks and discuss their evolution over a nine-year period. Next, in Chapter 4, we reconstruct and characterize the daily flow of drugs moving through the distribution networks; and we discuss possible methods for identifying prominent distributors where the flow is concentrated.

**Part II: formation, evolution, and growth.** Real-world distribution networks are large, self-organized, and heterogeneous in both types of economic actors and the type of material flow moving through the network. As a result, a highly complex structure emerges. Part II of this thesis poses the question: How do these complex structures originate, evolve, and grow?

We begin to address this question in Chapter 5. Here we develop a network growth model to explain the formation of such systems from a topological perspective. We aim to design link-formation rules that can explain the emergence of the observed structures. In Chapter 6, we use statistical tools to evaluate the growth of firms operating within the opioid distribution networks. Finally, in Chapter 7, we extend the model presented in Chapter 5 to describe how distribution networks can grow both in size (i.e., by increasing the number of links and nodes) and volume (i.e., by increasing the quantity shipped).

The models developed are calibrated and validated against the empirical data. It is worth noting that our goal is not to predict the system properties accurately. Instead, our goal is to understand the few but sufficient principles that govern the growth of these networks and, thus, their constituent firms. In doing so, we strive towards parsimonious models with few interpretable parameters that can be tuned and calibrated against firm-level data.

**Part III: cascade dynamics and network responses.** Cost-reduction strategies and just-in-time delivery principles have led to highly functional distribution networks. To what extent are their structures also resilient to



**Figure 1.2:** Roadmap for reading.

disruptions? How do supply shocks propagate on these networks? Also, what are the possible response strategies, and how can we model them?

These questions are addressed in Chapter 8 and Chapter 9. Specifically, in Chapter 8, we develop an Agent Based Model (Schweitzer, 2022a) to study how local supply shocks may propagate through the network and impact other firms not directly affected by the shock. Next, in Chapter 9, we investigate to what extent supply substitution policies can alleviate potential shortages and enhance system resilience.

In contrast to Part II, the focus of this part is not on validating the models using the data. Rather, the data are utilized to inform the models developed. Our objective is to incorporate various features of the real-world systems into the models to begin with the most realistic representation. Subsequently, through “what-if” scenarios, we simulate the proposed mitigation strategies and examine their efficacy.

**Roadmap for reading.** Fig. 1.2 shows different paths the reader can take when approaching this thesis. We recommend starting from Chapter 2, where we describe the data used throughout the thesis. Then, the thesis develops

along two main streams: network growth (Part II) and network resilience (Part III). We identify two separate reading paths because we provide different methodologies to tackle the two topics.

Path (1), shown in red on the left-hand side of the Figure, begins with Chapter 3, which provides a comprehensive characterization of the empirical system. The insights gained from this analysis can provide a foundation for understanding the growth model presented in the subsequent Chapter 5. It is recommended to continue reading Chapter 6, where we explore firm growth in distribution systems. The knowledge gained in both Chapter 5 and Chapter 6 is then necessary to understand Chapter 7.

Path (2), in blue on the right-hand side of the Figure, starts with Chapter 4. This Chapter reconstructs the empirical trajectories of drug packages moving through the opioid distribution system. These trajectory data serve as input of the model developed in Chapter 9. We recommend that the reader proceeds to Chapter 9 after reading Chapter 8, as the latter introduces the mathematical formalism required to understand the analysis presented in Chapter 9. Therefore, both Chapter 4 and Chapter 8 are prerequisites for reading Chapter 9.

Following either of the two paths depicted, the reader can eventually reach Chapter 10. In this Chapter, we summarize the main findings, draw conclusions from the three research streams, and outline possible directions for future research.

## Part I

# Distribution networks: an empirical analysis

“Without data you’re just another person with an opinion.”

*W. E. Deming*

## Chapter 2

# Data and network reconstruction

### Summary

This Chapter presents an overview of the ARCOS dataset, which contains approximately 500 million records of opioid drug shipments in the United States. These records span a period of nine years, from 2006 to 2014, and have a daily resolution. The dataset is a central component of all Chapters in this thesis and has not been thoroughly studied by the scientific community. Therefore, we provide a detailed description of the data's source, composition, and our procedure for reconstructing distribution networks from shipment records. Finally, we explore the local topological properties of the reconstructed networks using network motifs.

---

This Chapter has been written specifically for this thesis.

## 2.1 The ARCOS dataset

**On the data origin.** Starting from 1997, the US Drug Enforcement Administration (DEA) has introduced a data collection software for keeping track of all the shipments of opioids across the country\*. The software utilized to collect the shipping records is called ARCOS, an acronym for “Automation of Reports and Consolidated Orders System”. The system has been in place for many years, and it has collected billions of shipping records. For several years, these data were kept confidential. Only in 2019, the Federal Court ordered the DEA to release the data to the public.

The Court decision was the final stage of a protracted lawsuit between a set of drug companies and a civil union comprising 2,000 cities, towns and counties. The companies were accused of overselling the drugs for their economic interest. In the course of the litigation, the civil union gained the support from the owners of two leading US newspapers, namely “The Washington Post” and the “Charleston Gazette-Mail”, which played an essential role in the data release.

The shipments recorded between 2006 and 2014 have been made public and are now available on two platforms: the GitHub page of the Washington Post (Washington Post Investigative, 2019) and the website of the SLCG consulting group (SLCG, 2019)<sup>†</sup>. While the Washington Post published only a subset of the data (i.e., shipments of Hydrocodone and Oxycodone drugs), the SLCG made the complete raw data available on its website and provided a preprocessed version of the data. In this thesis we analyze the preprocessed version of the data downloaded from the SLCG website (SLCG, 2019). To the best of our knowledge, this project is the first work to utilize these data for scientific purposes.

---

\*For completeness, it tracks not only opioids but all “controlled substances”. These are drugs, or other substances, controlled by the government since they may cause addiction, e.g., opioids, stimulants, depressants, hallucinogens, and anabolic steroids. However, only the shipping transactions of opioids have been made publicly available. Hence, in the following, we will always refer to the opioid data.

<sup>†</sup>The SLCG was responsible for checking data anomalies during the opioid litigation.

**A complete dataset.** The ARCOS dataset comprises **499,534,836** shipping transactions recorded between 2006 and 2014. The dataset is complete, meaning it stores all national, yet legal, shipments confined to the US borders. Specifically, all the 50 US states, six US territories (i.e., district of Columbia, Puerto Rico, Guam, US Virgin Islands, American Samoa, Northern Mariana Islands), and four service lands (i.e., Armed Forces-Americas, Armed Forces-EMEA, Armed Forces-Pacific, and the Palau Republic) are considered.


Shipments involve about **300,000** economic actors: manufacturers (610), distributors (1,318) and actors who only receive packages (299,344), including for instance hospitals, pharmacies and doctors. According to the DEA regulations<sup>‡</sup>, all opioid manufacturers and distributors must report their shipments periodically, about once a year, via the ARCOS software. They must specify several details, including the shipping locations, times, and product type. Thus, the data collected boast a very high resolution concerning three main aspects: (i) the economic actors; (ii) the products shipped; (iii) and the shipping dates. We illustrate these three aspects below.

**Actors.** Each actor is identified by a unique alphanumeric code called *dea-number*, which is assigned by the DEA during registration. In addition to the *dea-number*, the data also contain the names, locations, and business types (e.g., manufacturers, distributors) of the companies involved. It is worth noting that the *dea-number* is assigned to a specific business location rather than the entire company. Thus, multiple business locations, even if they belong to the same company, are assigned to different *dea-numbers*.

**Drug products.** 18,677 different products are reported in the dataset. Each product is uniquely identified through a National Drug Code (*ndc*). This code usually appears on the package tag, as illustrated by Fig. 2.1. The *ndc* is an 11 digits number which comprises three segments encoding three different pieces of information: (i) the first five digits indicate the manufacturer or labeller; (ii) the second four digits indicate the drug type, i.e., its formulation, the active

---

<sup>‡</sup>More precisely, the “Controlled Substances Act of 1970 (§ 827)” created the requirement for the drug manufacturers and distributors to report their transactions to the Attorney General. Then the Attorney General delegated this authority to the DEA.



**Figure 2.1:** Illustrative picture of a drug product uniquely identified by a National Drug Code (*ndc*), 11 digits long. The first five digits encode the manufacturer (labeler); the second four digits encode the active ingredient, strength, and formulation (drug); the last two digits encode the package size (package). The code 59011-440-10 identifies Oxycodone tablets in 40 mg (segment 440) that are produced by the labeler Purdue Pharma (segment 59011) and packed in boxes of 100 tablets (segment 10). Credits: the photo of the drug package has been taken from an online article published by the journal *The Guardian*. Photograph: George Frey/Reuters.

ingredient, and strength; (iii) the last two digits indicate the package size, e.g., package of 100 tablets, package of 50 tablets.

Each product contains only one active ingredient or basic opioid, e.g., Fentanyl, Hydrocodone, or Oxycodone. There are fourteen different basic opioids in the dataset. The complete list is given in Table 2.1, along with the total number of shipping transactions and the units sold to consumers between 2006 and 2014.


We notice that two big players dominate the opioid market: Hydrocodone and Oxycodone. Jointly, their number of transactions encompasses more than 60% of the total. We obtain a similar result if we look at the number of units sold. We see that the total number of transactions of Oxycodone and Hydrocodone drugs is (at least) one order of magnitude bigger than the one associated with the other basic opioids.

The observed predominance of the two opioids is also reflected in the number of products. In Fig. 2.2, the total number of products (y-axis) per basic opioid (x-axis) is shown. We see that there are almost 6,000 products containing Hydrocodone and almost 3,000 products containing Oxycodone. The former corresponds to 32% of all the products, the latter to 15%. From Fig. 2.2, we also note a relatively large number of Fentanyl-based products (2,638), comparable to those containing Oxycodone. In contrast, we find a much lower number



Basic opioid	Transactions	Units sold	Transactions (%)	Units sold(%)
HYDROCODONE	174,285,342	4.153850e+10	0.3943	0.46115
OXYCODONE	105,477,472	1.557179e+10	0.2386	0.17287
FENTANYL	42,868,242	7.003753e+09	0.0969	0.07775
MORPHINE	39,066,352	1.213574e+10	0.0883	0.13473
BUPRENORPHINE	24,063,069	4.330039e+09	0.0544	0.04807
CODEINE	14,539,286	1.621039e+09	0.0328	0.01799
HYDROMORPHONE	14,466,475	4.812868e+09	0.0327	0.05343
METHADONE	14,069,036	8.150097e+08	0.0318	0.00904
OXYMORPHONE	5,013,407	1.318083e+09	0.0113	0.01463
MEPERIDINE	4,138,042	1.621182e+08	0.0093	0.00180
TAPENTADOL	2,901,272	3.388228e+08	0.0065	0.00376
DIHYDROCODEINE	752,273	4.141162e+08	0.0017	0.00459
OPIUM, POWDERED	297,356	6.878540e+06	0.0006	0.00007
LEVORPHANOL	63,414	5.518305e+06	0.0001	0.00006


**Table 2.1:** List of the fourteen basic opioids in the data, along with the number of shipping transactions and units sold to consumers. Absolute and percentage values are displayed. Opioids are sorted according to the number of transactions. The top two are highlighted in red.



**Figure 2.2:** Number of products (y-axis) per basic opioid (x-axis).

of shipping transactions for Fentanyl than for Oxycodone (see Table 2.1). The high potency of Fentanyl may explain this observation. Fentanyl is 100 times stronger than Morphine. Presumably, it cannot be prescribed regularly; therefore, it is shipped with much less frequency (i.e., fewer transactions).

Finally, as we observe a predominance of very few basic opioids in the market, we also observe a predominance of very few products. For Hydrocodone, only 60 products, corresponding to the 12% of products, appear in 80% of the shipments. Similarly, for Oxycodone, we find that only 60 products, corresponding



**Figure 2.3:** Number of new products (blue line) and products withdrawn from the market (red line) per year.

to 5% of products, appear in 83% of the shipments. The full list of products is reported in Table A.2 in Appendix A. These preliminary observations already reveal that the opioid market is organized so that very few players, in terms of products and basic opioids, dominate the shipments.

**Shipping dates.** The dataset covers nine years, from 2006 to 2014. Shipping dates are reported with a daily resolution, i.e., in the day-month-year format<sup>§</sup>. During the observation period, a quite pronounced turnover of drugs is observed. In Fig. 2.3, we show the number of new products entering the market every year (blue line) and the number of products exiting the market every year (red line)<sup>¶</sup>. We define as *new products* those that appear in the data, for the first time, in 2008 or later (i.e., no transactions recorded in the first two years). Similarly, we define as *withdrawn products* those that appear in the data, for the last time, in 2012 or earlier (transactions not recorded in the last two years). We find that approximately 1,000 products enter the market annually ( $\approx 5\%$  of products). A slightly decreasing trend is observed between 2008 and 2014. In contrast, an increasing trend characterizes the number of

<sup>§</sup>Note that the dates when orders are placed or packages delivered are not reported in the data and may differ from the shipping dates.

<sup>¶</sup>For these statistics, we do not distinguish between products with different packaging (i.e., by the last two *ndc* digits) but only with different compositions (i.e., by the first nine *ndc* digits).

products withdrawn from the market. The number of products that exit the market increases from about 300 in 2006 to about 1,000 in 2012.

The observed turnover may be due to various factors, such as the introduction of experimental drugs (on the market for short periods), or the effect of government interventions to reduce opioid prescriptions by favouring the use of others (with different compositions). This market dynamism may impact the structure of the underlying distribution systems. For example, the entry of a new drug may induce changes in the network structures. We investigate this aspect in Section 2.3 of the current Chapter.

**Preprocessing.** The data used for this thesis was obtained from the SLCG website (SLCG, 2019). The dataset is about 977 GB in size, which is too large to handle without proper management. To overcome this, we stored the data in a PostgreSQL database on our local server.

We used the preprocessed version of the data provided by SLCG, which had already been checked for anomalies during the opioid litigation. We did not encounter any additional anomalies during our analysis. Furthermore, we assumed that there were no missing records since all shipments are required by law to be reported.

When we analyzed the import and export transactions, we discovered that they accounted for only about 0.01% of the total records. This finding indicates that the opioid manufacturing and distribution system is a closed system, thereby eliminating concerns about missing information.

We divided the actors into three distinct groups: (i) manufacturers, (ii) distributors, and (iii) end consumers. The end consumer category includes all actors who are not distributors or manufacturers and who receive drugs, such as hospitals, clinics, practitioners, and pharmacies. Although they do not consume the drugs directly, we use the term “end consumers” since they represent the final destination of the distribution process. The data include 54 different business activities. They are reported in Table A.1.

Finally, we convert these shipping data into distribution networks. Given the importance of this step in our analysis, it is crucial to carefully consider the theoretical assumptions and simplifications involved, as they can significantly

impact the reconstructed distribution networks and subsequent analyses. We provide a detailed description of the data processing steps in the following Section.

## 2.2 Reconstructing distribution networks

### 2.2.1 From shipping transactions to networks

Throughout this thesis, we adopt a complex network approach and abstract distribution systems as networks. Although constructing a network representation of these systems may seem straightforward, it is not when the starting point is a large and multifaceted dataset. Given a set of shipping transactions, multiple representations of the same system can be constructed depending on how these transactions are aggregated and preprocessed. As Butts (2009) noted, a network representation is a theoretical act. As such, it always involves assumptions and simplifications that need to be made explicit (Peel *et al.*, 2022). This Subsection aims to explain the procedure we follow when converting the available data into networks and the limitations associated with such a network representation.

In our network representation, *nodes* represent firms, i.e., manufacturers and distributors, and *links* represent supply relations. We assume that a supply relation is established between two firms whenever at least one shipment has been observed. Specifically, we always consider an annual time-window to determine the existence of a supply link. Links are *directed* according to the direction of the shipments, from sender to receiver.

Note that we are interested in the distribution process that ensures that goods are moved from producers to consumers. We are not interested in the so-called “reverse distribution” process occurring in the opposite direction: expired or damaged products are returned back to manufacturers or destroyed. Therefore, we make two choices. First, we exclude distributors that are in charge of destroying packages. These are labelled as “reverse distributors” in the data, according to the *bus act*<sup>||</sup>. Second, we exclude back shipments. Given two

<sup>||</sup>Further, note that all the records are labelled by a code defying the transaction type, e.g., sale, return, package destroy. We exclude shipping transactions of destroyed or returned

firms, A and B, it is possible to observe shipments in both directions, i.e., from A to B and from B to A. To correctly identify the directionality of the supply link, the shipping volumes are compared, and a link is created in the direction of the larger shipment\*\*.

Since the dataset contains thousands of different products, millions of distribution networks could be constructed based on the products (or combinations of products) considered. In this thesis, we always consider distribution networks of multiple products. Specifically, we deal with two classes of distribution networks: what we call (i) *single-manufacturer* distribution networks and (ii) *single-opioid* distribution networks. The distribution networks of class (i) are reconstructed by aggregating transactions of products belonging to the same manufacturer. These products share the same label code, the first five digits of the *ndc*. The distribution networks of class (ii) are reconstructed by aggregating transactions of products containing the same basic opioid (even produced by different manufacturers).

Depending on the focus of the Chapter, we deal with one of the two classes. Specifically, in most Chapters, i.e., Chapter 3, Chapter 6, Chapter 8 and Chapter 9, we use the networks of class (ii), as we are interested in the large-scale analyses of nationwide distribution networks. In Chapter 7 and Chapter 5, instead, we use the networks of class (i) as we aim at modelling the formation of distribution systems with a single root, namely a single manufacturer.

Finally, we want to clarify how consumers are abstracted in our network representation. We first notice that the total number of consumers is 299,344. Instead, the total number of firms (distributors plus manufacturers) is 1,928, namely two orders of magnitude smaller. This represents both a computational and conceptual challenge for our network analysis. Using a network abstraction where all consumers are represented as single nodes would completely mask firms' behavior and only highlight the consumers' behavior.

---

packages and only consider transactions labelled by code "S" (i.e., almost 90% of the total number of transactions). This indicates an action of "sale, disposition, or transfer".

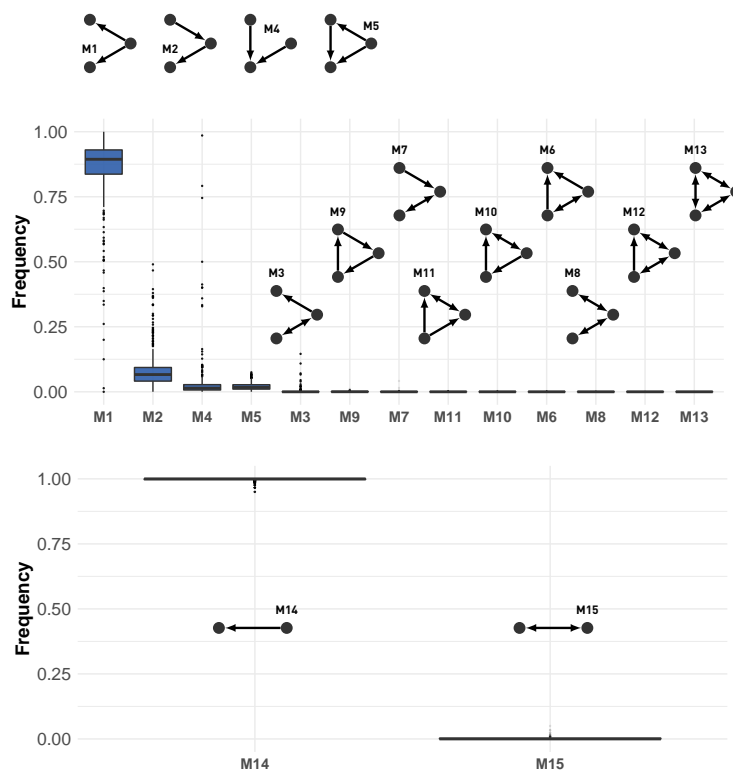
\*\*This is done when one of the two shipments occurs with a much lower frequency. Therefore, it is clear that one follows the direction of the distribution and the other represents a back shipment.

To solve this challenge, we use *representative* consumers. A representative consumer is obtained from aggregating the whole set of consumers to whom a distributor ships. Moreover, we are mainly interested in firms' rather than consumers' behaviors. Therefore, consumers will not be part of our network abstraction throughout the thesis. The sole exception are Chapter 5 and Chapter 7. In these Chapters, we investigate microscopic mechanisms to explain the evolution of the network topology and the goods flow. As we want to reproduce the flow that reaches consumers, we model representative consumers as nodes of the network.

### 2.2.2 Network motifs in directed tree-like structures

This Subsection presents a preliminary characterization of the distribution networks reconstructed from the ARCOS by analyzing their *local* topological properties. Specifically, we analyze small subgraphs of these networks often called network motifs (Milo *et al.*, 2002). These are recurring relational patterns between small groups of nodes, e.g., two nodes (dyadic motifs) or three nodes (triadic motifs). Network scientists have vastly studied motifs as they have been found to carry meaningful information on the functions of the underlying systems. They have been used, for example, to classify regulatory mechanisms in transcription regulation networks (Alon, 2007), or identify early-warning signals of the topological collapse of the Dutch inter-bank network (Squartini *et al.*, 2013).

Clearly, we do not expect that all the previously studied subgraphs that have been meaningful for other networks would also be meaningful in the context of distribution networks. We expect to observe patterns in the data that are consistent with a hierarchical tree structure, in which a manufacturer node serves as the root and is connected to child nodes (distributors) along the tree. Such structures are commonly found in distribution systems as they facilitate efficient deliveries and reduce costs (Perera *et al.*, 2017a). Direct loops or bidirectional connections, on the other hand, are unlikely to occur frequently. These patterns indicate backward shipments or shipments going in circles, which suggests inefficient distribution processes.



**Figure 2.4:** Box plots displaying the frequencies of the dyadic and triadic motifs. The sample comprises the distribution networks of 100 opioid producers in the US.

Yet, the question of the most frequent subgraphs in real-world distribution networks remains. To answer this question, we reconstruct the distribution networks of 100 (the largest) opioid producers. For each network, we count the number of dyadic and triadic motifs. In Fig. 2.4, we display the normalized count, i.e. frequency, of the dyadic and triadic motifs. The box plot shows the mean and variance of the motif frequencies obtained in our sample.

By construction, motifs representing bidirectional relations have zero frequency. On the contrary, all motifs with unidirectional links and not representing cycles are indeed observed. These are the four triadic motifs M1,

M2, M4 and M5 and the unidirectional link M14. The most observed triad is M1, followed by M2, M4 and M5. M1 represents a single firm that ships to two other firms. The more branched the network, the less frequent this pattern is. Think of a single long chain of distributors. In this case, M1 is absent. Conversely, the more centralized the network (i.e., a single distributor that ships to all others), the more frequent this pattern is. The very high frequency of these motifs suggests high centralization of the network.

M2 represents a path of length two in the network. M4 represents a firm receiving goods from two different partners. M5 has a similar interpretation but with a redundancy: the two partners are themselves connected. The configurations expressed by M4 and M5 are not possible in a perfect tree. By definition, nodes in a tree can have only one parent.

This preliminary investigation already provides us with two main insights. First, distribution networks exhibit structural patterns (i.e., M4 and M5) that distinguish them from perfect trees. Second, the pattern expressed by M1 (centralization) is predominant. We will return to these two insights in Chapter 5, while developing the mechanisms that drive the network formation.

## 2.3 Local topological changes under product introduction

**New products and topological changes.** In Section 2.1, we noted that the opioid market is quite dynamic: about 1,000 products enter the market each year. In this Section, we want to test the extent to which changes in market conditions due to a new product introduction can affect distribution networks and induce significant changes in their topologies.

According to Pero *et al.* (2010), introducing a new product can prompt manufacturers and distributors to reconfigure their business relations, thus, causing structural changes in the underlying distribution network (Caridi *et al.*, 2012; Zimmermann *et al.*, 2016). To test whether this phenomenon is actually observed in real-world networks, we reconstruct the distribution networks of the opioid producers that introduced a new drug between 2006 and 2014. Next, we use network motifs to capture topological properties of these networks and



apply the Differences-in-Difference (DiD) statistical technique to detect causal effects.

**Difference-in-Differences.** DiD is a statistical technique used to estimate the causal effect of a given event, or treatment, on some outcome variable (Gertler *et al.*, 2016). The statistical sample is split into two sub-samples: (i) the treatment sample, that is, the set of entities that do experience the event; (ii) the control sample, that is, the set of entities that do not experience the event.

DiD has been widely used in economics and the social sciences to estimate, for example, the effect of an economic policy or the passage of a law. In these fields, it is difficult to arrange experiments that control for differences between control and treatment samples, such as randomized control trials (RCTs)<sup>††</sup>. A DiD approach is a valid alternative to RCT to enable statistical control. It corrects for differences between treatment and control groups by comparing the *trends* of the outcomes rather than the outcomes directly. In other words, it measures the *changes* in outcomes over time (before and after the event) between the treatment and control samples.

**Testing the impact of new product introduction.** To assess the effect of the product introduction, we start constructing two data samples: (i) the treated sample (ii) and the control sample. The treated sample comprises distribution networks of the manufacturers that *did* introduce one new product between 2006 and 2014. The control sample comprises distribution networks of the manufacturers that *did not* introduce any new product between 2006 and 2014. For (i), we *do* expect local topological changes. For (ii), we *do not* expect local topological changes.

---

<sup>††</sup>RCT is a setting experiment used to control factors that differ across control and treatment groups. Individuals that join an experiment differ from one another, and these differences can influence study outcomes. One way to control for that is to set up an RCT experiment where individuals are randomly allocated among control and treatment groups. RCTs are used in clinical trials to test the effect of a drug or, in general, a medical treatment. If we think instead of a road repair program, operated at the district level, it is less feasible to set up RCT. One can not randomly assign a road repair program to districts, as they have the faculty to decide whether they want to enrol or not enrol in the program.

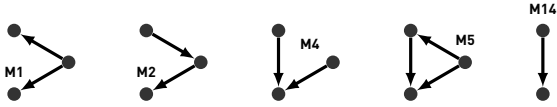
Note that to produce meaningful statistics, we ensure that: products introduced remain on the market for at least two years after their first appearance (in this way, we exclude any trial products that we do not expect to have significant effects on existing business relations); a minimum size (i.e., 20 nodes) characterizes the networks analyzed. The data composition is the following: the treatment sample includes 18 distribution networks; the control sample comprises 28 distribution networks. Among these 28 networks, 7 networks experience the introduction of the new product in 2009, another 7 networks in 2010, and the remaining 4 networks in 2012.

Next, we use network motifs to capture local properties of the empirical networks. Specifically, we examine the motifs M1, M2, M4, M5, and M14, discussed in Section 2.2.2, and run a DiD model for each of them. The outcome variable is the motif count. Significant increases, or decreases, in this outcome would signal (local) structural changes due to the product introduction. The following equation expresses the model:

$$Y_{igt} = \alpha_t + c_g + \beta D_{gt} + \theta X_{igt} + \varepsilon_{it} \quad (2.1)$$

In Eq. (2.1),  $Y_{igt}$  is the outcome variable (motif count); the subscript  $i$  labels the distribution network (unit of analysis);  $g$  indicates whether the network is either in the control sample or in the treatment sample; and  $t$  indicates whether the network is considered either before or after the event. Further,  $\alpha_t$  is the time fixed effect,  $c_g$  is the sample fixed effect,  $\beta$  is the treatment effect,  $X_{igt}$  captures network-level control variables and  $\varepsilon_{it}$  is the error term.

For each treated network  $i$ , we consider four observations: two observations at the two years before the event (new drug introduction) and two observations after it. Similarly, for each control network  $i$ , we consider four observations, two before the year of the event, and two after it. Specifically, the control sample is used repeatedly for each year of the event. According to the standard DiD model, both time and fixed effects are dummy variables. Hence,  $\alpha_t = 0$  if  $t$  is the time before the product introduction, and  $\alpha_t = 1$  if  $t$  is the time after it. Similarly,  $c_g = 0$  if  $g$  is the control sample, and  $c_g = 1$  if the  $g$  is the treatment sample.  $D_{gt}$  is also a dummy variable that equals one when the network is both (i) in the treatment sample and (ii) evaluated after the



	Motif 1	Motif 2	Motif 4	Motif 5	Motif 14
Treated	0.27 (0.20)	-0.69** (0.22)	-0.44 (0.28)	0.00 (0.19)	0.01 (0.04)
After	-0.25 (0.19)	-0.69*** (0.20)	-0.90*** (0.26)	-0.56** (0.17)	-0.13** (0.04)
DiD	0.10 (0.31)	0.83* (0.33)	0.79 (0.43)	0.48 (0.29)	0.10 (0.07)
Year 2010	-0.18 (0.15)	-0.05 (0.16)	0.19 (0.21)	-0.08 (0.14)	-0.01 (0.03)
Year 2012	-0.55** (0.19)	-0.56** (0.21)	-0.35 (0.27)	-0.52** (0.18)	-0.14** (0.04)
Size(log)	2.37*** (0.11)	1.62*** (0.12)	1.07*** (0.15)	1.80*** (0.10)	1.21*** (0.02)
R <sup>2</sup>	0.77	0.58	0.29	0.68	0.95
Adj. R <sup>2</sup>	0.76	0.56	0.27	0.67	0.95
Num. obs.	155	155	155	155	155

\*\*\*  $p < 0.001$ ; \*\*  $p < 0.01$ ; \*  $p < 0.05$ ;  $\cdot$   $p < 0.1$

**Table 2.2:** Results from the OLS regression to estimate the causal effect of the new product introduction. The model specifications are expressed by Eq. (2.1). The motifs' count is the outcome variable.

product introduction.  $D_{gt}$  is zero otherwise. The coefficient  $\beta$  quantifies the treatment effect according to the following equation:

$$\beta = (\langle Y \rangle_{\text{before}}^{\text{control}} - \langle Y \rangle_{\text{after}}^{\text{control}}) - (\langle Y \rangle_{\text{before}}^{\text{treated}} - \langle Y \rangle_{\text{after}}^{\text{treated}}) \quad (2.2)$$

We include two control variables in our analysis: (i) the year in which the new product is introduced, (ii) and the size of the network, namely the number of distributors (plus the manufacturer). Yet, we do not expect a linear relationship between the motif counts and the network sizes. Therefore, we consider a log transformation of both variables in Eq. (2.1).

We report the results from the OLS regression in Table 2.2. Each column shows the results obtained for a given motif. We remind the reader that the

DiD coefficient ( $\beta$  in Eq. (2.1)) quantifies the treatment effect. The obtained value is highlighted in Table 2.2.

We see that for all five motifs, the DiD coefficients are positive. This means that the new product introduction causes an overall increase in the number of motifs. Yet, the effect is significant only for motif M2 and slightly significant for motifs M4 and M5. Specifically, we observe a quite high increase in the count of M2, that is, a 129% increase (i.e.,  $100 \times [e^{0.83} - 1]$ ). This suggests that the number of paths of length two is increased or the length of the existing paths is increased.

The increase in the M2 count is significant and higher than that of the other motifs. It is worth noting that this is not due to a simple increase in the number of links, as no significant increase is observed in the M14 count (representing direct links).

Finally, we confirm our expectation of the non-linear relation between network size and motifs' count. Looking at the last row of the Table, we see that a 1% increase in the network size produces a 2.37% increase in the number of M1, a 1.86% increase in the number of M2, and a 1.80% increase in the number of M5. The motifs' count and the network size are correlated non-linearly by construction. Nevertheless, it is still interesting to comment on the differences between the regression coefficients obtained for the different motifs. In particular, we see that the correlation is stronger in the case of M1 than for the other motifs. This indicates that larger networks favor the centralization pattern (expressed by M1), compared to the other patterns, as it becomes much more abundant with the increasing of the network size.

It is safe to say that the obtained results are conditioned to the reliability of the model presented. This is characterized by a few, albeit important, limitations. We discuss them in the next paragraph.

**Limitations of the analysis.** The DiD model relies on two main assumptions: (i) the so-called parallel trend assumption; (ii) and the independency of the data records.

(i) According to the first assumption, the outcomes of both samples should display parallel trends in the absence of the treatment. DiD controls for dif-

ferences between the treatment and control samples that do not change over time. It does not control for those differences that do change over time. A good practice to check the validity of this assumption is to compare the two samples' trends before the event. For lack of data, we can not properly compare the two trends. Before 2009 (the first year of the event in our analysis), we can only compare the trend made up of three data points (i.e., 2006, 2007, and 2008) which we argue is not statistically meaningful.

(ii) The records in our regression analysis are likely to be correlated. Although most manufacturers work independently, their distribution networks may not be independent. Some networks may overlap as distributors ship products sent by different manufacturers. These correlations may bias the results. Single-manufacturer distribution networks are intrinsically connected. It is not possible to decouple them. All regression-based analyses would have this limitation as the distribution networks of different manufacturers may overlap through shared distributors. To account for this, a network-based null model could be considered in future studies, where firms' interdependencies are taken into account and single networks are evaluated independently. However, due to time constraints, we leave this exploration for future research.

## 2.4 Discussion

This Chapter provided a detailed description of the data used throughout this thesis. With half a billion shipping records, these data represent a unique source of empirical information into large-scale distribution systems and present exciting opportunities for future research.

We have explained the preprocessing steps required to transform these data into networks to be used for our analysis in later Chapters of this thesis. We also provided an overview of the local properties of the reconstructed networks by analyzing network motifs. Our findings suggested a tree-like topology, but we also identified variations from this structure. A more comprehensive characterization of the observed network structures is presented in the next Chapter.

## Chapter 3

# Opioid distribution networks: structure and evolution

### Summary

This Chapter presents a comprehensive empirical analysis of the opioid distribution networks in the United States. We examine the trend of opioid consumption and identify two opioids, Oxycodone and Hydrocodone, with substantial growth in demand from 2006 to 2014. Thus, we reconstruct their distribution networks; we evaluate the main structural properties of these networks; and trace their evolution over nine years. We find that, albeit the demand for opioids increased, the distribution networks did not expand. Interestingly, the topological properties of the Oxycodone distribution network remained largely unchanged, while the distribution network of Hydrocodone decreased in size. Upon further investigation, we conclude that the observed shrinking was likely due to cost-reduction strategies, rather than the network response to the increased demand.

---

This Chapter has been written specifically for this thesis. AA contributed to the design of the research questions, performed data processing and data analysis, and interpreted the results.

## 3.1 Introduction

The US opioid market has experienced substantial growth since the late 1990s, with the number of opioid prescriptions nearly tripling during this time. How did the distribution system evolve during this period?

Drawing on the ARCOS dataset, we reconstruct the distribution networks of different opioids in the United States. We assess their primary structural features and examine their stability over a nine-year period, spanning from 2006 to 2014. Our objective is to assess any structural shifts within the distribution system during a period of market expansion.


## 3.2 The opioid consumption

**On the opioid crisis.** Over the past two decades, there has been a significant surge in opioid prescriptions in the United States (Dasgupta *et al.*, 2018; DeWeerd, 2019; Guy Jr *et al.*, 2017). The number of prescriptions nearly tripled between 1991 and 2013, soaring from 76 million to approximately 207 million (Volkow, 2014). This surge in prescriptions has been accompanied by a rise in opioid-related abuses and overdose deaths, commonly referred to as the “opioid crisis” (Conrad, 2017; Eisenstein, 2019). Multiple factors have contributed to this crisis, including the inappropriate use of opioids, inadequate understanding of their adverse effects, and aggressive marketing strategies employed by certain opioid manufacturers\*.

Nowadays, opioid abuse is a serious problem for the healthcare system of the United States as well as for its social and economic welfare. For this reason, some US governments have implemented new regulations to tighten prescription requirements (Jones *et al.*, 2019). For instance, Florida introduced the “pill mill” law in September 2011, which declined opioid prescriptions by 1.4% after one year (Jones *et al.*, 2019). This corresponds to a monthly

---

\*In 2007, the pharmaceutical company Purdue Pharma was found guilty of making false claims about the low addictive power of its product. Similarly, in 2019, Insys Therapeutics pleaded guilty to inciting doctors to prescribe opioids for its financial interests (Dyer, 2016; Emanuel and Thomas, 2019; Whalen, 2018).



**Figure 3.1:** Consumption trend for the fourteen opioids in the ARCOS. The amount of consumption is reported in millions of grams.

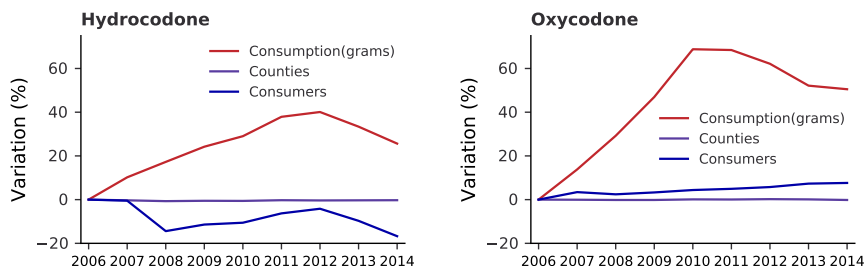
reduction of about half a million pills. Moreover, thanks to the new regulation, in 2012, Florida reported a 50% reduction in deaths from opioid overdose.

Clearly, studying the opioid crisis by embracing all its aspects (social, economic and health) presents significant challenges that are beyond this thesis's scope. Instead, the focus of this thesis is on the structural properties of the opioid distribution system and its potential changes due to the consumption shifts.

**Oxycodone and Hydrocodone: major players of the opioid market.** To analyze opioid consumption, we rely on the ARCOS dataset. We proxy annual consumption as the total amount (in grams) of opioids shipped to consumers in a given year.

In Fig. 3.1 we report the trend in consumption for each basic opioid. We notice two big players: Oxycodone and Hydrocodone. In 2006, Hydrocodone consumption was about 30 times higher than most other basic opioids, while Oxycodone consumption was almost 40 times higher. The two opioids show a similar consumption trend: since 2006, it follows a sharp increase that peaks in 2010 for Oxycodone, and in 2012, for Hydrocodone, which declines afterwards. These findings are in agreement with recent studies that, analyzing a different





**Figure 3.2:** Percentage variation of opioid consumption (red line), number of consumers (blue line), and number of counties supplied (violet line). On the left: Hydrocodone. On the right: Oxycodone.

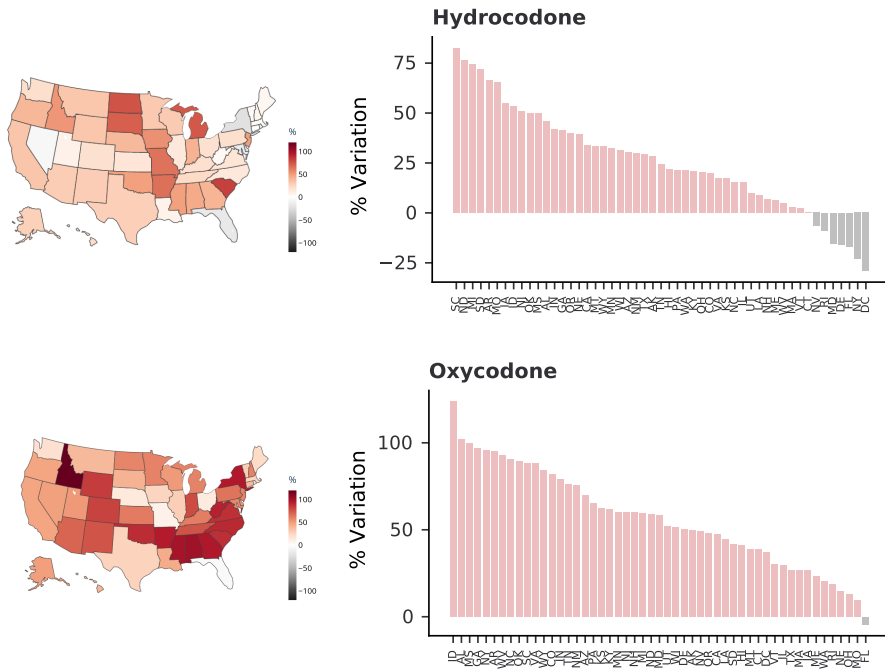
dataset, observed a slight decline in opioid consumption after 2010 (Guy Jr *et al.*, 2017; Jayawardana *et al.*, 2021).

For other basic opioids, however, we find very different trends. For example, we find (i) a downward trend that began as early as 2006 for Codeine, Morphine, and Fentanyl; (ii) an upward trend for Buprenorphine and Oxymorphone; and (iii) a new opioid, Tapentadol, that entered the market in 2009.

In the following, we focus our analysis on the two highest-selling opioids, Oxycodone and Hydrocodone, which display comparable consumption patterns. We defer the investigation of other basic opioids to future studies.

**Geography of the opioid markets.** The Hydrocodone and Oxycodone markets experienced substantial growth after 2006. Have the markets also expanded geographically? In other words, has the number of geographic areas supplied increased? We analyze the number of counties supplied every year. In Fig. 3.2, we report its variation from 2006 (violet line). For comparison, we also show the variation in the number of consumers (blue line) and the variation in consumption (red line).

Despite the substantial increase in consumption, we see that the number of counties supplied remains approximately the same each year. Our results suggest that both opioid markets did not expand geographically. Further, we see that the number of consumers shows a slight variation (about 10%), positive



**Figure 3.3:** Percentage variation of Oxycodone and Hydrocodone consumption across states between 2006 and 2014.

for Oxycodone and negative for Hydrocodone. This highlights some differences between the evolution of the two markets. After 2006, Oxycodone is still in a growth phase, as indicated by the increase in consumption and in the number of consumers. Hydrocodone, instead, is growing in terms of volumes sold (consumption) but decreasing in terms of the number of consumers. One possible explanation of these diverging observations is that, as early as 2006, government measures were taken to reduce the consumption of Hydrocodone, that mainly affected the number of consumers.

