

RÉPUBLIQUE ALGÉRIENNE DÉMOCRATIQUE ET POPULAIRE
MINISTÈRE DE L'ENSEIGNEMENT SUPÉRIEUR ET DE LA
RECHERCHE SCIENTIFIQUE

ÉCOLE NATIONALE POLYTECHNIQUE



المدرسة الوطنية متعددة التقنيات
Ecole Nationale Polytechnique



GE HealthCare

Département de Génie Industriel

End of Studies Project Dissertation

In partial fulfillment of the requirements for the State Engineer's Degree in
Industrial Engineering, specialization in Industrial Management

Improving Supply Chain Resilience and Delivery Performance in
the Medical Device Industry: A Ripple Effect Simulation and
AI-Based Decision Support Approach

BOUDJELLAL Zakaria Abdelbassat & BOUSLIMANE Housseem Eddine

Supervised by **Dr. ZOUAGHI Iskander**

Presented and publicly defended on June 13, 2026

Jury composition:

Chairman:	Dr. BELDJOUDI Samia	MCA	ENP
Examiner:	Dr. BARECHE Imene	MCB	ENP
Supervisor:	Dr. ZOUAGHI Iskander	MCA	ENP
Supervisor:	Mr. HAZI Mohamed	Logistics Manager	GEHC Algeria
Guest:	Mr. IZEM Samy	Country Manager	GEHC Algeria

RÉPUBLIQUE ALGÉRIENNE DÉMOCRATIQUE ET POPULAIRE
MINISTÈRE DE L'ENSEIGNEMENT SUPÉRIEUR ET DE LA
RECHERCHE SCIENTIFIQUE

ÉCOLE NATIONALE POLYTECHNIQUE



المدرسة الوطنية متعددة التقنيات
Ecole Nationale Polytechnique



GE HealthCare

Département de Génie Industriel

End of Studies Project Dissertation

In partial fulfillment of the requirements for the State Engineer's Degree in
Industrial Engineering, specialization in Industrial Management

Improving Supply Chain Resilience and Delivery Performance in
the Medical Device Industry : A Ripple Effect Simulation and
AI-Based Decision Support Approach

BOUDJELLAL Zakaria Abdelbassat & BOUSLIMANE Housseem Eddine
Supervised by **Dr. ZOUAGHI Iskander**

Presented and publicly defended on June 13, 2026

Jury composition:

Chairman:	Dr. BELDJOUDI Samia	MCA	ENP
Examiner:	Dr. BARECHE Imene	MCB	ENP
Supervisor:	Dr. ZOUAGHI Iskander	MCA	ENP
Supervisor:	Mr. HAZI Mohamed	Logistics Manager	GEHC Algeria
Guest:	Mr. IZEM Samy	Country Manager	GEHC Algeria

RÉPUBLIQUE ALGÉRIENNE DÉMOCRATIQUE ET POPULAIRE
MINISTÈRE DE L'ENSEIGNEMENT SUPÉRIEUR ET DE LA
RECHERCHE SCIENTIFIQUE

ÉCOLE NATIONALE POLYTECHNIQUE



المدرسة الوطنية متعددة التقنيات
Ecole Nationale Polytechnique



GE HealthCare

Département de Génie Industriel

Mémoire de projet de fin d'études

En vue de l'obtention du diplôme d'Ingénieur d'État en Génie Industriel,
spécialité Management Industriel

Amélioration de la Résilience de la Supply Chain et des
Performances de Livraison dans l'Industrie des Dispositifs
Médicaux : Une Approche par Simulation de l'Effet Ripple et Aide
à la Décision basée sur l'IA

BOUDJELLAL Zakaria Abdelbassat & BOUSLIMANE Housseem Eddine

Encadré par **Dr. ZOUAGHI Iskander**

Présenté et soutenu publiquement le 13 juin 2026

Composition du jury :

Président :	Dr. BELDJOUDI Samia	MCA	ENP
Examinatrice :	Dr. BARECHE Imene	MCB	ENP
Encadrant :	Dr. ZOUAGHI Iskander	MCA	ENP
Encadrant :	M. HAZI Mohamed	Logistics Manager	GEHC Algérie
Invité :	M. IZEM Samy	Country Manager	GEHC Algérie

الملخص

تُعد سلاسل إمداد الأجهزة الطبية عرضةً بشكل خاص للاضطرابات بسبب طبيعتها العالمية، والقيود التنظيمية التي تحكمها، واعتمادها الكبير على التوريد الدولي. تتناول هذه الدراسة ظاهرة انتشار الاضطرابات داخل سلسلة الإمداد الشاملة (Wing-to-Wing - W2W) الخاصة بشركة GE HealthCare الجزائر، وهي سلسلة إمداد تعتمد كلياً على الاستيراد، وتعمل وفق نظام الإنتاج حسب الطلب دون الاحتفاظ بمخزون احتياطي. وتعتمد الدراسة على منهجية مكونة من ثلاثة محاور رئيسية. يمثل المحور الأول في تطبيق منهجية Fuzzy ANP-Fuzzy TOPSIS لترتيب الاضطرابات المحددة مسبقاً وفقاً لتأثيرها واحتمالية حدوثها وقدرتها على الانتشار وإمكانية التحكم فيها، مع تطوير تمثيل احتمالي متمحور حول الطلبات من أجل نمذجة آليات انتشار الاضطرابات داخل السلسلة. أما المحور الثاني فيعتمد على تطوير نموذج محاكاة بالأحداث المنفصلة باستخدام AnyLogic لتحليل انتشار الاضطرابات وقياس تأثيرها على أداء التسليم. ويقدم المحور الثالث SC Copilot، وهو نظام دعم قرار قائم على الذكاء الاصطناعي يدمج نماذج اللغة الكبيرة مع مخرجات المحاكاة من أجل دعم تحليل السيناريوهات ومساعدة متخذي القرار. أظهرت النتائج وجود عدم تماثل واضح في تأثير الاضطرابات، حيث تسببت اضطرابات نقص المكونات في تأخيرات أكبر بكثير من تأخيرات المصانع نتيجة فرضها دورة استيراد كاملة إضافية. وتساهم هذه الدراسة في تقديم إطار منهجي لتحديد أولويات الاضطرابات والنمذجة الاحتمالية والتحليل القائم على المحاكاة لسلاسل الإمداد الحالية من المخزون، كما تبرز إمكانات الذكاء الاصطناعي الوكيلي في تعزيز مرونة سلاسل الإمداد وتحسين عملية اتخاذ القرار في قطاع الأجهزة الطبية.

الكلمات المفتاحية: مرونة سلسلة الإمداد؛ أثر التموج (Ripple Effect)؛ الأجهزة الطبية؛ Fuzzy ANP-TOPSIS؛ النمذجة الاحتمالية؛ المحاكاة بالأحداث المنفصلة؛ الذكاء الاصطناعي الوكيلي؛ دعم اتخاذ القرار.

Résumé

Les chaînes d'approvisionnement des dispositifs médicaux sont particulièrement vulnérables aux perturbations en raison de leur structure mondialisée, des contraintes réglementaires et de leur dépendance aux approvisionnements internationaux. Ce mémoire étudie la propagation des perturbations au sein de la chaîne logistique Wing-to-Wing (W2W) de GE HealthCare Algérie, un système à la commande, sans stock et entièrement dépendant des importations. Une méthodologie en trois piliers est développée. Le premier pilier applique une approche Fuzzy ANP–Fuzzy TOPSIS afin de hiérarchiser les perturbations préalablement identifiées selon leur impact, leur probabilité, leur effet de propagation et leur contrôlabilité, tout en établissant une représentation stochastique centrée sur les commandes pour formaliser les mécanismes de propagation. Le deuxième pilier développe un modèle de simulation à événements discrets sous AnyLogic afin de quantifier l'impact des perturbations sur la performance de livraison. Le troisième pilier introduit SC Copilot, un système d'aide à la décision basé sur l'intelligence artificielle intégrant des modèles de langage de grande taille et les résultats de simulation pour l'analyse de scénarios et la prise de décision managériale. Les résultats mettent en évidence une forte asymétrie du comportement des perturbations, les anomalies liées aux pièces manquantes générant des retards nettement plus importants que les retards d'usine en raison du déclenchement d'un second cycle complet d'importation. Cette étude propose un cadre méthodologique pour la priorisation des perturbations, la modélisation stochastique et l'analyse par simulation des chaînes sans stock, tout en démontrant le potentiel de l'intelligence artificielle agentique pour renforcer la résilience et la prise de décision dans le secteur des dispositifs médicaux.

Mots-clés : Résilience de la chaîne logistique ; Effet de propagation (Ripple Effect) ; Dispositifs médicaux ; Fuzzy ANP–TOPSIS ; Modélisation stochastique ; Simulation à événements discrets ; Intelligence artificielle agentique ; Aide à la décision.