To further investigate this point, we analyze the variation in consumption across states. Since states behaved very differently in the measures taken to tackle the crisis, we expect these differences to be reflected in their consumption (Jones *et al.*, 2019). In Fig. 3.3, we show the geographical heatmap of

consumption variation, between 2006 and 2014, for Oxycodone (lower side) and Hydrocodone (upper side). The bar plots report the percentage variation per state.

As expected, we observe considerable differences across states. Examining Oxycodone consumption, we find a more than 90% increase in Idaho and many southeastern states such as Alabama, Mississippi, and Georgia, and a minor increase (about 10%), for example, in Ohio and Missouri. In addition, we note that Florida is the only state where we observe a decrease in consumption. Unlike Oxycodone, the Hydrocodone consumption drops in many states, i.e., Washington DC, New York, Florida, Delaware, Maryland, Rhode Island, and Nevada. Our results suggest that regulations to limit opioid prescriptions in these states have been effective. Nevertheless, positive variations are still observed in many states, e.g., South Carolina, North Dakota, and Michigan.

**Direct and indirect distribution.** Opioids, like any other product, can be supplied either through direct distribution, i.e., manufacturers ship directly to consumers, or through indirect distribution, i.e., manufacturers use distribution networks to ship to consumers. We find that more than 99% of opioid supply is through indirect distribution, thus confirming the critical role of distribution networks in meeting demand and coping with its variations (see Fig. B.1 in Appendix B) . In the next Section we analyze the extent to which these changes in demand actually affected the structure of the opioid distribution system.

### 3.3 Network characterization

In Section 3.2, we showed that the consumption of Oxycodone and Hydrocodone drugs experienced significant growth before 2010, followed by a slight decline. This raises the question of how distribution networks evolve in response to such changes in demand. To address this question, we reconstruct the distribution networks of Oxycodone and Hydrocodone drugs using data from the ARCOS. We characterize them by evaluating various network features, including network size, density, and nodes' heterogeneity, and discuss

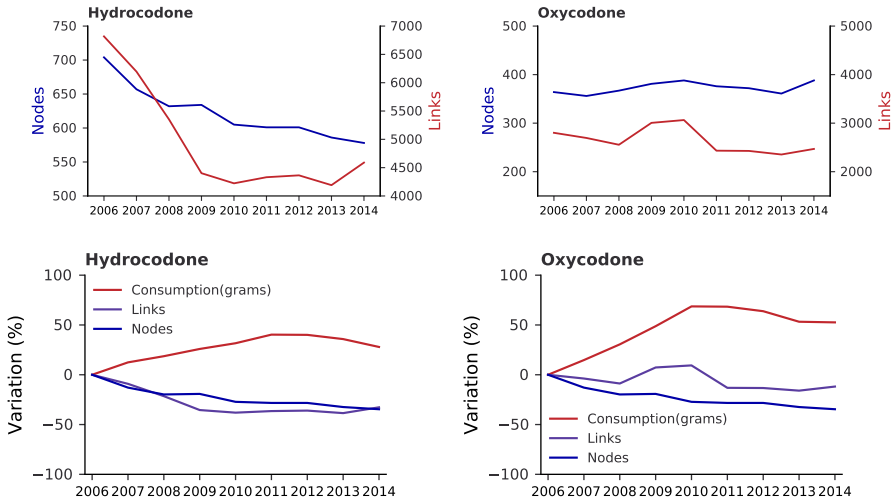
whether these features change over time. The findings are organized into three Subsections: (i) network size and density, (ii) firm heterogeneity, and (iii) firm position.

### 3.3.1 Network size and density

We start our analysis by monitoring the network size, i.e., the number of links and the number of nodes, and the network density. As discussed in Subsection 2.2.1, in our network representation nodes represent firms, i.e., manufacturers and distributors, and links represent supply relations. The latter are directed according to the direction of the shipments, from senders to receivers. The network density,  $D$ , is determined as the number of existing links divided by the number of possible links in a directed network, which can be expressed as  $D = \frac{L}{N(N-1)}$ , where  $L$  is the number of links and  $N$  the number of nodes.

In Fig. 3.4, we show the evolution of  $N$ ,  $L$  for the Oxycodone and Hydrocodone distribution networks. And in Table 3.1, we report the network density in the nine-year period. From Fig. 3.4, we see that the size of the Oxycodone distribution network is relatively stable over the years. We observe a small decrease in the number of links (about 10%) and a slight increase (about 5%) in the number of nodes between 2006 and 2014. A different pattern characterizes the Hydrocodone distribution network. Despite the increase in Hydrocodone demand, the underlying distribution network shrinks in size: we find an almost 20% nodes' decrease and more than 30% links' decrease.

These findings suggest that despite the growth in consumption between 2006 and 2014, both networks did not expand significantly during this period. This result can be better visualized in Fig. 3.4 (bottom plots), where we plot the variation in opioid consumption (red line), as well as the variation in the number of links (violet line) and the number of nodes (blue line). Variations are calculated by comparing the given year to the first year of observation. We observe that while the consumption variation is always positive with respect to the first observation year, the variation of the network size is always negative or slightly around zero.



**Figure 3.4:** Top: evolution of the number of nodes (blue line) and links (red line) of the Hydrocodone (left) and Oxycodone (right) distribution networks. Bottom: percentage variation in the number of nodes (blue line), number of links (violet line), and consumption (red line) for the two networks. Variations are measured with respect to the first year of observation.

Density	2006	2007	2008	2009	2010	2011	2012	2013	2014
Hydrocodone	0.014	0.014	0.013	0.011	0.012	0.012	0.012	0.012	0.014
Oxycodone	0.021	0.021	0.019	0.021	0.020	0.017	0.018	0.018	0.016

**Table 3.1:** Density of the Oxycodone and Hydrocodone distribution networks in the nine-year observation period.

Finally, monitoring network density, we find that both the Hydrocodone and Oxycodone networks exhibit consistently low-density values across all years, with average values of  $D=0.013$  for Hydrocodone and  $D=0.019$  for Oxycodone. These low-density values suggest that only a small portion of all possible supply relationships are observed. For the Hydrocodone network, only 1.3% of the supply relationships are observed, while for the Oxycodone network, only 1.9% of the supply relationships are observed. These results are consistent across the years, as we find that the annual density variations are quite small, with

an average of 0.3% for Hydrocodone and 9% for Oxycodone. This finding is in line with prior research that has identified low densities in other empirical supply networks (Brintrup *et al.*, 2015; Kim *et al.*, 2011; Wiedmer and Griffis, 2021). Having fewer connections in supply networks can promote efficiency and reduce costs.

### 3.3.2 Firm heterogeneity

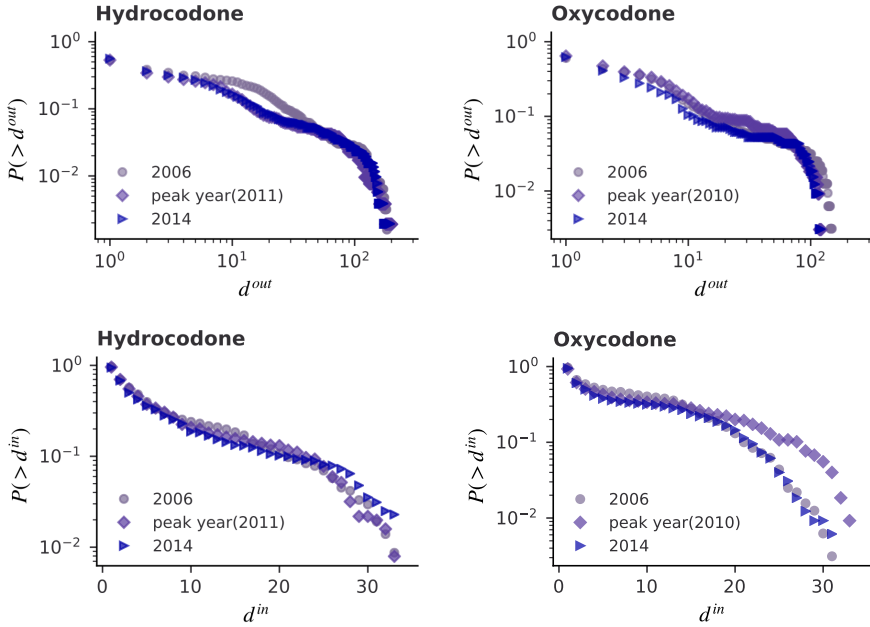
In this Section, we forward our empirical investigation and analyze the degree distributions of the networks under study. In network science, the term “degree” typically refers to the number of links that a node has. In the context of a supply network, where nodes represent firms, these firms can establish business relationships with two types of partners: (i) *source* partners who supply goods to them and (ii) *target* partners to whom they ship goods (i.e., clients). Therefore, we differentiate between in-degree, which denotes the number of source partners a firm has, and out-degree, which represents the number of target partners a firm has.

In Fig. 3.5, we show the in-degree and out-degree distributions for the Oxycodone and Hydrocodone networks. We evaluate them in three reference years: the first year (2006), the last year (2014) and the “peak” year. The peak year refers to the year in which the highest sales were recorded, i.e., 2010 for Oxycodone and 2011 for Hydrocodone. Note that we plot the complementary cumulative distribution function (CCDF) as it is more stable on the tails. The statistical properties of such distributions are reported in Table 3.2.

We find that the out-degrees exhibit very broad distributions. Specifically, the out-degree values for Oxycodone vary from 0 to 121, while those for Hydrocodone range from 0 to 195<sup>†</sup>. The positive skewness of the distributions implies the existence of right tails. The high kurtosis further indicates that these tails are remarkably pronounced, meaning that a few nodes act as “hubs” with many links, while many nodes have only a few links. The presence of hubs and the observed heterogeneity among nodes is common in economic networks. Degree distributions with heavy-tailed behavior are also characteristic of sev-

---

<sup>†</sup>It should be noted that nodes with  $d^{out} = 0$  represent firms that do not ship to other firms but rather they ship directly to consumers.



**Figure 3.5:** In-degree and out-degree distributions for the distribution network of Oxycodone and Hydrocodone in three years: the first year of observation (2006), the last year (2014), and the peak year (2010 for Oxycodone and 2011 for Hydrocodone).

eral other real-world networks, including financial networks (Boginski *et al.*, 2005), the world trade network (Li *et al.*, 2003), and supply networks (Perera *et al.*, 2017b)

In contrast to the out-degree distributions, the in-degree distributions exhibit narrower ranges. For instance, the maximum value is 31 for Oxycodone’s distribution network and 45 for Hydrocodone’s. Additionally, unlike the out-degree distributions, the in-degree distributions do not show heavy tails, as evidenced by the negative or small kurtosis values. Hence, in the networks under study, firms are very heterogenous in the number of target partners and much more homogeneous in the number of source partners.

	Oxycodone			Hydrocodone		
	2006	Peak Year	2014	2006	Peak	2014
$P(d_{out})$						
Mean	8.519	9.024	7.305	10.519	7.586	8.586
Std	23.802	21.702	20.777	25.791	21.451	25.228
Skewness	4.293	3.502	4.043	4.167	4.697	4.576
Kurtosis	18.420	12.005	15.700	19.483	25.021	22.237
$P(d_{in})$						
Mean	8.519	9.024	7.305	10.519	7.586	8.586
Std	8.295	10.518	8.436	14.509	9.534	11.486
Skewness	0.712	1.022	1.039	2.017	1.787	1.627
Kurtosis	-0.788	-0.371	-0.368	3.786	2.549	1.325

**Table 3.2:** Mean value, standard deviation, third and fourth moment of the out-degree and in-degree distributions for the Oxycodone and Hydrocodone distribution networks.

	Oxycodone		Hydrocodone	
	D(KS)	$p$ -value	D(KS)	$p$ -value
$d_{in}$				
2006 VS Peak	0.089	0.151	0.074	0.086
2006 VS 2014	0.118	0.021	0.073	0.097
$d_{out}$				
2006 VS Peak	0.065	0.494	0.109	0.002
2006 VS 2014	0.079	0.252	0.098	0.008

**Table 3.3:** Distances, D(KS), and  $p$ -values obtained from the KS-test that compares the out and in-degree distributions: (i) in the first year and in the peak year; (ii) and in the last year and in the peak year.

From a visual inspection, the statistical properties discussed so far look stable and robust across the years. We perform a Kolmogorov-Smirnov test (KS-test) to verify their robustness better. We report the distances and the  $p$ -values obtained from the test in Table 3.3. For the distribution network of



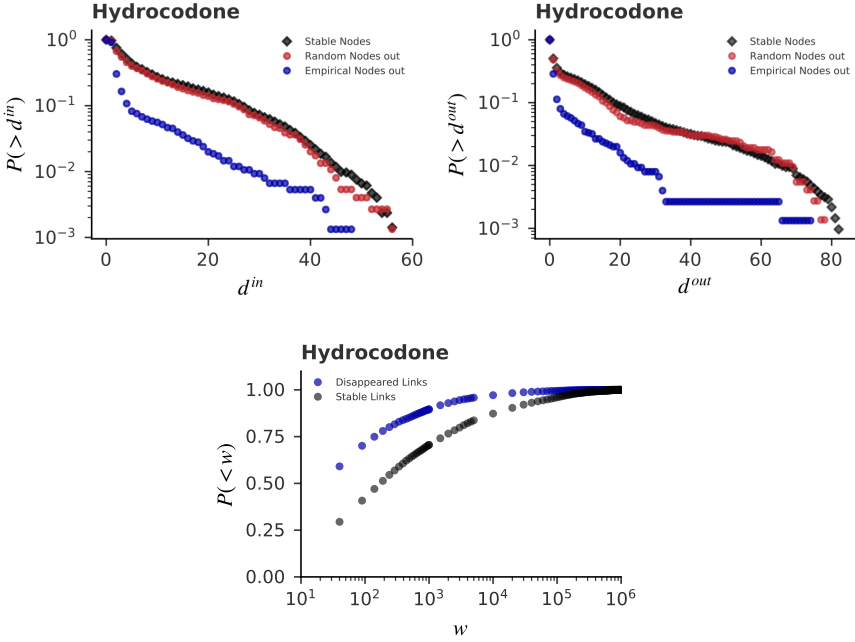
Oxycodone, we do not find any significant difference (for both the out and in-degree), suggesting that the structural properties of the network do not change significantly across the years. In the case of Hydrocodone, we find a similar result for the in-degree distribution. The out-degrees distribution, instead, significantly changes.

The distribution network of Hydrocodone has also experienced a great loss of nodes and links over the years, which may explain the significant differences in the out-degree distribution. This motivates us to investigate the part of the network that is shrinking. For this, we analyze the subset of nodes leaving the network and the subset of links disappearing over the years.

**Small nodes leave the network.** We determine the set of nodes leaving the network as the union of the sets of nodes disappearing every year (from 2006 to 2013). This definition makes our analysis independent of the specific year of observation. We compare this set of nodes with the nodes obtained from a random removal model. Precisely, in the random setting, we first remove  $n$  nodes uniformly at random from the empirical network, where  $n$  is the number of nodes that disappeared in the empirical one. Then, as in the empirical case, we collect the union of nodes randomly removed every year.

In Fig. 3.6, we show the CCDF of the in and out-degrees for the set of nodes leaving the network in the empirical case (blue dots) and in the random case (red dots). For comparison, we also plot the degree-distributions of the stable nodes (black dots), namely nodes that do not leave the network from one year to the next. As expected, under a random removal process, nodes leaving the network have a very similar degree distribution to those not leaving. Red and black curves overlap. Instead, we notice that the blue line (empirical case) diverges from the red line (random case) already for small values of both in-degree and out-degree.

Performing a KS-test, we confirm that the nodes leaving the network have significantly smaller degrees (both in and out-degrees) than expected at random, given the network topology. The distances and the  $p$ -values obtained from the test are reported in Table 3.4.



**Figure 3.6:** Top: in-degree (left plot) and out-degree distributions (right plot) of the nodes leaving the network in the empirical case (blue dots) and under the random model (red dots). Bottom: distribution of weights of links that disappear (blue dots) and that remain in the network (black dots).

	D(KS)	p-value
$d^{in}$	0.357	1.057e-42
$d^{out}$	0.231	4.096e-18

**Table 3.4:** Results from the KS-test evaluating the similarity between the degree distributions of nodes leaving the network in the empirical case and in the random model.

**Low-weight links disappear.** In Subsection 3.3.1, we observed that number of links in the Hydrocodone distribution network reduces by 30% between 2006 and 2014. These links are of two types: (i) links connecting one or two nodes leaving the network; and (ii) links connecting two stable nodes. While the first

type of disappearing links can be explained by the nodes leaving the network (see paragraph above), the second type can not. To understand this second type, we investigate their properties.

First, we look at the degree of the nodes that lost a link. By this, we aim to check whether the disappearing links are correlated to the degree of nodes. We find that this is not the case as the degree distribution of nodes losing a link is compatible with a random model. Results are reported in Fig. B.2 and Table B.1 in Appendix B.

Second, we consider an intrinsic property of the links, i.e., their weights. In Fig. 3.6, we show the cumulative distribution function of the weights of links disappearing (blue dots) and links stable on the network (black dots).

We find that about 70% of the links disappearing have low weight (lower than 100 grams), whereas only 40% of the stable links have such low weight. This result suggests that firms shipping low volumes are likely to cut their supply links in the following year. One possible explanation is that cost-saving strategies have been adopted, favoring higher shipping volumes and disfavoring lower ones.

### 3.3.3 Firms' positions

As the final step in our empirical investigation, we examine the topological position of firms within the supply network. In previous research, scholars have used the term *tier* to identify the position of a specific firm along the supply chain (Schwartz and Voß, 2007; Wiedmer and Griffis, 2021). The tier defines the distance of a firm from the focal manufacturer, and it has been applied to both upstream firms, such as suppliers (Brintrup *et al.*, 2015; Perera *et al.*, 2017a), and downstream firms, such as distributors (Jiang and Prater, 2002; Lan *et al.*, 2018). For example, a 1-tier supplier is the supplier directly connected to the manufacturer upstream, whereas a 1-tier distributor, e.g., a wholesaler, is the distributor to which the manufacturer directly ships its goods downstream.

In contrast to previous studies, our available data does not provide information on the tier of firms. We are only able to differentiate between manufacturers

and distributors. In the absence of additional information, we assume that all manufacturers are focal firms and assign them to tier 0. Precisely, given two nodes  $i$  and  $j$ , the the shortest path is defined as the length of the shortest sequence of nodes connecting  $i$  to  $j$ . Hence, we determine the tier,  $t_j$ , of the node  $j$ , as the shortest path connecting any manufacturer  $i$  to  $j$ , i.e.:

$$t_j = \min_{i \in \{m_1, m_2, \dots\}} \{\text{len}(p_{i \rightarrow j})\} \quad (3.1)$$

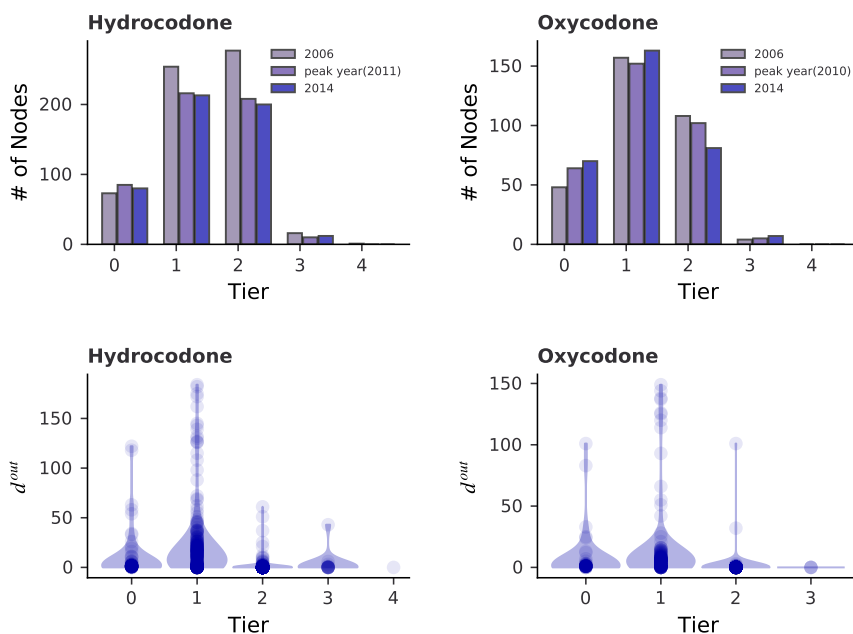
where  $p_{i \rightarrow j}$  indicates a path connecting  $i$  to  $j$ .

It is important to highlight that tier values may not be uniquely assigned, particularly in large-scale supply networks. A firm may have multiple tier values depending on the focal manufacturer considered or the product being shipped. The measure defined in Eq. (3.1) does not consider this variability and assumes that the nearest producer determines the tier of a firm. To address this limitation, we enhance our empirical investigation on firm positions in Chapter 4, where we improve Eq. (3.1) to account for all products shipped along all the empirical distribution paths.

Fig. 3.7, top row, displays the number of nodes per tier for the Oxycodone and Hydrocodone distribution network in 2006, in the peak years, and 2014. Fig. 3.7, bottom row, shows the correlation between the nodes' tier and their out-degree. First, we notice that there are only three tiers of distributors, and most distributors are located in tier 1 or tier 2. This suggests that the observed distribution networks are very short. By construction, in tier 0, producers are placed. They correspond to 10% of the total number of firms.

Furthermore, we observe that nodes with higher out-degrees tend to be closer to the focal firms and are typically positioned in tier 1. However, within this tier, we also identify many nodes with lower out-degree, indicating that manufacturers may distribute their products to both wholesalers (i.e., nodes with high out-degree) and retailers (i.e., nodes with low out-degree). In tiers 2 and 3, on the other hand, we mostly find small nodes with low out-degree, which likely represent individual retailers.

Finally, we note that the number of firms per tier varies little across the years. We observe only a slight increase in the number of manufacturers and a small



**Figure 3.7:** Top: number of nodes per tier in the first year, the last year, and the peak year, that is 2010 for Oxycodone and 2011 for Hydrocodone. Bottom: violin plot of out-degrees of nodes positioned at different tiers.

decrease in the number of distributors at tier 2 in Oxycodone’s distribution network. Similarly, we see a small decline in the number of distributors in Hydrocodone’s network, which occurs uniformly in both tier 1 and tier 2. The number of manufacturers remains stable.

### 3.4 Discussion

In this Chapter, we reconstructed the opioid distribution networks of the US from a collection of shipping records. We analyzed their structural properties and traced their evolution over nine years, from 2006 to 2014. Specifically, we

focused on the distribution networks of two top-selling opioids, i.e., Oxycodone and Hydrocodone.

We found that the consumption of the two opioids shows similar trends: it increased from 2006 to 2010 (or 2011 for Hydrocodone) and decreased slightly afterwards. However, the underlying distribution networks showed different structural evolutions.

Despite significant changes in the Oxycodone demand, its distribution network did not show significant changes from 2006 to 2014. We observed a slight increase in the number of firms (about 5%) and a small decrease in the number of supply links (about 10%). Next, we analyzed the out-degree and in-degree distributions and checked for similarity across years. We found no significant changes. Therefore, our results suggested that the topology of this distribution network did not experience significant reconfigurations in the face of shifts in demand.

In contrast, the distribution network of Hydrocodone showed a rather pronounced decline in the number of firms and supply links. Despite the growth in consumption between 2006 and 2010, the distribution network has shrunk in size. However, the further investigation suggested that the part of the network that shrunk is associated with small nodes (low out-degree) and weak links (i.e., low weight). One possible interpretation of this result is that, over the observed period, firms adopted cost-saving strategies that reduced the number of low-weight links in favour of high-weight links. Also, the changes observed (from 2006 to 2010) are unlikely to result from the enactment of new regulations limiting opioid prescribing for additional two reasons: (i) we still observe a growth in demand between 2006 and 2011; and (ii) most of the regulations have only been adopted since 2010.

Finally, this Chapter brings empirical insights to the supply chain domain. Summarizing our results, we found that large-scale distribution networks are characterized by: (i) very low density; (ii) high firm heterogeneity in the number of clients; (iii) firms' tendency to rely on more than one source partner; and (iv) short length. Based on previous literature, we classified firms into tiers and evaluated the resulting structure. However, applying the theoretical concept of "tiers" to real, large-scale distribution networks was not straightfor-

---

ward. Due to several focal firms, distributors may be assigned to different tiers depending on the product shipped. Therefore, the number of tiers cannot be uniquely assigned. In this Chapter, we analyzed firm position using a simple measure based on the shortest path concept, which captures only topological information. In Chapter 4, we improve our investigation on the number of tiers used in a large-scale distribution process. In the latter Chapter, we extract the flow of products with a resolution at the level of individual packages and use this information to assign the tier to distributors.

## Chapter 4

# Extracting trajectories of goods from shipping records

### Summary

In this Chapter, we move away from studying network evolution and instead analyze the small-scale dynamics on the network. Our focus is on reconstructing the daily distribution paths of individual packages as they move from production to consumption, which we refer to as trajectories. We reconstruct a large dataset of almost 40 billion trajectories using opioid shipping records. These trajectory data provide us with unique empirical insights into the distribution dynamic of goods within a nationwide distribution system. Moreover, we demonstrate how these data can be used to identify central distributors where the flow of packages is concentrated.

---

AA wrote this Chapter specifically for this thesis. Results are based on A. Amico, L. Verginer, F. Schweitzer “Reconstructing distribution paths from shipping records”, *Working Paper*. AA contributed to designing the research question and writing the code. AA performed the data analysis and interpreted the results.



## 4.1 Introduction

The previous Chapter investigated the main topological properties of the opioid distribution networks using a static network approach. The temporal information in the data was aggregated to construct yearly network snapshots. Thus, we investigated the network dynamic *across* years and discarded the dynamic *within* the single years.

Within a given year, distribution networks are intrinsically dynamic, as manufacturers and distributors interact daily to provide goods to numerous consumers. At a short time-scale (i.e., days), goods move from manufacturers to consumers through the network. To capture this daily flow, this Chapter goes beyond a static network perspective and considers the interactions of firms that occur on a daily basis. Starting from a set of shipping records in the AR-COS, we aim to reconstruct the complete trace of a single drug package moving from producer to consumer. As we do this, we aim at keeping track of the time sequence of all the dispatches in the data.

The Chapter is structured as follows. In Section 4.2, we discuss the limited applicability of existing techniques for extracting sequential data, i.e., distribution paths, in the context of distribution systems and propose a data-mining approach to overcome these limitations. In Section 4.3, we present large-scale descriptive statistics of these paths extracted, and in Section 4.4, we show how these data can be used to identify prominent distributors where the flow is concentrated. In Section 4.5, we close the Chapter by discussing the key findings.

## 4.2 A domain-driven approach

### 4.2.1 From shipments to trajectories

Supply chain data often lack comprehensive information, providing only limited details. Typically, the data available only reveal direct supply relations of firms while failing to capture the entire path of a good from its origin to its destination. Reconstructing complete distribution paths, starting from dyadic

data, is a challenging task that requires careful consideration. This issue is also known as the “transitivity problem” (Wichmann *et al.*, 2018).

The problem can be illustrated with the following example: suppose we have dyadic data that establish a supply relationship between firm  $C$  and firm  $D$ , denoted as  $C \rightarrow D$ . Additionally, we have information that firm  $A$  supplies to firm  $C$ , and firm  $B$  supplies to firm  $C$ . Can we conclude that both firm  $A$  and firm  $B$  supply to firm  $D$ ? In other words, can we identify two distinct distribution paths, one from  $A$  to  $C$  to  $D$  and another from  $B$  to  $C$  to  $D$ ? Unfortunately, this is not possible to infer directly, as we lack the ordered sequence of shipments between these firms.

Here, we address the question of reconstructing such paths from the available data in the ARCOS. In contrast to existing supply chain data, the ARCOS offers a significantly higher temporal resolution, with shipments recorded on a daily basis. Leveraging this temporal information, we can reconstruct the complete sequence of shipments required to move a package from its source to its destination.

In this thesis, we use the term *trajectories* to identify these distribution paths\*. We define a trajectory as the sequence of actors that goods travel through from their origin to their destination. In the next Sections, we will present our proposed method for reconstructing these trajectories. Before doing so, it’s important to note the limitations of previous techniques.

### 4.2.2 Limitations of temporal network techniques

As previously stated, data with a daily resolution are highly uncommon in the context of supply networks. Therefore, there is a lack of established techniques for analyzing temporal data within this domain. Yet, the question of how to extract sequential data from time-stamped data is a general one that scholars in various fields of applied network science have addressed. For example, researchers have tackled this problem to extract information flow in e-mail

---

\*Note that we prefer the term trajectory to the term “path”, widely used in network science. The term trajectory better conveys the idea of a finite quantity moving from a starting point to an ending point, as a package does when it moves from producer to consumer.

communication networks (Iribarren and Moro, 2009), to retrieve regulatory sequences in biological networks (Ko and Brandizzi, 2020), or to reconstruct mobility patterns of scientists (Vaccario *et al.*, 2020).

A prevalent approach involves using temporal networks (Holme, 2015; Lambiotte *et al.*, 2019; Peixoto and Rosvall, 2017). A temporal network  $G_T = (V, E_T)$  can be defined as a set of nodes  $v \in V$  and links  $e \in E_T$ , where nodes interact only at a given time,  $t$ . Thus, the interaction between two nodes,  $v_1$  and  $v_2$ , is represented by a time-stamped link,  $e = (v_1, v_2, t)$ , with  $t$  representing the time of the interaction. Scholars have developed methods to leverage these network representations and extract causal paths of interactions while respecting their underlying chronological order (Lambiotte *et al.*, 2019; Scholtes *et al.*, 2016, 2014).

Here we argue that these techniques are not directly applicable to distribution systems. Paths extracted using temporal networks would represent an oversimplified version of the underlying distribution process. The authors Mattsson and Takes (2021) made a similar criticism while studying other real-world processes, i.e., money flow in financial networks and ball passages in football games. We list below the main limitations temporal network techniques exhibit when applied to distribution systems.

**Conservation quantity not ensured.** In principle, multiple time-respecting paths exist on a distribution network, but only a subset is also quantity-respecting (as packages cannot multiply themselves). Consider the simple situation where two manufacturers,  $m_1$  and  $m_2$ , send one package each to a common distributor  $d$  at time  $t$ . At the time  $t + 1$ , the distributor  $d$  sends one package to pharmacy  $p_1$  and one package to pharmacy  $p_2$ . Thus there are a total of four time-respecting paths, namely  $m_1 \rightarrow d \rightarrow p_1$ ,  $m_1 \rightarrow d \rightarrow p_2$ ,  $m_2 \rightarrow d \rightarrow p_1$  and  $m_2 \rightarrow d \rightarrow p_2$ . However, only two can be observed as only two are also quantity-respecting (manufacturers originally sent only two packages).

**Stocks not modelled.** An additional complication to the point discussed above is that packages may be held in stock for some time and queued one after

the other. Then, they are shipped out according to some given stock-managing rules. This complication would not be captured by temporal network models, which do not account for the storage capacity of nodes.

**Critical selection of an optimal time-window.** The extraction of causal paths in temporal networks is usually based on the following assumption: given the interaction between node  $v_1$  and node  $v_2$  at time  $t_1$ , and the interaction between node  $v_2$  and node  $v_3$  at time  $t_2$  (with  $t_2 > t_1$ ), there exists a causal path  $v_1 \rightarrow v_2 \rightarrow v_3$ , if the time interval  $t_2 - t_1 < \delta$ , where  $\delta$  is a given time-window. Choosing the optimal time-window for a given system is a well-known challenge in itself (Caceres and Berger-Wolf, 2013), and quite problematic to address in our data. Opioid drugs are very diverse in terms of demand and shelf life. Some may move from manufacturers to consumers within a few days because of, for example, their high demand, while others may remain in storage longer. Thus, the optimal time-window selection would be product-specific, adding further complexity to the method.

### 4.2.3 Extraction method and boundary specification

**Method principles.** We aim to extract a set of trajectories from time-stamped shipping data. To this aim, we propose a data-mining algorithm that takes as input a set of shipping records and provides as output the observed trajectories.

The shipping records are processed in a time-respecting order. Next, we assign stocks to every firm, used to store the products received. Stocks are managed according to the first-in-first-out heuristic and updated based on the quantity recorded in shipping data. Note that we do not model the quantity to be shipped or received by a given firm. We take this information from the data while processing the shipping records.

Since we process the shipping records in a time-respecting order, we ensure that firms do not ship goods unless they have previously received them (volume constraint). This principle is valid for all firms labelled as distributors in the data. If, instead, firms are labelled as manufacturers, then we account for the possibility that such firms can produce new quantities, that is, they ship

products without previously receiving them. Again, we do not model the quantity produced by such manufacturers. This quantity is determined by the one recorded in the shipping data.

Finally, some rare cases can occur where some distributors appear in a given shipping record for the first time as a sender and not as a receiver. This happens because the previous shipments, where such distributors received those quantities, are not recorded in our data as before 2006, i.e., the first year of observation. Our algorithm adds the quantity to the distributors' stocks in these cases, assuming a previous shipment.

**Input: shipping records.** Let us consider a dataset  $D$  collecting  $S$  shipping records. We represent the record,  $\rho_i$ , of a given shipment  $i$ , as a six-tuple:  $\rho_i = (s_i, r_i, b_{s_i}, b_{r_i}, t_i, q_i)$ . In the tuple,  $s_i$  is the sender,  $r_i$  is the receiver,  $t_i$  is the shipping date, and  $q_i$  is the amount shipped (e.g., number of product units). Finally,  $b_{s_i}$  and  $b_{r_i}$  represent the business activities of the sender and receiver, respectively. We distinguish between manufacturers ( $m$ ), distributors ( $d$ ), and consumers ( $c$ ). Each element of the tuple is passed as input to the algorithm.

**The algorithm.** The extraction algorithm starts with an empty set of trajectories  $P := \emptyset$ . We assign an empty stock,  $\sigma_v := \emptyset$ , to each firm in the data. Then, for each shipping record,  $\rho_i$ , the algorithm performs the steps described below.

- 1 **Sender activity.** The sender,  $s_i$ , gets access to its stock,  $\sigma_{s_i}$ .
  - A Case:  $\sigma_{s_i} \geq q_i$ . The sender has enough quantity in its stock. Thus, it ships the package  $j$  and the trajectory is updated. Hence, we perform the following steps:
    - 1  $q_i$  units are taken out from the stock  $\sigma_{s_i}$ .
    - 2 The trajectory of the package  $j$  is updated:
      - a) if  $j$  does not have previous shipments, the trajectory is created, i.e.,  $p_j = (s_j^1 := s_i, r_j^1 := r_i, t_j^1 := t_i, q_j^1 := q_i)$ ;

- ▶ b) if  $j$  has  $k$  previous shipments, the trajectory is extended. The tuple  $(s_j^{k+1} := s_i, r_j^{k+1} := r_i, t_j^{k+1} := t_i, q_j^{k+1} := q_i)$  is appended to the existing trajectory  $p_j = (\mathbf{s}_j, \mathbf{r}_j, \mathbf{t}_j, \mathbf{q}_j)$ , where each element has length  $k$ .
- 3 The sender  $s_i$  updates its stock:  $\sigma_{s_i} = \sigma_{s_i} - q_i$ .
- B Case:  $\sigma_{s_i} < q_i$ . The sender does not have enough units in stock. Therefore, it produces  $q_i$  units. After production,  $\sigma_{s_i} = q_i$  and steps [1], [2] and [3] of Case [A] are performed.
- 2 **Receiver activity.** The receiver,  $r_i$ , gets access to its stock and adds the quantity received,  $\sigma_{r_i} = \sigma_{r_i} + q_i$ .
- 3 **Trajectory extraction.** If  $\rho_i$  is the last record of  $D$ , namely  $i = S$ , the algorithm performs its final step, hence, extracting the complete set of trajectories stored. We do not extract all the trajectories stored, but only those that satisfy some boundary conditions. In particular, we want to ensure that these trajectories begin on the production side and end on the consumption side. Therefore, we use the business activities,  $b_{s_i}$  and  $b_{r_i}$ , to identify the start and the end of the trajectory. For every trajectory the first shipping record has  $b_{s_i} = m$  and the last one has  $b_{r_i} = c$ , i.e.:

$$\rho_{i_{\text{start}}} = \{(s_i, r_i, b_{s_i}, b_{r_i}, t_i, q_i) \mid b_{s_i} = m\} \quad (4.1)$$

$$\rho_{i_{\text{end}}} = \{(s_i, r_i, b_{s_i}, b_{r_i}, t_i, q_i) \mid b_{r_i} = c\} \quad (4.2)$$

By this, we only select those trajectories that start in a manufacturer and end in a consumer. Finally, the nodes traversed  $\mathbf{n}_j$  are obtained as the union  $\mathbf{s}_j \cup \mathbf{r}_j$  and added to the tuple  $p_j$ .

**Stock-managing policy: FIFO.** In our analysis, we assume that firms within the opioid distribution system adhere to a first-in-first-out (FIFO) stock management policy. Under this policy, the first units to arrive are also the first to be sold. This approach minimizes the impact on the product's shelf-life and, therefore, it is commonly utilized in the supply chains of perishable goods, such as medicine (Harahap *et al.*, 2022; Vujanac *et al.*, 2016; Yadav *et al.*,

2021). Hence, our algorithm considers that firms arrange their inventory by queuing the units of products received. Every time, a newly acquired unit is added to the end of the queue. When fulfilling orders, the units at the front of the queue (i.e., the first to arrive) are dispatched first.

**Output: trajectories of goods.** The output of the algorithm is a set of trajectories  $P$ . A single trajectory belonging to the set  $P$  is represented as a five-tuple:  $p_j = (\mathbf{s}_j, \mathbf{r}_j, \mathbf{n}_j, \mathbf{t}_j, \mathbf{q}_j)$ , where  $j$  is the index of the package. Then,  $\mathbf{s}_j$  is the sequence of the senders,  $\mathbf{s}_j = (s_j^1, s_j^2, \dots, s_j^l)$ ;  $\mathbf{r}_j$  is the sequence of the receivers,  $\mathbf{r}_j = (r_j^1, r_j^2, \dots, r_j^l)$ ;  $\mathbf{n}_j$  is the sequence of the nodes traversed by the package  $j$ ; and  $\mathbf{t}_j$  denotes the sequence of the shipping dates,  $\mathbf{t}_j = (t_j^1, t_j^2, \dots, t_j^l)$ .

Note that a single trajectory in our analysis corresponds to the distribution of a single drug package. However, in reality, goods are shipped by manufacturers in big batches and often repackaged as they move along the supply chain to reach consumers. As a result, the number of units in each batch shipped at different distribution stages may vary depending on the firm’s practices. In fact, firms typically rely on “economic order quantity” models to determine the optimal amount of goods to procure to meet demand while minimizing inventory management costs (Berlec *et al.*, 2014; Huang and Wu, 2016). Our approach enables us to monitor the package’s status by reporting the number of units in each batch shipped. These are captured by the vector  $\mathbf{q}_j = (q_j^1, q_j^2, \dots, q_j^l)$  where  $q_j^l$  is the number of units shipped at step  $l$  of the distribution process.

Finally, note that the extracted trajectories are *time-respecting*. Hence, the condition  $t_j^k < t_j^{k+1}$  for all  $k$ , with  $1 \leq k \leq l$ , is always verified.

### 4.3 Descriptive data analysis

We use the set of shipping records in the ARCOS dataset as input for the extraction algorithm proposed above. Thus, applying the algorithm, we extract

nearly 40 billion trajectories<sup>†</sup>. These track the daily flow of individual packages moving from producers to consumers over a nine-year period (2006- 2014). In this Section, we perform extensive descriptive statistics to present the data obtained.

**Number of trajectories.** We start our empirical characterization by looking at the number of trajectories per product. The *ndc-number* identifies a single product in our data (see Chapter 2). More than 15,000 products reach consumers between 2006 to 2014. These products are very different from each other. Some products are used only for experimental trials and are shipped in small quantities; others are established products that are shipped in large quantities and reach thousands of consumers.

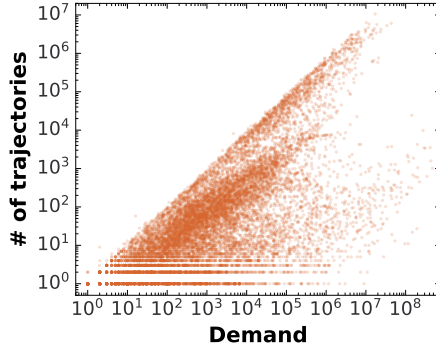
Since products differ in demand, we expect to observe differences in the number of trajectories extracted per product. In Fig. 4.1, we show the correlation between the number of trajectories per product (y-axis) and the product demand (x-axis). Demand is measured as the total volume (in product units) shipped to consumers within the observation period. Each orange dot marks a single product. Note that, by construction, the total number of trajectories can only be lower (or equal) than the total demand. The sharp diagonal cut in the plot marks this. Indeed, it is only possible to ship a minimum of one product unit along a single trajectory.