Abstract

Medical device supply chains are particularly vulnerable to disruptions due to their globalized structure, regulatory constraints, and reliance on international sourcing. This dissertation investigates disruption propagation within the Wing-to-Wing (W2W) supply chain of GE HealthCare Algeria, an import-dependent, zero-inventory, make-to-order system. A three-pillar methodology is developed. The first pillar applies a Fuzzy ANP–Fuzzy TOPSIS approach to prioritize previously identified disruptions according to their impact, likelihood, propagation effect, and controllability, while establishing an order-centric stochastic representation of the supply chain to formalize disruption propagation. The second pillar develops a discrete-event simulation model in AnyLogic to quantify disruption impacts on delivery performance. The third pillar introduces SC Copilot, an AI-powered decision-support system integrating large language models with simulation outputs for scenario analysis and managerial decision-making. Results reveal a strong asymmetry in disruption behavior, where missing-item disruptions generate substantially greater delays than factory delays by triggering a complete second import cycle. The study contributes a methodological framework for disruption prioritization, stochastic modeling, and simulation-based analysis in zero-inventory supply chains, while demonstrating the potential of agentic AI to enhance resilience and decision-making in the medical device sector.

Keywords: Supply Chain Resilience; Ripple Effect; Medical Devices; Fuzzy ANP–TOPSIS; Stochastic Modeling; Discrete-Event Simulation; Agentic AI; Decision Support.

Acknowledgements

“In the name of Allah, the Most Gracious, the Most Merciful”

All praise and gratitude belong first and foremost to **Allah**, the Almighty, for granting us the strength, patience, and clarity of mind to carry out and complete this work. Whatever good is found in these pages is by His grace alone.

Before presenting this work, we would like to express our sincere gratitude to all the people who supported us and contributed, in one way or another, to the completion of this final-year project.

First and foremost, we would like to thank our academic supervisor, **Dr. Iskander Zouaghi**, for his guidance, patience, availability, and constant support throughout this project. Beyond this internship, we have had the privilege of learning from him throughout our years in the department. His commitment to his students and his willingness to help whenever needed have left a lasting impression on us.

We would also like to thank all the professors of the **Industrial Engineering Department at the École Nationale Polytechnique**. Over the past three years, they have provided us with the knowledge, skills, and mindset that have shaped both our academic journey and our professional aspirations.

Our sincere thanks go to **Mr. Mohamed Hazi**, our supervisor at GE HealthCare, for welcoming us into the company, for his guidance, patience, and continuous encouragement, and for ensuring that our internship was a valuable learning experience. His guidance, encouragement, and continuous support played a major role in the success of this work.

We would also like to thank the entire **GE HealthCare Algeria team**. From our first day, we were welcomed into a dynamic, supportive, and highly professional environment that encouraged learning, initiative, and continuous improvement. We are especially grateful to **Mr. Malek Ouaret**, **Mr. Bilel Mellouk**, and **Mr. Nassim Lakabi** for their advice, kindness, and readiness to share their experience with us.

We would like to extend our appreciation to the members of the jury very dear professors to us for taking the time to review and evaluate our work. We are equally thankful to **Mr. Samy Izem** for honoring us with his presence and for the remarkable leadership through which he has built and managed such an outstanding team.

Finally, we would like to thank the **École Nationale Polytechnique**, which has been much more than a place of study for us. Over the years, it has become a second home where we grew academically, professionally, and personally.

Zakaria and Houssein

Dedication

I dedicate this work, which represents five years of learning, maturity, challenges, and unforgettable memories, to the following people:

*To my **family**, both near and extended, thank you for your constant support, encouragement, and unwavering belief in me throughout this journey.*

*To the memory of my beloved **grandfather**, and to you, "Foufou". I know you would have enjoyed this moment more than anyone else. May you rest in paradise. Your memory continues to inspire me every day.*

*To **Polytech**, thank you for being my second home, my comfort zone, and the place where many of my dreams took shape. You have been everything I could have hoped for in a university, and more.*

*To my very first friends in this journey, **Assil** and **Abderahmane**, thank you for your companionship from the very beginning. I was fortunate to start this adventure alongside such distinctly talented people.*

*To **GH**, where I took my first steps and met people I am grateful to have known. **Amira** and **Wassim**, "the binôme that almost was," you both have been very dear to me. Thank you for being present and valuable from the start until the very end. **Khalil** and **Nesrine**, you have been much more than a source of help and inspiration. I cannot imagine how these years would have unfolded without your guidance and friendship.*

*To the group that shifted my prépa and my entire Polytech experience, **Vniverse**. I deeply appreciate every person I met there. It was much more than a group; it was a source of comfort, meaning, and unforgettable memories. My sincere affection goes to my cherished friends **Ascem**, **Khaled**, **Ibtihal**, **Chichi**, **Haithem**, **Madina**, and **Lydia**. You made my prépa easier, my days at Polytech more joyful, and my journey more meaningful. I am especially grateful to **Ascem** (W1) for bringing us together with such purpose.*

*To **Tarek**, **Maria** (W2), and **Amina**, your early-morning breakfasts remain among my fondest memories, providing energy that carried me far beyond the day itself. Each of you has been a true source of inspiration and support. Thank you.*

*To all the wonderful people I met through **IEC** and the **2025 Committee**, and through **VIC**, particularly the members of **OT Ignite 2024**, thank you for bringing purpose and unforgettable memories to my years at Polytech. More broadly, I am equally grateful to the many **IEC alumni** who guided me and strengthened my belief in the value of our network.*

To **Issam**, thank you for being a voice of wisdom, for your care, and your constant support.

To the younger people I was fortunate enough to meet throughout my journey at Polytech: **Manar, Pulad, Ahlam, Malik, Hadjer, Ramzi, Mokri**, and my most recent acquaintance, **Noured-dine**. Thank you for making Polytech a better place simply by being yourselves. Your kindness and energy have left a lasting mark on my experience.

To my classmates, “**Promo Sérieuse**,” thank you for being such a unique and fun group, and for all the help and memories you contributed throughout these years.

And finally, to you, **Housseem**.

Last but certainly not least, thank you for being such a remarkable companion throughout this journey. Thank you for the countless discussions, the passionate debates, and the friendship we built together. You have been one of the smartest and most interesting people I have had the pleasure to meet. I am grateful for everything we accomplished and learned together.

To all those mentioned, and to those whom I may have unintentionally forgotten, thank you for being part of this remarkable chapter of my life.

Zakaria Abdelbassat

Dedication

To you, two thousand years from now.

*To my **family**, who supported me through everything: my **father**, who has never stopped working so that I could fully dedicate my time to myself; my **mother**, whose love and care guided me through the world; my **sisters**, who were always ready to help with everything they could, whether financially or with something as simple as a hug when I needed it most; my **brother**, who was always there for me whenever I needed him; and my **uncles**, **aunts**, and **cousins**, with whom I'm happy to share the same blood.*

*To **L'Haouch**, the place where I grew up, and to **Kechafa**, where I spent ten years learning, growing, and becoming the person I am today.*

To my high school friends, the special four with whom I shared my final year. For their priceless care and help, I am forever in their debt.

*To my second home, **ENP**. All the teachers and leaders who believed in me, and to everyone who voted for me, lost with me in competitions, followed me, guided me, or simply spent time with me, thank you.*

*My prepa years with **G3** and **G4**. The **Indus** community and **IEC**. To those who came before me and guided me through my internships and studies. To my classmates, **Promo Faible** and **DATAs**, with whom I shared countless projects and presentations, and to those who came after me Jma3et **Domines w Hackathons**. in addition to **IEEE** and les **ELNs**, my second speciality, which gave me memories I will always cherish.*

*My professional experiences and the people who shaped them, **Abdou**, **Chicha**, **Ribouh**, **Diaa**, **Madjid**, **Moumen**, and **Yamat Dar**; Major the **916** and **ENP Project** team; **KABAS** and **HEETCH**, who placed their trust in me.*

*My biggest family, the **VIC**. To those who guided me as a new member. To **Polymaze OT 2023** and my Co-PMs **Nihed** and **Ibrahim** and **Idara**. To those who lived my presidency with me, my committee board: **Amira**, **Yasser**, **Rouji**, **Amina**, and **Lyna**; all the **PMs**; the **Respos**; the **Top 100**; **RDS** with **Issam**; **Memers**; **Mentoring**; **Council-26** and every **Vicer** ever.*

All the groups that accompanied me through the years:

***Vniverse** ; **Haitham**, **Madina**, **Lydia** and the prepa days , **Raouf**, **Imed**, **Djouzi**, **Nazim**, **Mehdi**, **Dark**, **Fleet** and ALL the **Vs** ;*

***Neitherlands** and my final days at ENP with **Maria**, **Tarek**, **Amina** et al.;*

***Village**; **Heaven**; **Tassili**; **Groupat l3cha** ;**Dzo**, and the countless memories my words can't fully express.*

To the little “**Stairs**”, my younger siblings: **Hadjer-Yanir**, **Maroua**, **Meriem**, and **Lyna**. Thank you for following me until the end, and I will forever be proud to have you.

To **Basset**, with whom I studied for five years and shared every course and internship, thank you.

To **Ascem**, with whom I loved to play, plan, and dream. Despite our differences, we always found our way back. Thank you.

To **Khaled**, with whom I’ve shared every task, small or gigantic, for five years, thank you.

To the person who was my companion through all these years, thank you, **Ibtihal**.

To the young me, who made that decision back in 2018 he will never regret it.

Tch-nipa7 will live freely forever in the lives of the people he has met.

To the person I am today, and to whatever the future holds, thank you.

Housseem Eddine "Tch-Nipa7"

Table of contents

List of Tables

List of Figures

List of acronyms

General Introduction	19
1 State of Play	22
1.1 Introduction	22
1.2 The Medical Devices Sector: Global and Algerian Context	22
1.2.1 The Global Medical Devices Market	22
1.2.2 The Algerian Medical Devices Market	24
1.3 Presentation of GE HealthCare	25
1.3.1 GE HealthCare Worldwide	25
1.3.2 GE HealthCare Algeria	27
1.3.3 Supply Chain	29
1.4 W2W Process Mapping and International Network	30
1.4.1 BPMN Mapping of the W2W Process	30
1.4.2 Geographic Mapping of the Upstream Network	32
1.4.3 Disruption Taxonomy	34
1.5 Quantitative Diagnosis of Operational Vulnerabilities	36
1.5.1 Operational Performance of the W2W Process	36
1.5.2 Diagnosis of the Customs Clearance Process	39
1.5.3 Analysis of Geographic Sourcing Concentration	41
1.5.4 Temporal Analysis of Trends and Seasonality (2022–2025)	42
1.6 Diagnostic Summary and Problem Statement	45
1.6.1 Summary of Identified Vulnerabilities	45
1.6.2 Gap Analysis	45
1.6.3 Problem Statement	46
2 State of The Art	47
2.1 Supply Chain Management and Medical Device Supply Chains	47
2.1.1 Foundations of Supply Chain Management	47
2.1.2 Specificity and Structure of Medical Device Supply Chains	48
2.1.3 Import Dependence, Component Complexity, and Regulatory Constraints	49
2.1.4 Project-Based Delivery and Installation Synchronisation	50
2.2 Supply Chain Disruptions and the Ripple Effect	51
2.2.1 Nature, Definition, and Conceptual Properties of the Ripple Effect	51
2.2.2 Root Causes, Structural Dynamics, and Control Layers	52
2.2.3 The Bullwhip Effect, the Efficiency Trade-off, and Boundaries of the Ripple Effect	53
2.2.4 Ripple Effect in Healthcare and Medical Device Supply Chains	55
2.2.5 Methodological Approaches to Ripple Effect Analysis	55

2.3	Supply Chain Resilience	56
2.3.1	Conceptual Evolution of Supply Chain Resilience	56
2.3.2	Resilience Capabilities and Performance Indicators	56
2.3.3	From Resilience Capabilities to Managerial Action	57
2.4	Visibility, Digital Decision Support, and Agentic AI for Resilience	58
2.4.1	Visibility as a Condition for Resilience	58
2.4.2	From Passive Monitoring to Active Decision Support	58
2.4.3	Digital Twins, Agentic AI, and Governance of AI-Based Decision Support	58
2.5	Multi-Criteria Decision-Making and Discrete Event Simulation as Methodological Pillars	60
2.5.1	Positioning and Necessity of the Multi-Criteria Approach	60
2.5.2	Selection of the Fuzzy AHP–TOPSIS Hybrid	60
2.5.3	Discrete Event Simulation: Foundations and Paradigm Selection	61
2.5.4	DES Application and Software Platform	62
2.6	Research Gap and Conceptual Framework	63
2.6.1	Synthesis of the Literature	63
2.6.2	Identification of the Research Gap	64
2.6.3	Conceptual Framework of the Study	64
3	A Multi-Method Modelling Framework for Supply Chain Disruption Analysis	66
3.1	Introduction	66
3.2	Modelling GE HealthCare Algeria’s Supply Chain	67
3.2.1	Contextualizing the Ripple Effect Framework for GE HealthCare Algeria	67
3.2.2	The Structural Reframe: Orders as the Unit of Analysis	67
3.2.3	Network Representation of GE HealthCare Algeria’s Supply Chain	67
3.2.4	The Date-Based Performance Framework	68
3.3	Prioritization of Supply Chain Disruptions through a Multi-Criteria Decision Making Approach	69
3.3.1	Framing the Problem and Method Selection	69
3.3.2	Evaluation Criteria and Expert Panel	70
3.3.3	Results: Weights, Ranking, and Simulation Scope	72
3.3.4	Disruption Subnetwork Mapping and Network Reduction	75
3.4	Mathematical Formalisation	77
3.4.1	The Order as a Stochastic Flow Object	77
3.4.2	Pipeline Stage Formalisation	78
3.4.3	JIT Scheduling Logic	80
3.4.4	D3 Formalisation: Factory Delay	81
3.4.5	D1 Formalisation: Missing Parts and Recovery Loop	81
3.5	Conclusion	82
4	Simulation-Based Analysis of Disruption Propagation	83
4.1	Introduction	83
4.2	Simulation Model Architecture	83
4.2.1	Model Overview and Agent Definition	84
4.2.2	Resource Pool Specification	84
4.2.3	Stochastic Service Functions	84
4.2.4	Key Design Decisions	85
4.2.5	Output Collection and Performance Metrics	86
4.3	Design of Experiments	87
4.3.1	Experimental Framework	87
4.3.2	Design Principles	88
4.3.3	Super-Additivity Ratio	88
4.3.4	Run-Length and Sufficiency Justification	89

4.3.5	Experimental Plan	89
4.3.6	Response Variables	90
4.4	Simulation Results and Analysis	91
4.4.1	Overview of Experimental Outcomes	91
4.4.2	OTD Rate, Mean Delay, and Duration Sensitivity	91
4.4.3	Node Sojourn Times, Ripple Deltas, and Attenuation Factors	94
4.4.4	Disruption Intensity Effects and SAR	96
4.4.5	Modality-Level Analysis	98
4.4.6	Statistical Validation	100
4.4.7	Summary of Findings	101
4.5	Conclusion	102

5 SC Copilot: Design and Implementation of an AI-Powered Decision Support Layer for Supply Chain Simulation 104

5.1	Introduction	104
5.2	Contextual Positioning of Simulation and Agentic Decision Support	104
5.2.1	The Gap Between Simulation and Decision	104
5.2.2	Agentic AI Systems	105
5.3	Business Objectives and Design Requirements	105
5.3.1	Operational Context and Stakeholder Needs	105
5.3.2	Formalised Design Requirements	106
5.4	System Architecture	106
5.4.1	Overview	106
5.4.2	Technology Stack	107
5.4.3	File-Based Integration Layer	108
5.5	Agent Components	109
5.5.1	Component 1 – Natural Language Scenario Interpreter	109
5.5.2	Component 2 – Pipeline Orchestrator	112
5.5.3	Component 3 – Data Processing and KPI Engine	114
5.5.4	Component 4 – LLM Analysis and Conversation Engine	116
5.5.5	Component 5 – Executive Dashboard (Streamlit UI)	117
5.6	Implementation Considerations and Engineering Decisions	118
5.6.1	Prompt Engineering Constraints	118
5.6.2	API Rate Limiting and Fallback	119
5.6.3	Environment Configuration	119
5.6.4	Cross-Platform Considerations	119
5.7	Evaluation Against Design Requirements	119
5.7.1	Functional Assessment	119
5.7.2	Value Contribution	120
5.8	Conclusion	121

General Conclusion 122

Bibliography 125

Appendices 131

List of Tables

1.1	HHI Calculation by Country of Origin	41
2.1	Comparative analysis of MCDM methods (adapted from Hsu et al. 1)	60
3.1	Expert panel composition.	71
3.2	Criteria weights derived from Fuzzy ANP (Chang’s extent analysis).	72
3.3	Final ranking of disruptions by Fuzzy TOPSIS closeness coefficient (CC).	73
3.4	Operative simulation pipeline: 4 processing stages and 3 transfer segments.	77
3.5	Processing stage definitions: capacity and service distributions.	78
3.6	Transfer segment distributions: normal path vs. D1 recovery path.	79
4.1	Full experimental plan: 15 simulation runs	90
4.2	KPI summary across all 15 experiments	91
4.3	Estimated inter-stage attenuation factors $\hat{\alpha}_{ij}$ by scenario.	96
4.4	Statistical test results with effect sizes	100
5.1	SC Copilot technology stack.	108
5.2	Simulation parameter space managed by the NL Scenario Interpreter.	109
5.3	Key columns of the <code>orders_tab.csv</code> data artefact produced by AnyLogic.	115
5.4	Evaluation of SC Copilot against design requirements.	120
6.1	GE HealthCare, Key Historical Milestones	131
6.2	Core conceptual properties of supply chain disruption theory	132
6.3	Root causes of the ripple effect: structural and operational dimensions	133
6.4	Ripple effect vs. bullwhip effect	134
6.5	The four-layer disruption control framework [2]	135
6.6	Methodological streams for ripple effect analysis [2–4]	135
6.7	Extended fuzzy Saaty scale — conversion of verbal judgements into TFNs	137
6.8	Comparative properties of simulation paradigms for supply chain disruption analysis	142
6.9	Comparative analysis of DES software platforms for supply chain disruption modelling	142
6.10	Service duration functions and distribution families	145
6.11	Order agent variable definitions	145
6.12	Resource pool configuration	146
6.13	Structure of the <code>orders_tab</code> output CSV	147
6.14	Structure of the Fuzzy ANP + Fuzzy TOPSIS Python Script	149
6.15	AI agent contribution to supply chain visibility dimensions	150