We see that the demand varies by six orders of magnitude across products. For some products, billions of units are sold, and for others, less than one hundred. Remarkably, we find that less than 5% of the products have extremely high demand, i.e., more than one million units sold. As expected, this pattern is reflected in the number of trajectories extracted. Only a few products are associated with many trajectories: we count about one million trajectories for only 101 products ( less than 1% of the products). On the contrary, many products (1,519) are shipped only once. The correlation between the demand for a product and the number of trajectories observed is also evident from the

---

<sup>†</sup>In the ARCOS data products are identified via the *ndc-number*. Therefore, we preliminary process the data to group the shipping records by *ndc-number* . Next, we apply the extraction method to the shipping data of each product separately.






**Figure 4.1:** Total number of trajectories (x-axis) and product demand (y-axis). Demand is measured as the total volume sold (in product units) to consumers between 2006 and 2014. Products are sorted by demand (from lowest to highest).

scatter plot in Fig. 4.1. The Pearson correlation coefficient of 0.67 confirms this correlation.

The considerable differences in demand indicate that some products are commercial (i.e., sold in substantial quantities to consumers) while others are not. Non-commercial products may be subject to different distribution processes; for example, they are sold to specific consumers within a well-defined limited period. To reduce the noise, we decide to focus on commercial products only. Thus, in the rest of our analysis we consider only products with a total demand of at least 1,000 units.

**Trajectory frequency and concentration of flow.** A given sequence of nodes (e.g.,  $m \rightarrow d_1 \rightarrow d_2 \rightarrow c$ ) can be used several times to distribute goods. In other words, multiple packages can follow the same trajectory over nine years. We want to answer the questions: are some trajectories used more frequently than others? How is the total flow distributed among them?

We define the frequency  $f_p$  of a trajectory  $p$  as the number of times the trajectory is observed within the observation period. Thus,  $f_p$  represents the number of packages shipped through  $p$ . In Fig. 4.2, we show the probability distribution of  $f_p$  for a set of trajectories traversed by packages containing the same



**Figure 4.2:** Probability distribution of trajectory frequency  $f_p$  for the single product analysis (left plot), and the multiple-product analysis (right plot).

product (left plot) and different products (right plot)<sup>‡</sup>. We notice that the two probability distributions are strongly skewed. This observation indicates a rather high heterogeneity in the number of packages shipped per trajectory, suggesting a concentration of flow along some of them. In the single-product analysis, we find that more than 70% of the total flow passes through only 10% of the trajectories. In the multiple-product analysis, we find that about 70% of the total flow passes through only 8% of the trajectories.

Our results indicate large differences in the amount of goods passing through the observed trajectories and suggest a concentration of flow along a few trajectories. Various reasons, such as the presence of active consumers or predominant producers, can explain this result. Also, it may be due to consumer preferences toward specific producers or distributors.

**Trajectory length: number of distributor tiers.** We investigate the number of tiers of distributors involved in the shipping of a single product. As discussed in Chapter 3, the number of tiers indicates the number of sequential distributors connecting manufacturer to consumer.

<sup>‡</sup>For the single-product analysis, we select one of the products with the highest demand. It has the *ndc* equal to 00406-0357-05 and 10,646,008 trajectories extracted. For the multiple-product analysis, we select the top 200 products containing Oxycodone. For these, we extracted a total of 94,145,723 trajectories. Note that the top 200 products already cover 99% of all the transactions.

Let's consider the sequence of senders  $\mathbf{s}_j$  of the package  $j$ . The number of tiers is expressed by the cardinality of the set  $\mathbf{s}_j$  minus one. Note that the cardinality of  $\mathbf{s}_j$  would account for all senders including the first one, namely the manufacturer. As we want to measure the number of distributors, we use the minus one to exclude the manufacturer from the count. Thus, we compute the number of tiers  $\langle\theta\rangle_\mu$  traversed by a product  $\mu$  by averaging over all its packages, i.e.:

$$\langle\theta\rangle_\mu = \frac{1}{J^\mu} \sum_{j=1}^{J^\mu} \theta_{p_j} \quad (4.3)$$

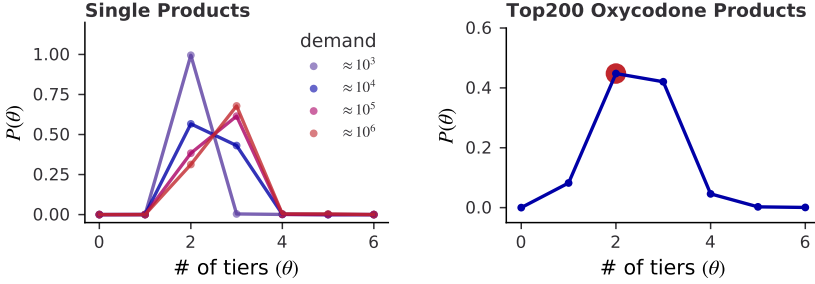
where  $J^\mu$  is the total number of packages containing product  $\mu$ . In Fig. 4.4 (a) we show  $\langle\theta\rangle_\mu$  per product (orange solid line). The orange-filled region displays the variance of the sample.

Note that nearly half of the products in the data (i.e., 2,913 out of 6,087) are shipped through a direct distribution process. This means that manufacturers ship directly to consumers without the intervention of any distributor. Because we are interested in the indirect distribution processes, we restrict our analysis to the subset of products with  $\langle\theta\rangle_\mu > 0$ , meaning that at least one tier of distributors is used.

For these products, we observe relatively short trajectories. Most products (56%) traverse only two tiers of distributors. And, at most, we observe three tiers of distributors for all products (97%).

In Fig. 4.3, we show the probability distribution of the number of tiers, namely how often a given number of tiers is used to ship packages of a single product (left plot) and multiple products (right plot). In the single-product analysis, we select four products with very different demands. In the multiple-product analysis, we select the 200 top products containing Oxycodone.

From the left plot in Fig 4.3, we see that for the two products with the lowest demand ( $\text{dmd}=10^3$  and  $\text{dmd}=10^4$ ), trajectories are very short. The probability distributions show a peak for  $\theta=2$ , which means that mainly two tiers of distributors are observed. For the two products with the highest demand ( $\text{dmd}=10^5$  and  $\text{dmd}=10^6$ ), the peak shifts slightly from  $\theta=2$  (two tiers) to  $\theta=3$  (three tiers). However, trajectories with two tiers of distributors are still quite



**Figure 4.3:** Probability distributions of the number of sequential distributors (i.e., number of tiers) traversed by packages containing the same product (left plot) and different products (right plot).

frequent. They correspond to 30% of the trajectories. Therefore, the difference between low-demand and high-demand products is minimal, and the trajectories are still very short in both cases. A similar result is obtained even when products are aggregated. In the right plot of Fig. 4.3, we find the same pattern: the probability distribution peaks at  $\theta=2$  and has a maximum at  $\theta=3$ . The number of trajectories with  $\theta=4$  is negligible. In conclusion, our analysis shows that many products pass through rather few tiers of distributors.

**Geodesic length: travel distance.** Opioids may, in principle, travel long distances before reaching end consumers. How far do they travel? To get this empirical insight, we define the travel distance  $d_{p_j}$  of the package  $j$  along the trajectory  $p_j$ , as the total geodesic distance<sup>§</sup> travelled by  $j$  while moving from the manufacturer to the consumer, i.e.:

$$d_{p_j} = \sum_{k=1}^l d(s_j^k, r_j^k) \quad (4.4)$$

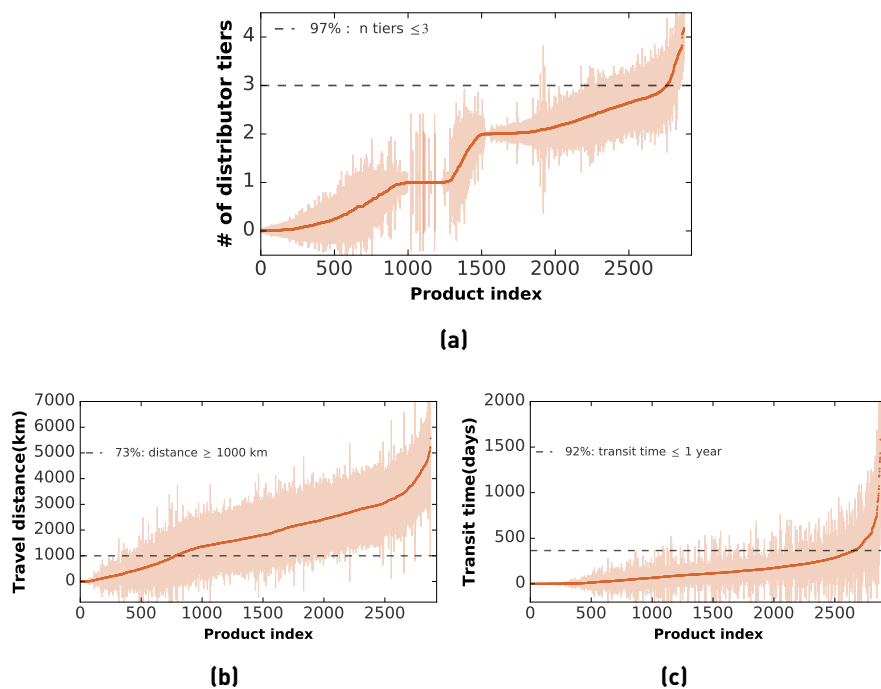
where  $d(s_j^k, r_j^k)$  indicates the geodesic distance between the sender and the receiver at step  $k$  of the distribution process. The sum runs over all the steps between the first sender (the manufacturer) and the last receiver (the con-

<sup>§</sup>The geodesic distances are computed through the python package GEOPY.

sumer). Next, we obtain the travel distance  $\langle d \rangle_\mu$  of a product  $\mu$ , by averaging over all packages containing  $\mu$ , i.e., :

$$\langle d \rangle_\mu = \frac{1}{J^\mu} \sum_{j=1}^{J^\mu} d_{p_j} \quad (4.5)$$

where  $J^\mu$  is the total number of packages of product  $\mu$ .



**Figure 4.4:** The orange line shows the average number of tiers (a), travel distance (b) and transit time (c) per product (x-axis). The orange-filled regions show the sample variance.

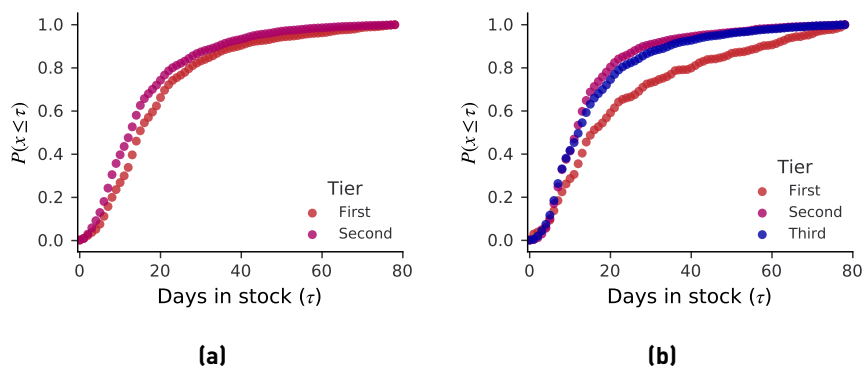
In Fig. 4.4 (b) we show  $\langle d \rangle_\mu$  per product. The orange-filled region displays the variance of the sample. On average, most products (73%) travel more

than 1000 km, and half of the products travel more than 2000 km. This result confirms our expectation that the distribution process under analysis operates nationwide, as rather large distances separate consumers from manufacturers through distributors.

**Trajectory duration: transit time and time in stock.** We now characterize the temporal aspect of the extracted trajectories and examine two features: (i) the transit time of the packages and (ii) their time in stock. Specifically, we define the time in transit  $T_{p_j}$  of the package  $j$  along the trajectory  $p_j$ , as the time elapsing between the first and last shipment, i.e.,  $T_{p_j} = t_j^l - t_j^1$ . Next, we define the time in stock  $\tau_{p_j}^k$  of the package  $j$  at the stock of the  $k^{\text{th}}$  sender, as the time elapsing between the current and the previous shipment, i.e.,  $\tau_{p_j}^k = t_j^{k-1} - t_j^k$ .

Before discussing the results, we want to clarify some aspects of the above measures. First, note that we cannot measure the time spent by packages at manufacturers' stocks because we do not have data on production times. As a result, below, we will only consider distributors' stocks. Also, in the absence of data on transportation times, we assume them to be zero. In other words, we assume that deliveries from one stock to another occur within one working day. We leave it to future studies to see how data on travel distances can be collected and merged with current data. Finally, we would like to clarify that our definition of stock time differs from the standard definition of "lead time". The latter indicates the time between the start and completion of a process. In the distribution process, the start is when the order is placed, while completion is when the delivery is made. Therefore, it measures the time between an order and a delivery. We cannot measure the lead times because we do not have data on the orders. In contrast, the time in stock defined above measures only the time spent in stock regardless of when orders are placed.

In Fig. 4.4 (c), we show the average time in transit  $\langle T_{p_j} \rangle_\mu$  per product. We find that nearly all products (92% of the products) have been in transit for less than a year. This confirms our expectation that distribution processes occur at short time scales, typically one year. Also, it is consistent with the observation that opioids (like many other drugs) are perishable products and, hence, are characterized by short shelf life. Only 5% of the products have much

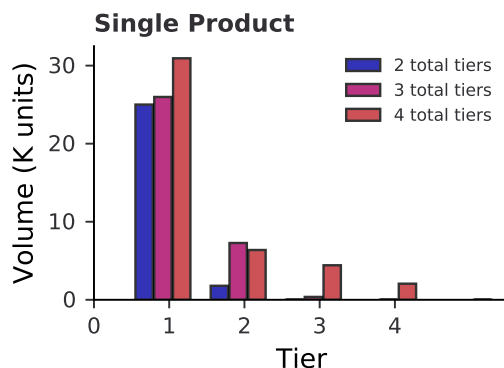


**Figure 4.5:** Probability distribution of time in stock of packages that traverse: two tiers of distributors (a), and three tiers of distributors (b). The colors specify the  $k^{\text{th}}$  tier the distributor belongs to.

longer transit times than one year: they range from 500 to 2,000 days (about five years). These products may have very low demand or large fluctuations in demand.

Next, we investigate the tier at which packages spend more time. In other words, given the sequence of distributors traversed, where do they stay longer? In Fig. 4.5 we plot the cumulative distribution of times in stock for packages that pass through two tiers of distributors (a), and three tiers of distributors (b). Note that the cumulative function,  $P(x \leq \tau)$ , measures the probability that a given package remains in stock for less (or equal) than a given time  $x$ .

First, we see that about 80% of the packages remain in stock for up to 40 days. This is true for all stocks, regardless of the total number of tiers along the trajectory and the tier the distributor belongs to. In addition, we observe that in the case of trajectories with lengths two (Fig.4.5 a), the two curves almost overlap, suggesting that the times in stock for the first and second-tier distributors are quite similar. In contrast, for longer trajectories (Fig.4.5 b), we observe that the first-tier distributor has longer times in stock than the second and third-tier distributors, as the red curve always remains below the other two.



**Figure 4.6:** Tier zero indicates the producer, tier one the first distributor, and so on. The height of the bars indicates the amount shipped by each actor at a given tier. Color distinguishes trajectories of different lengths.

**Quantity shipped.** The average number of units of a particular product shipped per tier is depicted in Fig. 4.6. We can observe a significant reduction in volume from tier 0, which represents the manufacturer, to tier 1, which refers to the first-tier distributor, as depicted in Fig. 4.6. In this diagram, we can see that the initial shipment of 30,000 units from the manufacturer undergoes a process of repackaging by the first distributor, resulting in smaller units that are eventually delivered to consumers in even smaller packages. This empirical finding provides strong evidence supporting the utilization of economic order quantity models by firms (Huang and Wu, 2016). According to saving-cost strategies, manufacturers typically ship large batches of products, which are subsequently reshaped and distributed by intermediaries. It is noteworthy that the number of units shipped is reduced by more than 60% as products pass from the manufacturer to the subsequent distributors. This consistent reduction in volume holds true for trajectories of varying lengths, indicating the robustness of this result.



## 4.4 Ranking nodes according to the empirical flow

In the Section above, we generated a high-resolution dataset tracking the empirical flow of goods from manufacturers to consumers, through distributors. In this Section, we show how these trajectories can be used to rank nodes based on the flow they handle.

Node ranking techniques are particularly valuable for supply chain managers because they help them identify the most prominent firms where the flow is concentrated (Craighead *et al.*, 2007; Mizgier *et al.*, 2013). In this way, they can make more informed decisions about how to allocate resources to mitigate possible disruptions. The techniques so far proposed in the literature are usually based on a static network approach: the network topology is used to inform the ranking algorithm. The actual flow of goods is discarded. The static approach is usually justified in the absence of more descriptive and granular data, i.e., sequential or temporal data.

Because we have sequential and high-resolution data (i.e., the set of extracted trajectories), we can move beyond the static network approach. Our goal is to revise established ranking measures to incorporate empirical flow and, thus, assess the error we make if we use only topological information. In particular, we focus on one of the most established measures for performing node ranking: Freeman’s betweenness centrality (Freeman, 1977; Scholtes *et al.*, 2016). This has been used by various scholars in the supply chain domain (Mizgier *et al.*, 2013).

Given a network,  $G(N, E)$ , comprising  $N$  nodes and  $E$  links, the (unnormalized) betweenness centrality of a node  $v$  is simply calculated as the total number of shortest paths passing through node  $v$ , i.e.,:

$$BC^p(v) := \sum_{i \neq v \neq j} |p(i, j; v)| \quad (4.6)$$

where  $p(i, j; v)$  denotes the set of shortest paths passing through node  $v$ , and connecting every pairs  $i$ - $j$ . According to Eq (4.6), the more shortest paths pass through a given node, the more flow it controls.

The flow is not directly measured in Eq (4.6). It is inferred from the network structure. We revise Eq (4.6) to include the actual flow. This is expressed in the form of trajectories, as defined in Section 4.2.3. Given a set of observed trajectories  $S$ , the revised centrality measure is written as:

$$BC^s(v) := \sum_{\substack{\forall i \in \{m\} \\ \forall j \in \{c\} \\ i \neq v \neq j}} |s(i, j; v)| \quad (4.7)$$

where  $s(i, j; v)$  denotes the set of trajectories that start from a manufacturer  $i$ , i.e.,  $i \in \{m\}$ , end in a consumer  $j$ , i.e.,  $j \in \{c\}$ , and pass through distributor  $v$ . According to Eq (4.7), the more trajectories pass through a node, the more flow it handles.

The two measures differ for a major aspect: the standard definition of betweenness centrality accounts for all possible (shortest) paths given the network topology, whereas the revised centrality measure directly accounts for the empirical flow.

Note that while the standard betweenness centrality is based on the transitive assumption, the revised one is not. According to the transitivity assumption, the presence of two paths, namely one from  $v_0$  and  $v_k$ , i.e.,  $v_0, v_1, \dots, v_k$ , and the other from  $v_k$  and  $v_l$ , i.e.,  $v_k, v_{k+1}, \dots, v_l$  implies that a path connecting  $v_0$  to  $v_l$  necessarily exists (Scholtes *et al.*, 2016).

On the contrary, the centrality measure introduced in Eq. (4.7) does not include such an assumption as it builds on the observed flow. The trajectories extracted are by construction *time-respecting* and *quantity-respecting* (see Section 4.2.3).

We question the validity of a transitivity assumption and investigate to what extent the two centrality measures produce different rankings. As a test case for our analysis, we consider the opioid distribution network, reconstructed from the ARCOS data<sup>¶</sup>. For every node  $v$  in this network, we compute (i) the standard centrality  $BC^p(v)$  and the revised centrality  $BC^s(v)$ . To assess the

<sup>¶</sup>All drugs containing Oxycodone have been considered for the current study.

Year	Kendall-Tau ( $p$ -value)	Pearson ( $p$ -value)
2006	0.74 (1.88e-47)	0.58 (1.83e-17)
2007	0.66 (3.45e-36)	0.63 (1.60e-19)
2008	0.66 (8.62e-36)	0.51 (3.79e-12)
2009	0.70 (2.13e-40)	0.59 (1.75e-16)
2010	0.66 (1.74e-35)	0.59 (7.30e-17)
2011	0.69 (7.89e-37)	0.64 (1.20e-19)
2012	0.71 (3.18e-38)	0.60 (4.88e-16)
2013	0.67 (7.54e-35)	0.60 (9.28e-17)
2014	0.75 (5.22e-43)	0.67 (1.50e-21)


**Table 4.1:** Pearson coefficients and the Kendall-Tau rank correlation coefficients, along with the corresponding  $p$ -values, which resulted from comparing the static centrality ranking with the one that incorporates the observed flow. As a test case for the comparison, we use the distribution network of Oxycodone drugs.

similarity between the paired values ( $BC^p(v), BC^s(v)$ ) (for every  $v$  in the network), we measure their correlation using the Pearson correlation coefficient.

Besides the absolute centralities' values, we are also interested in the pure ranking, namely the sequence of nodes ranked from the most central (highest flow controlled) to the lowest one (lowest flow controlled). Indeed the two rankings may be similar, even if the actual centrality values differ. To compare the two rankings, we use the Kendall-Tau rank correlation coefficient.

In Table 4.1, we report the Pearson coefficients and the Kendall-Tau rank correlation coefficients (with the corresponding  $p$ -value) obtained for every yearly network snapshot from 2006 to 2014. As the first observation, we notice that all the values obtained are distant from 1. They range between 0.51 and 0.75, suggesting that the two measures produce results that are indeed correlated yet quite different.

Further, we see that the Kendall-Tau correlation coefficients are higher than the corresponding Pearson coefficients. This means that the error we commit if we use a static approach (i.e., only using information from the network topology) is smaller if we want to assess the relative centralities, namely the ranking, rather than absolute centralities.



**Figure 4.7:** The y-axis shows the Pearson (in blue) and Kendall-tau coefficients (in dark violet) resulting from comparing the static centrality ranking with the one that incorporates the observed flow. The test has been performed on 200 different networks. They differ in the number of products shipped through them. The x-axis shows the number of products shipped by every network. This is a progressive number from 1 to 200.

Ultimately, the Pearson and the Kendall-Tau rank correlation coefficients do not change much over the years. Specifically, the Pearson coefficient varies from a minimum of 0.51 to a maximum of 0.67; the Kendall-Tau coefficient varies from a minimum of 0.66 to a maximum of 0.75, suggesting that the differences between the two approaches are stable within the observation period.

We advance our investigation and perform the same analysis on multiple distribution networks, which differ by products shipped. We start considering the distribution network of a single product. Then, we progressively add one product. We stop at 200 products. With these 200 products we already cover 99% of all transactions of Oxycodone drugs. For each product added, a different network is constructed. Thus, we compute the centrality values for each of these 200 networks and measure the Pearson and Kendall-tau coefficients.

In Fig. 4.7, we show the results from the test (y-axis) as a function of the number of products examined (x-axis). The similarity between the two measures is very high when considering only a few products. It is almost 1, precisely 0.99, for the Pearson test and 0.8 for Kendall-Tau. This means that for a network of a single product, the error we commit using only information about the network topology is low.

However, the approximation of a static network approach becomes less accurate if we consider larger distribution networks with more products shipped. We see that the correlation decrease as we increase the number of products. Then it stabilizes at  $\approx 0.65$  for Kendall-Tau and  $\approx 0.5$  for Pearson.

The previous findings suggest that a static approach can be inadequate for assessing prominent nodes in distribution networks. Albeit it can be accurate for simpler network realizations, where only a few products are considered, it starts to lose accuracy as soon as we enlarge the scale of investigation. The simple test conducted in this Section shows that ranking measures based only on topological information should be taken with caution, as the error committed may not be negligible.

## 4.5 Discussion

In summary, this Chapter presented a novel method for extracting trajectories of goods from shipping records. Unlike existing data-mining techniques, the proposed method respects system-specific constraints (e.g., time and volume).

Applying the method to the ARCOS data, we extracted nearly 40 billion trajectories. These track the daily flow of individual packages from production to consumption through the distribution network. Next, we performed comprehensive statistics to present to data. By this, we delineated the key properties of the flow observed. First, we found a quite pronounced concentration of flow along a few trajectories. About 70% of the flow pass through a few trajectories (less than 10%). Then, we noticed that while the travel distances are quite long for most packages (at least 1000 km), the topological length of the trajectories is relatively small. 60% of the products traverse only two tiers of distributors before reaching the end consumers. Almost all products are shipped through only three tiers, at most. Further, we confirmed two expectations. We observed a quite short time scale for the distribution process: 97% of the products remain in the distribution system for less than one year. And we observed fragmentation of volumes, namely, packages are reshaped along their pathways from producers to consumers.

---

Besides the empirical insights gained from this analysis, this Chapter's main contribution is to have generated a rich data with a unique high resolution, i.e., the trajectory data. These data can be used to inform models operating on short-time scales. For instance, we showed how to use these data to revise standard measures for node rankings. Then, in Chapter 9, we incorporate these data into the cascade model proposed in order to capture readjustment of day-to-day flow under substitution policies.

## **Part II**

# **Formation, evolution, and growth**

“It is in the character of growth that we should learn from both pleasant and unpleasant experiences.”

*N. Mandela*

## Chapter 5

# Modeling network formation and evolution

### Summary

In this Chapter, we develop a network growth model to explain the formation and evolution of distribution networks. Taking a firm's perspective, we consider two driving mechanisms for the emergence of the observed structures: centralization and multi-sourcing. The former indicates firms' preference to link to central firms in the network. The latter indicates firms' preferences to link to more than one source partner. We calibrate and validate the model against various statistical properties of a real-world distribution system. Our results reveal a good match between model predictions and real-world observations. Thus they suggest that the two mechanisms proposed, albeit simple, are sufficient to reproduce the observed topologies.

---

AA wrote this Chapter specifically for this thesis. Results are based on A. Amico, G. Vaccario, F. Schweitzer "Formation and evolution of large-scale distribution networks", Working Paper. AA contributed to designing the research question and deriving the model principles. AA wrote the code, performed the simulations, and interpreted the results.



## 5.1 Introduction

Manufacturers and distributors organize themselves into hierarchical structures to efficiently distribute goods to end consumers (Kito *et al.*, 2014). In this Chapter, we pose the question: how do these types of structures form and evolve?

In our view, distribution networks emerge spontaneously from the interaction and cooperation of several firms. Previous studies, instead, have viewed distribution networks as entirely designed by a focal manufacturer (Wang *et al.*, 2018). These studies typically assume that a single manufacturer is responsible for determining the structure of its distribution network and have developed methods to help manufacturers design optimal distributor locations such that total costs are minimized.

This idea of a focal manufacturer controlling the whole supply network has provided an appealing vision for managers, yet poorly suited to real-world systems (Choi *et al.*, 2001). Focal manufacturers can usually choose their direct suppliers and distributors. However, they have limited control over whom their direct distributors ship to (Brintrup and Ledwoch, 2018). This “eco-systemic” view, where firms are self-sufficient and autonomously join the network, is novel in the supply-chain domain and little discussed in the current literature.

Building on this reasoning, we propose an evolutionary growth model to describe the network formation of distribution systems. We explain the formation of such systems as the outcome of an evolutionary process rather than an accurate design strategy carried out by a single manufacturer. We aim to derive simple yet sufficient interaction rules at the micro-level, allowing us to reproduce real-world features at the macro-level. To validate our model, we use opioid distribution networks reconstructed from the ARCOS data.

The Chapter is organized as follows. We start introducing the network growth model: we derive the underlying interaction rules in Subsection 5.2.2 and clarify the role of the model parameters in Subsection 5.2.3. Next, we validate the model against real-world data using a two-steps procedure: we first fit the model to the data and, hence, obtain the optimal values for the parameters

in Subsection 5.3.1; we then test the goodness of the model in reproducing the main properties of real-world networks, including out and in-degree distributions, path length distributions in Subsection 5.3.2. Despite the model's simplicity, we find a good match between model prediction and real-world observations. Yet, simplicity comes with some model limitations discussed in Section 5.4 along with the Chapter summary.

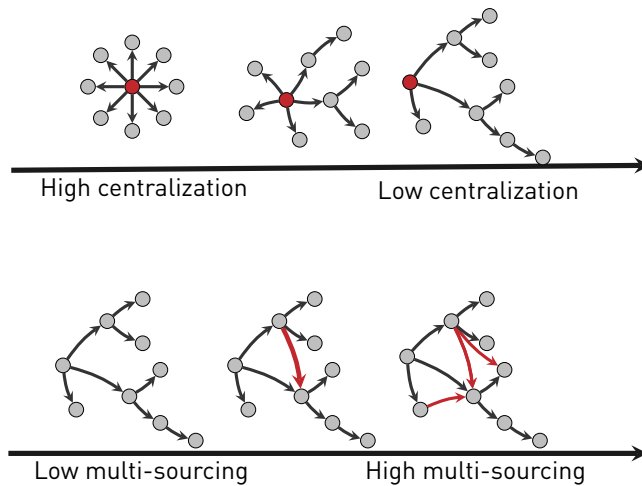
## 5.2 Model

### 5.2.1 Two driving mechanisms

**Centralization.** The term “network centralization” is commonly used in the supply chain domain to describe the degree to which a supply network is organized around a single firm (Schmitt *et al.*, 2015; Treiblmaier, 2018). A highly centralized network exhibits a star structure, where a central firm is connected to all others while the others have no connections. On the other hand, the lowest level of centralization occurs when all firms have an equal number of connections. Various studies have explored the concept of centralization in supply networks, emphasizing its implications for network performance (Kim *et al.*, 2015, 2011). Highly centralized supply networks have the benefit, for example, of increasing controllability in production planning. However, the drawback of such networks is a potential lack of system responsiveness, as firms in different tiers may not have direct contact with each other (Choi and Hong, 2002).

Scholars have empirically investigated the level of centralization of several supply networks. For example, some studies have examined the textile supply network in Prato, Tuscany (Paniccia, 1998) and automotive supply networks (Choi and Hong, 2002; Clark *et al.*, 1987). Additionally, centralization has been studied in industries such as chemical, electrical, cargo, electronic, and aircraft (Perera *et al.*, 2017b). These investigations have revealed a range of network centralization levels, with some networks exhibiting high centralization (Paniccia, 1998), while others display lower levels (Perera *et al.*, 2017b).

In our study, we examine “centralization” as a key mechanism in the formation of distribution networks. In our work, we refer to centralization as the firms’ tendency to connect to the most central firm in the network, i.e., the one with the most connections. If this tendency is strong, a star-like network structure emerges. Conversely, a low level of centralization will result in a more branched structure. A schematic representation of a high, medium and low centralized structure is shown in Fig. 5.1 (top row). The red node in the Figure is the central company, e.g., the focal manufacturer.



**Figure 5.1:** Schematic representation of network structures with high, medium and low centralization (top row); and with low, medium and high level of multi-sourcing (bottom row).

**Multi-sourcing.** Firms typically source products from multiple independent partners, resulting in numerous incoming connections (Burke *et al.*, 2004). The two primary sourcing strategies are multiple-sourcing, which involves having multiple source partners, and single-sourcing, which consists in having only one source partner (Inderst, 2008).

Single-sourcing can strengthen business relationships and foster cooperation, leading to efficient deliveries and successful “just-in-time” inventory strategies.

However, it is not always the optimal choice as it increases firms' exposure to the risk of supply interruption. To mitigate this risk, firms may implement multi-sourcing. A schematic representation of the network structure with low, medium and high multi-sourcing is shown in Fig. 5.1 (bottom row). The red links represent connections with additional source partner, which are created when multi-sourcing is implemented.

We propose that multi-sourcing, which represents firms' preferences to connect with multiple source partners, is the second key driver for network formation.

In the following Subsection, we formalize the two key mechanisms of network formation, i.e., centralization and multi-sourcing. Our work builds on the network growth model developed by Klemm *et al.* (2005). Based on simple interaction rules, the model well captures the emergence of tree networks\*. If all firms in the network implement single-sourcing, a perfect tree structure emerges (see discussion on motifs in tree-like networks in Subsection 2.2.2). As this is not always the case in reality, we use the model proposed by Klemm *et al.* (2005) as a starting point, and we extend it to capture structural deviations from trees consistent with real-world observations.

### 5.2.2 Deriving the interaction rules

Let's consider a distribution system comprising a single manufacturer, multiple distributors and consumers. We describe this system as a directed network where links represent supply relations and nodes represent firms and consumers. As this Chapter focuses on the network topology, we are not so much interested in the amount shipped, i.e., link weight. The latter is the object of study of the next Chapter.

Given a direct link,  $i \rightarrow j$ , we define the node  $i$ , the node from which the link departs, as *source* and the node  $j$ , the node in which the link ends, as *target*. Further, by  $d_i^{\text{out}}$  we denote the out-degree of  $i$ , that is, the number of outgoing links, and by  $d_i^{\text{in}}$  we denote its in-degree, that is, the number of incoming links.

---

\*Specifically, the model proposed by Klemm *et al.* (2005) was introduced to explain the formation of computer directories' trees. Nevertheless, it provides a general approach to describe the emergence of tree networks. The model was further investigated by Geipel *et al.* (2009) who derived an analytical solution for the parameters' calibration.

More explicitly,  $d_i^{\text{out}}$  represents the number of distribution channels  $i$  handles, whereas  $d_i^{\text{in}}$  represents the number of source nodes  $i$  relays on.

Initially, at  $t = 1$ , there are two nodes: the root, representing the manufacturer, and a child node, representing the distributor. These are connected through a direct link from root to child. At every time  $t$ , a new link is created between two nodes, a source node and a target node. Thus at time  $t$  the number of links is  $E(t) = t$ .

The source node  $i$  is selected among the existing nodes with a probability  $p_i(t)$ , given by:

$$p_i(t) = q_s \frac{d_i^{\text{out}}(t)}{t} + (1 - q_s) \frac{1}{N(t) - 1} \quad (5.1)$$

The first term describes a preferential attachment mechanism, where the probability of becoming a source node is proportionate to the number of outgoing links the node has. Note that according to a preferential attachment mechanism the normalization of the first term would be the sum of all nodes' out-degree, i.e.,  $\sum_i d_i^{\text{out}}(t)$ . Yet, we know that the sum of nodes' out-degrees in directed networks equals the total number of links  $E(t)$ . Since  $E(t) = t$ , we can simply write  $t$  as normalization factor.

Supported by empirical evidences, we consider a fully outsourced distribution process. This means that the manufacturer outsources its distribution activities to a single direct distributor (e.g., its warehouse) which then, in turn, links to other distributors. Hence, we have  $p_i = 0$  for the root  $i$ , meaning that the manufacturer has zero probability of establishing new links in addition to the one it has at time  $t = 1$ . Given this setting, we need to assume that the first distributor has out-degree different from zero, i.e.,  $d_i^{\text{out}}(t) = 1$ . Otherwise for  $t = 1$  and  $q_s = 1$  we have  $p_i = 0$  for the only existing nodes, i.e., the manufacturer and the first distributor. Note that this assumption does not change the normalization factor in Eq. (5.1). Indeed we have that the sum of out-degrees increases by one as the distributor node acquires a new link at time  $t = 1$ , but it also decreases by one as the manufacturer's out-degree is not included in the sum. Hence, the normalization factor is still  $t$ .

The second term describes the situation where all existing nodes have the same probability of being selected as source node. The minus one in the

normalization factor account for the fact that the manufacturer is excluded from the candidate nodes to be selected as source.

Now we come to the interpretation of both terms. The first term describes a preferential allocation of business opportunities, based on the “rich-get-richer” dynamic. The more *central* a firm is the more likely it grows. Note that by the term *central* we refer to the number of (outgoing) connections the node has.

The second term represents a random allocation of business opportunities. The chance for a firm to be selected as source and enlarge its business (by increasing the number of distribution channels), is subjected to random factors. These include, for instance, the stochasticity of demand, possible changes of market conditions (the entrance of new products), or possible changes of environmental conditions (e.g., the entrance of newcomer firms). Using the parameter  $q_s$  to interpolate between the two terms, we can tune the level of centralization. When  $q_s \approx 1$ , the first term is predominant, and we have high centralization. When  $q_s \approx 0$ , the second term is predominant, and we have low centralization.

Clearly, the mechanisms just proposed do not embrace all possible strategies for source selection. For instance, the geographical proximity of the distributor may play a role in selecting it as source partner, as well as its loyalty or efficiency in delivery. Including all these factors individually would add unnecessary complexity to the model, and limits its empirical validation when data, e.g., on shipping times or travel routes, are not available. Most importantly, a detailed description of all possible partner selection mechanisms is beyond the scope of this study. Our goal is to determine the set of simple and sufficient mechanisms, at micro-level, that allow to reproduce the main network properties observed at the macro-level.

Once the source node  $i$  is selected, the target node  $j$  is chosen. While  $i$  is always one of the existing nodes in the network,  $j$  can be either a newcomer node, or an existing node. Specifically, at a rate  $\alpha$   $j$  is a newcomer, and at a rate  $1 - \alpha$   $j$  is an existing node. In the latter case the target is selected among the existing nodes with a probability  $p_j$ , given by:

$$p_j(t) = q_t \frac{d_j^{\text{out}}(t)}{t - n_i(t)} + (1 - q_t) \frac{1}{N(t) - n_i(t) - 1} \quad (5.2)$$

$p_j$  is the probability for an existing node  $j$  to establish a link with a new source node in the network. Hence, the first term in Eq. (5.2) describes a preferential attachment mechanism, where the probability of a node to attach to a new source node is proportionate to the number of outgoing links the node has. As for the first term of Eq. (5.1), the normalization of the first term of Eq. (5.1) would be  $t$ . Nevertheless, differently from the source selection mechanism not all nodes in the network are eligible as targets. Non-eligible targets are nodes already connected to the source  $i$  and the source itself. In other words, we do not allow for the formation of self-loops, meaning links of nodes with themselves, and multi-edges, meaning links between the same source and target node. This choice implies that  $p_j(t)=0$  for  $j=i$  or  $j \in \text{nb}_i^{\text{out}}$ . The correcting factor  $n_i(t)$  in the denominator is used to exclude the non-eligible targets from the normalization, i.e.,  $n_i(t) = d_i^{\text{out}} + 1$ .

The second term describes the situation where all existing nodes have the same probability of linking to an additional source node. For the normalization term we account for all nodes in the network, yet excluding the non-eligible targets, by means of the factor  $n_i(t)$ , and the single manufacturer, by means of the minus one.

The mechanism expressed by Eq. (5.2) is the firm tendency to implement multi-sourcing. The current literature lacks a clear understanding on the factors driving firms to decide for multi-sourcing, rather than for single-sourcing. In absence of empirical insights, we propose two factors. The first factor is the “size” of the firm, evaluated on the number of distribution channels it handles. The bigger the firm is, the more likely it is to implement multi-sourcing, as expressed by the first term in Eq. (5.2). The second is a random factor: every firm are equally likely to implement multi-sourcing.

In addition to firms’ size and random effects, there may be other factors that drive strategic sourcing decisions. For example, products that are highly demanded and classified as essential goods may lead firms to implement multi-sourcing to ensure a reliable supply. Although our work does not capture these additional dimension, it is essential to explore it in future studies.


The parameter  $q_t$  is used to interpolate between the two factors. Fitting the value of  $q_t$  to real-data, we can better understand to what extent multi-sourcing

is actually preferred by big firms. We better clarify the role of the parameter  $q_t$ , along with the role of  $\alpha$  and  $q_s$  in the next Subsection.

### 5.2.3 Synthetic case

The model proposed is characterized by three global parameters:  $q_s$ ,  $q_t$ , and  $\alpha$ . To clarify their role we discuss a few explanatory networks obtained through model simulations.


We start by elucidating the role of  $q_s$ . Thus, keeping fix  $\alpha$  and  $q_t$ , we tune  $q_s$  to generate three different network realizations. Specifically, we set  $\alpha=1$ , and  $q_t=0$ , and consider the following  $q_s$  values:  $q_s=1$ ,  $q_s=0.8$ , and  $q_s=0.4$ . The resulting topologies are depicted in Fig. 5.2. Note that for the sake of clarity we show networks comprising a small set of nodes and links.



**Figure 5.2:** Synthetic networks generated through model simulations with  $\alpha=1$ ,  $q_t=0$  and:  $q_s=1$  (figure a);  $q_s=0.7$  (figure b),  $q_s=0.3$  (figure c). Nodes' size is proportionate to their out-degree. The red node is the root representing the manufacturer. Orange nodes represent distributors and consumers (leaves).

We see that by setting  $q_s=1$ , we obtain a star network where a single node, in the centre, is connected to all the others which, in turn, are not connected among themselves. By decreasing the value of  $q_s$  from 1 to 0.7 and further down to 0.3 we increase network branching while reducing the star-like shape.





**Figure 5.3:** Synthetic networks generated through model simulations with  $q_s=0.5$  and:  $\alpha=1$ ,  $q_t=0$  (figure a);  $\alpha=0.6$ ,  $q_t=0.8$  (figure b). Nodes' size is proportionate to their in-degree. The red node is the root representing the manufacturer. Orange nodes represent distributors and consumers (leaves).

Thus, from a mere topological perspective,  $q_s$  controls for the network branching and the departure from a star-like shape. We see that high  $q_s$  values signal high centralization of the underlying distribution system, whereas low  $q_s$  signal low centralization.

Next, we clarify the role of  $\alpha$  and  $q_t$ . For this, we fix  $q_s$  and vary both  $q_t$  and  $\alpha$ . In Fig. 5.3, we depict two networks generated through the model with the following parameters: (a)  $\alpha=1$  and  $q_t=0$ ; (b) and  $\alpha=0.7$  and  $q_t=0.8$ .  $q_s$  is kept fix to a medium value, i.e.,  $q_s=0.6$ . The size of the nodes is proportionate to their in-degree.

We see that for  $\alpha=1$  we obtain a perfect tree. All nodes in the tree have only one parent node. On the contrary, for  $\alpha \neq 1$  we see a departure from the tree structure. A few nodes have multiple parents.  $\alpha$  is an indicator of the number of times a multi-sourcing strategy is implemented. The higher the  $\alpha$  the less multi-sourcing strategies are implemented. While tuning  $\alpha$ , one controls for the number of times this strategy is implementing, tuning  $q_t$ , one increases (or decreases) the correlation between nodes' preference towards multi-sourcing and their out-degrees. Fig. 5.3 shows the case of  $q_t=0.8$ , i.e., high correlation. Note that nodes' size is proportionate their in-degrees. In this case we notice

that nodes with high out-degree also have high in-degree (i.e., multiple source nodes).

To briefly recap, the model introduced describes the link formation in distribution networks as driven by two main mechanisms: the firms' preference to link to central firms (centralization); and firms' tendency of increasing its number of suppliers (multi-sourcing)<sup>†</sup>. The former is controlled by the parameter  $q_s$ , whereas the latter by the parameters  $\alpha$  and  $q_t$ .

## 5.3 Model validation

In this Section, we validate the model against the real-world data. The validation procedure consists of two steps: (i) first we fit the model to the data to find the optimal parameters, namely the parameters for which we obtain the best fit; (ii) then feeding the model with the optimal parameters, we test it against stylized properties of a real-world distribution network. The first step is explained in Subsection 5.3.1, whereas the second step is explained in Subsection 5.3.2.

### 5.3.1 Optimal parameters' estimation

We start by fixing the network measures we can directly compute from the data. These are: (i) the number of nodes  $N$ ; (ii) the number of links  $E$ ; (iii) the entry rate  $\alpha$ .

The entry rate is computed analytically. From Subsection 5.2.2, we know that the number of links evolves as  $E(t)=t$ ; and the number of nodes grows, on average, as  $N(t)=\alpha \times t + 1$ <sup>‡</sup>. We can substitute the time  $t$  with  $E(t)$  in the latter equation and replace the number of edges  $N(t)$ , and links  $E(t)$

<sup>†</sup>Note that, by construction, the two mechanisms are not independent.

<sup>‡</sup>Note that the term  $+1$  accounts for the manufacturer node. Further, note that the expression for the evolution of  $N$  is always valid except for one case: very low  $\alpha$  values. In this case, to avoid forming multiple links between the source and target, one should consider the entrance of additional newcomer nodes. Nevertheless, very low  $\alpha$  values are not interesting for our study as they imply very high network density. High density is rarely detected in real-world distribution networks and tree-like networks. These are intrinsically sparse networks. Therefore,  $N(t) = \alpha \times t + 1$  is always valid for our case of interest.

with corresponding empirical values  $N$  and  $E$ . Thus, we obtain the following expression for the entry rate:

$$\alpha = \frac{N - 1}{E} \quad (5.3)$$

Differently from  $\alpha$ , the parameters  $q_s$  and  $q_t$  can not be calibrated analytically. Therefore, we adopt a Monte Carlo approach. We perform extensive computer simulations to explore the entire bi-dimensional space of the parameters and assess the best fit. Specifically, we consider values of  $q_s$  and  $q_t$  ranging from 0.1 to 0.9, with an interval of 0.02. For each pair  $(q_s, q_t)$ , we run the model and assign a “fitting score”. The optimal parameters are the ones for which we obtain the highest fitting score. Thus, we perform a total of 1,600 computer simulations. We stop every simulation when the synthetic network reaches the same number of links as the empirical one. We define  $\tau$  as the time at which every simulation ends.

We discard extreme values (i.e.,  $0.1 \leq q_s$  and  $q_s \geq 0.9$ ;  $0.1 \leq q_t$  and  $q_t \geq 0.9$ ) as we do not expect them to be optimal for the networks analyzed. Low  $q_s$  values produce branched networks characterized by long distribution paths from manufacturers to consumers. On the contrary, high  $q_s$  values generate star-like networks with short distribution paths, i.e., only one-step length. Further, high  $q_t$  values lead to networks where only a few nodes have an in-degree different from one. On the contrary, low  $q_t$  values produce networks where almost all nodes have in-degree different from one. As we know from the empirical insights gathered in Chapter 4 and Chapter 3, all the four situations mentioned above do not characterize our data sample. Hence, the choice of the parameters’ range.

We follow the approach proposed by Tomasello *et al.* (2014) to construct the fitting score. We start by considering a few macroscopic network features or observables  $\Omega$  (e.g., the average out-degree). For each  $\Omega$ , we measure the distance,  $\delta$ , between the empirical feature and the synthetic one, as expressed by:

$$\delta(q_s, q_t) = \frac{\Omega_e - \Omega_s(q_s, q_t)}{\Omega_e} \quad (5.4)$$

Manufacturer	<i>dea-number</i>	N	$\alpha$	$q_s$	$q_t$	Fit Score
MALLINCKRODT	RM0231821	719	0.62	0.64	0.56	0.75
WATSON LAB	RW0288933	634	0.63	0.64	0.28	0.59
BOEHRINGER	RR0112514	443	0.52	0.68	0.54	0.37
TEVA PHARM	RL0146274	428	0.87	0.78	0.84	0.41
NOVARTIS	PD0038667	355	0.64	0.62	0.24	0.58
ACTAVIS	RA0306490	314	0.86	0.82	0.90	0.17

**Table 5.1:** Name, *dea-number* of the top six opioid manufacturers, as well as the number of nodes in the distribution networks reconstructed, the optimal parameters ( $\alpha, q_s, q_t$ ), and the fitting score.

where the subscript  $e$  stands for empirical, and the subscript  $s$  stands for synthetic, or simulated.


If  $\delta$  in Eq. (5.4) is smaller than a relative error  $\varepsilon$ , we argue that the network generated by the model is close enough to the empirical one<sup>§</sup>. To translate this boolean measure, i.e., close or not close, into a score, for each parameter’s combination, we generate 100 synthetic networks. The frequency at which the condition  $\delta(q_s, q_t) \leq \varepsilon$  is met defines our fitting score.

By this, the fitting score is normalized between zero (none of the 100 networks meets the condition) and one (all the 100 networks meet the condition). Simply put, a fitting score close to one means a very good match between the model predictions and data, whereas a fitting score close to zero means a very bad match. The parameters with the highest fitting score are defined as optimal.

In our study, we select the following observables: the first and second moments of the distribution of out-degrees and in-degree distribution and average path length. We then consider a 5% relative error ( $\varepsilon = 0.05$ ) on the first moments and a 25% relative error on the second moments ( $\varepsilon = 0.25$ ). We select the  $\Omega$  such that they incorporate the less, albeit sufficient, information to reproduce real-world features. Indeed, besides the most straightforward choice of first moments, we need to include also second moments as the first moments alone are little informative about highly-skewed distributions.

<sup>§</sup>Note that if multiple observables are analyzed, the condition needs to be fulfilled under all the  $\Omega$  evaluated.

We fit the model on six different real-world distribution networks, reconstructed from the ARCOS data. We follow the network reconstruction method described in Section 2.2.1<sup>¶</sup>. The root node of each of the six networks represents a top opioid producer. In Table 5.1, we report the obtained values for the optimal  $q_s$  and  $q_t$  as well as the corresponding fitting score, the producer’s names and the total number of nodes in the network. To further illustrate the fitting procedure, in Fig. 5.4, we display the bi-dimensional parameter space for a single distribution network. The largest network is selected as a case example. We use both a 3-dimensional color map (left plot) and a 2-dimensional color map (right plot) to show the fitting scores for each pair  $(q_s, q_t)$  in the parameter space. The scores follow the color schema on the right-hand side of the Figure.



**Figure 5.4:** Parameter space illustrated via a 3-dimensional (left plot) and a 2-dimensional color map (right plot). The fitting score for each pair  $(q_s, q_t)$  in the parameter space is visualized through the color schema on the right-hand side.

From Fig.5.4, we see that the parameters with the highest scores are confined to a clear “optimal” region (colored in dark orange). We notice that this optimal region is very narrow along the  $q_s$  dimension, and broader along the  $q_t$  dimension, suggesting the  $q_t$  have more suboptimal values than  $q_s$ . The optimal combination of values (represented by the peak in the 3-dimensional

<sup>¶</sup>For the current analysis, we consider all drugs containing Oxycodone as a basic opioid.

color map) has coordinates:  $q_s = 0.64$  and  $q_t = 0.56$ . These values are also reported in Table 5.1, together with the optimal values obtained for the other networks.

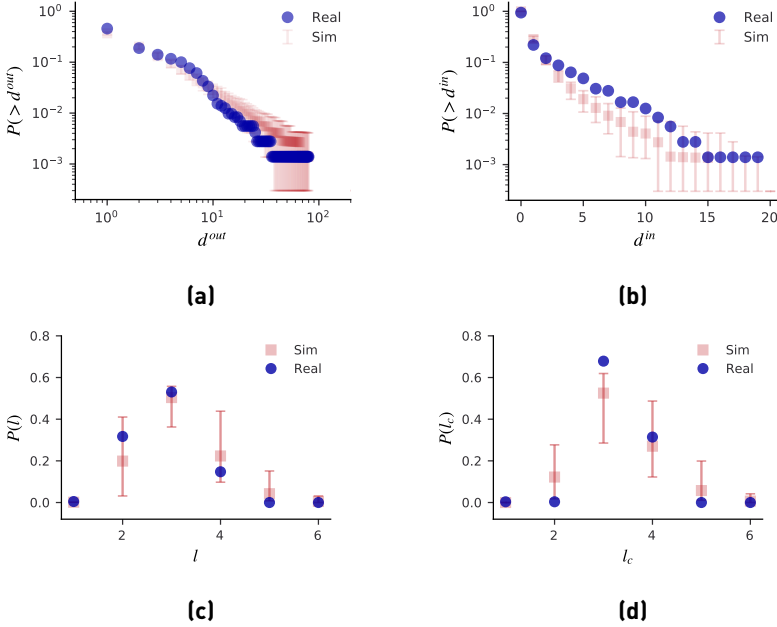
From the Table 5.1, we see that all the distribution networks examined exhibit high  $q_s$  values, ranging from 0.62 to 0.82. This suggests that centralization is predominant to the random mechanism: firms are more likely to connect to central firms than pick them randomly. As a result, the observed structures are very much star-like.

Next, we see that all  $\alpha$  values are quite different from 1, signalling that the observed topologies have a clear departure from perfect tree structures. Also, it confirms our expectation: firms decide for strategic multi-sourcing, yet at a lower frequency than single-sourcing. Note that  $\alpha$  does not quantify the proportion of firms implementing one strategy rather than the other. It instead indicates the frequency at which a single-source strategy has been adopted. And in turn, it indicates the frequency  $(1 - \alpha)$  at which a multi-sourcing strategy has been pursued even, for instance, by the same firm multiple times.

While we obtain similar values of  $q_s$  and  $\alpha$  across the six networks examined, the  $q_t$  values are more diverse. For two networks, we find that firm size, defined by its out-degree, is a predominant factor in identifying a firm's tendency towards multi-sourcing ( $q_t \geq 0.8$ ). The opposite is true in the other two networks ( $q_t \leq 0.2$ ). For the remaining two networks, firm size is comparable to the random factor ( $q_t \approx 0.5$ ). For the latter two cases, our results suggest that firm size is not a driving factor for multi-sourcing and that other factors could be relevant. These include, for instance, the reliability of existing suppliers, or their efficiency in terms of time-delivery, or even their risk exposure. What factors are the most relevant ones is still an open question that future studies can explore.

### 5.3.2 Statistical validation

In this Subsection, we perform the second step of our validation procedure. Feeding the model with the optimal  $\alpha$ ,  $q_s$ , and  $q_t$ , obtained in Subsection 5.3.1, we generate 200 network realizations to compare against the empirical network.



**Figure 5.5:** Probability distribution of out-degrees (a) in-degrees (b), path lengths from root to any node (c), and from root to leave nodes (d). The blue color indicates real-world data, whereas the red one indicates simulated data (mean and error bars from 200 runs).

The comparison is performed by evaluating the following network properties: (i) distribution of in-degrees,  $d^{in}$ ; (ii) distribution of out-degrees,  $d^{out}$ ; (iii) distribution of shortest paths from the root to any node in the network,  $l$ ; distribution of shortest paths from the root to only leave nodes in the network  $l_c$ . Leave nodes to represent end consumers.

In Fig. 5.5 we show the results obtained for a single distribution network, the largest one. We mark in blue the real-world data and in red the simulated one. Specifically, we report the median (red square) and the error bar computed at each bin for the simulated data. Note that the error bar is not obtained as the standard deviation of the data. Since the normality of the sample at each bin is not guaranteed, we discard 5% of the most extreme values and compute the

upper (and lower) limit of the error bars from the minimum (and maximum) value of the remaining 95% of data.

We see that the model reproduces the typical heavy tail behavior of the empirical distribution of out-degrees (plot a). It also leads to right-skewed distributions of in-degrees similar to the one observed (plot b). To properly test these observations and assess the statistical similarity between the empirical distribution and each distribution in the synthetic sample, we perform a KS-test. Considering a  $p$ -value of 0.001, we obtain compatibility between real-world data and simulations in 86% of the cases for the out-degree distribution and in 95% of the cases for the in-degree distribution.

Next, we analyze the distribution of path lengths from root to any node (plot c) and the distribution of path lengths from root to leaves (plot d). Note that in both figures, distance zero means the distance of the root node (i.e., the manufacturer) from itself. We see that the model can reproduce the very peaked shape of the two distributions. As most empirical data points fall within the respective error bars, we can confirm a good match between real-world and simulated data.

Further, note that the peak registered at  $l = 3$  and at  $l_c = 3$  suggests very short-depth networks, as most of the leaf nodes (and nodes in general) are at a distance 3 from the manufacturer. This confirms our findings discussed in Chapter 4 where we empirically reconstructed the distribution paths and highlighted the short topological distances connecting manufacturers to consumers along these paths.

Although we can overall confirm that that model well predicts the main topological features of an empirical distribution network, we still observe a few mismatches. First, we see a few deviations of the model prediction from the real-world data in correspondence of medium in-degree values, i.e.,  $5 \leq d^{\text{in}} \leq 8$ . The empirical distribution exhibits a slightly more complex behavior than a simple power law. The simple preferential mechanism based on nodes' out-degrees seems to not capture this additional complexity.

Second, we notice a small mismatch in the distribution of path lengths from root to leaves (plot d). Specifically, in correspondence of  $l_c = 3$ , the model underestimates the number of paths and overestimates it in correspondence of



$l_c = 5$ . As our major assumption is that consumers and distributors obey to the same attachment rules, a slight discrepancy in the distance of consumers from the root is observed. Improvement of the model should consider incorporating node heterogeneity in the interaction rules. We leave to future studies this investigation.

## 5.4 Discussion

In the present Chapter, we introduced a network growth model to explain the formation of large-scale distribution networks. Taking a firm's perspective, we proposed two mechanisms to explain the emergence of the observed structures: centralization and multi-sourcing. Centralization expresses firms' preferences in linking with central firms, whereas multi-sourcing represents firms' tendency to acquire additional suppliers or source nodes. The first strategy provides cost-saving advantages and improves operational efficiency, whereas the second strategy enhances resilience by guaranteeing alternate sources in case of a single supplier's failure.

The proposed mechanisms can be tuned using three model parameters:  $q_s$ ,  $q_t$ , and  $\alpha$ . Parameter  $q_s$  controls for network centralization, while  $q_t$  and  $\alpha$  control for multi-sourcing. We tested the model using the opioid distribution networks reconstructed from the ARCOS. Optimal parameters were identified using a Monte Carlo approach, revealing that the networks exhibit high centralization, frequent multi-sourcing, and a clear departure from perfect tree structures.

Next, we statistically validated the model and found a good match between model predictions and real-world observations. Specifically, we could reproduce the distinctive shapes of out and in-degree and path length distributions. Although the model already captures the main features of an empirical network, one could improve it on various fronts. Firstly, the model assumes identical linking strategies for consumers and firms. More realistic linking rules for consumers can be considered in future works at the cost of increased model complexity.

Secondly, the manufacturer's probability of linking to other distributors is assumed to be zero. The assumption is based on observations from the ARCOS

dataset. However, the model can be adapted to allow for different probabilities or calibrated to other real-world data.

Thirdly, the model disregards link weights, i.e., the amount of goods passing through each link. Yet, the evolution of link weights may influence topological growth and vice-versa. An extension of the model, presented in Chapter 7, addresses this limitation by exploring the growth of distribution systems as weighted networks.

Finally, the proposed model is applicable to distribution networks of various industrial sectors. The centralization and multi-sourcing strategies considered are widely discussed in the literature in industries as varied as automotive, electronic, and chemical industries (Inderst, 2008; Perera *et al.*, 2017b). Therefore, we expect these principles to be relevant driving mechanisms for other networks as well. However, it remains an open question to what extent the proposed model can replicate the stylized facts of different distribution networks and whether additional mechanisms need to be incorporated. Further research in this direction is encouraged.

## Chapter 6

# Firms' growth and their determinants

### Summary

This Chapter aims to empirically analyze the growth rates of firms operating within the opioid distribution network. We provide a brief overview of existing theories on firm growth and evaluate their relevance within the supply chain domain. Next, we explore the potential factors that influence the growth of opioid manufacturers and distributors and assess their effects through regression analysis.

---

This Chapter has been written specifically for this thesis. AA contributed to the design of the research question, performed the data analysis and interpreted the results.

## 6.1 Introduction

For many years, scholars have investigated the growth of business firms, drawing from the seminal work of Gibrat (1931). Gibrat's law of proportionate effect simply states that firm growth rates are independent of their sizes. When expressed in logarithmic form, these rates follow a Gaussian distribution.

Although extremely simplistic, Gibrat's law has become a good benchmark for comparing real-world observations. Scholars have used a variety of data to test the validity of this law including, for example, the Italian manufacturing firm data (Bottazzi and Secchi, 2006), the US pharmaceutical companies data (Fu *et al.*, 2005), and the Dutch manufacturing firms data (Zhou and de Wit, 2009). They found that the distribution of growth rates is "tent-shaped" and better described by a Laplacian distribution rather than a Gaussian one (as hypothesized by Gibrat). Interestingly, a similar tent-shape also characterizes the distribution of growth rates of other economic actors as diverse as countries and cities (Bottazzi *et al.*, 2011; Bottazzi and Secchi, 2006), becoming a stylized fact for economic growth processes.

In this Chapter, we empirically investigate the growth rates of the US opioid manufacturing and distribution firms. Our goal is to test the validity of the traditional theories against these novel data. Specifically, we start by testing the previously proposed functional forms, i.e., the Gaussian and Laplacian (Section 6.2). Next, we perform an econometric analysis on the determinants of growth rates (Section 6.3). Conclusions from the analysis are drawn in Section 6.4.

## 6.2 Growth rates and their functional form

As a first step to evaluating growth rates, we need to define the firm size and how we can measure it. Previously, scholars have proposed various measures as indicators of firm size, including their number of employees, total sales, and total assets (Bottazzi and Secchi, 2003; Coad, 2007; Zhou and de Wit, 2009). Also, physical measures of output have been used, e.g., tons of goods produced, when firms under study were manufacturers producing the same type of good.

From the ARCOS dataset, we do not have information on the total sales of firms, nor on their assets or number of employees, but we do have information about the amount of goods they ship. Assuming that the quantity shipped by a firm,  $i$ , is indicative of its sales, we proxy its size,  $s_i$ , as its total ship-out,  $W_i^{\text{out}}$ , i.e.:

$$s_i = W_i^{\text{out}} \quad (6.1)$$

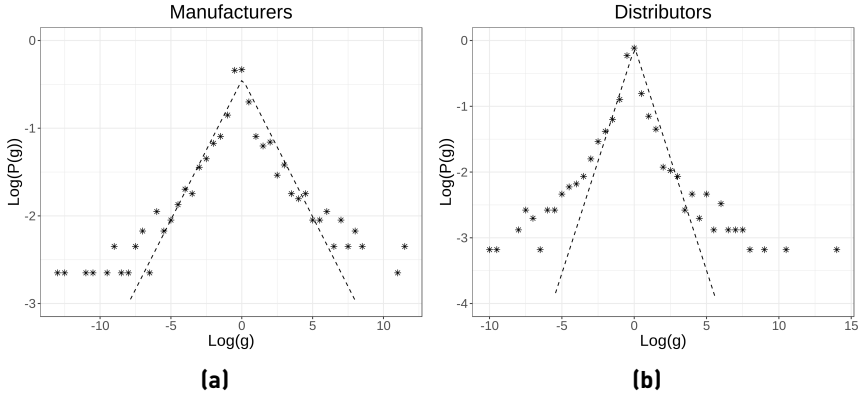
The total ship-out in Eq. (6.1) is measured in grams and represents the aggregate amount of basic opioids shipped, regardless of the specific drug product. A manufacturer can produce multiple drug products with varying sales that may offset each other or add up. Our goal is not to measure the sales of individual products but to determine a firm's overall annual sales in a specific opioid market, such as the Hydrocodone one. Therefore, we use the total grams shipped as a proxy for the product sale and sum them across all the packages shipped by the firm.

Moreover, sales of products from different opioid markets may also influence each other. For instance, an increase in the sales of Oxycodone products may lead to a decrease in sales of Hydrocodone products. To better assess firms' total sales, one should aggregate the sales from products containing even different basic opioids. However, using grams shipped as a proxy for sales, in this case, would be unreliable. Opioids have varying potencies, resulting in different grams per product. Hence, a separate measure should be used when aggregating products belonging to various opioid markets. For instance, a possible approach would be to convert all grams to morphine equivalents and use that measure as a proxy for sales. As this analysis requires further investigation, it is left to future work.

Thus, following previous works, we determine the growth rate of firm  $i$  as the logarithm of its *1-year* size variation, i.e.,

$$g_i(y) = \log \left[ \frac{s_i(y)}{s_i(y-1)} \right] \quad (6.2)$$

where  $y$  indicates the year under study. Note that the standard definition of  $g_i$ , as expressed by Eq. (6.2), assumes zero dependencies among firm growth rates. This assumption does not hold in distribution networks. Distributors



**Figure 6.1:** Empirical growth rates,  $g$ , of manufacturers (a) and distributors (b) in the opioid distribution network. The star symbols indicate the empirical distribution of growth rates. The dashed line indicates the Laplacian fit. We binned the data for visualization purposes only and fitted the theoretical distribution to the non-binned data.

can increase their shipments only if they, in turn, have received enough goods. In the current Chapter, we assume firm growth rates to be independent. In the next Chapter, we revise this critical assumption and propose a network growth model to account for the inter-firm dependencies.

Based on Eq. (6.2), we compute the annual growth rates of the opioid manufacturers and distributors from the year 2006 to the year 2014. Fig. 6.1 displays the distribution of growth rates of manufacturers (a) and distributors (b). Note that both plots have been obtained by pooling the growth rates of the nine observation years and considering all the shipments of a single opioid, i.e., Hydrocodone. The star symbols indicate the empirical values, whereas the dashed line indicates the Laplacian fit against the data.

The fit is based on the maximum likelihood estimation of the theoretical parameters. We report in Table 6.1 the location parameter,  $\mu$  and the scale parameter,  $b$ , estimated from the fit, as well as the goodness of the fit, i.e., the log-likelihood.

Manufacturer growth rates					
	$\mu$	$b$	Log-likelihood	D (KS)	p-value
Gauss	0.13 (0.08)	2.34 (0.05)	-2026	0.15	2.2e-16
Laplace	0.05 (0.03)	1.36 (0.04)	-1789	0.09	7.1e-08

Distributor growth rates					
	$\mu$	$b$	Log-likelihood	D (KS)	p-value
Gauss	-0.01 (0.02)	1.23 (0.02)	-4643	0.20	2.2e-16
Laplace	0.03 (0.01)	0.62 (0.01)	-3759	0.11	2.2e-16

**Table 6.1:** On the right: Optimal parameter values ( $\mu$  and  $b$ ) resulted from the Gaussian and Laplacian fit the empirical growth rates. On the left: distances and  $p$ -values from the KS-test compare the similarity between the theoretical and empirical distributions.

From a visual inspection, we see that the empirical values do not align with the theoretical ones. This is true for both data on manufacturers and distributors, albeit it is particularly evident in the case of distributors. The statistical test (KS-test) confirms this observation. Table 6.1 shows the distances and the  $p$ -values obtained from the KS-test. Given the low  $p$ -values ( $\ll 0.01$ ), we reject the hypothesis of similarity between the empirical distribution and the Laplacian one for both manufacturers and distributors data. In Table 6.1, we also show the results from comparing the empirical distribution with a Gaussian one fitted to the data. As for the Laplacian distribution, we reject the hypothesis of similarity between the empirical and the theoretical distribution. We obtain even worse results with the Gaussian fit than the Laplacian fit.

We conclude that the functional forms previously proposed in the literature (Laplacian and Gaussian) are not valid for the growth rates observed in our data. We see from Fig. 6.1 that the major mismatch between the empirical growth rates and those predicted by the Laplacian model is on the tails of the probability distribution. The empirical distribution shows fatter tails than those predicted by the model. Correlations between growth rates may be at the origin of such fat tails (Bottazzi *et al.*, 2001). As mentioned above, the outflows of firms in distribution networks are tightly coupled, possibly explaining the large values of growth rates observed. However, identifying the

correct functional form of the observed growth rate distribution is beyond this study's scope and is left for future research.

## 6.3 Assessment of the determining factors

### 6.3.1 An econometric approach

This Section investigates the determinants of growth rates of firms in distribution networks. We consider firm-specific characteristics such as age, size, and the number of partners. Further, we also consider the firm position in the network. Specifically, the tier the firm belongs to, defined in Chapter 3 and expressed by Eq. (3.1), identifies its position.

We take an econometric approach and use a multivariate model in which the dependent variable is the firm growth rate, and the independent variables are the determinants. In addition, we include control variables to ensure that the model does not capture variations due to external factors, such as a general increase in shipments in a given year. Then, we perform an ordinary least squares (OLS) regression to assess the statistical significance of the selected independent variables and the sign of their correlation with the growth rates. The independent and the control variables are explained in detail below.

**Age and size.** The role of age and size as determinants of firm growth has been vastly investigated in previous studies (Becchetti and Trovato, 2002; Wynarczyk and Watson, 2005; Zhou and de Wit, 2009). Younger firms show higher growth rates. Smaller firms are more likely to grow since they have to achieve a minimum size efficiency (Zhou and de Wit, 2009). Building on previous studies, we include these factors as independent variables in our analysis. Firm size is measured according to Eq. (7.1). The age,  $a_{i,y}$ , is given by the number of years since the first appearance of the firm:  $a_{i,y} = y - y_i$  where  $y_i$  is the year of the first appearance of  $i$  in the data. Note that we only have nine observation years (from 2006 and 2014). Therefore, the range of age values thus obtained is rather narrow. Section 6.3.2 discusses how this data limitation impacts our results.



**Number of partners.** Based on previous studies, we test the hypothesis that the number of partners plays a role in firm growth (Wynarczyk and Watson, 2005). In light of these studies, we expect that the higher the number of partners a firm has, the higher its growth rate. In particular, we consider separately the number of source partners, i.e., partners from whom the firm receives goods, and the number of target partners, i.e., partners to whom the firm ships (i.e., clients). Note that clients include both end consumers and other firms in the network.

**Tier.** To our knowledge, no previous empirical studies have shown how a firm's position in a distribution network can impact its growth rate. In this study, we aim to test this hypothesis and determine the sign of the effect using a regression model. The effect can be either positive or negative, with a positive effect meaning that growth rates increase moving downstream and a negative effect meaning that growth rates decrease moving downstream. For instance, a positive effect may arise given that firms positioned downstream, i.e., retailers, tend to have smaller sizes than those positioned upstream, e.g., wholesales. And it has been observed that small firms have higher growth rates than big ones (Zhou and de Wit, 2009). Conversely, a negative effect may occur in the presence of demand changes. In this case, last-tier distributors may be subjected to lower demand fluctuations, given their proximity to consumers, and hence to lower growth rates than first-tier distributors. We will incorporate the tier value as a determinant of firm growth in the regression model to assess which effect is more prevalent.

**Controls: year, state.** In Chapter 3, we discussed the evolution of the opioid demand across years and US states. We observed considerable changes in the opioid demand over the years. Further, we also highlighted considerable differences across states. These variations and differences can be reflected in the amount shipped by the firms in the network. We do not want to capture such variations with our model. We rather want to explain the growth rates of firms independently of their specific location or the specific year of observation. Therefore, we use two control variables: one for the year and one for the state where the firm is located. Moreover, our proxy of firm age depends strictly on

	<b>Model A</b> Manufacturers	<b>Model A</b> Distributors	<b>Model B</b> Distributors
(Intercept)	<b>2.70</b> (0.49) <sup>***</sup>	0.08 (0.37)	0.58 (0.38)
age	0.09 (0.09)	-0.06 (0.03) <sup>*</sup>	-0.06 (0.03) <sup>*</sup>
log(size)	<b>-0.23</b> (0.02) <sup>***</sup>	<b>-0.07</b> (0.01) <sup>***</sup>	<b>-0.10</b> (0.01) <sup>***</sup>
N clients	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
N sources		<b>0.01</b> (0.00) <sup>***</sup>	<b>0.01</b> (0.00) <sup>**</sup>
tier 1			<b>-0.33</b> (0.06) <sup>***</sup>
tier 2			<b>-1.09</b> (0.31) <sup>***</sup>
tier 3			-1.05 (0.66)
Controls: year, state			
R <sup>2</sup>	0.17	0.09	0.10
Num. obs.	730	2776	2776

\*\*\* $p < 0.001$ ; \*\* $p < 0.01$ ; \* $p < 0.05$

**Table 6.2:** Results from the OLS regressions. The coefficients with  $p$ -values smaller than 0.01 are reported in bold character.

the year it is calculated. Trivially, in 2007 all firms may have an age of zero or one (i.e., their first year is 2006). Therefore, we include the interaction effect between  $a_{i,y}$  and  $y$ .

### 6.3.2 Results from the OLS regression

To build our panel dataset, we consider all shipments of a single opioid. As a case example, we use Hydrocodone. The unit of analysis is the growth rate of a single firm. A single firm appears in the panel dataset (at maximum) 8 times, i.e., once per year. Thus, we analyze two models: *Model A* employs the following independent variables: age, size, number of source and target partners; *Model B* employs the same variables of *Model A* plus the tier. As in the Subsection above, we evaluate the growth rates of manufacturers and distrib-

utors separately. In Table 6.2, we show the results from the OLS regressions.

We start discussing the results obtained from the *Model A* and *Model B* run on the growth rates of distributors. We find that the factor age has a negative coefficient in both *Model A* and *Model B*. This confirms our expectation that younger distributors are more likely to have higher growth rates. However, the effect is not strong ( $p$ -value  $< 0.05$ ). As mentioned above, the validity of our proxy of firm age is limited by the lack of data. We found a very skewed distribution of age values, where old distributors (i.e., the ones appearing over the whole observation period) are the majority and few of them have younger ages ( $< 9$  years). Therefore, the low significance of the age may depend on the small range of values obtained from the data.

As expected, we find that the size of a distributor is significantly and negatively correlated with its growth rate: small distributors are more likely to exhibit higher growth rates. Further, we found that the number of source partners of a distributor has a positive and significant effect on its growth rate. Particularly, having one more source partner produces a 10% increase in its annual ship-out.

In contrast, we note that the number of clients does not play an important role. Some distributors, particularly those at the last tiers, ship to many consumers (about a thousand). Other distributors ship to other distributors in the network (about a dozen). We argue that our linear model does not capture these large differences in the number of clients, making the independent variable not significant. Further studies could test this argument by including, for example, an interaction effect between tier and number of clients or by considering a possible re-scaling of the variable.

Furthermore, we observe that the tier variable exhibits a significant negative correlation with firm growth rates. This finding implies that growth rates tend to decrease as we move downstream in the network. Specifically, second-tier distributors have lower growth rates than first-tier distributors, and third-tier distributors have lower growth rates than second-tier distributors. As we discussed in Chapter 3, Hydrocodone consumption increased during the period of observation. Our results suggest that the observed demand changes mostly af-

affected the first-tier and second-tier distributors, leading them to exhibit higher growth rates compared to the final tiers.

The results from the analysis conducted on the growth rates of the manufacturers reveal that most determinants considered, i.e., age and number of target partners, are not significant in explaining growth rates. We observe a significant effect only for the size. This, as expected, is negatively correlated to manufacturer growth. The non-significance of the other variables is related to data issues. Most manufacturers appear in the data from the beginning of our observation period. Therefore, they are few diversified in age, making our proxy of age unsuitable for this model. Further, similarly to distributors, manufacturers are very diversified in the number of clients. Some manufacturers ship directly to consumers, thus having a much higher number of clients than others that ship via distributors. For the latter, the number of clients is almost two orders smaller than the former. As for distributors, we argue that our linear model does not capture these big differences, making this variable not significant.

## 6.4 Discussion

In this Chapter, we empirically investigated the growth rates of firms in a distribution network. For this aim, we used the ARCOS dataset and analyzed the size and the growth rate of opioid manufacturing and distribution firms in the US. Our goal was to test previous theories against these novel data to deepen our understanding of firm growth in the context of distribution networks.

We found that the stylized fact of a Laplacian shape for the growth rate distribution does not hold for the firms in our data. We argued that the discrepancy between the theoretical prediction and the empirical observation lies in the tails of the growth rate distribution. These may arise from the strong correlations between the growth rates of firms in distribution networks.

Next, we forwarded our analysis and defined possible determinants of firm growth rates. Building on previous studies, we considered the firm age, size, and number of partners as predictors in our model. In addition, we considered

the firm position, as expressed by the tier the firm belongs to. We assessed their statistical significance with an OLS regression.

Our results revealed that the distributors that are: (i) smaller in size, (ii) younger, (iii) and with more source partners are more likely to exhibit higher growth rates. These results aligned with our expectations and previous findings obtained on other firm data (Zhou and de Wit, 2009). In addition, we found that first-tier distributors are more likely to have higher growth rates than second-tier distributors, and second-tier distributors exhibit higher growth rates than third-tier distributors. On the contrary, we did not find any significant effect for the age and number of clients in the manufacturers' data. The latter result may depend on the method we used to compute manufacturer ages and the considerable differences in the number of clients manufacturers ship to.

It is fair to say that the validity of the above results is limited by the low predictive power of our model:  $R^2 \approx 0.1$ . Several reasons may explain this observation. First, the OLS regression may not be well suited to investigate growth determinants within this context. The linear model may not capture the big differences in the growth rates of firms in our data. Choosing a type of regression that is not linear can be the object of future studies.

Further, we could not include other import determinants of firm growth due to data limitations. For example, the financial capital of a firm and its marketing or management strategy could be relevant to predict growth rates improving the actual model.

Finally, our model has a central critical assumption: the independency of the data records. In other words, we assumed that firm growth rates in the distribution network are independent. However, they are not. Distributors can increase their ship-out only if their source partners have done the same. These correlations would limit any regression-type analysis. To properly overcome this limitation and better investigate firm growth in distribution networks, we propose a network-based approach in the next Chapter.

## Chapter 7

# Modeling network growth: topology and goods flow

### Summary

This Chapter expands on exploring the factors influencing network and firm growth. Building on the analysis presented in Chapter 5, we introduce a network growth model that accounts for the co-evolution of network topology and goods flow. Considering inter-firm dependencies, we overcome the limitations of the regression analysis presented in Chapter 6. Specifically, our model employs a proportionate growth to explain the evolution of flows, which is subject to demand fluctuations and upstream fluctuations due to supply availability. To test and validate our model, we apply it to the opioid distribution system. Our findings suggest that upstream fluctuations significantly impact growth dynamics and are crucial for reproducing the observed patterns of firm growth.

---

This Chapter has been written specifically for this thesis. Results are based on A. Amico, G. Vaccario, F. Schweitzer “The growth of business firms in distribution networks”, Working Paper. AA contributed to designing the research question and deriving the interaction rules. AA wrote the code, performed the simulations, and interpreted the results.

## 7.1 Introduction

In Chapter 5, we proposed a growth model to describe the formation of distribution systems. Validating the model against real-world data, we gained new insights into the complex architecture of such systems and the mechanisms driving their evolution. Nevertheless, our understanding was limited to the topological aspect of this evolutionary process, namely how new firms enter the network and how new links are formed between them. Distribution networks grow both in size by increasing the number of firms and links, *and* in volume, by increasing the amount of goods flowing through each link.

To comprehensively understand the mechanisms driving the growth of distribution networks, this Chapter aims to derive the microscopic rules describing the evolution of goods' flows. Specifically, we want to explain how firms' inflows and outflows evolve over the years and how the topological inter-firm dependencies affect such process. In accounting for these dependencies, we overcome the limitations of the econometric approach presented in Chapter 6.

The Chapter is structured as follows. We start by evaluating the statistical properties of the goods flows in an empirical system (Section 7.2). Guided by the empirical insights, we derive the microscopic rules for the evolution of goods flow (Section 7.3). Next, we propose an analytic approach to calibrate the model against real-world data (Section 7.3.2). Feeding the model with the optimal parameters obtained from the calibration, we test the model against various stylized facts of the empirical system, e.g., growth rate distribution and growth rate volatility (Section 7.4). We close the Chapter by summarizing the main findings and outlining the limitations of the present study (Section 7.5).

## 7.2 Empirical evidences

As a case study, we analyze the distribution network of the top-selling opioid manufacturer (Mallinckrodt), reconstructed from the ARCOS data. By this, we use the same data of the analysis carried out in Chapter 5, as well as the same procedure to reconstruct the network. Nodes represent firms and consumers,

and links represent supply relations observed in a given year. Every link has a weight,  $w$ , to account for the annual amount shipped, measured in grams.

In Fig. 7.1, we show the probability distribution of link weights in the year 2008 (blue circles) along with the log-normal fit performed on the data (red dashed line). For the fit, we use the method proposed by Clauset *et al.* (2009), which performs better on probability distributions that have heavy tails\*.

The distribution exhibits a rather pronounced heavy tailed character, indicative of a great heterogeneity in link weights. Large weights (about a million grams) characterize very few links, while most links have small weights (about a thousand grams). The observed heterogeneity should be traced back to the different roles firms play in a distribution network. The largest weights represent the outflows of large distributors, e.g., warehouse distributors. Smaller weights, instead, represent the outflows of smaller distributors, e.g., retail distributors. Furthermore, we see that a log-normal functional form describes our data well. Log-normal distributions may originate from generative processes based on proportional growth dynamics (Schweitzer, 2020). This empirical hint will be used in the formulation of the model in the Section 7.3.

Next, we study the growth rates of firms. As explained in Chapter 6, we determine the growth rate  $g_i$  of firm  $i$  as the logarithm of its 1-year size variation, expressed by Eq. (6.2). Moreover, we proxy firm size  $s_i$  by its total ship-out (see Chapter 6). Thus, we have:

$$s_i(y) = \sum_{j \in \text{nb}_i} w_{ij}(y) \quad (7.1)$$


where  $\text{nb}_i$  identifies the set of neighbors to whom  $i$  ships,  $w_{ij}$  indicates the amount shipped through each of them, and  $y$  is the year of analysis.

As the ARCOS dataset spans over nine years, i.e., from 2006 to 2014, we could potentially compute the annual growth rates at eight different years. Then we could study the probability distribution of the pooled values to have a

---

\*Due to their large fluctuations on the tails, power-law distributions are difficult to detect by standard least-squared methods. The method proposed by Clauset *et al.* (2009) combines a maximum-likelihood approach and a KS-based fitting method, which is more robust on the tails. The method is available within the python package called “power-law” (Alstott *et al.*, 2014).





**Figure 7.1:** Left side: probability distribution of link weights in the year 2008 (blue circles), and log-normal fit (red dashed line). Right side: probability distribution of growth rates, pooling the years 2007, 2008, and 2009, and Laplacian fit (blue dashed line).

$\mathbf{P(g)}$	<b>KS-test p-values</b>			
	2007	2008	2009	Pool
2007	1.000	0.961	0.163	0.893
2008		1.000	0.169	0.892
2009			1.000	0.458

**Table 7.1:**  $p$ -values obtained from the KS-test that compare the statistical similarity of  $P(g)$  in different years. Not all the years of observation appear in this Table, but only the ones for which a  $p$ -values  $> 0.01$  has been obtained.

more comprehensive statistic. However, we find that only some growth rates computed at different years are statistically compatible, according to the KS statistics. Therefore, we restrict our analyses to those years we find a statistical agreement. These are the years 2007, 2008, and 2009. The results from the KS-test on these years are reported in Table 7.1.

In Fig. 7.1, we show the empirical distribution of growth rates together with the Laplacian fit performed on the data (dashed line). This distribution exhibits a symmetric exponential character that, plotted in log-log, appears as a “tent-shape”. This result is similar to the one of Chapter 5, obtained under aggregation of multiple products manufactured by different companies. Similar to what we found in Chapter 5, the Laplacian functional form does not

characterize our data. Despite a quite good visual agreement between data (circles) and fit (dashed line), the KS-test returns a  $p$ -value of  $2.91 \times 10^{-04}$ . Hence, we reject the null hypothesis of statistical similarity between the real data and the theoretical predictions.

As in Chapter 6, we can explain this result by noticing that the major mismatch between the two distributions occurs on the tails. The empirical tails are “fatter” than the ones predicted by the Laplacian model. Pronounced tails may result from dependencies between the growth of firms (Bottazzi *et al.*, 2011). We keep in mind this empirical insight, which, together with the log-normality of link weights, forms the basis for the formulation of our model.