List of Figures

1.1	Medical Devices Market Size 2025 to 2035 (USD Billion) [5]	23
1.2	Global Top 10 Medical Device Manufacturers by Revenue (2025) [6]	23
1.3	Global Medical Devices Market Share by Type [7]	24
1.4	GE HealthCare Product Segments	26
1.5	Geographic Presence of GE HealthCare (From GE HealthCare Website)	27
1.6	Company Organizational Chart (Internal Document)	27
1.7	GEHC Algeria Market Share Segmentation (Internal Document)	28
1.8	BPMN – Upstream Phase of the W2W	31
1.9	BPMN – Downstream Phase of the W2W	32
1.10	Geographic Mapping of GE HealthCare Algeria’s Upstream Network	33
1.11	GE HealthCare Algeria’s Distribution Network from the Khemis El Khechna Warehouse	34
1.12	Disruption Interaction Heatmap of the Supply Chain	36
1.13	Diagram of the Backward Calculation Logic of the W2W Process	38
1.14	Median Lead Time by Modality (in weeks)	38
1.15	Lead Time Distribution by Modality (Box Plot)	39
1.16	Coefficient of Variation of Lead Time by Modality	39
1.17	Distribution of Customs Clearance Lead Times	40
1.18	Breakdown of Transport Modes for Equipment Importation	40
1.19	Sourcing Share and HHI Contribution by Country	42
1.20	Temporal Evolution of Delivery Lead Times by Modality (2022–2025)	44
2.1	Medical devices supply chain. [8]	49
2.2	Disruption propagation in the supply chain [3]	51
2.3	Bullwhip and ripple effects [9]	54
2.4	Functional architecture of an LLM-based autonomous agent. [10]	59
2.5	Framework for investigating the ripple effect on the supply chain with the help of simulation research methodology [9]	63
3.1	Hierarchical organization of candidate criteria derived from the literature review.	71
3.2	Final criteria weights from Fuzzy ANP.	72
3.3	Closeness coefficients (CC) for the thirteen disruptions.	73
3.4	Distribution of disruptions in the d^+/d^- space. Disruptions in the upper-left zone (low d^+ , high d^-) are the most critical.	74
3.5	Heatmap of averaged expert evaluations.	74
3.6	Disruption subnetwork for D1 and D3 mapped onto the full GE Algeria network.	76
4.1	AnyLogic model canvas: full process flow for the normal path and the D1 recovery sub-path.	84
4.2	KPI heatmaps across all 15 experiments: OTD rate (top) and mean delay in days (bottom).	92
4.3	OTD rate versus disruption intensity across scenario families.	93
4.4	Sensitivity of OTD and mean delay to simulation duration under baseline and compound disruption conditions.	94

4.5	Mean sojourn time (days) per stage per scenario.	94
4.6	Ripple delta heatmap: signed deviation of mean sojourn time from No Disruption baseline, per stage and scenario.	95
4.7	Delay distributions $\Delta_k = \text{ATD}_k - \text{COSD}_k$ by disruption type.	97
4.8	Delay distribution by equipment modality and disruption scenario.	99
5.1	SC Copilot high-level architecture	107
5.2	Decision flowchart for the Natural Language Scenario Interpreter	111
5.3	UML Activity Diagram of the SC Copilot pipeline orchestrator across three execution domains	113
5.4	Layout of the SC Copilot executive dashboard	118
6.1	Stochastic input distributions for DES supply chain modelling [11]	141
6.2	An example of triangular probability density function [12]	141
6.3	Full supply chain network of GE HealthCare Algeria.	144

List of acronyms

- **ABM** : Agent-Based Modelling
- **AHP** : Analytic Hierarchy Process
- **AI** : Artificial Intelligence
- **ANP** : Analytic Network Process
- **ANPP** : Agence Nationale des Produits Pharmaceutiques
- **AOSD** : Actual Order Ship Date
- **API** : Application Programming Interface
- **BMD** : Bone Densitometry
- **BOM** : Bill of Materials
- **BPMN** : Business Process Model and Notation
- **CAGR** : Compound Annual Growth Rate
- **CI** : Consistency Index
- **COG** : Centre of Gravity
- **COSD** : Confirmed Order Ship Date
- **COVID-19** : Coronavirus Disease 2019
- **CR** : Consistency Ratio
- **CRD** : Customer Request Date
- **DES** : Discrete Event Simulation
- **ENP** : École Nationale Polytechnique
- **ERP** : Enterprise Resource Planning
- **FE** : Field Engineers
- **FMI** : Field Modification Instruction
- **FNIS** : Fuzzy Negative Ideal Solution
- **FPD** : Factory Promised Date
- **FPIS** : Fuzzy Positive Ideal Solution
- **GE** : GE HealthCare

- **GEHC** : General Electric HealthCare
- **HHI** : Herfindahl-Hirschman Index
- **IDE** : Integrated Development Environment
- **JIT** : Just-In-Time
- **JSON** : JavaScript Object Notation
- **KPI** : Key Performance Indicator
- **LLM** : Large Language Model
- **LSI** : Local Supplier Items
- **MAMMO** : Mammography
- **MCDM** : Multi-Criteria Decision Making
- **MRI** : Magnetic Resonance Imaging
- **MTO** : Make-to-Order
- **NM** : Nuclear Medicine
- **OECD** : Organisation for Economic Co-operation and Development
- **OFS** : Operations Factory Specialist
- **OTD** : On-Time Delivery
- **PCH** : Pharmacie Centrale des Hôpitaux
- **PCS** : Patient Care Solutions
- **PDR** : Parts & Depot Replenishment
- **PET** : Positron Emission Tomography
- **PJM** : Project Manager
- **PLE** : Personal Learning Edition
- **PMI** : Project Manager Installation
- **RAD** : Radiology
- **RAG** : Retrieval-Augmented Generation
- **RI** : Random Index
- **SAR** : Super-Additivity Ratio
- **SC** : Supply Chain
- **SCM** : Supply Chain Management
- **SIS** : Susceptible-Infected-Susceptible
- **SOSD** : Scheduled On-Site Date
- **SSSD** : Scheduled Ship Distribution Date

- **TFN** : Triangular Fuzzy Number
- **TOPSIS** : Technique for Order of Preference by Similarity to Ideal Solution
- **UI** : User Interface
- **UML** : Unified Modeling Language
- **USD** : United States Dollar
- **VASC** : Vascular
- **W2W** : Wing-to-Wing

General Introduction

The global medical devices industry has become one of the most strategically sensitive sectors of the modern economy. In developing markets where import dependence is structural and regulatory environments are still maturing, the uninterrupted supply of medical equipment is not merely a commercial concern. Any disruption to the supply chain translates directly into delayed diagnoses, deferred treatments, and degraded healthcare service delivery.

Algeria illustrates this challenge clearly. The country has invested significantly in expanding its public healthcare infrastructure over the past decade, driving sustained demand for advanced imaging equipment. Yet the market remains almost entirely import-dependent, with no local manufacturing capacity for modalities such as CT scanners, MRI systems, or PET machines. Every unit sold must travel from a global manufacturing site through a multi-stage international logistics chain, across Algerian customs, to a hospital site where field engineers complete installation. The absence of any inventory buffer means that a disruption at any single node propagates directly and visibly to the end customer.

It is within this context that GE HealthCare Algeria operates. As the local subsidiary of one of the world's leading medical technology companies, GE HealthCare Algeria manages the full wing-to-wing (W2W) delivery cycle from order receipt to on-site installation. The supply chain it operates is structurally different from inventory-holding chains. There is no finished-goods stock to absorb shocks, no production capacity to accelerate. Each order is an individual project processed through a sequence of actors spanning multiple countries and institutional environments. When something goes wrong, the consequences travel through the entire pipeline with no dampening mechanism.

This dissertation addresses how such disruptions propagate, how they can be prioritised and quantified, and how decision-makers can be equipped to manage them. The phenomenon at the centre of this analysis is the **ripple effect**: the downstream propagation of supply chain disruptions beyond their point of origin, amplifying delays and threatening network resilience. While extensively studied in manufacturing and retail supply chains, its dynamics in zero-inventory, make-to-order, import-dependent chains have received very limited attention. This study addresses that gap.

Research question

How do supply chain disruptions propagate through GE HealthCare Algeria's import-dependent, zero-inventory pipeline, which disruptions are the most critical, and how can decision-makers be supported in assessing and managing their impact on delivery performance?

Methodology

The thesis develops a three-pillar framework. The first pillar is a disruption modelling layer that begins by representing GE HealthCare Algeria's supply chain as a stochastic order-flow network adapted to its zero-inventory, import-dependent architecture, then applies a Fuzzy ANP–Fuzzy TOPSIS multi-criteria model – drawing on four GE Algeria expert panels and evaluating thirteen identified disruptions across four dimensions: global impact, probability and frequency, propagation effect, and controllability – to prioritise the scenarios that warrant simulation-based analysis. The second pillar is a discrete-event simulation built in AnyLogic, formalising the four-stage W2W pipeline as a stochastic order-flow system and tested across a fifteen-run Design of Experiments covering isolated and compound disruption scenarios to quantify ripple propagation, delivery reliability degradation, and super-additivity effects. The third pillar is SC Copilot, an agentic AI decision-support layer integrating a large language model with the simulation engine to deliver natural-language scenario interpretation, automated KPI analysis, and executive-grade recommendations to non-expert managers.

Thesis structure

This thesis is organised into five chapters, as illustrated in Figure ??.

Chapter 1 establishes the operational context through a state of play at GE HealthCare Algeria, mapping the W2W process, characterising the upstream global network, and presenting a quantitative diagnosis of operational vulnerabilities based on the company's and sales dataset.


Chapter 2 reviews the relevant academic literature, covering supply chain resilience, the ripple effect, MCDM methods, and discrete-event simulation, and identifies the research gaps motivating the present work.

Chapter 3 develops the full analytical apparatus: the order-based network representation, the Fuzzy ANP–Fuzzy TOPSIS prioritisation model, and the mathematical formalisation of the stochastic simulation pipeline.

Chapter 4 presents the simulation model, the Design of Experiments, and the empirical results, deriving eight thesis-level findings on asymmetric disruption sensitivity, compound effects, and modality-level vulnerability.

Chapter 5 introduces SC Copilot, detailing its architecture, five-component pipeline, LLM prompting strategy, and executive dashboard, evaluated against formal design requirements.

The thesis closes with a general conclusion summarising contributions, limitations, and directions for future research.



Structure of the dissertation.

Chapter 1

State of Play

1.1 Introduction

This chapter establishes the analytical foundation of the study. It opens with a contextual overview of the global and Algerian medical devices market, before presenting GE HealthCare as a company and its Algerian subsidiary. It then maps the Wing-to-Wing (W2W) supply chain process and the upstream sourcing network, identifies the thirteen disruptions affecting the chain, and closes with a quantitative diagnosis covering lead time performance, customs clearance behavior, geographic concentration, and temporal trends. Together, these elements converge on a structured problem statement that frames the analytical work developed in the subsequent chapters.

1.2 The Medical Devices Sector: Global and Algerian Context

1.2.1 The Global Medical Devices Market

The medical devices sector is one of the pillars of the global healthcare economy. Valued at approximately 640 billion US dollars in 2024, the global market is projected to exceed 1,140 billion dollars by 2034, growing at a CAGR of around 6% over the period [5], as illustrated in Figure 1.1. This growth is driven by the rising prevalence of chronic diseases, population aging, technological advances in medical devices, and increasing demand for minimally invasive procedures, as well as the integration of artificial intelligence and robotics into healthcare technologies [13].

Geographically, North America dominates the market, accounting for nearly 40% of global revenues in 2024 due to its advanced healthcare infrastructure and rapid technology adoption [5]. Europe represents the second-largest hub, supported by the presence of major manufacturers. In 2023, the United States remained the leading exporter of medical devices, with exports valued at 34.8 billion dollars and destined mainly for the Netherlands, China, Germany, Mexico, and Japan [14].




Figure 1.1: Medical Devices Market Size 2025 to 2035 (USD Billion) [5]

The market is highly concentrated and dominated by large American and European multinationals. Major companies include Medtronic, Johnson & Johnson, Siemens Healthineers, Abbott, Stryker, GE HealthCare, BD, Philips, Cardinal Health, and Boston Scientific [15]. As shown in Figure 1.2, the United States hosts several of the world's largest medical technology firms, with Abbott, Medtronic, and Johnson & Johnson among the leading companies by revenue [16].




Figure 1.2: Global Top 10 Medical Device Manufacturers by Revenue (2025) [6]

By product type, the market includes orthopedic, cardiovascular, diagnostic imaging, in vitro diagnostics (IVD), minimally invasive surgery, wound management, diabetes management, ophthalmic, dental, nephrology, and general surgery devices, among other categories. Figure 1.3 presents the relative weight of each segment, with IVD holding the largest share in 2025 [7].




Figure 1.3: Global Medical Devices Market Share by Type [7]

1.2.2 The Algerian Medical Devices Market

The Algerian medical devices market is characterized by a strong dependence on imports, which cover virtually all domestic needs. It results from the interaction between sustained growth in healthcare demand, a largely import-driven supply, and an evolving regulatory framework. In this context, the supply chain is heavily structured around a central public entity, the Pharmacie Centrale des Hôpitaux (PCH), which controls access to a significant share of the market, estimated at over 40% [17, 18].

Demand Dynamics and Evolving Healthcare Needs

Demand for medical devices in Algeria is underpinned by structural trends driven by demographic and epidemiological shifts. Life expectancy rose from 52.6 years in 1970 to 79.6 years in 2023, reflecting a relative aging of the population [19]. This shift is accompanied by an increase in chronic diseases, particularly cardiovascular, renal, and oncological conditions. According to the TAHINA national survey, hypertension and diabetes have prevalence rates of 25.87% and 12.89%, respectively [19]. This epidemiological shift generates growing demand for diagnostic, monitoring, and treatment devices, such as medical imaging equipment, cardiac monitors, dialysis devices, and diabetes management systems. The market is estimated at 1.39 billion dollars in 2024, with a CAGR of 3.83% through 2029, when it is expected to reach 1.68 billion dollars [20]. Public investment reinforces this dynamic: under the 2025 Finance Bill, an envelope of 1,004 billion dinars (approximately 7.52 billion dollars) was allocated to the health sector, representing a 19% increase over 2024, and 134 new public hospitals have been planned for 2024 [21, 22].

Supply Structure and Regulatory Framework

The supply of medical devices in Algeria relies primarily on imports, which account for approximately 98% of domestic needs, covering more than 129,000 types of devices with an estimated value of 630 million dollars in 2025 [23]. The main supplier countries are the United States (18%), China (14.5%), Germany (12.5%), France (11%), and Japan (10.5%) [17]. The local landscape is predominantly composed of more than 200 national operators, acting mainly as importers and distributors [24]. Public facilities procure primarily through centralized tendering processes managed by the PCH, while the private sector absorbs nearly 60% of volumes through more decentralized mechanisms [17, 18].

The regulatory framework, anchored in Law 18-11 of July 2, 2018, requires prior registration for any medical device issued by the Agence Nationale des Produits Pharmaceutiques (ANPP), with a statutory processing period of 240 days [25]. From September 7, 2025, all imported medical devices must undergo prior review by the ANPP before being placed on the market [26, 27]. A ministerial order adopted in October 2025 further makes the serialization of imported medical devices mandatory from 2027, through the use of Data Matrix codes compliant with international standards [28, 29].

1.3 Presentation of GE HealthCare

1.3.1 GE HealthCare Worldwide

Built on over a century of history in medical imaging, GE HealthCare was spun off from General Electric on January 3, 2023, and began trading on Nasdaq under the ticker symbol GEHC, becoming a fully independent entity incorporated in the State of Delaware and headquartered in Chicago, Illinois. This independence marked a strategic turning point, driving a transformation in the company's positioning: from equipment supplier to integrated healthcare solutions provider. In 2024, GE HealthCare generated revenues of 19.7 billion dollars, employed approximately 53,000 people across more than 160 countries, and continued to advance its *precision care* strategy, centered on personalized, connected, data-driven care.

GE HealthCare's operations are organized around four main segments. The Imaging segment encompasses MRI, CT scanning, digital radiology, mammography, and nuclear medicine, representing the company's historical core with over 8.9 billion dollars in annual revenues. The Ultrasound segment covers radiology and primary care equipment, women's health, cardiology, and point-of-care solutions. The Patient Care Solutions (PCS) segment includes patient monitoring equipment, anesthesia devices, cardiology equipment, and ventilators dedicated to critical care. The Pharmaceutical Diagnostics (PDx) segment encompasses contrast media and molecular imaging agents that complement the equipment of the other three segments. These four segments are illustrated in Figure 1.4.



Diagnostic Imaging (DI)



Ultrasound (US)



Patient Care Solutions (PCS)



Pharmaceutical Diagnostics (PDx)

Figure 1.4: GE HealthCare Product Segments

Global Supply Chain Organization

GE HealthCare conducts R&D activities in 18 countries and manufactures its products across 20 countries worldwide, serving customers in more than 160 countries, as reflected by the geographic footprint shown in Figure 1.5. For the MENEAT region (Middle East, Northeast Africa, and Turkey), Dubai serves as the regional commercial and logistics hub. For the VEGA zone, which covers Francophone African countries including Algeria, the consolidation hub of Saint-Witz in the Paris region handles order grouping, pre-shipment quality control, and documentation preparation before dispatch.