## 7.3 Model

### 7.3.1 Demand and supply fluctuations

We consider a distribution network comprising  $N$  nodes, i.e., a single manufacturer, multiple distributors and multiple consumers. These are connected through direct links used to ship goods. Links have weights to characterize the amount of goods shipped. We denote with  $W_i^{\text{out}}$  the total outflow of node  $i$  and  $W_i^{\text{in}}$  its total inflow, i.e.:

$$W_i^{\text{out}}(t) = \sum_{j \in \text{nb}_{\text{out}}} w_{ij}(t); \quad W_i^{\text{in}}(t) = \sum_{j \in \text{nb}_{\text{in}}} w_{ji}(t) \quad (7.2)$$

where  $\text{nb}_{\text{out}}$  is the set of neighbors to whom  $i$  ships, and  $\text{nb}_{\text{in}}$  is the set of neighbors from whom  $i$  receives. The differences  $\Delta W_i(t) = W_i^{\text{out}}(t) - W_i^{\text{in}}(t)$  are captured by the inventory.

At time zero, one manufacturer and one distributor populate the distribution network. Over time, the network evolves: new nodes enter the network, and new links are created. This process follows the rules for link formation derived in Section 5.2.2 and expressed by Eq. (5.1) and Eq. (5.2).

While the topology is evolving, outflows and inflows change. In our model, nodes have no control over their inflows, only over their outflows. Specifically, we make two assumptions: (i) nodes ship-out proportionate to their last ship-

ments; (ii) nodes ship-out conditioned to their supply availability. Specifically, nodes experiencing shortages, i.e.,  $\Delta W_i \leq 0$ , reduce their ship-out; nodes experiencing surplus, i.e.,  $\Delta W_i > 0$ , increase their ship-out. We name the latter condition as the *upstream constraint*.

Translating the above principles into equations, we write the link weights dynamic as:

$$\frac{w_{ij}(t+1)}{w_{ij}(t)} = \begin{cases} \frac{1}{k_i} e^\varepsilon & \Delta W_i(t) \leq 0 \\ k_i e^\varepsilon & \Delta W_i(t) > 0 \end{cases} \quad (7.3)$$

where  $w_{ij}$  is the outflow of  $i$  towards  $j$ ,  $k_i \geq 1$  is a rescaling factor, and  $\varepsilon$  is a random Gaussian noise with finite mean and variance,  $\varepsilon \sim \mathcal{N}(\mu, \sigma^2)$ . The dynamic expressed by Eq. (7.3) combines two processes.

On the one hand, outflows grow according to a stochastic proportionate dynamic, where the stochastic term captures demand fluctuations. This builds on the traditional law of proportionate effect, introduced by Gibrat in 1931 (Fu *et al.*, 2005; Riccaboni and Schiavo, 2010; Schweitzer, 2020). Note that, given a variable  $w$ , the original expression of the law has the form:

$$w(t+1) = w(t)\eta(t) \quad (7.4)$$

where  $\eta$  is the stochastic growth factor. We choose  $\eta$  to be the *exponent* of a Gaussian noise, i.e.,  $\eta = e^\varepsilon$ . By this, we ensure that positive values are obtained for the outflows. Moreover, the proportionate growth dynamic implies that the quantity under study is log-normally distributed. This is in line to what we observe for the link weights in our data (Fig. 7.1).

On the other hand, outflows are subject to the upstream constraint whose effect is controlled by the sign of  $\Delta W_i$  and whose strength is controlled by the rescaling factor  $k_i$ . This accounts for dependencies between firms: firm growths are constrained by a given finite quantity, i.e., the total supply (Amaral *et al.*, 1998; Bottazzi *et al.*, 2011). Note that, in principle,  $k_i$  is firm-dependent, as it relates to the firm internal organization. It may also be time-dependent: product life-cycle leads firms to experience periodic phases of growth or recession. To keep our model simple, we assume  $k_i$  to be homogenous across firms and constant over time,  $k_i = k$ .

Finally, we consider a logistic growth for the production dynamic at the manufacturer, i.e.,:

$$\Omega(t+1) = \left[ 1 + r \left( 1 - \frac{\Omega(t)}{K} \right) \right] \Omega(t) \quad (7.5)$$

where  $\Omega(t)$  is the production at time  $t$ ,  $r$  is the growth parameter, and  $K$  is the production limiting value. In other words, according to Eq. (7.5), we assume that the production grows exponentially at a rate  $r$ , saturating afterwards at the limiting value  $K$ . While it is clear that a logistic dynamic can not capture various realistic aspects of production growth (e.g. possible fluctuations), it is also worth noticing the main advantage: its simplicity allows us to determine the model parameters analytically.

### 7.3.2 Parameter calibration: an analytic solution

**Mean and variance of demand fluctuations.** The model introduced above is characterized by two free parameters:  $\mu$  and  $\sigma$ . Here we derive an analytic solution to estimate their values.

Let's consider the case  $k = 1$ . The growth dynamic in Eq. (7.3) reduces to the simple proportionate growth as:

$$w(t+1) = w(t)e^\varepsilon \quad (7.6)$$

The stochastic process expressed by Eq. (7.6) is a geometric Brownian motion with a close-form solution. Specifically, the solution  $w(t)$  is log-normally distributed with a finite mean given by:

$$\mathbb{E}(w(\tau)) = w(0)e^{(\mu + \frac{\sigma^2}{2})\tau} \quad (7.7)$$

and a finite variance given by:

$$\text{Var}(w(\tau)) = w(0)^2 e^{2\mu\tau} \left( e^{\sigma^2\tau} - 1 \right) \quad (7.8)$$

where  $\tau$  is the final time, assuming that the initial time is zero.

Our model assumes that each time a new link is created. Once created, the link remains in the network for  $\tau - t$  time steps, where  $t$  is the time when

the link is created, and  $\tau$  is when the simulation ends. Therefore, considering  $t' = \tau - t$  the time of the link weight evolution, we obtain the mean value of the link weights by averaging over all times, i.e., :

$$\langle w(\tau) \rangle_{\text{sim}} = \frac{1}{\tau} \sum_{t'=0}^{\tau} w(t') \quad (7.9)$$

Next, we approximate the variable  $w(t')$  in Eq. (7.9) to its expected value, i.e.,  $w(t') \approx \mathbb{E}(w(t'))$ . Thus, we substitute the expression of  $\mathbb{E}(w(t'))$  from Eq. (7.7) into Eq. (7.9) and obtain the following expression for the mean value of the link weights:

$$\langle w(\tau) \rangle_{\text{sim}} \approx \frac{1}{\tau} \sum_{t'=0}^{\tau} \mathbb{E}(w(t')) = \frac{w_0}{\tau} \sum_{t'=0}^{\tau} e^{(\mu + \frac{\sigma^2}{2})t'} \quad (7.10)$$

The right term in Eq. (7.10) is a geometric series. Substituting its closed form into Eq. (7.10), we arrive to the the final expression of the first moment (the mean value):

$$\langle w(\tau) \rangle_{\text{sim}} \approx \frac{w_0}{\tau} \frac{\left(1 - e^{(\mu + \frac{\sigma^2}{2})\tau+1}\right)}{\left(1 - e^{(\mu + \frac{\sigma^2}{2})}\right)} \quad (7.11)$$

Following the same approach, we derive the expression for the second moment  $\langle w^2(\tau) \rangle_{\text{sim}}$  and we have:

$$\langle w^2(\tau) \rangle_{\text{sim}} \approx \frac{w_0^2}{\tau} \frac{\left(1 - e^{(2\mu + \sigma^2)\tau+1}\right)}{\left(1 - e^{(2\mu + \sigma^2)}\right)} \quad (7.12)$$

The coupled Eq. (7.11) and Eq. (7.12) can be solved numerically after fixing the first and the second moments to their empirical values:  $\langle w(\tau) \rangle_{\text{sim}} = \langle w \rangle_{\text{emp}}$ ; and  $\langle w^2(\tau) \rangle_{\text{sim}} = \langle w^2 \rangle_{\text{emp}}$ . By this, we obtain the optimal values for the parameters  $\mu$  ad  $\sigma$ . These are:  $\mu = 0.011$ , and  $\sigma = 0.049$  for the network in the year 2008.

**Parameters of the production dynamic.** The production dynamic in Eq. (7.5) is characterized by two parameters:  $K$ , the limiting production value, and  $r$ , the growth rate. We set  $K$  to the maximum value observed in the data, i.e.:

$$K = \max_{\{y\}}[\Omega(y)] \quad (7.13)$$

where  $\{y\}$  is the set of years under analysis, and  $\Omega(y)$  is the empirical production. Hence, we obtain:  $K = \Omega(2009) \approx 60$  million grams.

Then, we determine the production growth rate  $r$  by imposing that the simulated production reaches the empirical value within the period of the simulations,  $\tau$ . Given the logistic dynamic, the simulated production at time  $\tau$  is:


$$\Omega(\tau) = \frac{K}{1 + \frac{K - \Omega_0}{\Omega_0} e^{-r\tau}} \quad (7.14)$$

Thus, enforcing the condition  $\Omega(\tau) = \Omega(y)$ , and inverting (7.14), we obtain the expression for  $r$ , i.e.,:

$$r = -\frac{1}{\tau} \log \left[ \left( \frac{K}{\Omega(y)} - 1 \right) \left( \frac{\Omega_0}{K - \Omega_0} \right) \right] \quad (7.15)$$

where  $\Omega_0$  is the initial production. We obtain  $r = 0.013$ .

**Simulation time shift.** The proposed model does not ensure a direct match between simulation time steps and empirical times. One time step of the simulation is the time needed for one link to be created, and it does not correspond to one year (empirical time). As we are interested in reproducing the annual growth rate, we need to determine the simulation time shift,  $\nu$ , corresponding to the one-year shift. We determine  $\nu$  such that the simulated production at time  $\tau - \nu$  matches the empirical one at the  $y - 1$ :  $\Omega(\tau - \nu) = \Omega(y - 1)$ . By this, we are fixing the production to its empirical value in both the year  $y$  and the  $y - 1$ . We obtain  $\nu = 49$ .



**Figure 7.2:** Distance (y-axis) between the empirical distribution of growth rates and the one obtained from computer simulations for different  $k$  values (x-axis) and assessed using the KS-test.  $k \in [1, 4]$  with an interval of 0.25.

## 7.4 Model validation

### 7.4.1 Role of the upstream constraint

In this Section, we aim to test the model against the real-world data. To this aim, we feed the model with the optimal parameter values. The parameters  $q_s$ ,  $q_t$ , and  $\alpha$  (that control for the topological growth) are calibrated according to the procedure discussed in Section 5.3; the values of the parameters  $\mu$ ,  $\sigma$ ,  $r$  and  $K$  are obtained from the analytic solutions derived in the Section above.

Then, to determine the optimal  $k$  (upstream constraint), we consider a set of  $k$  values, ranging from 1 to 4, with an interval of 0.25. We run the model for each  $k$  value and measure the distance between the empirical distribution of growth rates and one obtained through computer simulations. The optimal  $k$  is the one that minimizes such distance. The statistical distances are evaluated through the KS-test and reported in Fig. 7.2. From the Figure, we see that a set of  $k$  values, ranging from 1.25 to roughly 2, could potentially suit the data. We select the minimum observed:  $k^{\text{opt}} = 1.50$ .


In Fig. 7.3, we show the distribution of growth rates obtained from the model simulations (red diamonds) with  $k = 1$ , that means without upstream constraint, in the left plot; and with  $k = k^{\text{opt}}$ , that means with the optimal upstream constraint, in the right plot. The blue circles indicate the empirical

Model	mean	std	skewness	kurtosis	KS	p-value
Emp	0.05	1.41	0.20	10.28		
Sim $k=k^{\text{opt}}$	0.18	0.88	2.00	19.03	0.11	$0.77 \times 10^{-5}$
Sim $k=1$	0.57	0.21	0.14	0.34	0.59	$\ll 10^{-10}$

**Table 7.2:** On the left side: mean, standard deviation, skewness, and kurtosis of three distributions of growth rates: empirical (first row), simulated with  $k = k^{\text{opt}}$  (second row); simulated with  $k = 1$  (third row). On the right side: distances and  $p$ -values obtained from the KS-test between empirical distribution and simulated ones.

distribution. For a better comparison, in Table 7.2, we report their statistical properties (i.e., mean, standard deviation, skewness, and kurtosis).

We see that the model with  $k=k^{\text{opt}}$  can reproduce the characteristic tent-shape of the empirical distribution together with its broad variance (as indicated by the standard deviation) and pronounced tails (as noted in the kurtosis value). Instead, the model with  $k=1$  does not perform well: it produces a narrower distribution with low variance and no tails. This finding suggests that fluctuations generated by the upstream constraint are needed to reproduce the broadness and the tails of the observed distribution.



**Figure 7.3:** Empirical (blue circles) and simulated (red diamonds) distribution of growth rates. Left side: results from the model fed with  $k=1$ , i.e., no upstream constraint. Right side: results from the model fed with  $k = k^{\text{opt}} = 1.50$ , namely with the upstream constraint at its optimal value.




Albeit we observe a good match between real-world and simulated data with  $k=k^{\text{opt}}$ , we need to reject the hypothesis of statistical similarity of the two distributions. The KS-test returns a  $p$ -value  $\approx 10^{-5}$ . This result may originate from the mismatch between the quite asymmetric character of the simulated distribution (skewness=2) and the very symmetric empirical one (skewness=0.2). The higher and positive value of the skewness of the simulated distribution indicates that the model leans towards positive values, underestimating the number of negative growth rates. This may be due to the fact that, in our model, we are considering only the new entry of firms and the formation of new links. We are not accounting for the opposite process, i.e., the exit of firms or cutting links, which may cause negative growth rates. In this thesis, we don't have time to explore the combination of the entry and exit process. We leave this exploration to future studies.

#### 7.4.2 Volatility of growth rates

Since we used the empirical growth rates to calibrate and validate the model, we perform additional tests to evaluate the model performance. To this end, we consider the volatility of firm growth rates. In line with previous studies (Calvino *et al.*, 2018), we measure volatility as the standard deviation of growth rates,  $\sigma_g$ .

In Fig. 7.5, we plot  $\sigma_g$  ( $y$ -axis) as a function of firm size ( $x$ -axis). We investigate the correlation between these two quantities in both real data (left plot) and simulated data (right plot). From the empirical data (left plot), we see that growth rate volatility and firm size are linked by a negative relationship. This implies that the smaller the firm, the more volatile its growth is. To put it simply, small distributors experience greater fluctuations in their growth than large ones. This finding is in agreement with previous studies conducted on manufacturing firms (Calvino *et al.*, 2018; Riccaboni and Schiavo, 2010; Sutton, 2002), where a similar volatility-size relation has been found. It is also in agreement with the regression analysis conducted in Chapter 6.

From the visual inspection of the simulated data, we notice that the model can reproduce such negative relation. To better quantify the match between



**Figure 7.4:** Growth rate as function of firm size for empirical data (left plot) and simulated data (right plot). The dashed line indicates the linear fit as expressed by Eq. (7.16) and performed via an ordinary least squares method.


simulated and real-world data, we fit the data points assuming a power-law functional form (linear in log-log scale), i.e.:

$$\log(\sigma_g) = \alpha + \beta \log(s) \quad (7.16)$$

Using an ordinary least squares method, we fit the model expressed by Eq. (7.16) and estimate the slope of the curve  $\beta$ . We obtain  $\beta = -0.42$  for the real-world data and  $\beta = -0.12$  for simulated ones.

The obtained values indicate that the model is good at reproducing the negative nature of the observed volatility-size relation. But it has limitations in reproducing its strength. The model predicts a less strong relation between size and growth volatility. Specifically, we notice that the model predicts less volatile growth rates for small firms ( $\log(s) \leq 6$ ) and more volatile growth rates for big firms ( $\log(s) \geq 12$ ). In other words, the model does not capture the full heterogeneity of the data. The observed mismatch could be related to the choice of a homogeneous value for the model parameter  $k$ . Since this parameter controls the strength of the fluctuations in the outflows, the choice of a heterogeneous  $k$  could produce more diverse growth volatility. The testing of this hypothesis is left to future studies.

As the last step in our validation procedure, we analyze the relationship between the volatility of growth rates and the firm's position in the network. In



**Figure 7.5:** Growth rate of firms as a function of their tier in empirical data (blue dots) and synthetic data (red dots). The error bar in each box indicates the standard deviation of the sample representing the growth rate volatility of each tier.

particular, we consider the tier the firm belongs to. As expressed by Eq. (3.1), tier is measured as the length of the shortest path connecting the manufacturer to the distributor.

In Fig. 7.5, we use a scatter plot to show firm growth rates (y-axis) as function of their tier (x-axis). In addition, we use box blots to highlight the spread of the sample and mark its standard deviation  $\sigma_g$  (volatility per tier) through the error bar. We have two side-by-side box plots: the blue color indicates real-world data, the red color indicates the simulated data.

We find that volatility increases with the tier as indicated by the increase of the amplitude of the error bars. Specifically, third-tier distributors exhibits a higher spread than second-tier and first-tier distributors. In short, downstream distributors experience more volatile growths than upstream ones. Interestingly, these findings are obtained in real-world data (blue dots) and simulated data (red dots), suggesting a good performance of the model in reproducing the empirical feature under study.

It's worth commenting on a few mismatches that we still observe. The model underestimates the spread of data points at tier 3 and overestimates the spread of data points at tier 2. As a result, the model predicts a smaller difference in volatility of firms at tier 2 and 3, producing growth rates more uniform

across samples. As mentioned above, the model does not capture the full heterogeneity of the data. In addition, the model predicts a higher number of level firms than is actually observed. This is actually a mismatch related to topological growth, which we have already commented on in the Chapter 5. In the next Section, we outline directions for future improvements.

## 7.5 Discussion

In summary, this Chapter extended the network growth model presented in Chapter 5 to account for the growth of firms' inflows and outflows. Our goal was to delineate a few, albeit sufficient, principles to describe such process and test their validity against real-world data.

Driven by empirical evidence, such as the log-normality of firm outflows and the tent-shaped distribution of their growth rates, we proposed a proportionate growth dynamic to describe the evolution of goods flows. Specifically, the proposed dynamic comprises two main ingredients: (i) a stochastic growth motivated by the demand stochasticity and (ii) an upstream constraint inspired by the reasoning that a limited amount of resources is available and can be distributed. Both ingredients generate fluctuations in the firm outflows, thus affecting their growth.

We found that the main contribution to extreme values of the growth rates comes from the upstream constraint. The fluctuations originating upstream rather than downstream (i.e., from the demand) may produce a high jump in firm growth. Using the ARCOS dataset, we tested the predictive quality of the model against several features of an empirical distribution system. These include the distribution of firm growth rates, their volatility, and their relation with firm size and firm tier. The validation revealed that the model reproduces the analyzed features. However, it has two main limitations.

First, we had to reject the hypothesis of similarity between the empirical distribution of growth rates and the one obtained from the simulations. We argued that this mismatch originated on the left tails of the distribution, i.e., on the negative values of the growth rates. The model predicts less extreme negative values than the ones observed. Indeed, we explained that the model only

---

accounts for the entry process of firms and links. It does not account for the opposite process, namely the exit of the firms or the link deletions. For time constraints, we leave this investigation to future studies.

Second, our analysis of the volatility of growth rates revealed that the model underestimates the differences between firms characterized by different sizes and positions in the network. We pointed out that this mismatch could be due to a (too) simplistic choice of the model parameter  $k$  and could be improved by incorporating further heterogeneity. For instance, a non-symmetric character of  $k$  could be considered, producing more substantial fluctuations in case of shortages compared to the surplus phase. Future studies could consider the above research directions to advance our understanding of business growth.

## Part III

# Cascade dynamics and network responses

“Everything should be made as simple as possible, but not simpler.”

*A. Einstein*

## Chapter 8

# Upstream and downstream shock propagation

### Summary

This Chapter explores the cascade effects produced by localized supply-side shocks in distribution networks. We propose a cascade model that considers a bidirectional proportion of the shock, with firms rationing their shipment downstream and increasing orders upstream. Using the empirical distribution network of opioids as a test case, the results highlight the nontrivial dependencies among opioid producers due to upstream propagation. Under the assumption of perfect substitution, these dependencies can help mitigate supply shortages. However, in the long run, indirect losses occur due to faster depletion of other producers' inventories, adding to the direct losses.

---

This Chapter has been written specifically for this thesis. AA contributed to the development of the research question. AA designed the model principles, wrote the code and interpreted the results.

## 8.1 Introduction

This Chapter investigates the amplification of local disruptions in supply networks. The term disruption usually refers to an unexpected event, or shock, that interrupts the normal flow of goods (Craighead *et al.*, 2007). These shocks may have various origins, including, for instance, natural disasters (e.g., hurricanes, earthquakes, floods), socioeconomic factors, labour issues, or terrorist attacks (Bode and Wagner, 2015; Li and Zobel, 2020). An initial shock may directly damage only a few firms in a supply network. Yet, it may have ripple effects and *indirectly* damage several other firms (Craighead *et al.*, 2007; Zhao *et al.*, 2010). In other words, the shock's impact can be amplified due to the interconnectedness of the economic activities in the network (Carvalho *et al.*, 2021). Various terminologies have been used in the academic literature to identify this amplification effect: disruption propagation (Scheibe and Blackhurst, 2018); ripple effect (Dolgui *et al.*, 2018); risk diffusion (Basole and Bellamy, 2014); and cascading failures (Hearnshaw and Wilson, 2013). In the following, we use the term “cascade”.

Recently, cascades have gained attention from academics and practitioners in the supply chain domain. A series of real events has shown that (even small) localized shocks can have devastating consequences also for firms not directly affected by the shock, e.g., lower revenues, delivery delays, loss of market share and reputation (Craighead *et al.*, 2007; Dolgui *et al.*, 2018; Scheibe and Blackhurst, 2018).

Examples are the earthquake in Japan in 2011 and the hurricane Katrina in 2005, which disrupted the performance of several supply networks causing significant economic losses (Acemoglu *et al.*, 2015; Jabbarzadeh *et al.*, 2016). Besides disruptions due to natural hazards, accidents related to single firms can be mentioned. For example, in 2000, a fire accident hit one supplier of the telecommunications company Ericsson. The local accident brought about 400 million dollars in loss to the company as it could not rely on alternative suppliers (Latour, 2001). A similar situation occurred at Toyota when a fire hit its brake valve supplier in 1997. This caused a two-week shutdown of several Toyota plants and a loss of 195 million dollars (Tomlin, 2006).



The recent COVID-19 pandemic adds to a long list of events that have strained several supply networks over the past two decades, highlighting their fragility. Above all, the pandemic has shown that disruptions harm not only companies but also people directly (Diem *et al.*, 2022). Many nations suffered shortages of primary supplies (e.g., food and medical equipment) due to quarantine policies, which profoundly affected the well-being of millions of people.

While the vulnerability of supply networks to cascade effects is widely recognized, there remains a lack of investigation into the actual propagation dynamics within real-world systems. Moreover, existing studies have focused primarily on small-scale networks involving a single focal firm.

This Chapter addresses this gap by providing a model for the inter-firm propagation of supply shocks in large-scale distribution systems. Our objective is to discuss this process in a general context, drawing insights from the specific case study of the opioid distribution system. Through this analysis, we aim to evaluate and discuss the dependencies among firms that arise from cascade effects and to assess the possible indirect losses generated.

We describe the model principles and introduce the mathematical formalism in Section 8.2. Next, we use a stress-test approach and quantify the cascading impact for the Oxycodone distribution network in Section 8.3. In Section 8.4, we conclude the Chapter by outlining the limitations of the model and directions for possible improvements.

## 8.2 Modeling cascades in distribution systems

### 8.2.1 Related works

Cascade processes have been widely studied in the field of network science with applications to various economic networks, including financial networks (Elliott *et al.*, 2014), food trade networks (Burkholz and Schweitzer, 2019), and international trade networks (Kali and Reyes, 2010). After the 2008 financial crisis, scholars began to develop cascade models to explain how systemic failure can be generated by even small, localized shocks (Battiston *et al.*, 2012a,b; Lorenz *et al.*, 2009; Roukny *et al.*, 2013). These models usually consider only

one propagation direction. Instead, in supply networks, shocks may propagate downstream, from suppliers to consumers, and upstream, from consumers to suppliers (Diem *et al.*, 2022; Inoue and Todo, 2019). The downstream propagation would affect the shipments' dynamic, whereas the upstream propagation would affect the orders' dynamic. The two dynamics, in principle, follow different rules. Therefore, the previously proposed models are not directly applicable.

General equilibrium models have been developed to study shock propagation in production networks (Carvalho *et al.*, 2021). They describe how the output of one production unit affects the output of another production unit in the network. Since firms in distribution networks only ship finished goods, these models do not suit our case. In addition, they neglect inventories. Inventories are an essential tool for distributors to mitigate shortages.


A common approach in the supply chain field to model shock propagation is the system thinking/dynamics approach (Ghadge *et al.*, 2022; Ivanov, 2017; Sterman, 2010). This approach considers a limited number of “representative” agents, such as a single manufacturer, first-tier and second-tier suppliers, a distribution center, and a retailer. Then, feedback effects between these agents are evaluated using diagramming tools, such as causal loop diagrams and stock-flow diagrams (Dolgui *et al.*, 2018). In contrast, agent-based models (ABM) emphasize interactions between agents as well as their heterogeneity (Ding *et al.*, 2018). As we want to capture the impact of the observed network structure (how firms are connected via supply links) on the cascade dynamics, we employ an ABM approach. This is described in the next Section.

### 8.2.2 Model: underlying principles

We consider the situation in which a shocked manufacturer reduces its shipments to its distributors, who, in case of insufficient inventories, must shrink their shipments too. This triggers a cascade of supply shortages which may lead to unmet consumer demand. While this cascade is happening from upstream to downstream, a second dynamic happens from downstream to upstream: distributors with insufficient inventories increase their orders to manufacturers. Receiving more orders, these manufacturers may bring to exhaustion

their inventories and need to shrink their shippings too, thus triggering another cascade moving from upstream to downstream. A schematic representation of the proposed process is drawn in Fig. 8.1.

To model cascading effects in distribution systems and account for the role of inventories, we draw on previous ABM models that incorporate inventory dynamics (Hallegatte, 2014). The model proposed by Hallegatte has been used to analyze the cascade effect and the economic impact of major shocks resulting from, e.g., the Japanese earthquake (Inoue and Todo, 2019) and the hurricane Katrina (Hallegatte, 2008). However, unlike the work presented in (Hallegatte, 2008; Inoue and Todo, 2019), we do not consider the production side of the supply network, but only the distribution side. This results in significant differences in cascading dynamics.



**Figure 8.1:** Schematic representation of the upstream and downstream shock propagation in a distribution network. The bars on the side of the nodes represent indicative stock levels. Entirely blue bars represent inventories that are at their target level. Blue-red bars represent inventories that are affected by shortages. Entirely red bars represent empty inventors. Firms experiencing shortages (i.e., affected by the shock) are colored in red, as well as the links through which the shock propagates. The illustrative example considers two manufacturers,  $A$  and  $B$ , and three distributors,  $C$ ,  $D$  and  $E$ , in the network. At a given time  $t^*$ , a shock hits the inventory of manufacturer  $A$ . A supply shortage propagates downstream (red line), affecting distributor  $C$ . As a result, distributor  $C$  increases its orders and triggers an upstream propagation (red line). This upstream propagation leads the inventory of  $B$  to a faster depletion, thus generating a second cascade that moves downstream to consumers.

While firms in production networks receive intermediate goods as inputs, firms in distribution networks only receive finished goods. Therefore, in production networks, firms experiencing a shortage tend to reduce orders to their suppliers, as they do not require additional supplies if their production is interrupted or slowed down. In distribution networks, however, firms facing a shortage tend to increase orders to their suppliers (manufacturers and distributors upstream) to replenish their stocks. As a result, our work extends previous studies by incorporating an upstream propagation mechanism specifically suited to distribution systems.

### 8.2.3 Model: system dynamics

Consider a distribution system composed of  $N$  firms, i.e., manufacturers and distributors. We represent it as a network, where nodes indicate firms connected through directed links used to ship goods. The direction of the links follows the direction of the shipments. Each link has a weight indicating the volume shipped at a given time  $t$ . The weighted links are represented as elements of the shipping matrix  $W$ . The shipping vector  $\omega$  captures shipments from firms to consumers. For instance, the element  $W_{ij}$  indicates the amount shipped by firm  $i$  to firm  $j$ , whereas the element  $\omega_i$  indicates the amount shipped by firm  $i$  to consumers.

Every firm  $i$  holds an inventory or stock,  $s_i$ , where goods can be stored. This is updated according to its total ship-out,  $W_i^{\text{out}}$ , and ship-in,  $W_i^{\text{in}}$ , as below:

$$s_i(t) = s_i(t-1) + W_i^{\text{in}}(t-1) - [W_i^{\text{out}}(t-1) + \omega_i(t-1)] \quad (8.1)$$

Distributors use inventories to store goods received from manufacturers or upstream distributors, whereas manufacturers use inventories to store their own production. We assume that the production phase takes place on a much larger time-scale than the distribution phase. Specifically, we consider a *daily* time-scale for the distribution dynamics and an annual time-scale for the production dynamics. As a result, over the one-year analysis period, production is fixed. Instead, shipments and inventories are updated at each time of the simulation.

Following the principles outlined in the ARIO-inventory model, firms place orders to (i) meet demand and (ii) avoid empty inventories by keeping them at a constant target level,  $s^T$ , or safety buffer. Safety buffers are used to better manage supply shortages and sudden fluctuations in demand and reduce lead times (products are already in stock).

Translating the above principles into an equation, we describe the order placed by firm  $i$  as:

$$o_i(t) = d_i(t-1) + \frac{1}{\tau} [s_i^T - s_i(t)] \quad (8.2)$$

where  $d_i$  is the demand faced by  $i$  and  $s_i$  is its inventory level. The first term represents the orders needed to meet the received demand, whereas the second represents the orders needed to make the inventory converge towards its target level. The parameter  $\tau$  indicates how quickly the firm wants to restore its inventories. To keep our model simple, we consider  $\tau$  homogenous across firms and constant over time. In our study, we set it equal to one working week, i.e.,  $\tau = 5$  days.

The total demand  $d_i$  in Eq. (8.2) takes into account two terms: demand received from (i) consumers and (ii) demand from other firms. Following Hallegatte (2014), we model the two terms separately. We define  $c$  as the vector capturing consumer demand and  $O$  as the order matrix capturing orders placed by other firms. For example, the element  $c_i$  indicates the demand that firm  $i$  receives from consumers, while the element  $O_{ji}$  indicates the demand that firm  $i$  receives from firm  $j$ . In other words,  $O_{ji}$  indicates the order placed by firm  $j$  towards firm  $i$ . Hence, the total demand faced by  $i$  is equal to the following:

$$d_i(t) = \sum_j O_{ji}(t) + c_i \quad (8.3)$$

Note that  $c_i$  has no time dependence. We assume constant daily demand to avoid adding further complexity to the model. Future studies may explore more realistic settings for demand dynamics. The elements of the order matrix,  $O$ , in Eq. (8.3) are given as:

$$O_{ij}(t) = o_i(t)T_{ij} \quad (8.4)$$

where the transition matrix  $T$  captures firms' preferences towards their neighbors. Specifically, the element  $T_{ij}$  expresses the proportion of goods  $i$  orders to  $j$ . More explicitly, it is the number of units  $i$  orders from  $j$  normalized by the total number of units  $i$  orders to all its upstream neighbors.

In the absence of shortages, firms meet their demand and ship the requested quantity. Therefore, at each time  $t$ , shipments are equal to the orders placed at the previous time, that is:

$$W(t) = O^\top(t-1) \quad (8.5)$$

where  $O^\top$  is the transpose of the order matrix  $O$ . Similarly, at each time  $t$ , shipments to consumers are equal to the orders placed by them, that is:

$$\omega(t) = c(t-1) = c \quad (8.6)$$

where  $\omega$  and  $c$  are two vectors whose length is equal to the number of firms in the network.

#### 8.2.4 Model: system response to shocks

**Supply-side shock.** Suppose, now, that an external shock hits the distribution network at day  $t^*$ , and the manufacturer  $i$  is directly damaged. As discussed in the Introduction, the types of shocks affecting real-world supply networks are manifold, including natural hazards, human-made attacks, or internal labor issues. Modeling the direct damage of each type of shock can be challenging and goes beyond this thesis's scope. Since we are interested in the *indirect* damages due to the cascade, we make a simple assumption about the size of the direct shock. Specifically, we simulate a  $\sigma$  percentage reduction in the manufacturer's inventory as:

$$s_i(t=t^*) = (1-\sigma)s_i(t-1) \quad (8.7)$$

where  $\sigma$  is the size of the shock and  $s_i$  denotes the manufacturer's inventory level (used to store its production). Note that no recovery of production is allowed. We argue that this would require a longer period than the one

considered. However, the generality of the proposed model allows for improvements in this direction if necessary. Since production cannot be increased, the shock triggers a supply shortage that propagates downstream and, in turn, upstream through the network.

**Downstream propagation: rationing scheme.** When shortages occur, stocks may not be sufficient to meet demand, and firms ship what is left in their inventories. In other words, they try to meet demand as much as possible. Thus, the actual total ship-out of a firm  $i$ , at any given time, is limited by:

$$\Omega_i^{\text{act}}(t) = \min\{s_i(t), d_i(t-1)\} \quad (8.8)$$

where the term  $\Omega_i^{\text{act}}$  includes shipments directed towards firms and consumers, i.e.,:

$$\Omega_i^{\text{act}}(t) = W_i^{\text{out,act}}(t) + \omega_i^{\text{act}}(t) \quad (8.9)$$

When stocks are lower than demand, firms ration their shipments. Rationing is done in proportion to orders received. The rationing scheme first involves the division of shipments between consumers and firms as:

$$\omega_i^{\text{act}}(t) = \Omega_i^{\text{act}}(t) \frac{c_i}{d_i(t-1)}; \quad W_i^{\text{out,act}}(t) = \Omega_i^{\text{act}}(t) \frac{\sum_j O_{ji}(t-1)}{d_i(t-1)} \quad (8.10)$$

and then between firms, as:

$$W_{ij}(t) = W_i^{\text{out,act}}(t) \frac{O_{ji}(t-1)}{\sum_j O_{ji}(t-1)} \quad (8.11)$$

Note that the case  $\Omega_i^{\text{act}}(t) = d_i(t-1)$  and  $W_i^{\text{out,act}}(t) = \sum_j O_{ji}(t-1)$  represents the situation when no shortage is observed. In this case, Eq (8.11) reduces to Eq. (8.5) that expresses the system dynamics in a no-shock scenario.

**Upstream propagation: increase of orders.** When shortages occur, firms with insufficient inventories ration their shipping downstream and increase orders upstream. Firms increase orders according to the dynamics expressed by Eq. (8.2) and Eq. (8.4).

Assuming that firms keep their original preferences towards their neighbors, the transition matrix in Eq. (8.2) is given by:

$$T_{ij} = \frac{O_{ij}(0)}{\sum_j O_{ij}(0)} \quad (8.12)$$

This assumption is valid as long as one wants to evaluate the “short-term” effects of the cascade, that is, within a given year. In the long run, the network can adapt to cope with the shortage by reconfiguring itself. Firms may establish new relationships or modify their preferences with existing partners. Additionally, manufacturers can invest in new production to recover from shocks. As our model focuses on an annual time-frame, it does not consider the above mentioned aspects. The only immediate response mechanism for firms is to request substitute products from their source partners. The study of other possible mechanisms is left for future research.

### 8.2.5 Measuring the cascade impact

In our analysis, we evaluate the impact of the cascade by considering two aspects. On the one hand, we assess the effects of the cascade on manufacturers not directly harmed by the shock. On the other hand, we assess the total supply that can not be distributed, producing unmet demand among consumers. Besides this, one might be interested in assessing the costs that shortages impose on distributors. As they adjust operations to replenish inventories and mitigate shortages, they may face costs due to business changes. We are not interested in this aspect in the current Chapter. We will propose a measure for this in Chapter 9.

Because of the upstream propagation of the cascade, manufacturers not directly affected by the shock may also be harmed. In particular, they may experience an increase in orders from distributors, which leads to a faster depletion of their inventories. This, in turn, can cause management issues as production schedules need to be adjusted. To track this indirect impact, at each step of



the simulation, we measure the reduction in manufacturers' inventories,  $\delta(t)$ , in a shock scenario compared to a no-shock one, i.e. :

$$\delta_i(t) = \tilde{s}_i(t) - s_i(t) \quad (8.13)$$

where the subscript  $i$  indicates the manufacturer not directly shocked,  $\tilde{s}_i$  is its inventory level in a shock scenario, and  $s_i$  is its inventory level in a no-shock one.

In addition, we estimate the “indirect” supply loss produced by the cascade. By indirect loss, we mean the total amount of supply that could not be distributed because of the indirect damage. Thus, we are not interested in the loss generated directly by the shock. Instead, we are interested in the loss resulting from the shock propagation. To this end, we define the change in supply,  $\Delta$ , as the difference between the supply distributed when the system suffers the shock and when it does not, i.e.,:

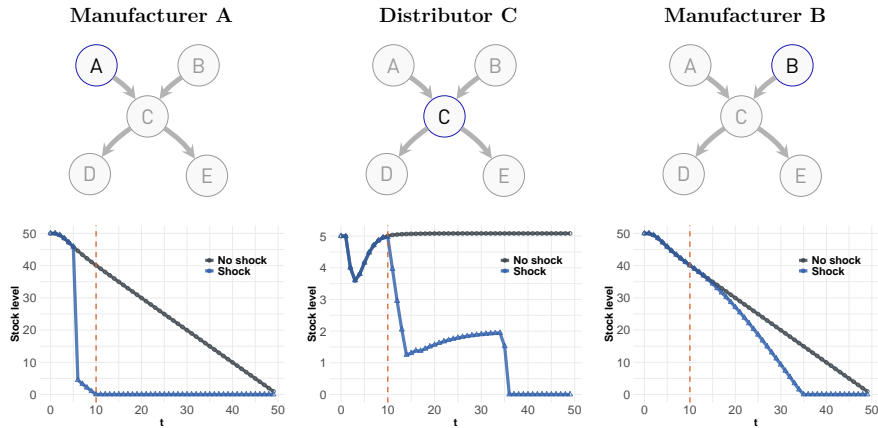
$$\Delta(t) = \sum_{\substack{i \in \{m_1, m_2, \dots\} \\ i \neq \tilde{i}}} \left( \tilde{W}_i^{\text{out}}(t) - W_i^{\text{out}}(t) \right) \quad (8.14)$$

where the sum runs over all the manufacturers in the network, except for the one directly shocked,  $\tilde{i}$ .

Positive  $\Delta$  values indicate that the supply shortage has been compensated through substitution. Manufacturers ship more than they used to in a no-shock scenario because the goods they produce are used to offset the shortage. On the other hand, negative  $\Delta$  values signal that the upstream propagation has triggered additional shortages. Manufacturers ship less than they used to in a no-shock scenario because they run out of stock sooner. In the latter case, the indirect supply loss arises, which adds to the direct one.

Thus, we determine the indirect supply loss,  $\lambda$ , as the total negative supply variation observed, i.e.:

$$\lambda(t) = \sum_{t'=t^*}^t |\Delta(t')| \times \Theta[-\Delta(t')] \quad (8.15)$$



**Figure 8.2:** Inventory dynamics for manufacturer *A*, distributor *C*, and manufacturer *B*. The network and the shock scenario are explained in Fig. 8.1. The shock hits the manufacturer *A*, propagates downstream, affecting the distributor *C*, and upstream, affecting the manufacturer *B*. In the plots on the right-hand side, the blue lines indicate the inventory dynamic in a shock scenario and the black lines in a no-shock scenario. The dashed red lines mark when the manufacturer *A* goes into stock-out in the shock scenario.

where  $\Theta[\cdot]$  denotes the Heaviside function, and  $t$  is the time of the simulation.

## 8.3 Results

### 8.3.1 Illustrative example for system dynamics

Before discussing the real-world application, we first illustrate our approach using a small synthetic network. We use the network depicted in Fig. 8.1 as an example case and run the model in the absence and the presence of the shock.

For simplicity, we consider target stocks and manufacturers' productions to be homogenous across distributors and manufacturers, respectively. Specifically, we set  $s_A(0) = s_B(0) = 50$  units and  $s^T = 5$  for the three distributors. Then, consumers order one unit at a constant rate, and the distributor *C* orders from the two manufacturers in equal proportions. Fig. 8.2 shows the inventory dynamics in the absence and presence of the shock.

As governed by Eq. (8.2), distributors increase their orders to meet the demand and replenish their inventories. The replenishment speed is controlled by the parameter  $\tau$  and represented by the positive slope of the curve in Fig. 8.2. The smaller the  $\tau$ , the faster the inventories get back to their target levels and the steeper the slope is. In the absence of external shocks, distributors replenish their inventories already in a few time steps, as  $\tau = 5$ . Thus, the inventories are stable to their target values, and the system reaches its “equilibrium”: the amount shipped equals the amount demanded.

Let us now consider the following shock scenario. At the time  $t^* = 6$ , a shock hits the manufacturer  $A$ , leading to a 90% reduction of its stocks. This produces a supply shortage downstream, affecting distributor  $C$  and further down distributors  $D$  and  $E$ , and finally, consumers. Thus, distributor  $C$  increases its orders to manufacturer  $B$ , triggering the upstream propagation of the shock. As depicted in Fig. 8.2, manufacturer  $B$  is harmed as it depletes its stock faster than in a no-shock scenario.

Hence, the external shock that directly affects manufacturer  $A$  also affects manufacturer  $B$  indirectly due to the upstream propagation. In other words, the two manufacturers are interdependent due to the cascade effects.

### 8.3.2 Firms' dependencies under cascade effects

We now turn our attention to a real-world application. We use the nationwide distribution network of opioids as a test-bed for our analysis\*. We adopt a stress-test approach: we start from an observed network and model its distortion due to a simulated shock.

Within this approach, we want to start with a realistic representation of the empirical system. Therefore, we incorporate as much empirical information as possible into the model. These include (i) the supply dependencies among distributors or manufacturers in the network; (ii) the consumer demand; (iii) the distributors' target stocks; and the (iv) annual productions.

---

\*The drugs that contain Oxycodone as a basic opioid are used for this analysis.

First, we reconstruct the empirical distribution network following the procedure described in Subsection 2.2.1. In this network, nodes represent firms, and links represent supply relations observed in a given year.

Next, we initialize the consumer demand, expressed by the vector  $c$ , to its empirical value. Note that we do not have data about demand. We only have data about shipments. Therefore, we reconstruct the daily demand from the observed daily shipments. Specifically, the demand is obtained by averaging daily shipments over the year. In other words, it is assumed that firms have met demand in a given year and that consumers have received the quantity ordered. The same reasoning applies to the orders placed by the firms. Therefore, we initialize the elements in the order matrix to the corresponding empirical values. Then, from Eq. (8.4), we reconstruct the transition matrix,  $T$ , expressing the distributors' preferences towards their neighbors.

Finally, we set the target stocks to the empirical buffers recorded at the end of the year. This holds under the assumption that distributors are left with the stock buffers they planned<sup>†</sup>. Inventories are then initialized to their target levels. Then, we proxy manufacturers' production as the annual total ship-out.

From the data, we find 56 entities registered as manufacturers. However, since most of them are small laboratories producing very few product units, we restrict our analysis to the top 20 manufacturers. In Table 8.1, we report their name, geographical location, and total production for the year 2008.

We run multiple independent computer simulations to evaluate how an external shock that affects a single manufacturer can indirectly harm other manufacturers in the network. In each simulation, a single manufacturer is shocked. Then, the shock propagates through the network according to the rules discussed in Subsection 8.2.4. We evaluate the shock's impact on the other manufacturers (not directly shocked) using the  $\delta$  indicator defined by Eq. (8.13).

We run a total of 20 independent simulations, one for each shock-affected manufacturer. The severity of the shock is  $\sigma=0.99$ , representing a 99% reduction in

---

<sup>†</sup>There are cases where distributors ship more than they receive in a year. This suggests that older stocks, i.e., those from previous years, have been used. It also suggests little need for these distributors to replenish their stocks. Therefore, without further information, we set a minimum stock target for these distributors.

Name	Short	Production	City	dea-num
MALLINCKRODT	MA	16,719,342	HOBART	RM0231821
PURDUE PHARM.	PU	5,643,717	WILSON	RP0257938
NOVARTIS	NO	3,888,791	LINCOLN	PD0038667
WATSON LAB.	WA	3,697,629	CARMEL	RD0118150
ACTAVIS	AC	1,179,084	TOTOWA	RA0306490
BOEHRINGER.	BO	929,319	COLUMBUS	RR0112514
SHARP CORP.	SH	529,071	ALLENTOWN	RS0308317
DSM PHARM.	DS	323,072	GREENVILLE	RC0240692
VINTAGE PHARM.	VI	209,722	HUNTSVILLE	RV0359299
BARR LAB.	BA	138,079	FOREST	RB0234005
GENERIC	GE	132,245	CHARLOTTE	RG0360115
THE PF LAB.	TH	122,349	TOTOWA	PT0164587
METRICS	ME	113,839	GREENVILLE	RM0282638
RX OF MCKESSON	RX	102,618	MEMPHIS	RR0276837
AMNEAL PHARM.	AM	56,261	BROOKHAVEN	RA0370712
INTERPHARM	IN	45,845	BROOKHAVEN	RI0347179
GENUS LIF.	GE	17,972	ALLENTOWN	RL0287385
MIKART	MI	12,213	ATLANTA	RM0197497
IMPAX LAB.	IM	12,015	PHILADELPHIA	RI0259300
JANSSEN	JA	10,384	GURABO	RJ0255453

**Table 8.1:** Name, production (product units), city, and *dea-number* of the top 20 opioid (Oxycodone) producers in 2008.

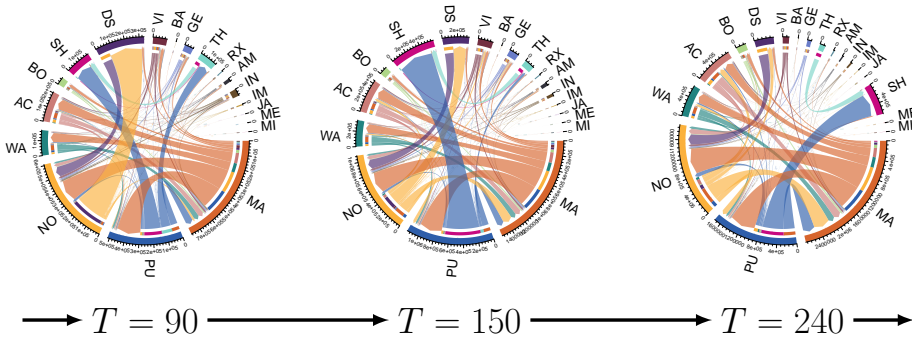
total production. We stop each simulation after 360-time steps, corresponding roughly to one year of analysis. It is not informative to study the system later, because a supply deficit occurs anyhow as annual stocks deplete.

In Fig. 8.3, we show the results from the simulations using chord diagrams. In these diagrams a colored fragment, or “sector”, placed on the outer part of the circular layout, represents a manufacturer. Specifically, it represents the manufacturer that is shocked during the simulation. There are, in total, 20 colored sectors, one per manufacturer (i.e., one per simulation). An arc between two sectors represents the cascade impact transmitted by a shocked manufacturer,  $j$ , to a second manufacturer,  $i$ , in the network. The direction of the arc follows the direction of the cascade: from the shocked manufacturer to the manufacturer indirectly impacted. The color of the arc matches the color assigned to the shocked manufacturer. The arc width is proportional to the

cascade impact  $\delta_i$ . The numerical values of the cascade impact are reported on the upper side of each sector.

Note that the diagrams do not display the direct impact of the shock. They only show the indirect one. The three diagrams display the results obtained at three different time steps of the simulations:  $T = 90$ , i.e., roughly 3 months after the shock;  $T = 150$ , i.e., roughly 5 months after the shock; and  $T = 240$ , i.e., roughly 8 months after the shock.

Note also that the 20 manufacturers are not directly connected in the network via supply links, i.e.,  $W_{ij} = 0 \quad \forall i, j \in \{m_1, m_2, \dots\}$ . Therefore the dependencies shown in Fig. 8.3 are obtainable only as a result of the model simulations.




**Figure 8.3:** Chord diagrams showing the dependencies between opioid producers under cascading supply shocks. The three diagrams display the results obtained at three different time steps of the simulations, i.e.,  $T = 90$ ,  $T = 150$ ,  $T = 240$ .

First we see that, a few manufacturers, when shocked, lead to large cascades. The arcs departing from manufacturer MA (in orange), from manufacturer PU (in blue), and from manufacturer NO (in yellow) are much wider than all the others. From Table 8.1, we see that the three manufacturers dominate the opioid market, as their annual production is much greater than that of the others (over 3 million units). Simply put: the biggest producers trigger the largest cascades.

Further, we notice that the diagrams appear rather sparse; there are few predominant connections between pairs of manufacturers. This observation suggests that cascades originating at one manufacturer are mainly transmitted to a few others. In other words, some manufacturers are much more closely interdependent than others under cascade effects. In fact, the cascade propagates in a non-linear manner. In our model, the propagation dynamics depend on two key elements: the network topology and the preferences of distributors toward their neighbors, as expressed by Eq. (8.4).

Comparing the three diagrams, we observe a few changes in the cascade pattern over time. For example, the cascade triggered by the shock at the manufacturer NO (yellow arc) is mainly transmitted to manufacturer DS, at time  $T = 90$ . At later times, this pattern changes, as MA is also quite exposed to the cascade. Similarly, at time  $T = 90$ , both manufacturers SH and TH are exposed to the shock affecting PU (blue arc). However, at time  $T = 150$ , SH became considerably more exposed than TH. In short, we observe a few variations of the cascade patterns depending on when the cascade is evaluated. These variations can be traced back to the stock availability of the manufacturers. If manufacturers run out of stock, distributors increase orders to alternative ones, thus redirecting the cascade.

As a final, albeit important, remark we note that the interpretation of cascade depends strictly on the reader's perspective. On the one hand, cascades can harm other firms in the network, not directly damaged by the shock. They expose other manufacturers to the shock as their inventories are put under stress and experience a faster depletion. This, in turn, can generate management issues because the production schedules need to be adjusted. In the long run, stock-outs can occur, resulting in multiple disadvantages, including loss of revenue or customer dissatisfaction. On the other hand, from a systemic perspective, dependencies among different producers can be used to leverage additional resources and mitigate supply shortages to consumers. Thus, the revealed dependencies can be good or bad depending on the point of view.



**Figure 8.4:** Evolution of the variation in supply (black solid line) during the simulation. On the left y-axis we show its absolute value and on the right y-axis we show its value normalized to the average direct loss. The blue-filled area indicates the quantity shipped to offset the shortage. The red-filled area indicates the supply loss due to the cascade.

### 8.3.3 Indirect loss estimation

In the Subsection above, we have revealed strong dependencies between opioid producers due to the cascade effects. Here we address the questions: to what extent do these dependencies help mitigate supply shortages? And to what extent can they trigger further losses?

To address the question, we simulate a shock (magnitude  $\sigma = 0.99$ ) hitting the inventory of the top producer in the network and run the cascade model prosed above. During the simulation we monitor the supply variation, measured according to Eq. (8.14). In Fig. 8.4 we show its evolution (black solid line). We report its absolute value (left y-axis) and the value normalized to the average direct loss (right y-axis), i.e., the loss resulting from the 99% reduction in the top producer's stock.

We see that during the first part of the simulation,  $\Delta$  is positive. This means that the direct loss is partially mitigated (about 25%) through supply substitution. Manufacturers who are not directly harmed provide substitute goods to compensate for the shortage. The total amount of substituted goods is represented by the blue-filled area in Fig. 8.4. Note that not all manufacturers contribute to substitutions, but only those connected to the shocked



manufacturer through the dependencies highlighted by the chord diagrams in Fig. 8.3.

Interestingly,  $\Delta$  goes from positive to negative at a certain point. Some manufacturers go into stock-out faster than in a no-shock situation. This causes an indirect supply loss to consumers. The total indirect loss,  $\lambda(t=T)$  in Eq. (8.15), computed at the end of the simulation<sup>‡</sup>, is represented by the red-filled area in Fig. 8.4. We see that the indirect loss became comparable to the direct one, as indicated by the percentage values on the right y-axis. Therefore, if no other mitigation strategies, e.g., increased production, are adopted during the period under consideration, indirect losses are generated, adding to the direct ones.

## 8.4 Discussion

In summary, our goal was to evaluate the cascade effects produced by localized shocks in real-world distribution networks. To this aim, we used the nationwide opioid distribution network as a testbed for our analysis. We do not observe big supply shocks in this data. Small shocks, instead, are challenging to be detected due to a lack of information. We only have access to the empirical shipments, not to the orders placed. Hence, small reductions in supply may be associated with small supply shortages or small negative fluctuations in demand. Distinguishing the two causes is a data challenge beyond this thesis's scope. In fact, we are interested in those shocks that are large enough to produce a significant effect on the entire network, not small fluctuations in the flow.

In the absence of empirical evidence, we proposed a cascade model to describe shock propagation in distribution networks. Hence, the magnitude of the shock, its location, and the time of its occurrence are all part of our modeling assumptions. Building on previous studies (Hallegatte, 2014; Inoue and Todo, 2019), we considered a bidirectional propagation of the shock: an upstream propagation that affects the shipping's dynamic and a downstream propagation that affects the orders' dynamic.

---

<sup>‡</sup> $T$  = marks the final time

Thus, we use the empirical network as starting point and model its distortion due to a simulated shock. We assumed that the observed supply relations are held fixed during the cascade. Consequently, we considered the change in orders placed and shipments as the only response of the network to the shock. A possible reconfiguration of the existing links is not allowed.

Our results revealed nontrivial dependencies among producers in the empirical network due to cascade effects. Under the assumption of perfect substitution, these dependencies can be leveraged to mitigate supply shortages. However, in the long run, indirect losses are generated due to the faster depletion of the inventories of other producers. From a certain point in time, the indirect losses start adding up to the direct ones. When planning redundancies in the system, managers and policymakers should consider not only the individual exposure of firms to possible shocks but also the cascade effects they produce. Our approach can help to understand where, given the network topology, more resources could be allocated to mitigate these effects.

Nevertheless, our approach is limited by some critical assumptions we made. First, our model runs on a daily time-scale where the demand is constant. Demand seasonality effects may speed up or slow down the cascade and underestimate or overestimate the indirect loss computed at a given time. Further studies could incorporate more realistic settings for the demand distribution into the model.

Second, price does not play any role in the proposed model. However, supply and demand are linked through price. Supply shortages may increase the price of the goods, reducing their demand. While this is true for most goods on the market, we argue that price may play a minor role in the case of pharmaceutical products. In this case, the demand elasticity is small since the demand does not change much as the price change, because drugs are essential goods. Further, the prices of opioid drugs are under the control of the government and its institutions, e.g., the DEA.

Third, the proposed model is valid for evaluating short-term effects. The network response may differ for long-term effects, i.e., longer than one year). New business contracts may be formed, and new supply links created. As a result, a different network topology can be observed. Since we limited our

---

analysis to an annual time-window, we believe that this assumption is still reasonable.

Finally, our model works under the assumption of *perfect substitution*. Substitution of goods supplied by different manufacturers implies the willingness of (i) distributors to change their usual business; (ii) and consumers to accept substitute products, which in some cases, especially for pharmaceuticals, may not be an option. More information needs to be incorporated into the model to relax this assumption. In the next Chapter we will discuss how to address this problem.

## Chapter 9

# Flexibility: a pillar of supply chain resilience

### Summary

This Chapter examines how supply substitution can serve as a response mechanism to supply-side shocks. We refer to the system's capacity to adapt to shocks via supply substitution as "flexibility," which we formalize as a system parameter using higher-order network models. We show that flexibility can indeed alleviate shortages by reducing the demand deficit. However, it also introduces costs due to flow readjustments after substitution. As a result of this trade-off, we establish an efficient frontier that policymakers can use to select the optimal flexibility value, balancing deficit reduction and system costs.

---

AA wrote this Chapter specifically for this thesis. It is based on the research findings presented in A. Amico, L. Verginer, G. Casiraghi, G. Vaccario, F. Schweitzer, "Adapting to Disruptions: Flexibility as a Pillar of Supply Chain Resilience", *arXiv:2304.05290*. AA contributed to designing the research question, to developing the model, performing the simulations, and interpreting the results.

## 9.1 Introduction

In the previous Chapter, we presented a cascade model to describe how supply shortages may propagate in large-scale distribution networks. Perfect supply substitution was a central assumption of the proposed model. Specifically, we assumed that (i) distributors were willing to reroute packages, and (ii) consumers were willing to accept goods produced by alternative suppliers (i.e., manufacturers). In a more realistic setting, this assumption may not hold. Distributors are interested in making a few changes to their operations to keep costs low; and consumers have preferences towards their suppliers that they may maintain during shortages. Also, in some application cases perfect substitution is limited by suitable conditions, e.g., the absence of allergens in substitute goods compared to usual ones. Drugs and food fall into this category.

We define *flexibility* as the degree of supply substitution the distribution system is willing, and capable, to accept during supply shortages. In a fully flexible system, firms respond to shortages by rerouting substitutable goods. In a non-flexible one, firms keep their business as usual.

This Chapter investigates the role of flexibility in mitigating disruptions and enhancing network resilience. Generally, resilience is defined as a system's ability to withstand and recover from shocks (Hollnagel *et al.*, 2006; Schweitzer, 2022b; Schweitzer *et al.*, 2021). It combines two fundamental aspects: *robustness* against shocks themselves and *adaptivity* to overcome the states that result from a shock (Helfgott, 2018).

Proactive and reactive mitigation strategies improve supply network resilience by focusing on robustness and adaptivity, respectively (Aldrighetti *et al.*, 2021). While proactive approaches build robustness by planning redundancies in the system (e.g., increasing inventory buffers, devising just-in-case production capacity) before the event occurs; reactive approaches aim at adjusting the operational processes when the shock hits the system. Our study contributes to the understanding of the latter approaches.

We view flexibility as the primary contributor to the adaptivity of a distribution network. Flexibility represents the ability of the system to adapt to

the shock by leveraging alternative (substitutable) resources already in the system. Thus, in the short run, manufacturers are not required to ramp up production, nor are distributors required to create new distribution channels.

To model flexibility we extend the method presented in the previous Chapter. Leveraging the higher-order network representation, we model the readjustment of flows after substitution. We formalize system flexibility as a model parameter,  $\phi$ , controlling for the degree of flow readjustment. We use the opioid distribution network as test bed for our analysis. Tuning  $\phi$ , we investigate and discuss the impact of flexibility: (i) on the consumers, measuring the demand deficit and (ii) on the distribution network, quantifying changes to its usual operations. We conclude the Chapter with an outlook for future works.

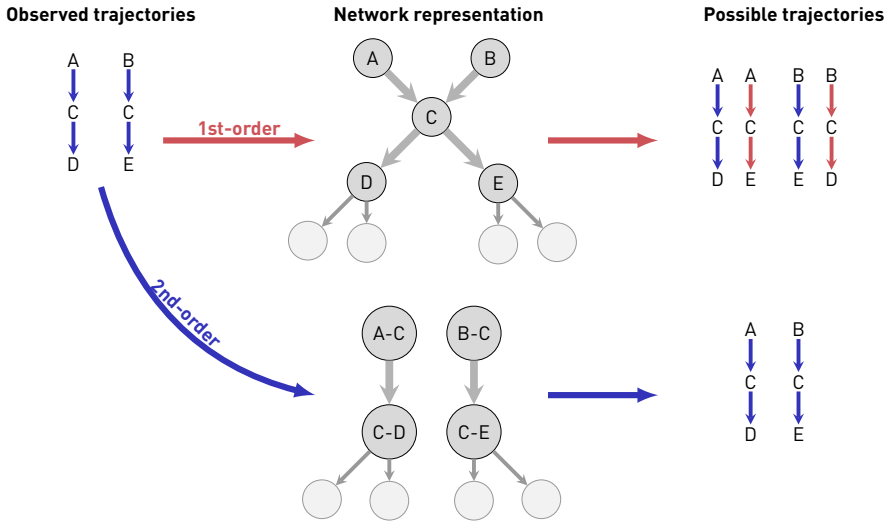
## 9.2 Modeling flexibility through higher-order networks

### 9.2.1 Capturing the empirical flow

Flexibility represents the system's ability to leverage alternative resources by readjusting the usual flow observed in day-to-day business operations. In a data-driven approach, the usual flow corresponds to the empirical flow observed in the system under study (assuming no empirical shock). To model the empirical flow, we move from a standard network representation to a higher-order network representation (Scholtes *et al.*, 2016). Let us explain the difference between the two representations with an illustrative example.

Suppose we have a distribution network comprising five firms: two manufacturers,  $A$  and  $B$ , and three distributors,  $C$ ,  $D$ , and  $E$ . These are connected through direct links used to ship goods. Further, suppose that the empirical flow of goods is provided in the form of trajectory data. As defined in Chapter 4, a trajectory is a sequence of nodes traversed by goods on their way from producers to consumers, via distributors. Specifically, we observe two trajectories: (i)  $A \rightarrow C \rightarrow D$ ; (ii)  $B \rightarrow C \rightarrow E$  (Fig. 9.1).

If we now use a standard network representation of our system, four trajectories would be possible, not two. These are: (i)  $A \rightarrow C \rightarrow D$ , (ii)  $A \rightarrow C \rightarrow E$ ,



**Figure 9.1:** First-order and second-order network representation of a distribution system comprising two manufacturers,  $A$  and  $B$ , and three distributors,  $C$ ,  $D$  and  $E$ . The empirical trajectories are two and depicted in blue (left side of the picture). These are preserved in the second-order representation. In contrast, in the first-order representation, we obtain two more trajectories, not observed in the data. These are depicted in red.

(iii)  $B \rightarrow C \rightarrow D$ , (iv)  $B \rightarrow C \rightarrow E$ . Based on the transitivity assumption, a standard network representation accounts for all possible trajectories given the network topology, thus discarding the constraints acquired from the trajectory data (Lambiotte *et al.*, 2019).

In contrast, higher-order network representations capture possible constraints of the data by recording information on previous steps along the trajectories. This is done by constructing networks where nodes represent trajectories and links represent connections between these trajectories (Lambiotte *et al.*, 2019; Scholtes, 2017; Scholtes *et al.*, 2014). Specifically, in a representation of order  $k$ , nodes represent trajectories of length  $k-1$ .

In Fig. 9.1, we depict a second-order ( $k = 2$ ) representation of the schematic distribution network considered above. In this representation, two firms traversed by an observed trajectory are aggregated in a single second-order node.

A link between second-order nodes represents a trajectory of length two. For example, the second-order node  $v_1 = (A, C)$  is obtained aggregating  $A$  and  $C$ , traversed by the trajectory  $A \rightarrow C \rightarrow D$ . The link between the second-order nodes  $v_1 = (A, C)$  and  $v_2 = (C, D)$  represents the trajectory  $A \rightarrow C \rightarrow D$ . Notice that there is no link between  $v_1 = (A, C)$  and  $v_3 = (C, E)$ , as the trajectory  $B \rightarrow C \rightarrow E$  is not observed in the data.

In short, we showed how higher-order network representations could be leveraged to capture empirical flows on networks in a compact form. As we learned from Chapter 4, most trajectories in the opioid distribution network have length two: goods move from manufacturers to retailers via warehouses (i.e., two tiers of distributors). Thus, we argue that a second-order representation ( $k = 2$ ) is already sufficient to capture the empirical flow of the system under study. We leave to future works the exploration of models having  $k > 2$ .

### 9.2.2 Higher-order representation

Let's consider a distribution network consisting of  $N$  firms, i.e., manufacturers and distributors, and a set of observed trajectories of goods  $P := \{p_1, p_2, \dots, p_S\}$ . Each trajectory is represented as a tuple,  $p_s = (k \rightarrow j \rightarrow i)$ , indicating the sequence of firms used to ship goods, e.g.,  $k =$  manufacturer,  $j =$  first-tier distributor, and  $i =$  second-tier distributor.

The first-order network representation  $G^{(1)}$  is defined as a tuple  $(V^{(1)}, \mathcal{E}^{(1)})$ .  $V^{(1)}$  is the set of firms, and  $\mathcal{E}^{(1)}$  is the set of supply links connecting them. We define  $w_{j,i}^{(1)}$  as the shipping *frequency* associated to the supply link  $j \rightarrow i$ , with  $(j, i) \in \mathcal{E}^{(1)}$ , namely the number of times a shipment occurs.

The second-order network representation  $G^{(2)}$  is defined as a tuple  $(V^{(2)}, \mathcal{E}^{(2)})$ .  $V^{(2)}$  is the set of second-order nodes, representing supply links between two firms in the underlying first-order network, and  $\mathcal{E}^{(2)}$  is the set of second-order links representing trajectories of length two belonging to the set  $P$  of the observed trajectories. For example, the link  $e = (v_1, v_2)$  between two second-order nodes  $v_1$  and  $v_2$ , i.e.,  $v_1 = (k, j)$  and  $v_2 = (j, i)$ , represents the trajectory  $k \rightarrow j \rightarrow i$  of length two. We define  $w_{v_1, v_2}^{(2)}$  as the *frequency* associated to the



second-order link. More explicitly,  $w_{v_1, v_2}^{(2)}$  is the number of times the trajectory is observed.

For each second-order link  $e$ , we consider a second-order link  $\tilde{e}$  in the opposite direction. While the second-order link  $e$  has the direction of the shipments, the second-order link  $\tilde{e}$  follows the direction of the demand, or orders\*. For instance, from the second-order link  $e = (v_1, v_2)$ , representing the shipping trajectory  $k \rightarrow j \rightarrow i$ , we obtain the second-order link  $\tilde{e} = (\tilde{v}_2, \tilde{v}_1)$ , representing the trajectory of the orders  $i \rightarrow j \rightarrow k$ . Note that also second-order nodes have a transposed representation, i.e., from  $v_1 = (k, j)$  we obtain  $\tilde{v}_1 = (j, k)$ , and from  $v_2 = (j, i)$  we obtain  $\tilde{v}_2 = (i, j)$ . Assuming that, during usual business, the number of shipments equals the number of orders, we can write the following equivalence:

$$\tilde{w}_{\tilde{v}_2, \tilde{v}_1}^{(2)} = w_{v_1, v_2}^{(2)} \quad (9.1)$$

The same idea applies to the first-order links. For each shipping link  $j \rightarrow i$ , we consider a link in the direction of the orders, i.e.,  $i \rightarrow j$ , with equal frequency, i.e.,:

$$\tilde{w}_{i, j}^{(1)} = w_{j, i}^{(1)} \quad (9.2)$$

Finally, firms in distribution networks hold inventories, or stocks (see discussion in Chapter 8) to store goods coming from various source partners. In a second-order representation, stocks have an interesting interpretation. The stock of a second-order node  $s_{v_1}$ , with  $v_1 = (k, j)$ , indicates the amount of good shipped by  $k$  and put in stock by  $j$ . In other words, it represents the part of the stock, or sub-stock, of  $j$  used to store goods coming from supplier  $k$ , i.e.,  $s_{(k, j)} = s_{(j|k)}$ . Note the advantage of the second-order representation: it allows for keeping track of sub-stocks held by firms in the first-order distribution network. On the contrary, this information was discarded in a first-order representation, where goods shipped by different source partners were aggregated in a single stock (see Chapter 8).

---

\*Note that the term “higher-order” does not connect with the dynamic of goods orders or the demand. It is the common term for models capturing flows in complex networks. (Scholtes *et al.*, 2016).

### 9.2.3 Formalizing flexibility

**Second-order transition matrices.** When a shock hits a distribution network, and supply substitution is implemented, firms and consumers adapt their orders towards their suppliers. To model how orders are placed, we use the second-order transition matrix  $T^{(2)\dagger}$  (Scholtes *et al.*, 2014).

We construct  $T^{(2)}$  such that it stores information on the probability that a given firm places orders to a given supplier via a given intermediary. Formally, it represents the “transition” probability of an order moving between two second-order nodes, i.e., along a trajectory of length two. For example, the entry  $T_{\tilde{v}_2, \tilde{v}_1}^{(2)}$ , with  $\tilde{v}_2 = (i, j)$  and  $\tilde{v}_1 = (j, k)$ , indicates the probability that firm  $i$  orders goods to firm  $k$ , via firm  $j$ . In other words, it represents the transition probability that order moves from firm  $i$  to firm  $k$ , via firm  $j$ , i.e.:

$$T_{(\tilde{v}_2, \tilde{v}_1)}^{(2)} = T_{(ij), (jk)}^{(2)} = P(i \rightarrow j \rightarrow k) \quad (9.3)$$

where  $i \rightarrow j \rightarrow k$  is the trajectory along which the *order* moves. We leverage the structure of this transition matrix to model system response to shocks. Let us consider two extreme cases.

**Zero flexibility: Non-Markovian transition matrix.** The system has zero flexibility. It does not implement supply substitution and keeps its business as usual. More explicitly, consumers keep their preferences towards their suppliers, and distributors keep their usual operations. The observed flow is kept unchanged.

This translates into constructing a second-order transition matrix that preserves the observed trajectories. We name this transition matrix “Non-Markovian” since it captures more than one step memory. It captures two steps memory along a given trajectory. Given the frequencies of orders,  $\tilde{w}^{(2)}$ ,

---

<sup>†</sup>In a second-order formalism, the second-order transition matrix,  $T^{(2)}$ , stores information on the transition probabilities between second-order nodes.

along the observed trajectories, we determine the entries of the second-order transition matrix,  $T^{(2)\text{Non-Mkv}}$ , as:

$$T_{\tilde{v}_2, \tilde{v}_1}^{(2)\text{Non-Mkv}} = \frac{\tilde{w}_{\tilde{v}_2, \tilde{v}_1}^{(2)}}{\sum_{\tilde{v}'} \tilde{w}_{\tilde{v}_2, \tilde{v}'}^{(2)}} \quad (9.4)$$

where the sum runs over all second-order nodes  $\tilde{v}$ .

Recalling that  $\tilde{v}_2 = (i, j)$  and  $\tilde{v}_1 = (j, k)$ , Eq. (9.4) ensures the dependency between the (probability of) orders placed by  $i$  towards  $j$  and the (probability of) orders placed by  $j$  towards  $k$ , namely  $P(i \rightarrow j \rightarrow k)$ .

**Full flexibility: Markovian transition matrix.** The system has full flexibility. It implements supply substitution and adapts to mitigate shortages. This means that consumers relax their preferences towards suppliers and distributors adapt their usual operations by rerouting goods on new trajectories.

To capture the orders' dynamic during substitution, we construct a transition matrix that stores no information about the usual operations. We define this transition matrix as the ‘‘Markovian’’ transition matrix,  $T^{(2)\text{Mkv}}$ , as it captures only information from the network topology discarding information about the empirical flow. Using the frequencies in the first-order network,  $\tilde{w}^{(1)}$ , we determine the entries of the second-order transition matrix,  $T^{(2)\text{Mkv}}$ , as:

$$T_{\tilde{v}_2, \tilde{v}_1}^{(2)\text{Mkv}} = \frac{\tilde{w}_{\tilde{v}_1}^{(2)}}{\sum_{\tilde{v}_1'} \tilde{w}_{\tilde{v}_1'}^{(2)}} = \frac{\tilde{w}_{j,k}^{(1)}}{\sum_{k'} \tilde{w}_{j,k'}^{(1)}} \quad (9.5)$$

where the last term indicates the probability that  $j$  places an order towards  $k$ . Recalling that  $\tilde{v}_2 = (i, j)$  and  $\tilde{v}_1 = (j, k)$ , Eq. (9.5) assumes that the (probability of) orders placed by  $j$  towards  $k$  and the (probability of) orders placed by  $i$  towards  $j$  are independent:  $P(i \rightarrow j \rightarrow k) = P(i \rightarrow j)P(j \rightarrow k)$  (see derivation in C). Specifically, the orders placed by  $i$  to the supplier  $k$  via the intermediary  $j$  do not depend on  $i$  anymore:  $i$  has relaxed its preferences towards  $k$ , aligning them to the intermediary  $j$ .

**Flexibility as model parameter.** The two limit cases described above represent two extreme, and quite unrealistic, system responses: zero or perfect supply substitution. To explore the full spectrum of responses lying between those two extreme cases, we consider the general form for the transition matrix of the model as:

$$T^{(2)}(\phi) = (1 - \phi)T^{(2)\text{Non-Mkv}} + \phi T^{(2)\text{Mkv}} \quad (9.6)$$

where  $\phi$  is a model parameter used to interpolate between two limit cases. We interpret  $\phi$  as the system flexibility that controls the degree of supply substitution implemented in response to the shock. Its value ranges from zero to one. When flexibility equals zero, the transition matrix is reduced to the Non-Markovian transition matrix,  $T^{(2)}(\phi = 0) = T^{(2)\text{Non-Mkv}}$ . When flexibility equals one, the transition matrix is reduced to the Markovian transition matrix,  $T^{(2)}(\phi = 1) = T^{(2)\text{Mkv}}$ .

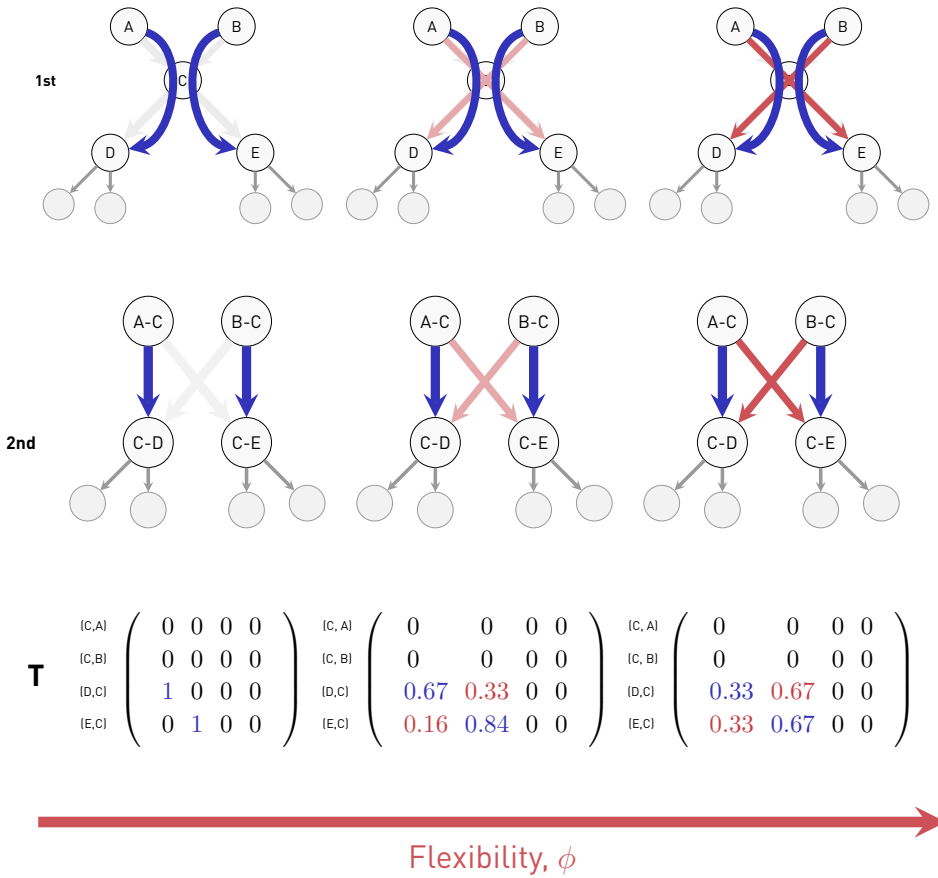
#### 9.2.4 An illustrative example

To illustrate our approach, we consider a simple example. We assume that there are five firms ( $A, B, C, D, E$ ) and two trajectories observed in a day-to-day business: (i)  $A \rightarrow C \rightarrow D$ , with a frequency  $f = 50$ , and (ii)  $B \rightarrow C \rightarrow E$  with frequency  $f = 100$ .

In Fig. 9.2, we draw three representations of this system under three different responses:  $\phi = 0$  (left column);  $\phi = 0.5$  (middle column); and  $\phi = 1$  (right column). For all three cases, we show the first-order network representation (top row), the second-order representation (middle row), and the transition matrix underlying the system response (bottom row).

For  $\phi = 0$ , the system keeps its usual business. From the the second-order transition matrix (bottom row, left column) we see that  $D$  only orders goods from  $A$ , via  $C$ , and  $E$  only order goods from  $B$ , via  $C$ . Notice that the transition matrix  $T^{(2)}$  indicates the probability of placing orders, not shipments.

When the system accepts a medium ( $\phi = 0.5$ ) degree of supply substitution,  $D$  relaxes its preferences towards  $A$  aligning them to the availability of  $C$ . From the transition matrix (bottom row, middle column) we see that now 33% of the



**Figure 9.2:** Three representations of the same distribution network under three different response strategies governed by (i) zero flexibility,  $\phi = 0$  (left column), (ii) medium flexibility  $\phi = 0.5$  (middle column), (iii) full flexibility  $\phi = 1$  (right column). First-order network representation (top row), second-order representation (middle row), second-order transition matrix (bottom row). The trajectories observed in a day-to-day business are colored in blue. The trajectories formed as result a of flexibility are colored in red.

orders placed by  $D$  are towards  $B$  (versus the 0% during usual business) and 67% of them are towards  $A$ . In other words, when flexibility is implemented  $D$  receives goods from both supplier  $A$  and  $B$ . The proportion of these orders

is dictated by the proportion of orders placed by  $C$ . Hence, on average, it indicates the availability of goods (from both suppliers) in the stock of  $C$ .

Moving from medium substitution to plain substitution ( $\phi = 1$ ), we see that 67% of orders placed by  $D$  are directed towards  $B$  (versus the 0% during a day-to-day business), and 33% are directed towards  $A$  (see bottom row, right column). In this scenario,  $D$  switches its preferences in favor of a perfect alignment with the (proportions of) orders placed by  $C$ .

Given the new proportion of orders, existing shipping trajectories are readjusted and new ones are formed. These new shipping trajectories are depicted in red in Fig. 9.2. For instance, for  $\phi = 1$ , the new trajectory  $B \rightarrow C \rightarrow D$  is used more often than the usual one  $A \rightarrow C \rightarrow D$ . Notice that our model does not account for any link rewiring in the distribution network. It only considers flow readjustment on the existing system.

### 9.2.5 System dynamics

The principles underlying the system dynamics in the first-order network and presented in Chapter 8 are still valid in a second-order representation. Distributors place orders to (i) restore their sub-stocks and (ii) meet demand, i.e.:

$$o_{(i,j)}^{(2)}(t) = d_{(i|j)}^{(2)}(t-1) + \frac{1}{\tau} \left[ s_{(i|j)}^{(2)T} - s_{(i|j)}^{(2)}(t) \right] \quad (9.7)$$

where  $o_{(i,j)}^{(2)}$  is the order placed by  $i$  towards  $j$  and  $d_{(i|j)}^{(2)}$  is the demand  $i$  faces on the goods received from  $j$ . As in the previous Chapter we set  $\tau = 5$ , corresponding to roughly five working days. In Eq. (9.7)  $s_{(i|j)}$  represents the sub-stock of  $i$  used to store goods received by  $j$ , whereas  $s_{(i|j)}^{(2)T}$  represents the target level (see Section 8.2.3). Sub-stocks are updated according to the total ship-out and the total ship-in:

$$s_{(i|j)}^{(2)}(t) = s_{(i|j)}^{(2)}(t-1) + W_{(j,i)}^{(2)\text{in}}(t-1) - \left[ W_{(j,i)}^{(2)\text{out}}(t-1) + \omega_{(i|j)}^{(2)}(t-1) \right] \quad (9.8)$$

where  $\omega_{(i|j)}^{(2)}$  indicates the amount of goods  $i$  receives from  $j$  and that ships to consumers.

$W^{(2)}$  is the shipping matrix in a second-order model. It defines the amount shipped along trajectories of length two. Specifically, the element  $W_{(kj)(ji)}$  indicates the amount of goods shipped by  $k$  to  $i$  via  $j$ , namely the amount shipped along the trajectory  $k \rightarrow j \rightarrow i$ . Similarly, the order matrix in a second-order model,  $O^{(2)}$ , defines the amount ordered along trajectories of length two. Specifically, the element  $O_{(i,j)(j,k)}^{(2)}$  indicates the amount of orders placed by  $i$  and directed towards  $k$ , via  $j$ .

Next, we obtain the orders along a given trajectory  $i \rightarrow j \rightarrow k$  as:

$$O_{(i,j)(j,k)}^{(2)}(t) = o_{(i,j)}^{(2)}(t)T_{(i,j)(j,k)}^{(2)}(\phi) \quad (9.9)$$

where  $T^{(2)}(\phi)$  is the transition matrix governing the system response, and  $o_{(i,j)}$  is obtained from Eq. (9.7).

Finally, we make a simple assumption about the direct supply shock. We consider a  $\sigma$  percentage reduction in the inventory level of the shocked manufacturer. Thus, by tuning the parameter  $\phi$  from zero to one we can study multiple responses to the shock depending on the system flexibility.

### 9.2.6 Measuring the shock impact

To evaluate the effect of flexibility, we consider two indicators: (i) demand deficit at consumers and (ii) changes to the usual operations. While the first indicator informs us about the effectiveness of the substitution policy in mitigating shortages among consumers, the second indicator is used to quantify the impact of the policy on the distribution system.

Specifically, we determine the demand deficit,  $dd(t)$ , at a given time  $t$ , as the percentage of the (cumulative) unfulfilled demand of consumers, i.e.:

$$dd(t) = \frac{\sum_{t'=0}^t \sum_v \omega_v(t') - c_v}{t \times \sum_v c_v} \quad (9.10)$$

where  $v$  is a second-order node,  $c_v$  is the (constant) daily demand from consumers, and  $\omega_v(t)$  is the quantity shipped to consumers at time  $t$ . Our indicator

is built assuming that goods demanded are shipped within the next working day.

The second indicator,  $\gamma(t)$ , measures the changes to the distribution system resulting from the readjustment of the usual flow. Hence, we consider the amount shipped,  $W_{(kj)(jk)}^{(2)}$ , along a given trajectory (of length two) in two scenarios: when supply substitution is implemented, i.e.,  $\phi \neq 0$ , and when it is not implemented, i.e.,  $\phi = 0$ . The difference between those two quantities gives the number of goods rerouted in case supply substitution is taken in place, i.e.,:

$$\gamma(t) = \frac{\sum_v \left| W_v^{(2)}(\phi, t) - W_v^{(2)}(\phi = 0, t) \right|}{\sum_v \left| W_v^{(2)}(\phi = 1, t) - W_v^{(2)}(\phi = 0, t) \right|} \quad (9.11)$$

where  $v$  is a second-order node. Since we are interested in the relative, or percentage, difference we chose to normalize the absolute difference, in the numerator, with the maximum possible difference, i.e., for  $\phi = 1$ , in the denominator.