Figure 1.5: Geographic Presence of GE HealthCare (From GE HealthCare Website)

1.3.2 GE HealthCare Algeria

GE HealthCare operates in Algeria as a Société à Responsabilité Limitée (SARL) under Algerian law, headquartered in El Biar, Algiers. Its primary activity is the importation and distribution of medical imaging equipment, together with installation and after-sales services. The Algerian subsidiary provides direct representation of the group to institutional clients, including public hospitals and private clinics, and manages commercial, technical, and after-sales relationships across the national territory. GE HealthCare has been present in Algeria since 2016.

The organizational chart below (Figure 1.6) illustrates the company’s hierarchical structure and its main departments:




Figure 1.6: Company Organizational Chart (Internal Document)

The company’s operations span ten departments covering Finance, Human Resources, Legal, Logistics, Marketing, Project Management (PJM), Quality, Sales, After-Sales Service, and Treasury,

each playing a defined role in the execution of the W2W process.

Product Portfolio Marketed in Algeria

The portfolio marketed by GE HealthCare Algeria reflects the company’s global segments, with a focus on high-technology products meeting the diagnostic needs of Algerian healthcare facilities. According to COCIR data (April 2025 to March 2026 period), GE HealthCare holds an overall market share of **43%** in the Algerian high-technology medical devices market, making it the sector’s leading player.

In the Imaging segment, GE HealthCare Algeria offers the full range of diagnostic imaging modalities: CT scanners, MRI systems, digital radiology systems, mammography systems, and molecular imaging equipment. In this segment, GE HealthCare holds a **45%** market share, facing direct competition from Siemens Healthineers, Philips Healthcare, Canon Medical, and Mindray.

The Ultrasound segment represents GE HealthCare’s most prominent stronghold in Algeria. With a **55%** market share, the company holds an unchallenged leadership position, covering cardiac ultrasound, women’s health, primary care, and surgical guidance.

The Patient Care Solutions segment covers patient monitors, ICU monitoring equipment, defibrillators, and ventilation systems. With a **12%** market share in this segment, GE HealthCare remains behind its positions in Imaging and Ultrasound, making it a priority commercial development axis for the coming years. The segmentation of GE HealthCare Algeria’s market positions across these three product families is summarized in Figure 1.7.




Figure 1.7: GEHC Algeria Market Share Segmentation (Internal Document)

Customers and Markets Served

GE HealthCare Algeria operates across two structurally distinct market segments. The public sector constitutes the company’s primary market. It includes university hospital centers (CHU), specialized hospital facilities (EHS), public hospital establishments (EPH), and polyclinics, which

procure primarily through centralized public tenders issued by the PCH or directly by the facilities themselves. This channel is characterized by long decision cycles, extensive documentation requirements, strict technical acceptance procedures, and extended payment terms. It nonetheless represents significant volumes due to the hospital equipment programs supported by the Algerian state, notably the program to build 134 new public hospitals launched in 2024. The private sector comprises private medical clinics, diagnostic centers, specialist practices, and private polyclinics, which are growing rapidly in major Algerian cities, offering shorter decision cycles and more favorable payment terms.

1.3.3 Supply Chain

The Wing-to-Wing (W2W) Operating Model

GE HealthCare Algeria's operational process is built around the *Wing-to-Wing* (W2W) concept, which refers to the comprehensive management of the order lifecycle from initiation through to installation and commissioning at the end customer's site. In the Algerian context, this W2W process unfolds across several sequential and interdependent phases. The commercial phase covers customer needs identification, responses to public tenders or direct negotiation, solution configuration, and contract signing. The order management phase consists of entering and validating the order in the global information system, triggering production planning or shipment scheduling from the relevant manufacturing plants. The international logistics phase encompasses transportation from the manufacturing plant to a regional hub in France, followed by sea or air freight to Algeria, customs clearance, and temporary storage in the local warehouse. The installation phase mobilizes GE HealthCare Algeria's field engineering teams for on-site delivery, installation, technical qualification, and end-user training. The after-sales service phase covers preventive maintenance, service contracts, and ongoing technical support.

Inventory Policy and Transit Warehouse Management

GE HealthCare Algeria's inventory management policy is structurally determined by the nature of the equipment it markets. Heavy equipment, including CT scanners, MRI systems, and radiology systems, is produced under a make-to-order (MTO) model: each unit is manufactured to order, configured according to customer specifications, and shipped upon receipt of a firm purchase order. This model entails long cycle times, ranging from several weeks to several months depending on equipment complexity and manufacturing site, with no finished goods inventory held in Algeria for these product categories.

In contrast, GE HealthCare Algeria maintains a transit warehouse used to store delivered equipment awaiting installation, as well as spare parts, consumables, and accessories required for maintenance and after-sales service activities. This warehouse plays a critical buffer role within the chain, absorbing the gaps between equipment delivery and actual on-site installation. Warehouse operations are fully outsourced to ARAMEX, a third-party logistics (3PL) provider operating a dedicated facility in Khemis El Khechna, east of Algiers.

The facility comprises four specialized storage aisles: one dedicated to diagnostic imaging equipment, one to ultrasound systems, one to patient care solutions, and two to spare parts and maintenance components, providing a combined storage capacity structured across nine shelving levels. An on-site repair center, also operated by ARAMEX, supports refurbishment and maintenance activities for returned or defective units. ARAMEX operates under standard operating procedures defined by GE HealthCare and is bound by service level agreements governing handling quality,

response times, and end-to-end traceability.

This configuration, MTO for equipment and outsourced transit stock management for accessories and parts, exposes GE HealthCare Algeria to specific risks of service disruption and capital tie-up, which form the core subject of the analysis developed in the following chapters.

1.4 W2W Process Mapping and International Network

This section analyzes the Wing-to-Wing (W2W) supply chain process of GE HealthCare Algeria, maps the upstream network, and identifies the main disruptions. The approach is based on triangulation between eight semi-structured interviews conducted with key company stakeholders, including the Logistics Manager, Service & Installation Manager, Bank Facilitator, Technical Pharmacist, Commercial Head, Sales Executor, and Project Manager, and the analysis of internal documents. This enabled the modeling of the W2W process, the geographic mapping of the upstream network, and the construction of a taxonomy of disruptions affecting GE HealthCare Algeria.

1.4.1 BPMN Mapping of the W2W Process

The W2W process represents an end-to-end view of the supply chain, covering the entire operational cycle from the formalization of the customer need through to the final installation of the equipment. In the case of GE HealthCare Algeria, it is structured into two main phases, upstream and downstream, modeled using Business Process Model and Notation (BPMN) to represent the actors, flows, and decision points of the process in a standardized way.

Upstream Phase

The upstream phase, depicted in Figure 1.8, covers the process from the initial commercial opportunity through to the shipment of equipment from the manufacturing plant.




Figure 1.8: BPMN – Upstream Phase of the W2W

The process begins with a commercial opportunity, when a healthcare facility expresses a need for medical equipment. The Sales team qualifies the request, verifies technical feasibility, and sets a target date referred to as the *Customer Request Date*, that's after the finance team verifies the customer solvency as well.

Following contract signature, the Project Manager conducts a contract review and registers the order in the global information system. A Global Order Number (GON), a unique identifier assigned to each order within the GE HealthCare system, is generated, organizing the order into logical groupings known as the Equipment Set and the Fulfillment Set.

The order is then routed to production based on its type and transmitted to the relevant manufacturing plant. Two fulfillment modes are possible: a direct plant-to-customer flow, or routing through an intermediate distribution center. Manufacturing is launched and tracked against a defined Manufacturing Lead Time. Once the equipment is ready, it is shipped from the plant and shipping documents are transmitted to the local Algerian teams.

Downstream Phase

The downstream phase, shown in Figure 1.9, begins upon receipt of the shipping documents and extends through to customer acceptance and order closure.




Figure 1.9: BPMN – Downstream Phase of the W2W

The downstream phase begins upon receipt of the shipping documents, triggering the import process in Algeria. A regulatory check is performed regarding the product's registration with the Agence Nationale des Produits Pharmaceutiques (ANPP); if registration is absent, a registration sub-process is initiated, which can hold the order for up to eight months, representing one of the most critical bottlenecks in the chain.

In parallel, financial operations are initiated: bank domiciliation and the establishment of a letter of credit. Once these steps are validated, the customs file is submitted through the ALCES platform. Following documentary compliance review, duties and taxes are paid, enabling issuance of the Bon à Enlever (BAE) and physical release of the goods.

The equipment is then transferred to GE HealthCare Algeria's local warehouse, received, inspected, and entered into the information system. Compliance with contractual timelines is tracked by comparing the Confirmed Order Ship Date (COSD) with the Actual Order Ship Date (AOSD).


Finally, project teams organize on-site delivery, verify site readiness, and carry out installation, performed by Field Engineers (FE). The process concludes with customer acceptance, revenue recognition, and formal order closure.

1.4.2 Geographic Mapping of the Upstream Network


GE HealthCare Algeria's upstream network spans four manufacturing countries, reflecting the geographic specialization of the group's production sites by product line. Heavy imaging equipment, including CT scanners and radiology systems, is primarily produced in the United States, while MRI systems are manufactured in Austria. Ultrasound equipment and part of the patient care solutions originate from China, and certain equipment items and accessories are produced or consolidated in France.

All flows converge at the consolidation hub of Saint-Witz, in the Paris region, which handles order grouping, pre-shipment quality control, and documentation preparation before dispatch to Algeria. As illustrated in Figure 1.10, shipments from Saint-Witz to Algeria are made via two complementary modes of transport: **air freight**, preferred for urgent equipment, spare parts, and small-volume accessories; and **sea freight**, used for heavy and bulky equipment such as CT scanners and MRI

systems.



(a) Global view of international flows toward France (Saint-Witz hub)



(b) Flow view at the French hub level

Figure 1.10: Geographic Mapping of GE HealthCare Algeria's Upstream Network

Once on Algerian territory, equipment and spare parts are transported to the transit warehouse located in Khemis El Khechna, in the wilaya of Boumerdès. This warehouse, whose logistics management is outsourced to Aramex, serves as the reception, temporary storage, and dispatching point for all goods imported by GE HealthCare Algeria. Deliveries to end customers across the national territory are organized from this site, as shown in Figure 1.11.




Figure 1.11: GE HealthCare Algeria's Distribution Network from the Khemis El Khechna Warehouse

1.4.3 Disruption Taxonomy

Field interviews allowed the identification of a set of disruptions affecting GE HealthCare Algeria's supply chain. To structure these rigorously, they were organized following the MECE principle (Mutually Exclusive, Collectively Exhaustive), ensuring distinct and comprehensive categories. Thirteen disruptions were identified and grouped into four families.

1. Upstream Supply Disruptions

This group covers disruptions occurring upstream of goods arriving in Algeria, at the production and consolidation stages:

- **D1 — Missing Items:** absence of accessory items due to an internal error or a plant-level stockout, blocking the installation process.
- **D2 — LSIs Governance:** disruptions linked to Local Supplier Items governed by the local entity, where delays, lack of qualification, non-compliance, or supplier failures lead to the absence or inadequacy of critical accessories, resulting in operational delays and potential financial losses.
- **D3 — Factory delays:** failure to meet production deadlines without sufficient advance notice to the local teams.
- **D4 — Sales/Delivery Order gaps:** commercial commitments that are incompatible with industrial planning constraints, generating tension throughout the chain.

2. Regulatory and Financial Interface Disruptions

These disruptions occur during banking and customs processes:

- **D5 — Financial hold:** order suspension linked to uncertainty over customer solvency or payment terms.
- **D6 — Banking hold:** documentary inconsistencies or foreign exchange constraints blocking the financial transaction.
- **D7 — Customs hold:** non-compliance of the import file resulting in goods being held at customs.

3. Local Operational Disruptions

These concern operations conducted on Algerian territory:

- **D8 — Dispatch Management of Field Engineers (FE):** Human constraints (unavailability of Field Engineers or difficulties in dispatching them efficiently when multiple installations are scheduled simultaneously) and technical constraints (unprepared premises) that delay installation activities.
- **D9 — Post-signature cancellation:** late order cancellation resulting in significant financial and logistics losses.
- **D10 — Blocked FMIs:** inability to apply Field Modification Instructions on site.

4. Structural Market Disruptions

This group covers systemic constraints specific to the Algerian context:

- **D11 — Spare Mobility Constraints:** Regulatory and authorization-related restrictions preventing the re-export of refurbished spare parts and the transfer of unused new spare parts from Algeria to GE HealthCare's global network, limiting inventory reallocation opportunities and affecting after-sales service continuity.

- **D12 — ANPP registration hold:** regulatory delays of up to eight months, limiting or delaying the commercialization of new equipment.
- **D13 — Upstream visibility gap:** absence of advance information on production disruptions, amplifying the cascading effects of delays.

To analyze the interdependencies between these disruptions, an interaction matrix was built in the form of a heatmap. It measures, on a scale of 0 to 3, the impact of each disruption D_i on the other disruptions D_j , enabling the identification of propagation effects and critical nodes in the system. As shown in Figure 1.12, D12 (ANPP registration hold) and D5 (Financial hold) emerge as the highest-influence nodes, with the broadest propagation reach across the taxonomy. This analysis is the subject of an in-depth treatment in the diagnostic chapter.




Figure 1.12: Disruption Interaction Heatmap of the Supply Chain

1.5 Quantitative Diagnosis of Operational Vulnerabilities

1.5.1 Operational Performance of the W2W Process

GE HealthCare's Wing-to-Wing process rests on backward calculation logic: all intermediate milestones are derived from the customer's requested date, rather than built sequentially from the order

date. This approach places the customer at the center of operational planning and requires each stage of the chain to conform to a time window derived from the final commitment.

The anchor point of this system is the **Customer Request Date (CRD)**, which represents the date by which the customer expects the equipment to be delivered and operational on site. From this date, several planned milestones are calculated by subtracting the estimated durations of each phase:

- The **Factory Promised Date (FPD)** is the date by which the plant commits to ship the equipment:

$$FPD = CRD - (L_D + L_S + L_M) \quad (1.1)$$

where L_D is the local delivery lead time, L_S the international transit lead time, and L_M the manufacturing lead time.

- The **Scheduled Ship Distribution Date (SSSD)** is the planned shipment date from the distribution hub:

$$SSSD = CRD - (L_D + L_S) \quad (1.2)$$

- The **Scheduled On-Site Date (SOSD)** is the planned arrival date at the customer site:

$$SOSD = CRD - L_D \quad (1.3)$$

- The **Commitment On-Site Date (COSD)** is the firm customer delivery commitment date, corresponding in the system to the confirmed CRD.

The system distinguishes three types of dates that coexist throughout project tracking. **R-dates** (Requested) are customer-requested dates, of which the CRD is the sole representative. **S-dates** (Scheduled) are dates planned and committed by GE, including FPD, SSSD, SOSD, and COSD. **A-dates** (Actual) are real execution dates, recorded after the fact: AFD (Actual Factory Date), ASSD (Actual Ship Distribution Date), and AOSD (Actual On-Site Date).

The gap between A-dates and S-dates forms the basis for measuring process performance. The **OTD (On-Time Delivery)** indicator is defined as:

$$OTD = \text{Actual Date} - \text{Promised Date} \quad (1.4)$$

A positive value indicates a delay; a negative or zero value indicates on-time delivery. This indicator, calculable at each process milestone, forms the backbone of the quantitative performance analysis of the W2W process, and in particular the measurement of delay propagation from one stage to the next. The backward calculation logic underpinning this date system is represented schematically in Figure 1.13.




Figure 1.13: Diagram of the Backward Calculation Logic of the W2W Process

To establish a quantitative diagnosis of the current-state process, a sample of 104 installations was assembled, covering the 2023–2025 period for the Imaging segment. This sample includes eight equipment modalities: MRI (MR), CT Scanner (CT), Mammography (MAMMO), Radiology (RAD), Nuclear Medicine (NM), Vascular (VASC), PET (PET), and Bone Densitometry (BMD).

Lead Time Analysis by Modality

An examination of lead time medians (Figure 1.14) reveals significant structural differences across modalities. MRI presents the longest lead time in the portfolio, with a median of 20 weeks, reflecting the intrinsic complexity of the equipment, including weight, installation constraints, and transport requirements. At the other end of the spectrum, radiology and mammography systems fall between 5 and 10 weeks, indicating a globally more responsive supply chain for lighter equipment. The remaining modalities (PET, VASC, NM, BMD, CT) fall within an intermediate range of 10 to 15 weeks.




Figure 1.14: Median Lead Time by Modality (in weeks)

Beyond medians, the lead time distribution (Figure 1.15) highlights broad variability across all modalities, reflecting a lack of lead time standardization from one project to the next. CT scanner stands out as the notable exception: it shows the tightest distribution in the portfolio, with few extreme outliers. Its high volume (46 projects in the sample) contributes to this consistency, lending the process a level of maturity that low-volume modalities cannot achieve. CT can thus serve as an internal benchmark for assessing achievable performance levels.




Figure 1.15: Lead Time Distribution by Modality (Box Plot)

The coefficient of variation (CV, Figure 1.16) reinforces this reading: low-volume modalities exhibit high CVs, indicating that observed variability is not fully controlled and that lead times on these lines remain difficult to predict.




Figure 1.16: Coefficient of Variation of Lead Time by Modality

COSD Revision Rate Analysis

A complementary indicator was constructed from the analysis of customer commitment date revisions. Out of the 104 lines in the sample, 19 COSD dates were modified during the course of a project, at the initiative of the upstream Project Manager (PjM) or the installation Project Manager (PMI). This represents a revision rate of 19%, signaling that nearly one project in five requires an adjustment to the date promised to the customer, with the resulting implications for customer satisfaction, commercial credibility, and internal coordination.

1.5.2 Diagnosis of the Customs Clearance Process

The customs clearance phase is a critical step in the W2W process, positioned between international receipt and delivery to the end customer. Its analysis is based on customs processing durations recorded from the historical shipment data.