Although we do not directly measure costs in our study, we use Eq. (9.11) as a proxy for that costs that may arise from increased flexibility. As flexibility increases, products are likely to flow through more distributors, leading to higher handling costs. Furthermore, using new distribution paths may result in additional costs due to labor and increased complexity. Hence, with our measure, we assess the so-called transportation and transshipment costs as defined by Aldrichetti *et al.* (2021). These costs include penalties for not using optimal transportation routes and the costs associated with utilizing alternative distributors' inventories. However, disruptions may also incur other costs, such as damage costs resulting from damage to machines, buildings, or inventories (Fattahi and Govindan, 2018; Turnquist and Vugrin, 2013); backlog costs due to loss of demand (Shukla *et al.*, 2011); or delay penalties for late deliveries (Elluru *et al.*, 2019). Our model does not capture the latter costs. Future studies can explore these aspects further.



## 9.3 Stress-testing the opioid distribution system

To study to what extent supply substitution helps distribution networks mitigate supply shocks, we implement a stress-test approach. We use the opioid distribution network as the starting point of our simulations and model its distortion due to the system response to the shock. Specifically, we consider the distribution network of Oxycodone in 2012.

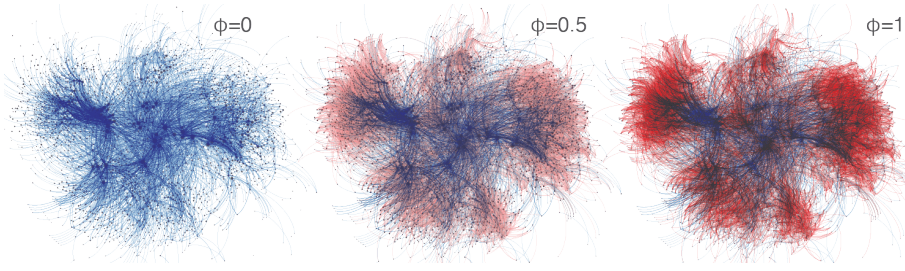
### 9.3.1 Model initialization

Because we want to start with the closest representation of the real system, our model incorporates various empirical information gathered from ARCOS data. We use these information to: construct an higher-order representation of the system; initialize stock levels and target stock levels; and determine consumers' demand.

**Higher-order representation.** As discussed in Chapter 4, our trajectory data comprise roughly two billion trajectories. We restrict our analysis to the trajectories of drugs containing Oxycodone and observed in a given year (2012). Given the set  $P := \{p_1, p_2, \dots, p_S\}$  of empirical trajectories, we compute their frequencies, in the selected year. Based on these frequencies, we create the second-order transition matrices according to Eq. 9.5 and Eq. 9.4<sup>‡</sup> and build the second-order network representation of the empirical system. In Fig. 9.3, we show the second-order representation of the Oxycodone distribution system for different flexibility values. In this representation, nodes represent links between two firms. Blue links represent observed trajectories of length two. Red links represent trajectories that are not observed, yet they are possible given the network topology. We see that by increasing flexibility, from  $\phi=0$  to  $\phi=1$ , we obtain new trajectories, depicted in red, that can be used to source and distribute goods.

---

<sup>‡</sup>For this computation we use the *pathpy* tool available in *python* (Scholtes and pathpy contributors, 2019).



**Figure 9.3:** Second-order representation of the Oxycodone distribution system in 2012. In this representation, a node represents a supply link between two firms, and an edge represents a trajectory of length two. Blue edges indicate trajectories that were empirically observed. Red edges indicate trajectories that were not observed, yet possible, given the network topology. Increasing the value of the parameter  $\phi$  in this representation, from  $\phi=0$  to  $\phi=1$ , increases the probability that these red trajectories become available for the distribution system.

**Target and initial stocks.** The target stock defines the amount of goods distributors plan as a safety buffer. Safety buffers are not recorded in the data. Assuming that all distributors met their planning within the observation year,  $y$ , we obtain target stocks as the buffer they remain with at the end of the year. From the yearly ship-in,  $W_i^{\text{in}}(y)$ , and the yearly ship-out to other firms,  $W_i^{\text{out}}(y)$ , and to consumers,  $\omega_i(y)$ , we compute the buffer of distributor  $i$  as:

$$s_i^T = W_i^{\text{in}}(y) - [W_i^{\text{out}}(y) - \omega_i(y)] \quad (9.12)$$

Note that, in some cases, the ship-out is bigger than the ship-in. This suggests that: (i) their inventories were not empty at the beginning of the given year, or (ii) they did not plan a target (safety) stock. For these firms, we set a minimum buffer equal to one. Since our model runs on the second-order representation, we determine the stocks in such representation. Respecting the proportion of volumes observed in the data, we compute the target sub-stock of  $i$  containing goods shipped by  $j$  as:

$$s_{(i|j)}^T = s_i^T \frac{W_{ji}(y)}{\sum_j W_{ji}(y)} \quad (9.13)$$

All stocks are initialized to their target values at the beginning of the simulation.

**Daily Demand.** Assuming that firms perfectly met demand within the observation year, we determine the constant daily demand faced by firm  $i$  as:

$$c_i = \frac{\omega_i(y)}{365} \quad (9.14)$$

where  $\omega_i(y)$  indicates the total ship-out of  $i$  directed towards consumers. As for the stock level, the second-order representation of  $c_i$  is obtained by respecting the proportion of volume observed, i.e.:

$$c_{(i|j)} = c_i \times \frac{W_{ji}(y)}{\sum_j W_{ji}(y)} \quad (9.15)$$

where  $c_{(i|j)}$  indicates the demand  $i$  receives from consumers on goods coming from  $j$ .


### 9.3.2 Tracing the efficient frontier

In our stress-test approach, we start from the opioid distribution network and run the cascade model on its second-order representation. As a supply-side shock, we consider a 30% reduction in *global* production, meaning that all manufacturers' stocks are affected by a reduction of 30%.

We remind the reader that the system response is, then, captured by the second-order transition matrix,  $T^{(2)}(\phi)$ , as function of the parameter  $\phi$ . Therefore, we vary  $\phi$  to investigate the role of flexibility as system response strategy. Specifically, we consider values of  $\phi$  ranging from 0 to 1, with an interval of 0.01, i.e.,  $\phi \in [0, 1]$ . Hence, we run in total 100 independent simulations, one for each  $\phi$  value, which indicates 100 possible system responses to shocks induced by a different degree of system flexibility each.

In Fig. 9.4 we show and the demand deficit  $dd$  (a), the system change  $\gamma$  (b) as a function of the system flexibility,  $\phi$ . Both indicators are evaluated after 45 times-steps (i.e., roughly one month and a half) from the time the shock hits the network.

From plot (a), we observe an initial decrease of  $dd$  for increasing  $\phi$  values, suggesting that flexibility effectively reduces consumer deficit. In other words, the



**Figure 9.4:** Change to the system (figure a) and demand deficit (figure b) as a function of system flexibility,  $\phi$ . The efficient frontier (figure c) is drawn from a collection of 100 points, representing a  $\phi$  value each. Low values are on the bottom-right, high values on the upper-right.

more the network adapts by rerouting goods leveraging additional, yet substitutable, sources, the less the impact on consumers. Interestingly, we see that  $dd$  increases again after reaching a minimum, in correspondence of  $\phi = 0.88$ . This finding indicates that while low or medium values of flexibility help reduce the deficit at consumers, high values of flexibility do not move the system towards a further reduction of demand deficit. A possible explanation is that a flexible system, leveraging alternative sources, may deplete the inventories of both distributors and manufacturers faster, which can not be used to mitigate the deficit further.

In short, we show that flexibility, up to some point, helps mitigate the deficit at consumers. Yet, what is the cost flexibility brings to the distribution system?

As discussed above, we evaluate those costs by measuring changes to the usual operations. As we show in Fig. 9.4 (b), the changes to the system increase with  $\phi$ . As expected, the more the system adapts, the more goods are rerouted, and

the more changes to the usual business are observed. Further, we find that the increasing trend is not linear. Specifically, we see that changes to the usual business are kept low, i.e., less than 10%, up to a given value of flexibility, i.e.,  $\phi \approx 0.25$ . For  $\phi > 0.25$  the changes increase almost exponentially with  $\phi$ . As we argued for the demand deficit, this non-linear behavior can be due to a faster stock depletion in case of high flexibility. Depleted stocks may produce major rerouting of goods, translating into pronounced changes to the usual business as observed for high  $\phi$  values. This hypothesis would enrich the interpretation of our results. However, we leave this investigation to further studies.

To combine the observations made so far, we study the role of flexibility in the  $dd - \gamma$  space, as reported in Fig 9.4 (c). Every point  $(dd, \gamma)$  in this space represents the system state reached with a given degree of flexibility  $\phi$ . The collection of points for various  $\phi$  values defines the shape of the observed curve. As before, both  $dd$  and  $\gamma$  are evaluated after  $t = 45$  time steps from the moment of the shock. In Fig. 9.4 (c) low  $\phi$  values are located on the bottom-right side of the curve, whereas high  $\phi$  values are located on the upper-right side. Hence, tuning  $\phi$  from low to high values corresponds to advance from bottom-right to bottom-left on the curve and further up from up-left to up-right.

Given the hyperbola shape of this curve, we determine two sets of points: an *efficient* set located on the lower side of the curve and defined by  $0 \leq \phi \leq 0.88$ ; and *inefficient* set located on the upper side of the curve and defined by  $0.88 < \phi \leq 1$ . The efficient set is colored in blue, while the inefficient one is colored in red.

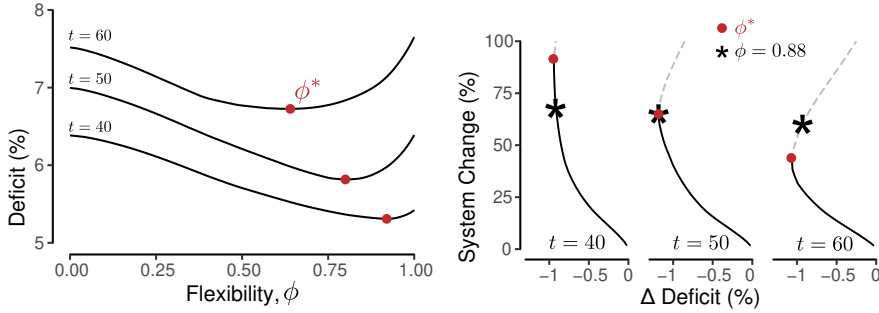
Within the efficient set, increasing flexibility allows to reduce demand deficit from 7% ( $\phi = 0$ ) to 5.5% ( $\phi = 0.88$ ). Still, this benefit to consumers comes with costs to the distribution system, which needs to change its usual operations of  $\approx 60\%$ . Hence, within this set of  $\phi$  values, we observe a trade-off between reducing demand deficit and keeping costs low for distributors. The decision on the optimal  $\phi$  value is left to a policymaker. The latter should balance the interests of the distributors in keeping low changes to their usual business with the general interest of governments and consumers in keeping the supply deficit low.

Within the inefficient set, instead, flexibility is not effective anymore. Interestingly, increasing flexibility further (more than  $\phi = 0.88$ ) generates big changes to the distribution system, which are not balanced by a further reduction in demand deficit. Indeed we see that high values of flexibility (located on the upper side of the curve) lead the system to have the same demand deficit as lower values of flexibility (located on the lower side of the curve), but with fewer changes to the usual operations.

The curve just discussed can be used by a policymaker to assess the degree of supply substitution needed while balancing two objectives: (i) keep costs low for distributors, (ii) keep deficit low for consumers. By borrowing the terminology from the field of finance, we use the term *efficient frontier* to define the observed curve. The Modern Portfolio Theory (MPT) shows a trade-off between portfolio risk and return. The term frontier identifies the optimal portfolio set that represents the best possible combination of expected return and investment risk (Fabozzi *et al.*, 2002; Mangram, 2013). We stress here that the analogy between MPT and our work only lies in the existence of a trade-off between two objectives, and we do not want to highlight additional similarities.

Finally, we further investigate the role of flexibility by studying the efficient frontier for multiple time horizons. In Fig. 9.5, we draw three frontiers (right plot) and three curves of deficit (left plot), one for each time-horizon considered:  $t = 40$ ,  $t = 50$ , and  $t = 60$ . In other words, each frontier indicates the system state, concerning demand deficit and operation changes, as observed after 40, 50, and 60 time steps by the time the shock hits the network.

We observe a reduction in deficit for all three horizons and a consequent trade-off: increasing flexibility reduces  $dd$  while increasing  $\gamma$ . Hence, for a given demand deficit that can be mitigated through flexibility, there is always an associated cost that the system needs to bear. More interestingly, the value of flexibility that minimizes the deficit,  $\phi^*$  (corresponding to the red point in Fig. 9.5), decreases over time. This implies that the inefficient set depends on the time horizon selected. Indeed, we see that the bending of the upper side of the frontier is more pronounced as the time-horizon increases. As time passes, more points move to the inefficient set (located on the upper side),



**Figure 9.5:** Demand deficit (left plot) and efficient frontier (right plot) as obtained after 40, 50, and 60-time steps since the simulated shock hits the distribution network. The red point indicates the flexibility value that leads to the minimum deficit. The star symbol indicates a given flexibility value that moves from the efficient to the inefficient set.

where the demand deficit is not further reduced. Indeed, the figure shows that the star symbol, corresponding to  $\phi = 0.88$ , moves from the efficient set to the inefficient one over the three-time horizons considered. As the time horizon indicates the duration of the shortage, we find that as the duration is longer, fewer flexibility values are efficient. Thus, as time goes on, choices about optimal flexibility are limited to a few values.

## 9.4 Discussion

In summary, we investigated the role of supply substitution as a system response to supply-side shocks. Within a short period, supply substitution is a quick and low-cost strategy against shortages. Communicating substitutable goods to consumers and distributors can rapidly reduce the shock's impact without carrying additional costs or times due, for instance, to a topological reconfiguration of the network. Indeed, we do not consider new link creation or rewiring but only redistribution of flows. Further, this policy does not demand that manufacturers ramp up production in a short period. By not requiring the production of new supplies, the policy takes advantage of the supply already available in the system. More generally, the proposed policy leverages existing

resources such as goods, business relations, and infrastructure, making them immediately available after a supply shock.

Yet, supply substitution implies (i) the willingness of distributors to reroute goods; and (ii) the willingness of consumers to accept them. We named this tendency *flexibility*. Using higher-order network models, we formalized flexibility as a single model parameter and investigated its impact on consumers and distributors.

We discussed the existence of two competitive objectives: (i) the government prefers to have a minimum deficit for consumers during shortages; (ii) distributors prefer to keep low changes to their operations to reduce costs. Using the opioid distribution network as a testbed, we quantified the trade-off between those two objectives.

We showed that increasing flexibility alleviates the deficit at consumers while generating changes, thus costs, to the distribution operations. The observed trade-off implies the existence of an efficient frontier traced in the space of demand deficit and changes to system operations. Interestingly, we identified two sets of flexibility values: an efficient set for low and medium flexibility values; and an inefficient set for high flexibility values. Within this inefficient set, big changes in the system operations are not balanced by a further reduction in demand deficit. Flexibility reduces demand deficit. Yet, too high flexibility produces no further benefit.

The optimal  $\phi$  selection is left to a policymaker since it depends on the practical implications (context-specific) that deficit and costs have for consumers and distributors, respectively. As we have characterized the frontier region along which a policymaker can select the optimal  $\phi$ , we have also revealed the existence of a region where a policymaker should *not* select the  $\phi$  since there is no optimal value.

In short, the main contribution of our study is to have characterized the efficient frontier along which a policymaker can balance the reduction in demand deficit against the increasing costs associated with flexibility. Our approach is versatile and applicable to a wide range of products, extending beyond pharmaceuticals. The only condition is that these products are substitutable and have partially overlapping distribution systems.



Besides this main contribution, the methodological contribution of our work needs to be mentioned. We proposed a novel method to study supply substitution as a quick response strategy of distribution networks toward supply shocks. Going beyond the standard network approach, we exploit the powerful abstraction of higher-order network models to capture flow readjustments of substitutable goods.

The proposed model incorporates several characteristics of an empirical distribution network (i.e., the goods' flow, network topology, inventories levels, and demand from consumers) while being parsimonious in the number of parameters. These characteristics make it well-suited for both future theoretical investigations and further empirical applications. This is particularly important since there are several aspects that our research has left open.

As we characterized the efficient frontier following the implementation of supply substitution, we did not investigate under which conditions this frontier is observed. Flexibility is likely to generate changes to the system operations under various circumstances. Instead, the behavior of demand deficit as a function of flexibility is not as trivial and calls for deeper investigations. Future studies can take multiple directions to enrich and extend the presented findings. First, we know that flexibility is linked to the number and the size of buffers (i.e., sub-stocks) distributors have and that they leverage during supply shortages. In our model, these buffers strictly depend on the (i) target stock levels and the (ii) topological dependencies between distributors and manufacturers. We took both this information from the data. Studying, even theoretically, how a change in one of these two aspects leads to change in the observed findings would enrich the presented work.

Second, we made a simple assumption about consumer demand: constant in time. Fluctuations of demand may speed up (or slow down) the cascading shortages, shifting up (or down) the flexibility values along the frontier. Hence, it may also affect the curve's shape at a given time horizon. Future research can investigate this point further.

Third, we mentioned that higher-order models could be generalized to any given order  $k$ , but then we focused only on the inclusions of the trajectories of length two. This choice makes sense for our application case since short-

length trajectories are used to ship goods from manufacturers to consumers in the opioid distribution network. However, it is up to further investigations on how the presence of (even very few) longer trajectories influences the presented results.

# Chapter 10

## Summary and Conclusions

### 10.1 Overview

In this thesis, we studied the formation, growth and resilience of large-scale distribution networks. Adopting a complex system perspective, we abstracted each constituent element (i.e., manufacturers, distributors, and consumers) and studied their interactions at the micro-level.

We set the goal to understand, on the one hand, the interaction rules driving link formation and firm growth and, on the other hand, the interaction mechanisms that enhance network resilience. For this aim, we developed a set of agent-based models to describe such interaction mechanisms at the micro-level and test the network properties at the macro-level. To test the effectiveness of our proposed models, we used a large-scale and real-world distribution system as a case study. Specifically, we examined the opioid distribution system that serves over 200,000 consumers and encompasses more than 1,000 firms.

Below, we summarize our main findings. These unfold along three lines: purely empirical, model-driven on network growth, and model-driven on network resilience. Next, we discuss our contributions to different scientific fields, from supply chains to network science. Finally, we provide an outlook on future research.

## 10.2 Summary in perspective

**Empirical characterization: network evolution.** In Chapter 3, we started our empirical investigation by analyzing the main structural properties of the nationwide distribution networks of two top-sold opioids, i.e., Oxycodone and Hydrocodone. According to our findings, these networks are characterized by (i) low density; (ii) firm heterogeneity in the number of distribution channels they operate; (iii) firms' tendency to link to more than one source partner; (iv) short depth, i.e., a maximum of three tiers of distributors. The properties (i), (ii), and (iv) can be explained as the result of cost-saving strategies and lean practices: the number of links is set close to the minimum; a few hub distributors handle most of the distribution channels; topological distances between manufacturers and consumers are kept short using a few tiers of distributors. Instead, property (iii) relates to system resilience: firms rely on multiple source partners to reduce their risk exposure to link failures.

We found these properties to be robust across the years. Interestingly, despite the substantial growth of opioid consumption between 2006 and 2011, the two distribution networks did not expand. Specifically, we did not detect any significant change in the topology of the distribution network of Oxycodone. Instead, we observed a shrinking in the size of the distribution network of Hydrocodone. Our further investigation revealed that the observed shrinking was due to the exit of small firms (i.e., low-degree firms) and the disappearance of weak links (i.e., low-weight links). Therefore, we interpreted the observed topological change as the system's adoption of cost-saving policies rather than the system response to demand growth.

**Empirical characterization: goods' flow.** In Chapter 4, we shift the perspective from a macroscopic time-aggregated perspective to a microscopic time-respecting one. From daily shipping records, we extracted more than 40 billion trajectories, tracking the flow of individual packages of opioids from production to consumption. We presented a novel data-mining technique that mimics a real-world distribution process and ensures domain-specific constraints (e.g., conservation of distributed quantity and compliance with shipping times).

This allowed us to investigate the system at a novel (high) resolution, never explored before.

Our main results revealed that the flow of goods is very heterogeneously distributed among the various distribution pathways connecting producers to consumers. Very few pathways are used ten thousand times, while many others are used only a few times. The observed concentration of flow along specific paths may be due to various reasons, e.g., the presence of very active consumers, their preferences for a particular producer, the dominance of a few large producers, and the existence of more efficient distribution chains.

Moreover, this analysis, conducted at the micro-level, allowed us to confirm the results obtained from a macroscopic topological analysis (in Chapter 3). In particular, we confirmed that only a few tiers of distributors operate in the distribution process. Opioids travel several thousand kilometres before reaching consumers. The opioid distribution networks span an entire nation. Nevertheless, the topological length of the observed trajectories is relatively short.

**Centralization and multi-sourcing: driving mechanisms for network formation.** Inspired by the empirical insights gained in Chapter 3, we developed a network growth model to explain the emergence of large-scale distribution networks in Chapter 5. Unlike previous works, we adopted a “eco-systemic” perspective. We viewed the distribution network as the outcome of an evolutionary process rather than the result of a careful design strategy of a single manufacturer.

Taking a firm’s perspective, we proposed two mechanisms to explain the emergence of an observed structure: centralization and multi-sourcing. Centralization expresses firms’ preferences in linking with central distributors, whereas multi-sourcing represents firms’ tendency to acquire additional source partners. The first strategy increases efficiency. The second one fosters local resilience to supply shocks.

We tested the model against the distribution networks of six opioid producers. We took a two-step validation procedure. First, through computer simulations, we identified the optimal set of parameters that generate the closest networks

to the empirical ones. Also, the fine-tuning of the parameters allowed us to gain empirical insights into the network structure and firms' strategies. Specifically, we found that the observed networks exhibit predominant centralization and frequent multi-sourcing. With a frequency of about 80%, firms connect to central distributors. Moreover, with a frequency of about 30%, firms implement multi-sourcing. Further, given the parameters obtained, such networks' topologies are star-like and exhibit a clear departure from perfect trees.

Next, we forwarded the validation procedure and tested the extent to which the model can reproduce a set of network properties of an empirical distribution network. By feeding the model with optimal parameters, we could reproduce the distinctive shapes of out and in-degree and path length distributions. The good match between model predictions and real-world observations suggested that the mechanisms proposed, despite their simplicity, were effective in reproducing the observed topology.

**The role of the upstream constraint in firm growth.** The network growth model of Chapter 5 was extended in Chapter 7 to account for the evolution of goods flow. Driven by empirical evidence (i.e., the log-normality of the distribution of outflows), we proposed a proportionate growth dynamic to describe the evolution of goods flow. Further, to explain the observed fluctuations of firms' growth rates, we proposed two factors: stochastic demand and upstream constraint. The first factor considers fluctuations due to the randomness of demand, i.e., downstream fluctuations. The second factor considers fluctuations due to supply availability, i.e., upstream fluctuations. Our results revealed that the second factor is necessary to reproduce the observed growth rates, especially the most extreme values. In short, upstream fluctuations have a predominant effect on firm growth.

**Shock propagation and upstream effects.** In Part III of this thesis, we shifted the focus from network growth to network resilience. In particular, we studied the propagation of supply shocks in large-scale distribution systems. As with network growth, we found that upstream effects also play an essential role in cascading dynamics.

In Chapter 8, we presented a cascade model that accounts for a bidirectional propagation of the shock: upstream and downstream. In the downstream propagation, distributors and manufacturers ration their shipments generating supply shortages. In the upstream propagation, distributors increase orders to mitigate shortages and meet the demand. We use the distribution network of Oxycodone as a test-bed for our analysis. Specifically, we implemented a stress-test approach: we started from the empirical network and modelled its distortion due to the shock propagation.

Our results uncovered non-trivial dependencies between producers under cascading shortages. Due to this upstream propagation, manufacturers not directly shocked may be harmed. They may deplete their stock earlier to offset shortages.

In addition, we noted that shortages harm not only firms but also consumers. We estimated the impact on consumers due to the cascade effects. If, at an early stage, stocks of alternative suppliers are leveraged to mitigate shortages (under the assumption of perfect substitution). From a certain point onward, alternative producers may run out of supplies, and indirect losses may be generated, adding to the direct ones.

**The role of supply substitution in mitigating shortages.** In Chapter 9, we explored the role of supply substitution as the system response to supply-side shocks. In the immediate aftermath of a shock, asking manufacturers to ramp up production may be neither feasible nor sufficient. Instead, communicating substitutable goods to consumers and distributors can quickly reduce the impact of the shock. Moreover, it does not entail additional costs due, for instance, to a topological reconfiguration of the network.

We introduced the term *flexibility* to indicate the degree of supply substitution the system is willing and capable of accepting. It identifies (i) the tendency of distributors to reroute packages after substitution; and (ii) the tendency of consumers to accept substitute goods.

Leverage the powerful abstraction of higher-order network models we captured flow readjustments after substitution and formalized flexibility as a model parameter. Tuning flexibility, we investigated the effect of various degrees of

supply substitution. The nationwide distribution network of Oxycodone drugs has been used as a test-bed for our analysis.

We showed that increasing flexibility alleviates the deficit among consumers. However, it also brings changes to usual operations that, in turn, may carry additional costs. We found that flexibility mitigates demand deficit up to a given value. High flexibility produces no further benefit.

Our analysis has highlighted the existence of a trade-off between shortage mitigation and costs after substitution and provided a novel tool that can help managers and policymakers better delineate mitigation strategies.

**From local to global resilience.** We want to point out to the reader that the effectiveness of supply-substitution policies in mitigating shortages depends closely on the topology of the network and its local properties. In Parts I and II of this thesis, we highlighted the tendency of firms to implement multi-sourcing. Firms create “redundant” connections with multiple source partners to reduce their exposure to disruptions of single links. Hence, this strategy increases network resilience at the local level. In Part III, it was found that it also promotes resilience at the global level. Supply substitution is not an option if all firms rely on one source partner. Only because firms have access to supplies from multiple sources, the system can leverage alternative distribution paths and mitigate shortages through substitution.

### 10.3 Scientific contribution

**Supply chain domain.** The main contribution of our work falls within the supply chain domain. Abandoning the traditional focal firm perspective, we have provided new methodologies for studying large-scale supply networks, in general, and distribution networks, in particular. The proposed models apply to large-scale dynamics and, at the same time, are grounded in detailed firm-level data.

Enhancing supply chain resilience is one of the biggest challenges of today’s economies and a top priority for several governments. In Part III of this thesis, we approached this demanding task. First, we proposed a firm-level



shock propagation model tailored to distribution networks. In this regard, we extended previous models limited to production networks. We highlighted the differences between shock propagation on the production side of a supply network and its distribution side. Next, we proposed supply substitution as a rapid response strategy to shortages. We provided a novel methodology (based on higher-order network techniques) to simulate it on a real-world network. Finally, we quantified the pros and cons this strategy carries.

In addition to network resilience, much of this thesis has been devoted to understanding the mechanisms underlying emerging network structures and firm growth. On the one hand, we have provided a simple methodology for generating network topologies similar to those observed. On the other, we have generalized traditional theories of firm growth to transfer them to the supply chain domain.

**Network science.** Network science is still in its infancy regarding supply chain applications. Our contribution to this field is to provide a new area of application on which old theories have been tested, and standard tools have been extended. For example, in Chapter 5, we extended previous network growth models to account for system-specific characteristics. Similarly, the cascade model proposed in Chapter 8 can be seen as an extension of threshold models previously used to study shock propagation in economic networks, e.g., in food trade networks (Burkholz and Schweitzer, 2019) or financial networks (Lorenz *et al.*, 2009). Ultimately, in Chapter 9, we exploited the framework of higher-order models, recently developed by network scientists, to describe more complex system responses to shocks.

**Data mining and Data Science.** One of the goals of these fields is to develop methods and algorithms to handle large-scale datasets and extract meaningful information from them. In addition, much attention has been paid to methods aimed at extracting sequential information (e.g., items flow) from time-stamped data. The method presented in Chapter 4 had the same objective: extracting goods flow from time-stamped shipping records. The method was built explicitly for distribution networks. However, it could inspire other

fields of application as long as the application concerns the flow of an item that is conserved in quantity.

**Generalizability.** In this thesis, we proposed modeling approaches that are not limited to the pharmaceutical industry. Specifically, the growth model we presented is rooted in general principles that encompass multi-sourcing, centralization, and the availability of supplies. These principles hold true across various industries, as they are derived from previous studies that have examined firm behavior and growth in sectors as diverse as automotive, chemical, and electronics.

Similarly, the cascade model we proposed has broad applicability. Shortages of goods are experienced across all industrial sectors. The approach proposed to mitigate such shortage can apply to various products, such as food, gas and electronics. The only requirement is that such products can be substituted and rely on even partially overlapping distribution systems.

Therefore, although our proposed models have been validated against a specific distribution system, we argue that our modeling framework is general enough to remain valid, with minor extensions, in other industrial sectors or geographic contexts. We encourage future research in this direction.

## 10.4 Outlook

Like any research project, the one presented in this thesis is not without its limitations. Nevertheless, limitations are opportunities to advance research. We outline below the main shortcomings of our work to better delineate future research directions.

**Network formation.** As part of model validation, we explored the entire two-dimensional parameter space to obtain the optimal values of  $q_s$  and  $q_t$ . To this end, we had to perform several time-consuming computer simulations. In the future, we could provide an analytical solution for the model parameters. The model is simple enough to support this investigation. In addition, Geipel *et al.* (2009) have already presented an analytical solution for the parameters

of a tree growth model. Based on their approach, the derivation of the parameter  $q_s$  should be straightforward. More challenging may be the derivation of  $q_t$ . However, it is worth exploring this direction since an analytical solution provides an undoubted advantage in terms of time gain.

From a conceptual point of view, the first possible extension would be to incorporate node heterogeneity into the model. For simplicity, we treated distributors and consumers in the same way: they all obey the same rules of interaction. This also leads to a small discrepancy between model prediction and real-world observation, discussed in Chapter 5. In the future, labels could be provided to nodes to distinguish between manufacturers, distributors, and consumers. Thus, linking rules could be designed to be consistent with node labelling, at the price of requiring more parameters.

In addition, the model could be further extended by allowing the entry of new manufacturers. From a broader perspective, one might be interested in replicating the topology of distribution networks of multiple manufacturers (e.g., operating in the same industry) rather than of a single manufacturer. Then, labelling nodes and designing different interaction rules might be required.

**Firm growth.** The study regarding firm growth is the one that left us with the most open questions. We had to reject the hypothesis of similarity between the empirical distribution of business growth rates and that obtained from the proposed model. Below we list several options for improving the model in the future.

First and foremost, an exit process for firms should be designed in addition to the entry process currently in the model. The primary mismatch between the model predictions and observations concerns the negative values of growth rates. The model predicts fewer negative values than those observed. Therefore, exit dynamics would not only describe a more realistic firm's behavior but could also produce more negative growth rates, thus reducing the discrepancy between simulated and real-world data.

Second, we could consider heterogeneous values of the upstream constraint,  $k$ , and make it firm-dependent. In other words, distributors could act differently in the case of shortage or surplus. This would produce more diverse

growth rates and reduce the observed differences between model predictions and observations. However, correctly defining the type of heterogeneity to be included is a challenging task. Information on the factors driving distributors' strategies during shortages or surpluses is needed. Usually, these data are private. Therefore, careful model design may be required, supported by existing theories. Alternatively, a data-driven approach could be taken: one could propose various distributions of  $k$  values and find the one that best fits the data, which leads to a minor mismatch between model predictions and data.

Finally, an attractive research direction to explore in the future is the investigation of how network topology affects growth rates. Previous work has studied firm growth, assuming that interactions between firms are negligible. Yet, firms' growth can be correlated due to competition. In distribution networks, firms' growth is correlated due to competition and collaboration. Investigating how these firms' dependencies affect their growth rates could advance the traditional studies of firm growth in a new direction. Does a less centralized structure lead to more extreme values of growth rates? Do perfect tree structures produce more homogenous values? Our modeling framework already provides a possible way to start addressing these questions. In future studies, we could vary the parameters  $q_s$  (centralization) and  $\alpha$  and  $q_t$  (multi-sourcing) to study their influence on firm growth.

**Cascades.** Some critical assumptions of the proposed cascade model could be revised and improved in future studies. The timing of the shock, its severity and its location are part of our modeling assumptions. Defining more realistic shock scenarios is quite a challenging task. It requires the joint effort of several research areas. In addition, it requires detailed information about the system, such as the company's exposure to natural hazards or occupational accidents. Future research should address this task to better design models with immediate practical implications.

A second central assumption is constant daily demand. Model extensions should consider a more realistic time dependence for consumer demand. To this end, data from ARCOS are a good starting point to find inspiration on the type of demand seasonality to include.

In the short run, it is reasonable to assume that supply relations are fixed. This may not be true in the long run: new supply links may be created, and others may be destroyed. Therefore, future work could extend the current model to account for the long-term effects of the cascade.

**Flexibility.** This research direction is promising to continue. We provided a new tool to study the impact of supply substitution policies with important practical implications. We showed the existence of a trade-off (between demand deficit reduction and cost after substitution) that a policymaker is confronted with when defining mitigation strategies. We expressed this trade-off through what we called an efficient frontier. We believe that of particular importance is to investigate the conditions under which this frontier is observed.

The first step to deepen our understanding is to perform synthetic tests. In the analysis presented in this thesis, the information about stock buffers and topological dependencies between firms is taken from the data. Instead, one could test different distributions of stock levels or network topologies to understand to what extent they give rise to similar frontier curves.

Next, the assumption we made in the simple cascade model is also present in this analysis: consumer demand is constant in time. Fluctuations of demand may speed up (or slow down) the cascading shortages, shifting up (or down) the flexibility values along the frontier. Hence, it may also affect the curve's shape at a given time horizon. Future research can investigate this point further.

We have pioneered an approach that improves system resilience by leveraging the resources already in place, such as infrastructure, business relations, and goods, making them immediately available without the need of building new ones. By striking a careful balance between flexibility-promoting policies and cost considerations, supply chains can enhance their resilience to disruptions and effectively adapt to them.

# Bibliography

- Acemoglu, D.; Ozdaglar, A.; Tahbaz-Salehi, A. (2015). Systemic risk and stability in financial networks. *American Economic Review* **105**(2), 564–608.
- Aldrighetti, R.; Battini, D.; Ivanov, D.; Zennaro, I. (2021). Costs of resilience and disruptions in supply chain network design models: a review and future research directions. *International Journal of Production Economics* **235**, 108103.
- Alon, U. (2007). Network motifs: theory and experimental approaches. *Nature Reviews Genetics* **8**(6), 450–461.
- Alstott, J.; Bullmore, E.; Plenz, D. (2014). powerlaw: a python package for analysis of heavy-tailed distributions. *PloS one* **9**(1), e85777.
- Altıparmak, F.; Gen, M.; Lin, L.; Karaoglan, I. (2009). A steady-state genetic algorithm for multi-product supply chain network design. *Computers & industrial engineering* **56**(2), 521–537.
- Amaral, L. A. N.; Buldyrev, S. V.; Havlin, S.; Salinger, M. A.; Stanley, H. E. (1998). Power law scaling for a system of interacting units with complex internal structure. *Physical Review Letters* **80**(7), 1385.
- Basole, R. C.; Bellamy, M. A. (2014). Supply network structure, visibility, and risk diffusion: a computational approach. *Decision Sciences* **45**(4), 753–789.

- Battiston, S.; Gatti, D. D.; Gallegati, M.; Greenwald, B.; Stiglitz, J. E. (2012a). Default cascades: when does risk diversification increase stability? *Journal of Financial Stability* **8(3)**, 138–149.
- Battiston, S.; Puliga, M.; Kaushik, R.; Tasca, P.; Caldarelli, G. (2012b). Debtrank: too central to fail? financial networks, the fed and systemic risk. *Scientific reports* **2**, 541.
- Becchetti, L.; Trovato, G. (2002). The determinants of growth for small and medium sized firms. The role of the availability of external finance. *Small business economics* **19(4)**, 291–306.
- Berlec, T.; Kušar, J.; Žerovnik, J.; Starbek, M. (2014). Optimization of a product batch quantity. *Strojniški vestnik-Journal of Mechanical Engineering* **60(1)**, 35–42.
- Bode, C.; Wagner, S. M. (2015). Structural drivers of upstream supply chain complexity and the frequency of supply chain disruptions. *Journal of Operations Management* **36**, 215–228.
- Boginski, V.; Butenko, S.; Pardalos, P. M. (2005). Statistical analysis of financial networks. *Computational statistics & data analysis* **48(2)**, 431–443.
- Bottazzi, G.; Coad, A.; Jacoby, N.; Secchi, A. (2011). Corporate growth and industrial dynamics: evidence from French manufacturing. *Applied Economics* **43(1)**, 103–116.
- Bottazzi, G.; Dosi, G.; Lippi, M.; Pammolli, F.; Riccaboni, M. (2001). Innovation and corporate growth in the evolution of the drug industry. *International journal of industrial organization* **19(7)**, 1161–1187.
- Bottazzi, G.; Secchi, A. (2003). Common properties and sectoral specificities in the dynamics of US manufacturing companies. *Review of Industrial Organization* **23(3-4)**, 217–232.
- Bottazzi, G.; Secchi, A. (2006). Explaining the distribution of firm growth rates. *The RAND Journal of Economics* **37(2)**, 235–256.

- Brintrup, A.; Ledwoch, A. (2018). Supply network science: emergence of a new perspective on a classical field. *Chaos: An Interdisciplinary Journal of Nonlinear Science* **28(3)**, 033120.
- Brintrup, A.; Wang, Y.; Tiwari, A. (2015). Supply networks as complex systems: a network-science-based characterization. *IEEE Systems Journal* **11(4)**, 2170–2181.
- Burke, G. J.; Carrillo, J. E.; Vakharia, A. J. (2004). *Sourcing decisions with stochastic supplier reliability and stochastic demand*. Tech. rep., Citeseer.
- Burkholz, R.; Schweitzer, F. (2019). International crop trade networks: the impact of shocks and cascades. *Environmental Research Letters* **14(11)**, 114013.
- Butts, C. T. (2009). Revisiting the foundations of network analysis. *science* **325(5939)**, 414–416.
- Caceres, R. S.; Berger-Wolf, T. (2013). Temporal scale of dynamic networks. In: *Temporal networks*, Springer. pp. 65–94.
- Calvino, F.; Criscuolo, C.; Menon, C.; Secchi, A. (2018). Growth volatility and size: a firm-level study. *Journal of Economic Dynamics and Control* **90**, 390–407.
- Caridi, M.; Pero, M.; Sianesi, A. (2012). Linking product modularity and innovativeness to supply chain management in the Italian furniture industry. *International Journal of Production Economics* **136(1)**, 207–217.
- Carvalho, V. M.; Nirei, M.; Saito, Y. U.; Tahbaz-Salehi, A. (2021). Supply chain disruptions: evidence from the great east japan earthquake. *The Quarterly Journal of Economics* **136(2)**, 1255–1321.
- Choi, T. Y.; Dooley, K. J.; Rungtusanatham, M. (2001). Supply networks and complex adaptive systems: control versus emergence. *Journal of operations management* **19(3)**, 351–366.



- Choi, T. Y.; Hong, Y. (2002). Unveiling the structure of supply networks: case studies in Honda, Acura, and DaimlerChrysler. *Journal of Operations Management* **20**(5), 469–493.
- Clark, K. B.; Chew, W. B.; Fujimoto, T.; Meyer, J.; Scherer, F. (1987). Product development in the world auto industry. *Brookings Papers on economic activity* **1987**(3), 729–781.
- Clauset, A.; Shalizi, C. R.; Newman, M. E. (2009). Power-law distributions in empirical data. *SIAM review* **51**(4), 661–703.
- Coad, A. (2007). Firm growth: a survey .
- Conrad, K. (2017). The Opioid Epidemic. *Current Emergency and Hospital Medicine Reports* **5**(4), 119–120.
- Craighead, C. W.; Blackhurst, J.; Rungtusanatham, M. J.; Handfield, R. B. (2007). The severity of supply chain disruptions: design characteristics and mitigation capabilities. *Decision Sciences* **38**(1), 131–156.
- Dasgupta, N.; Beletsky, L.; Ciccarone, D. (2018). Opioid crisis: no easy fix to its social and economic determinants. *American journal of public health* **108**(2), 182–186.
- DeWeerd, S. (2019). Tracing the US opioid crisis to its roots. *Nature* **573**(7773), S10.
- Diem, C.; Borsos, A.; Reisch, T.; Kertész, J.; Thurner, S. (2022). Quantifying firm-level economic systemic risk from nation-wide supply networks. *Scientific reports* **12**(1), 1–13.
- Ding, Z.; Gong, W.; Li, S.; Wu, Z. (2018). System dynamics versus agent-based modeling: a review of complexity simulation in construction waste management. *Sustainability* **10**(7), 2484.
- Dolgui, A.; Ivanov, D.; Sokolov, B. (2018). Ripple effect in the supply chain: an analysis and recent literature. *International Journal of Production Research* **56**(1-2), 414–430.

- Dong, Y.; Skowronski, K.; Song, S.; Venkataraman, S.; Zou, F. (2020). Supply base innovation and firm financial performance. *Journal of Operations Management* **66(7-8)**, 768–796.
- Dyer, O. (2016). Opioid manufacturer bribed doctors to prescribe fentanyl inappropriately, US says.
- Eisenstein, M. (2019). Treading the tightrope of opioid restrictions. *Nature* **573(7773)**, S13.
- Elliott, M.; Golub, B.; Jackson, M. O. (2014). Financial networks and contagion. *American Economic Review* **104(10)**, 3115–53.
- Elluru, S.; Gupta, H.; Kaur, H.; Singh, S. P. (2019). Proactive and reactive models for disaster resilient supply chain. *Annals of Operations Research* **283**, 199–224.
- Emanuel, G.; Thomas, K. (2019). Top executives of INSYS, an opioid company, are found guilty of racketeering. *The New York Times. New York* .
- Fabozzi, F. J.; Gupta, F.; Markowitz, H. M. (2002). The legacy of modern portfolio theory. *The journal of investing* **11(3)**, 7–22.
- Fattahi, M.; Govindan, K. (2018). A multi-stage stochastic program for the sustainable design of biofuel supply chain networks under biomass supply uncertainty and disruption risk: a real-life case study. *Transportation Research Part E: Logistics and Transportation Review* **118**, 534–567.
- Freeman, L. C. (1977). A set of measures of centrality based on betweenness. *Sociometry* , 35–41.
- Fu, D.; Pammolli, F.; Buldyrev, S. V.; Riccaboni, M.; Matia, K.; Yamasaki, K.; Stanley, H. E. (2005). The growth of business firms: theoretical framework and empirical evidence. *Proceedings of the National Academy of Sciences* **102(52)**, 18801–18806.

- Geipel, M. M.; Tessone, C. J.; Schweitzer, F. (2009). A complementary view on the growth of directory trees. *The European Physical Journal B* **71(4)**, 641–648.
- Gereffi, G. (2014). Global value chains in a post-Washington Consensus world. *Review of international political economy* **21(1)**, 9–37.
- Gereffi, G. (2019). Global value chains, development, and emerging economies 1. In: *Business and Development Studies*, Routledge. pp. 125–158.
- Gertler, P. J.; Martinez, S.; Premand, P.; Rawlings, L. B.; Vermeersch, C. M. (2016). *Impact evaluation in practice*. World Bank Publications.
- Ghadge, A.; Er, M.; Ivanov, D.; Chaudhuri, A. (2022). Visualisation of ripple effect in supply chains under long-term, simultaneous disruptions: a system dynamics approach. *International Journal of Production Research* **60(20)**, 6173–6186.
- Gibrat, R. (1931). *R and Les Inégalités Économiques*, Libraire du Recueil. Paris, France .
- Guy Jr, G. P.; Zhang, K.; Bohm, M. K.; Losby, J.; Lewis, B.; Young, R.; Murphy, L. B.; Dowell, D. (2017). Vital signs: changes in opioid prescribing in the United States, 2006–2015. *MMWR. Morbidity and mortality weekly report* **66(26)**, 697.
- Hallegatte, S. (2008). An adaptive regional input-output model and its application to the assessment of the economic cost of Katrina. *Risk Analysis: an International Journal* **28(3)**, 779–799.
- Hallegatte, S. (2014). Modeling the role of inventories and heterogeneity in the assessment of the economic costs of natural disasters. *Risk analysis* **34(1)**, 152–167.
- Harahap, F. A. U.; Siregar, R. N. I.; Sihotang, W. Y.; Nasution, A. N. (2022). Drug Management On Availability Of Drugs In Pharmaceutical

- Installations Pabatu General Hospital Pt Pmn Using Fifo & Fefo Methods. *International Journal of Health and Pharmaceutical (IJHP)* **3(1)**, 72–81.
- Hearnshaw, E. J.; Wilson, M. M. (2013). A complex network approach to supply chain network theory. *International Journal of Operations & Production Management* .
- Helfgott, A. (2018). Operationalising systemic resilience. *European Journal of Operational Research* **268(3)**, 852–864.
- Hollnagel, E.; Woods, D. D.; Leveson, N. (2006). *Resilience engineering: concepts and precepts*. Ashgate Publishing, Ltd.
- Holme, P. (2015). Modern temporal network theory: a colloquium. *The European Physical Journal B* **88(9)**, 1–30.
- Huang, B.; Wu, A. (2016). EOQ model with batch demand and planned backorders. *Applied Mathematical Modelling* **40(9-10)**, 5482–5496.
- Inderst, R. (2008). Single sourcing versus multiple sourcing. *The Rand journal of economics* **39(1)**, 199–213.
- Inoue, H.; Todo, Y. (2019). Firm-level propagation of shocks through supply-chain networks. *Nature Sustainability* **2(9)**, 841–847.
- Iribarren, J. L.; Moro, E. (2009). Impact of human activity patterns on the dynamics of information diffusion. *Physical review letters* **103(3)**, 038702.
- Ivanov, D. (2017). Simulation-based ripple effect modelling in the supply chain. *International Journal of Production Research* **55(7)**, 2083–2101.
- Jabbarzadeh, A.; Fahimnia, B.; Sheu, J.-B.; Moghadam, H. S. (2016). Designing a supply chain resilient to major disruptions and supply/demand interruptions. *Transportation Research Part B: Methodological* **94**, 121–149.

- Jayawardana, S.; Forman, R.; Johnston-Webber, C.; Campbell, A.; Berterame, S.; de Joncheere, C.; Aitken, M.; Mossialos, E. (2021). Global consumption of prescription opioid analgesics between 2009-2019: a country-level observational study. *EClinicalMedicine* **42**, 101198.
- Jiang, B.; Prater, E. (2002). Distribution and logistics development in China: the revolution has begun. *International Journal of Physical Distribution & Logistics Management* **32(9)**, 783–798.
- Jones, M. R.; Novitch, M. B.; Sarrafpour, S.; Ehrhardt, K. P.; Scott, B. B.; Orhurhu, V.; Viswanath, O.; Kaye, A. D.; Gill, J.; Simopoulos, T. T. (2019). Government legislation in response to the opioid epidemic. *Current pain and headache reports* **23(6)**, 1–7.
- Kali, R.; Reyes, J. (2010). Financial contagion on the international trade network. *Economic Inquiry* **48(4)**, 1072–1101.
- Kim, Y.; Chen, Y.-S.; Linderman, K. (2015). Supply network disruption and resilience: a network structural perspective. *Journal of operations Management* **33**, 43–59.
- Kim, Y.; Choi, T. Y.; Yan, T.; Dooley, K. (2011). Structural investigation of supply networks: a social network analysis approach. *Journal of Operations Management* **29(3)**, 194–211.
- Kito, T.; Brintrup, A.; New, S.; Reed-Tsochas, F. (2014). The structure of the Toyota supply network: an empirical analysis. *Saïd Business School WP* **3**.
- Klemm, K.; Eguíluz, V. M.; San Miguel, M. (2005). Scaling in the structure of directory trees in a computer cluster. *Physical review letters* **95(12)**, 128701.
- Ko, D. K.; Brandizzi, F. (2020). Network-based approaches for understanding gene regulation and function in plants. *The Plant Journal* **104(2)**, 302–317.

- Lambiotte, R.; Rosvall, M.; Scholtes, I. (2019). From networks to optimal higher-order models of complex systems. *Nature physics* **15(4)**, 313–320.
- Lan, Y.; Li, Y.; Papier, F. (2018). Competition and coordination in a three-tier supply chain with differentiated channels. *European Journal of Operational Research* **269(3)**, 870–882.
- Latour, A. (2001). Trial by fire: a blaze in Albuquerque sets off major crisis for cell-phone giants. *Wall Street Journal* **1(29)**, 2001.
- Li, X.; Jin, Y. Y.; Chen, G. (2003). Complexity and synchronization of the world trade web. *Physica A: Statistical Mechanics and its Applications* **328(1-2)**, 287–296.
- Li, Y.; Zobel, C. W. (2020). Exploring supply chain network resilience in the presence of the ripple effect. *International Journal of Production Economics* , 107693.
- Lorenz, J.; Battiston, S.; Schweitzer, F. (2009). Systemic risk in a unifying framework for cascading processes on networks. *The European Physical Journal B* **71(4)**, 441.
- Mangram, M. E. (2013). A simplified perspective of the Markowitz portfolio theory. *Global journal of business research* **7(1)**, 59–70.
- Martínez-Olvera, C.; Davizon-Castillo, Y. A.; Tozan, H.; Erturk, A. (2015). Modeling the supply chain management creation of value—a literature review of relevant concepts. *Applications of Contemporary Management Approaches in Supply Chains* , 57–82.
- Mattsson, C. E.; Takes, F. W. (2021). Trajectories through temporal networks. *Applied Network Science* **6(1)**, 1–31.
- Milo, R.; Shen-Orr, S.; Itzkovitz, S.; Kashtan, N.; Chklovskii, D.; Alon, U. (2002). Network motifs: simple building blocks of complex networks. *Science* **298(5594)**, 824–827.

- Mizgier, K. J.; Jüttner, M. P.; Wagner, S. M. (2013). Bottleneck identification in supply chain networks. *International Journal of Production Research* **51(5)**, 1477–1490.
- Paniccia, I. (1998). One, a hundred, thousands of industrial districts. Organizational variety in local networks of small and medium-sized enterprises. *Organization studies* **19(4)**, 667–699.
- Pathak, S. D.; Day, J. M.; Nair, A.; Sawaya, W. J.; Kristal, M. M. (2007). Complexity and adaptivity in supply networks: building supply network theory using a complex adaptive systems perspective. *Decision sciences* **38(4)**, 547–580.
- Peel, L.; Peixoto, T. P.; De Domenico, M. (2022). Statistical inference links data and theory in network science. *Nature Communications* **13(1)**, 1–15.
- Peixoto, T. P.; Rosvall, M. (2017). Modelling sequences and temporal networks with dynamic community structures. *Nature communications* **8(1)**, 1–12.
- Perera, S.; Bell, M. G.; Bliemer, M. C. (2017a). Network science approach to modelling the topology and robustness of supply chain networks: a review and perspective. *Applied network science* **2(1)**, 1–25.
- Perera, S.; Perera, H. N.; Kasthurirathna, D. (2017b). Structural characteristics of complex supply chain networks. In: *2017 Moratuwa Engineering Research Conference (MERCOn)*. IEEE, pp. 135–140.
- Pero, M.; Abdelkafi, N.; Sianesi, A.; Blecker, T. (2010). A framework for the alignment of new product development and supply chains. *Supply Chain Management: An International Journal* .
- Potter, A.; Wilhelm, M. (2020). Exploring supplier–supplier innovations within the Toyota supply network: a supply network perspective. *Journal of Operations Management* **66(7-8)**, 797–819.

- Reisch, T.; Heiler, G.; Diem, C.; Klimek, P.; Thurner, S. (2022). Monitoring supply networks from mobile phone data for estimating the systemic risk of an economy. *Scientific reports* **12**(1), 1–10.
- Riccaboni, M.; Schiavo, S. (2010). Structure and growth of weighted networks. *New Journal of Physics* **12**(2), 023003.
- Roukny, T.; Bersini, H.; Pirotte, H.; Caldarelli, G.; Battiston, S. (2013). Default cascades in complex networks: topology and systemic risk. *Scientific reports* **3**(1), 1–8.
- Scheibe, K. P.; Blackhurst, J. (2018). Supply chain disruption propagation: a systemic risk and normal accident theory perspective. *International Journal of Production Research* **56**(1-2), 43–59.
- Schmitt, A. J.; Sun, S. A.; Snyder, L. V.; Shen, Z.-J. M. (2015). Centralization versus decentralization: risk pooling, risk diversification, and supply chain disruptions. *Omega* **52**, 201–212.
- Schoenherr, T. (2010). Outsourcing decisions in global supply chains: an exploratory multi-country survey. *International Journal of Production Research* **48**(2), 343–378.
- Scholtes, I. (2017). When is a network a network? multi-order graphical model selection in pathways and temporal networks. In: *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*. pp. 1037–1046.
- Scholtes, I.; pathpy contributors (2019). pathpy. <https://www.pathpy.net/>. Python package version 2.2.0.
- Scholtes, I.; Wider, N.; Garas, A. (2016). Higher-order aggregate networks in the analysis of temporal networks: path structures and centralities. *The European Physical Journal B* **89**(3), 1–15.
- Scholtes, I.; Wider, N.; Pfitzner, R.; Garas, A.; Tessone, C. J.; Schweitzer, F. (2014). Causality-driven slow-down and speed-up of diffusion in non-Markovian temporal networks. *Nature communications* **5**(1), 1–9.



- Schueller, W.; Diem, C.; Hinterplattner, M.; Stangl, J.; Conrady, B.; Gerschberger, M.; Thurner, S. (2022). Propagation of disruptions in supply networks of essential goods: a population-centered perspective of systemic risk. *arXiv preprint arXiv:2201.13325* .
- Schwartz, F.; Voß, S. (2007). Distribution network design with postponement. *Wirtschaftsinformatik Proceedings 2007* , 78.
- Schweitzer, F. (2020). The law of proportionate growth and its siblings: applications in agent-based modeling of socio-economic systems. In: *Complexity, Heterogeneity, and the Methods of Statistical Physics in Economics*, Springer. pp. 145–176.
- Schweitzer, F. (2022a). Agents Networks Evolution: a quarter century of advances in complex systems.
- Schweitzer, F. (2022b). Group relations, resilience and the I Ching. *Physica A: Statistical Mechanics and its Applications* , 127630.
- Schweitzer, F.; Casiraghi, G.; Tomasello, M. V.; Garcia, D. (2021). Fragile, yet resilient: adaptive decline in a collaboration network of firms. *Frontiers in Applied Mathematics and Statistics* **7**, 634006.
- Shukla, A.; Agarwal Lalit, V.; Venkatasubramanian, V. (2011). Optimizing efficiency-robustness trade-offs in supply chain design under uncertainty due to disruptions. *International journal of physical distribution & logistics management* **41(6)**, 623–647.
- SLCG (2019). Opioid Data. <https://www.slcg.com/opioid-data>.
- Squartini, T.; Van Lelyveld, I.; Garlaschelli, D. (2013). Early-warning signals of topological collapse in interbank networks. *Scientific reports* **3**, 3357.
- Sterman, J. (2010). *Business dynamics*. Irwin/McGraw-Hill.
- Sutton, J. (2002). The variance of firm growth rates: the “scaling” puzzle. *Physica a: statistical mechanics and its applications* **312(3-4)**, 577–590.

- Tomasello, M. V.; Perra, N.; Tessone, C. J.; Karsai, M.; Schweitzer, F. (2014). The role of endogenous and exogenous mechanisms in the formation of R&D networks. *Scientific reports* **4**(1), 1–12.
- Tomlin, B. (2006). On the value of mitigation and contingency strategies for managing supply chain disruption risks. *Management science* **52**(5), 639–657.
- Treiblmaier, H. (2018). Optimal levels of (de) centralization for resilient supply chains. *The International Journal of Logistics Management* .
- Turnquist, M.; Vugrin, E. (2013). Design for resilience in infrastructure distribution networks. *Environment Systems & Decisions* **33**, 104–120.
- Vaccario, G.; Verginer, L.; Schweitzer, F. (2020). The mobility network of scientists: analyzing temporal correlations in scientific careers. *Applied Network Science* **5**(1), 1–14.
- Volkow, N. D. (2014). America’s addiction to opioids: heroin and prescription drug abuse. *Senate Caucus on International Narcotics Control* **14**.
- Vujanac, R.; Miloradovic, N.; Vulovic, S. (2016). Dynamic storage systems. *Annals of the Faculty of Engineering Hunedoara* **14**(1), 79.
- Wang, G.; Gunasekaran, A.; Ngai, E. W. (2018). Distribution network design with big data: model and analysis. *Annals of Operations Research* **270**(1), 539–551.
- Washington Post Investigative (2019). ARCOS API. <https://github.com/wpinvestigative/arcos-api>.
- Whalen, J. (2018). Purdue Pharma to stop promoting OxyContin to US doctors. *The Wall Street Journal* .
- Wichmann, P.; Brintrup, A.; Baker, S.; Woodall, P.; McFarlane, D. (2018). Towards automatically generating supply chain maps from natural language text. *IFAC-PapersOnLine* **51**(11), 1726–1731.

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

Wycisk, C.; McKelvey, B.; Hülsmann, M. (2008). “Smart parts” supply networks as complex adaptive systems: analysis and implications. *International Journal of Physical Distribution & Logistics Management* .

Wynarczyk, P.; Watson, R. (2005). Firm growth and supply chain partnerships: an empirical analysis of UK SME subcontractors. *Small Business Economics* **24(1)**, 39–51.

Yadav, A. S.; Sharma, V.; Agarwal, P.; Swami, A.; Yadav, P. K. (2021). Pharmaceutical drug two-warehouse inventory model under FIFO dispatching policy using ant colony optimization for travelling salesman problem. *Linguistics and Culture Review* **5(S2)**, 1148–1171.

Zhao, K.; Kumar, A.; Yen, J. (2010). Achieving high robustness in supply distribution networks by rewiring. *IEEE Transactions on Engineering Management* **58(2)**, 347–362.

Zhou, H.; de Wit, G. (2009). Determinants and dimensions of firm growth. *SCALES EIM Research Reports (H200903)* .

Zimmermann, R.; Ferreira, L. M. D.; Moreira, A. C. (2016). The influence of supply chain on the innovation process: a systematic literature review. *Supply Chain Management: An International Journal* .

# Appendix

# Appendix A

## Supplementary Material to Chapter 2

---

### Business Activities

---

RESEARCHER (II-V)	MLP-MILITARY	PHARMACY- MIL
CENTRAL FILL PHARMACY	REVERSE DISTRIB	PRACTITIONER-DW/100
CHEMICAL EXPORTER	MLP-NATUROPATHIC PHYSICIAN	HOSP/CLINIC FED
CHAIN PHARMACY	MLP-AMBULANCE SERVICE	MANUF (BULK)
HOSP/CLINIC-VA	DETOXIFICATION	HOSP/CLINIC- MIL
MLP-OPTOMETRIST	MLP-NURSE PRACTITIONER-DW/100	MLP-REGISTERED PHARMACIST
MLP-ANIMAL SHELTER	PHARMACY - FED	PRACTITIONER-MILITARY
IMPORTER	EXPORTER	PRACTITIONER
MLP-NURSE PRACTITIONER	CHEMPACK/SNS DISTRIBUTOR	ANALYTICAL LAB
M/O PHARMACY	MLP-DEPT OF STATE	RESEARCHER (I)
CHEMICAL MANUFACTURER	MLP-NURSE PRACTITIONER-DW/30	MAINTENANCE
COMPOUND/DETOX	MLP-PHYSICIAN ASSISTANT-DW/30	CHAIN HOSP/CLINIC
HOSPITAL/CLINIC	MLP-PHYSICIAN ASSISTANT	COMP/MAINT/DETOX
CHEMICAL DISTRIBUTOR	MLP-PHYSICIAN ASSISTANT-DW/100	MLP-NURSING HOME
MAINT DETOX	RETAIL PHARMACY	COMPOUND/MAINT
MLP-EUTHANASIA TECHNICIAN	PRACTITIONER-DW/275	AUTOMATED DISPENSING SYSTEM
PRACTITIONER-DW/30	TEACHING INSTITUTION	MANUFACTURER
DISTRIBUTOR	CANINE HANDLER	IMPORTER (C, LI)

**Table A.1:** Business activities listed in the ARCOS dataset. We group them into three categories: manufacturers, distributors, and “consumers”. The consumers are those who only receive goods and are not classified as either manufacturers or distributors. Importers and exporters are not included in our analysis since they only appear in less than 0.01% of the transactions.


Hydrocodone			Oxycodone		
ndc	count	(%)	ndc	count	(%)
00406035705	10,910,767	0.062	00406051201	6,232,985	0.059
00591034905	6,812,588	0.039	00406055201	4,471,392	0.042
00406036601	5,999,312	0.034	00406052301	4,100,955	0.038
00406035801	4,775,253	0.027	60951071270	3,995,670	0.037
00406036701	4,550,834	0.026	00591093201	3,667,308	0.034
00406036501	4,448,582	0.025	59011042010	2,526,600	0.023
00406036301	4,061,444	0.023	00406853001	2,494,118	0.023
00591320201	4,053,442	0.023	10702005601	2,449,938	0.023
00591054005	3,913,581	0.022	00406052201	2,370,002	0.022
00591085305	3,749,532	0.021	00406051205	2,347,109	0.022
53014054867	3,569,557	0.020	59011044010	2,233,023	0.021
00591320301	3,353,134	0.019	00591093301	2,212,098	0.020
00591038505	3,325,400	0.019	59011010510	2,120,228	0.020
00591038705	3,322,107	0.019	00406851501	1,993,107	0.018
00406036001	3,284,559	0.018	59011010310	1,911,144	0.018
00603388721	3,193,247	0.018	53746020401	1,855,454	0.017
00406036005	3,127,428	0.017	59011041010	1,839,704	0.017
00406036101	3,103,683	0.017	59011048010	1,831,662	0.017
00603129558	2,747,435	0.015	59011010710	1,820,031	0.017
00603388728	2,491,943	0.014	00228287911	1,525,280	0.014
53746010901	2,480,188	0.014	00591082501	1,505,802	0.014
00472103016	2,381,028	0.013	00228287811	1,435,455	0.013
00406036705	2,354,949	0.013	60951079770	1,424,639	0.013
00406036201	2,354,279	0.013	60951070070	1,399,898	0.013
00093516101	2,292,773	0.013	00591074905	1,386,795	0.013
00406035805	2,237,094	0.012	00603499121	1,258,779	0.011
00591050301	2,094,122	0.012	52152021402	1,239,981	0.011
60432045516	2,035,496	0.011	59011043010	1,210,664	0.011
00603389021	2,034,056	0.011	59011010010	1,171,686	0.011
00603389121	1,957,034	0.011	00603499221	1,152,195	0.010
00591054001	1,915,460	0.010	10702001801	1,143,682	0.010
53746014501	1,817,595	0.010	52152021502	1,111,036	0.010
62175049064	1,805,494	0.010	59011046010	1,060,151	0.010

00121065516	1,718,466	0.009	00603499021	996,659	0.009
00406035701	1,641,846	0.009	00603499821	932,652	0.008
00591085301	1,568,519	0.009	00591082401	904,828	0.008
62037052401	1,455,052	0.008	00406051262	851,748	0.008
00591038501	1,429,288	0.008	00406055401	783,627	0.007
00603388821	1,417,980	0.008	53746020301	782,929	0.007
00591050305	1,264,323	0.007	63481062970	764,743	0.007
00591261205	1,108,056	0.006	60951060270	707,572	0.006
00603388732	1,107,293	0.006	10702005701	690,618	0.006
50383004316	1,084,475	0.006	58177004104	684,372	0.006
00603388128	1,083,364	0.006	00406058201	632,841	0.006
49884023533	1,079,695	0.006	57664022388	613,833	0.005
00406036505	1,066,333	0.006	00406053201	588,304	0.005
00406035901	1,041,860	0.005	00406055262	586,295	0.005
00406036305	981,502	0.005	60951079670	579,048	0.005
00074197314	849,987	0.004	00093003101	573,648	0.005
00591038701	849,143	0.004	00054368663	559,198	0.005
53746011001	844,058	0.004	10702000801	558,788	0.005
00603389128	825,323	0.004	10702000901	557,748	0.005
58177088107	776,063	0.004	63481062370	542,343	0.005
64376064801	744,735	0.004	00093003201	527,752	0.005
00603389028	716,042	0.004	68382079401	517,357	0.004
53746011005	709,098	0.004	00093003301	516,448	0.004
00591034901	699,897	0.004	00228298311	488,727	0.004
58177087707	665,867	0.003	00054465025	436,924	0.004
00591038801	655,307	0.003	59011041510	432,445	0.004
00591320205	653,703	0.003	59011086010	431,103	0.004
Tot	140,590,671	0.81		87,741,121	0.83

**Table A.2:** List of the top 60 Hydrocodone (left side) and Oxycodone (right side) products ranked by the number of transactions, along with the total number of transactions reported in both absolute and percentage values. These 60 products constitute approximately 80% of the total number of transactions for both opioids


## Appendix B

# Supplementary Material to Chapter 3



**Figure B.1:** The red line on the chart shows the annual volume of Hydrocodone (left) and Oxycodone (right) drugs supplied to consumers through the distribution network (indirect process). In comparison, the blue line represents the volume of these drugs sold directly from manufacturers to consumers without the use of a distribution network. The majority of the volume is seen to pass through the distribution network.





**Figure B.2:** The figure shows the in-degree and out-degree distributions of nodes that lost a link in the empirical case (represented by the blue line) and in the random model (represented by the red line). As we discussed in Chapter 3, there are no significant differences between the empirical data and the simulation data points. We have confirmed this using a KS-test, and the results are reported in the table below.

Distribution	D(KS)	p-value
In-degree	0.016	0.087
Out-degree	0.021	0.010

**Table B.1:** Results from the KS-test that estimates the similarity between the (in and out) degree distributions of the nodes that lost a link in the empirical case and under the random model.

## Appendix C

# Supplementary Material to Chapter 9

**Markovian transition matrix under supply substitution.** In condition of perfect substitution we assume that the order placed by two interacting firms,  $i$  and  $j$  are independent. Specifically, we assume that the orders placed by  $i$  to  $j$  are independent from the the orders place by  $j$  to  $k$ , i.e.,  $P(i \rightarrow j \rightarrow k) = P(i \rightarrow j)P(j \rightarrow k)$ . Assuming such independency of the orders, we obtain for the second-order Markovian transition matrix :

$$T^{(2)\text{Mkv}}_{(i,j)(j,k)} \propto T_{(i,j)}T_{(j,k)} \quad (\text{C.1})$$

where  $T_{(i,j)}$  indicates the probability that  $i$  places orders to  $j$  and  $T_{(j,k)}$  indicates the probability that  $j$  places orders to  $k$ . To make  $T^{(2)\text{Mkv}}$  row-stochastic matrix, as  $T^{(2)\text{Mkv}}$ , we have to normalize it. We write  $T^{(2)\text{Mkv}}_{(i,j)(j,k)} = AT_{(i,j)}T_{(j,k)}$  and we compute the normalizing constant:

$$\sum_{(j,k)} T^{(2)\text{Mkv}}_{(i,j)(j,k)} = 1 = AT_{(i,j)} \sum_{(j,k)} T_{(j,k)} \quad (\text{C.2})$$

$$\Rightarrow A = \frac{1}{T_{(i,j)} \sum_{(j,k)} T_{(j,k)}} \quad (\text{C.3})$$

and hence, we have

$$\mathbb{T}^{(2)\text{Mkv}}_{(i,j)(j,k)} = \frac{T_{(j,k)}}{\sum_{(j,k)} T_{(j,k)}} = T_{(j,k)} \quad (\text{C.4})$$

From Eq. (C.4), we notice that the element  $\mathbb{T}^{(2)\text{Mkv}}_{(i,j)(j,k)}$  does not contain information on the node  $i$ . In other words, the orders placed by  $i$  to  $k$ , via  $j$ , do not depend on  $i$ , but only on the orders placed by  $j$  to  $k$ , i.e.,  $T_{(j,k)}$ .  $i$  has relaxed its preferences towards supplier  $k$  aligning them to its intermediary firm  $j$ .

# List of Figures

1.1	Schematic representation of a supply network. A single supply chain, consisting of one supplier, one manufacturer and two distributors, is highlighted by the red links in the figure. . . . .	4
1.2	Roadmap for reading. . . . .	9
2.1	Illustrative picture of a drug product uniquely identified by a National Drug Code ( <i>ndc</i> ), 11 digits long. The first five digits encode the manufacturer (labeller); the second four digits encode the active ingredient, strength, and formulation (drug); the last two digits encode the package size (package). The code 59011-440-10 identifies Oxycodone tablets in 40 mg (segment 440) that are produced by the labeller Purdue Pharma (segment 59011) and packed in boxes of 100 tablets (segment 10). Credits: the photo of the drug package has been taken from an online article published by the journal <i>The Guardian</i> . Photograph: George Frey/Reuters. . . .	15
2.2	Number of products (y-axis) per basic opioid (x-axis). . . . .	16
2.3	Number of new products (blue line) and products withdrawn from the market (red line) per year. . . . .	17

---

2.4	Box plots displaying the frequencies of the dyadic and triadic motifs. The sample comprises the distribution networks of 100 opioid producers in the US. . . . .	22
3.1	Consumption trend for the fourteen opioids in the arcos. The amount of consumption is reported in millions of grams.. . . .	31
3.2	Percentage variation of opioid consumption (red line), number of consumers (blue line), and number of counties supplied (violet line). On the left: Hydrocodone. On the right: Oxycodone. . . . .	32
3.3	Percentage variation of Oxycodone and Hydrocodone consumption across states between 2006 and 2014. . . . .	33
3.4	Top: evolution of the number of nodes (blue line) and links (red line) of the Hydrocodone (left) and Oxycodone (right) distribution networks. Bottom: percentage variation in the number of nodes (blue line), number of links (violet line), and consumption (red line) for the two networks. Variations are measured with respect to the first year of observation. . . .	36
3.5	In-degree and out-degree distributions for the distribution network of Oxycodone and Hydrocodone in three years: the first year of observation (2006), the last year (2014), and the peak year (2010 for Oxycodone and 2011 for Hydrocodone). . . .	38
3.6	Top: in-degree (left plot) and out-degree distributions (right plot) of the nodes leaving the network in the empirical case (blue dots) and under the random model (red dots). Bottom: distribution of weights of links that disappear (blue dots) and that remain in the network (black dots). . . . .	41
3.7	Top: number of nodes per tier in the first year, the last year, and the peak year, that is 2010 for Oxycodone and 2011 for Hydrocodone. Bottom: violin plot of out-degrees of nodes positioned at different tiers. . . . .	44

4.1	Total number of trajectories (x-axis) and product demand (y-axis). Demand is measured as the total volume sold (in product units) to consumers between 2006 and 2014. Products are sorted by demand (from lowest to highest). . . . .	56
4.2	Probability distribution of trajectory frequency $f_p$ for the single product analysis (left plot), and the multiple-product analysis (right plot). . . . .	57
4.3	Probability distributions of the number of sequential distributors (i.e., number of tiers) traversed by packages containing the same product (left plot) and different products (right plot). . . . .	59
4.4	The orange line shows the average number of tiers (a), travel distance (b) and transit time (c) per product (x-axis). The orange-filled regions show the sample variance. . . . .	60
4.5	Probability distribution of time in stock of packages that traverse: two tiers of distributors (a), and three tiers of distributors (b). The colors specify the $k^{\text{th}}$ tier the distributor belongs to. . . . .	62
4.6	Tier zero indicates the producer, tier one the first distributor, and so on. The height of the bars indicates the amount shipped by each actor at a given tier. Color distinguishes trajectories of different lengths. . . . .	63
4.7	The y-axis shows the Pearson (in blue) and Kendall-tau coefficients (in dark violet) resulting from comparing the static centrality ranking with the one that incorporates the observed flow. The test has been performed on 200 different networks. They differ in the number of products shipped through them. The x-axis shows the number of products shipped by every network. This is a progressive number from 1 to 200. . . . .	67

5.1	Schematic representation of network structures with high, medium and low centralization (top row); and with low, medium and high level of multi-sourcing (bottom row). . . . .	74
5.2	Synthetic networks generated through model simulations with $\alpha=1$ , $q_t=0$ and: $q_s=1$ (figure a); $q_s=0.7$ (figure b), $q_s=0.3$ (figure c). Nodes' size is proportionate to their out-degree. The red node is the root representing the manufacturer. Orange nodes represent distributors and consumers (leaves). . . . .	79
5.3	Synthetic networks generated through model simulations with $q_s=0.5$ and: $\alpha=1$ , $q_t=0$ (figure a); $\alpha=0.6$ , $q_t=0.8$ (figure b). Nodes' size is proportionate to their in-degree. The red node is the root representing the manufacturer. Orange nodes represent distributors and consumers (leaves). . . . .	80
5.4	Parameter space illustrated via a 3-dimensional (left plot) and a 2-dimensional color map (right plot). The fitting score for each pair $(q_s, q_t)$ in the parameter space is visualized through the color schema on the right-hand side. . . . .	84
5.5	Probability distribution of out-degrees (a) in-degrees (b), path lengths from root to any node (c), and from root to leave nodes (d). The blue color indicates real-world data, whereas the red one indicates simulated data (mean and error bars from 200 runs).. . . . .	86
6.1	Empirical growth rates, $g$ , of manufacturers (a) and distributors (b) in the opioid distribution network. The star symbols indicate the empirical distribution of growth rates. The dashed line indicates the Laplacian fit. We binned the data for visualization purposes only and fitted the theoretical distribution to the non-binned data. . . . .	93

7.1	Left side: probability distribution of link weights in the year 2008 (blue circles), and log-normal fit (red dashed line). Right side: probability distribution of growth rates, pooling the years 2007, 2008, and 2009, and Laplacian fit (blue dashed line). . . . .	104
7.2	Distance (y-axis) between the empirical distribution of growth rates and the one obtained from computer simulations for different $k$ values (x-axis) and assessed using the KS-test. $k \in [1, 4]$ with an interval of 0.25. . . . .	110
7.3	Empirical (blue circles) and simulated (red diamonds) distribution of growth rates. Left side: results from the model fed with $k=1$ , i.e., no upstream constraint. Right side: results from the model fed with $k = k^{opt}=1.50$ , namely with the upstream constraint at its optimal value. . . . .	111
7.4	Growth rate as function of firm size for empirical data (left plot) and simulated data (right plot). The dashed line indicates the linear fit as expressed by Eq. (7.16) and performed via an ordinary least squares method. . . . .	113
7.5	Growth rate of firms as a function of their tier in empirical data (blue dots) and synthetic data (red dots). The error bar in each box indicates the standard deviation of the sample representing the growth rate volatility of each tier. . . . .	114



- 8.1 Schematic representation of the upstream and downstream shock propagation in a distribution network. The bars on the side of the nodes represent indicative stock levels. Entirely blue bars represent inventories that are at their target level. Blue-red bars represent inventories that are affected by shortages. Entirely red bars represent empty inventories. Firms experiencing shortages (i.e., affected by the shock) are colored in red, as well as the links through which the shock propagates. The illustrative example considers two manufacturers, *A* and *B*, and three distributors, *C*, *D* and *E*, in the network. At a given time  $t^*$ , a shock hits the inventory of manufacturer *A*. A supply shortage propagates downstream (red line), affecting distributor *C*. As a result, distributor *C* increases its orders and triggers an upstream propagation (red line). This upstream propagation leads the inventory of *B* to a faster depletion, thus generating a second cascade that moves downstream to consumers. . . . . 122
- 8.2 Inventory dynamics for manufacturer *A*, distributor *C*, and manufacturer *B*. The network and the shock scenario are explained in Fig. 8.1. The shock hits the manufacturer *A*, propagates downstream, affecting the distributor *C*, and upstream, affecting the manufacturer *B*. In the plots on the right-hand side, the blue lines indicate the inventory dynamic in a shock scenario and the black lines in a no-shock scenario. The dashed red lines mark when the manufacturer *A* goes into stock-out in the shock scenario. . . . . 129
- 8.3 Chord diagrams showing the dependencies between opioid producers under cascading supply shocks. The three diagrams display the results obtained at three different time steps of the simulations, i.e.,  $T = 90$ ,  $T = 150$ ,  $T = 240$ . . . . 133

- 8.4 Evolution of the variation in supply (black solid line) during the simulation. On the left y-axis we show its absolute value and on the right y-axis we show its value normalized to the average direct loss. The blue-filled area indicates the quantity shipped to offset the shortage. The red-filled area indicates the supply loss due to the cascade. . . . . 135
- 9.1 First-order and second-order network representation of a distribution system comprising two manufacturers, *A* and *B*, and three distributors, *C*, *D* and *E*. The empirical trajectories are two and depicted in blue (left side of the picture). These are preserved in the second-order representation. In contrast, in the first-order representation, we obtain two more trajectories, not observed in the data. These are depicted in red. . . . . 142
- 9.2 Three representations of the same distribution network under three different response strategies governed by (i) zero flexibility,  $\phi = 0$  (left column), (ii) medium flexibility  $\phi = 0.5$  (middle column), (iii) full flexibility  $\phi = 1$  (right column). First-order network representation (top row), second-order representation (middle row), second-order transition matrix (bottom row). The trajectories observed in a day-to-day business are colored in blue. The trajectories formed as result a of flexibility are colored in red. . . . . 148
- 9.3 Second-order representation of the Oxycodone distribution system in 2012. In this representation, a node represents a supply link between two firms, and an edge represents a trajectory of length two. Blue edges indicate trajectories that were empirically observed. Red edges indicate trajectories that were not observed, yet possible, given the network topology. Increasing the value of the parameter  $\phi$  in this representation, from  $\phi=0$  to  $\phi=1$ , increases the probability that these red trajectories become available for the distribution system. . . . . 153

- 9.4 Change to the system (figure a) and demand deficit (figure b) as a function of system flexibility,  $\phi$ . The efficient frontier (figure c) is drawn from a collection of 100 points, representing a  $\phi$  value each. Low values are on the bottom-right, high values on the upper-right. . . . . 155
- 9.5 Demand deficit (left plot) and efficient frontier (right plot) as obtained after 40, 50, and 60-time steps since the simulated shock hits the distribution network. The red point indicates the flexibility value that leads to the minimum deficit. The star symbol indicates a given flexibility value that moves from the efficient to the inefficient set.. . . . 158
- B.1 The red line on the chart shows the annual volume of Hydrocodone (left) and Oxycodone (right) drugs supplied to consumers through the distribution network (indirect process). In comparison, the blue line represents the volume of these drugs sold directly from manufacturers to consumers without the use of a distribution network. The majority of the volume is seen to pass through the distribution network. . . . . 191
- B.2 The figure shows the in-degree and out-degree distributions of nodes that lost a link in the empirical case (represented by the blue line) and in the random model (represented by the red line). As we discussed in Chapter 3, there are no significant differences between the empirical data and the simulation data points. We have confirmed this using a KS-test, and the results are reported in the table below. . . . 192

# List of Tables

2.1	List of the fourteen basic opioids in the data, along with the number of shipping transactions and units sold to consumers. Absolute and percentage values are displayed. Opioids are sorted according to the number of transactions. The top two are highlighted in red. . . . .	16
2.2	Results from the OLS regression to estimate the causal effect of the new product introduction. The model specifications are expressed by Eq. (2.1). The motifs' count is the outcome variable. . . . .	26
3.1	Density of the Oxycodone and Hydrocodone distribution networks in the nine-year observation period. . . . .	36
3.2	Mean value, standard deviation, third and fourth moment of the out-degree and in-degree distributions for the Oxycodone and Hydrocodone distribution networks. . . . .	39
3.3	Distances, $D(KS)$ , and $p$ -values obtained from the KS-test that compares the out and in-degree distributions: (i) in the first year and in the peak year; (ii) and in the last year and in the peak year. . . . .	39

3.4	Results from the KS-test evaluating the similarity between the degree distributions of nodes leaving the network in the empirical case and in the random model. . . . .	41
4.1	Pearson coefficients and the Kendall-Tau rank correlation coefficients, along with the corresponding $p$ -values, which resulted from comparing the static centrality ranking with the one that incorporates the observed flow. As a test case for the comparison, we use the distribution network of Oxycodone drugs. . . . .	66
5.1	Name, <i>dea-number</i> of the top six opioid manufacturers, as well as the number of nodes in the distribution networks reconstructed, the optimal parameters $\{\alpha, q_s, q_t\}$ , and the fitting score.. . . .	83
6.1	On the right: Optimal parameter values ( $\mu$ and $b$ ) resulted from the Gaussian and Laplacian fit the empirical growth rates. On the left: distances and $p$ -values from the KS-test compare the similarity between the theoretical and empirical distributions. . . . .	94
6.2	Results from the OLS regressions. The coefficients with $p$ -values smaller than 0.01 are reported in bold character. . . . .	97
7.1	$p$ -values obtained from the KS-test that compare the statistical similarity of $P(g)$ in different years. Not all the years of observation appear in this Table, but only the ones for which a $p$ -values $> 0.01$ has been obtained. . . . .	104
7.2	On the left side: mean, standard deviation, skewness, and kurtosis of three distributions of growth rates: empirical (first row), simulated with $k = k^{\text{opt}}$ (second row); simulated with $k = 1$ (third row). On the right side: distances and $p$ -values obtained from the KS-test between empirical distribution and simulated ones. . . . .	111

---

8.1	Name, production (product units), city, and <i>dea-number</i> of the top 20 opioid (Oxycodone) producers in 2008. . . . .	132
A.1	Business activities listed in the arcoss dataset. We group them into three categories: manufacturers, distributors, and “consumers”. The consumers are those who only receive goods and are not classified as either manufacturers or distributors. Importers and exporters are not included in our analysis since they only appear in less than 0.01% of the transactions. . . . .	188
A.2	List of the top 60 Hydrocodone (left side) and Oxycodone (right side) products ranked by the number of transactions, along with the total number of transactions reported in both absolute and percentage values. These 60 products constitute approximately 80% of the total number of transactions for both opioids . . . . .	190
B.1	Results from the KS-test that estimates the similarity between the (in and out) degree distributions of the nodes that lost a link in the empirical case and under the random model. . . . .	192