The customs clearance lead time distribution (Figure 1.17) follows a log-normal law, confirmed by a Kolmogorov-Smirnov test ($KS = 0,068$, $p = 0,107$, the only distribution not rejected at the 5% significance level). This distribution is characterized by strong right-skewness: a skewness of 3,80 and a kurtosis of 22,40, well above warning thresholds (1 and 3, respectively). In practice, the majority of files are processed within reasonable timeframes, but a small number of cases accumulate extreme delays that pull the average upward and obscure the true performance of the standard flow.




Figure 1.17: Distribution of Customs Clearance Lead Times

This finding has an important operational implication: the customs clearance problem is not systemic. The process functions well in the majority of cases. Improvement efforts should therefore focus on the proactive identification of at-risk files (incomplete documents, complex customs regimes, specific transport modes) rather than an overhaul of the entire process.

In addition, the breakdown by transport mode (Figure 1.18) shows that approximately 65 % of imports travel by sea, with the remainder by air. This modal profile directly influences lead times and the window of exposure to customs delays: sea shipments, being longer in transit, accumulate more customs waiting risk than air shipments.




Figure 1.18: Breakdown of Transport Modes for Equipment Importation

Overall, customs clearance behaves as an absorbing node within the W2W: it functions normally

in most situations, but as soon as a degradation occurs, whether from poor documentary preparation, volume peaks, or a regulatory dispute, it is fully absorbed at this stage, with non-negligible downstream propagation effects on lead times.

1.5.3 Analysis of Geographic Sourcing Concentration

To objectively assess the degree of geographic dependency in the sourcing base, the Herfindahl-Hirschman Index (HHI) was calculated from a dataset covering 785 active orders, spread across 13 manufacturing sites identified by their internal GE HealthCare plant codes. Sites were aggregated by country of origin, in line with the objective of measuring geographic concentration. The index is defined as:

$$HHI = \sum_{i=1}^n s_i^2 \quad (1.5)$$

where s_i represents the percentage share of country i in total order volume. The calculation yields the distribution presented in Table 1.1.

Table 1.1: HHI Calculation by Country of Origin

Country	Orders	Share (%)	HHI Contribution
France	393	50.1 %	2,510
United States	221	28.2 %	795
China	171	21.7 %	475
Total	785	100 %	3,780

The resulting HHI score is **3,780 / 10,000**, which unambiguously classifies the sourcing base in the high concentration zone (critical threshold: 2,500). Several factors compound this finding. France's individual contribution (2,510 points) alone exceeds the moderate concentration threshold, meaning a localized disruption affecting French sourcing would be sufficient to destabilize the entire chain. The three identified countries collectively cover 100 % of total volume, leaving no capacity to redirect supply to alternative geographies. This profile also exposes GE HealthCare Algeria to cumulative risks of diverse natures: logistics risks tied to dependence on the Europe-Mediterranean corridor, geopolitical risks linked to the Sino-American concentration, and regulatory risks inherent to any excessive dependence on a single supplier territory. These proportions and their HHI contributions are visualized in Figure 1.19.




Figure 1.19: Sourcing Share and HHI Contribution by Country

1.5.4 Temporal Analysis of Trends and Seasonality (2022–2025)

This analysis covers the evolution of delivery lead times between 2022 and 2025, based on 90 projects. Four modalities were selected, RAD, MAMMO, CT, and MR, on the basis of a sufficient volume of observations to draw statistically valid conclusions. The BMD, NM, VASC, and PET modalities were excluded, as their sample sizes ($n \leq 5$) do not allow reliable statistical interpretation. The analysis crosses three reading angles: overall trend across the period, month-by-month seasonality profile, and year-over-year evolution.

Trend Analysis

The most significant finding is that all modalities show increasing lead times over the period, each at its own pace. This convergence points toward common causes, likely external to the local organization: component shortages, pressure on manufacturing capacity, or tightening customs procedures.


MAMMO records the sharpest rise, with lead times increasing by 2.5 weeks between 2023 and 2025, and several breaches of the upper control limit (UCL) visible on the control chart. RAD shows a more moderate but real increase. CT remains the most stable modality, with a slightly negative year-over-year delta (−1,0 week), indicating a marginal improvement consistent with the process maturity identified earlier. MR presents the most concerning case: the deterioration does not manifest through visible spikes but through a slow, continuous drift in the baseline level, with a year-over-year delta of +12 weeks between 2022 and 2025, representing a degradation of approximately 60 % relative to the starting point. These modality-level trajectories are captured in the control charts presented in Figure 1.20.

Seasonality Analysis


The monthly analysis reveals that pressure points do not occur randomly throughout the year: they concentrate on relatively stable periods, which means they are in principle foreseeable. RAD and

MAMMO experience a systematic peak in March, corresponding to the acceleration of hospital orders at the start of the Algerian budget year. CT peaks in September, linked to the post-summer restart. MR stands out with a peak in July, associated with the summer shutdowns of European manufacturing plants: orders are brought forward to avoid these closures, but this artificial concentration generates congestion similar to what it was meant to prevent.


The fact that these peaks recur each year without having been meaningfully attenuated reflects an essentially reactive response, in the absence of structured demand planning.




(a) CT Scanner



(b) MRI (MR)



(c) Mammography (MAMMO)



(d) Radiology (RAD)

Figure 1.20: Temporal Evolution of Delivery Lead Times by Modality (2022–2025)

1.6 Diagnostic Summary and Problem Statement

1.6.1 Summary of Identified Vulnerabilities

All analyses conducted in this chapter, including the W2W mapping, disruption taxonomy, quantitative lead time diagnosis, customs analysis, geographic concentration, and temporal evolution, converge on the same finding: GE HealthCare Algeria's supply chain is exposed to a set of structural vulnerabilities whose effects accumulate and reinforce one another. Four main fragility factors emerge.

1. Variability and unpredictability of W2W lead times. Analysis of the 104 Imaging segment installations reveals high coefficients of variation across almost all modalities, a progressive drift in lead times over the 2022–2025 period, particularly pronounced for MRI at +12 weeks or +60 %, and a customer commitment date revision rate of 19 %. These indicators reflect a process whose behavior remains difficult to predict and reliably plan.

2. Geographic concentration of sourcing. With an HHI of 3,780, the sourcing base is concentrated across three countries covering 100 % of total volume. The complete absence of any active alternative geography makes the supply chain vulnerable to any geographically localized shock, whether logistical, geopolitical, or industrial in nature.

3. Asymmetric behavior of the customs clearance process. The log-normal distribution of customs lead times, with a skewness of 3,80 and a kurtosis of 22,40, reveals a process that is globally functional but exposed to extreme cases whose impact is disproportionate. This node absorbs and amplifies any upstream degradation without any buffering capacity.

4. Lack of visibility and anticipatory management. Seasonal peaks recur each year without a structured response. The drift in MRI lead times built up over three years without triggering a corrective action. The current management system detects drifts once they are already established, rather than preventing them.

These four vulnerabilities are not independent: geographic concentration amplifies the impact of a customs delay; the absence of a buffer turns a manufacturing delay into a customer disruption; the visibility gap prevents any proactive response. It is this interlocking nature that gives the problem its systemic character.

1.6.2 Gap Analysis

The diagnosis established in this chapter allows precise identification of the supply chain's fragility sources. It does not, however, answer operational questions that are nonetheless essential: which disruptions deserve priority attention given their combined effects? What is the actual system behavior when a disruption propagates through the W2W? What is the chain's capacity to survive and recover from different scenarios?

These questions remain unanswered not from a lack of data, but from the absence of an adequate analytical framework. GE HealthCare Algeria currently lacks a tool to model the stochastic behavior of the W2W process, simulate the impact of disruptions on performance, or quantitatively assess supply chain resilience. This is the gap the present work seeks to fill.

1.6.3 Problem Statement

The diagnosis carried out highlights that GE HealthCare Algeria's supply chain is exposed to multiple disruptions whose effects propagate non-uniformly through a heavily constrained Wing-to-Wing (W2W) process. Against this backdrop, the central question structuring the remainder of this work is as follows:

How can disruption propagation within GE HealthCare Algeria's W2W process be represented, analyzed, and operationalized in order to assess supply chain resilience and support decision-making?

This question breaks down into three research sub-questions:

- SQ1** — Which disruptions affecting the W2W process should be prioritised for simulation-based analysis, and how can GE HealthCare Algeria's supply chain be represented as a stochastic order-flow model adapted to a zero-inventory, import-dependent architecture?
- SQ2** — How can discrete-event simulation be used to model and quantify the propagation of priority disruptions through the supply chain pipeline, what ripple effects do they generate on operational performance?
- SQ3** — How can simulation findings be transformed into an operational decision-support capability that assists managers in interpreting disruption outcomes, exploring alternative scenarios, and identifying targeted mitigation actions?

These three sub-questions provide the analytical structure of the study. The first addresses the representation of the supply chain and the identification of the most influential disruption sources. The second focuses on understanding the mechanisms through which disruptions propagate and affect supply chain performance. The third explores how the resulting knowledge can be leveraged to support operational decision-making and resilience improvement. Collectively, they establish a logical progression from problem characterization to impact assessment and managerial implications.

Conclusion

This chapter has established a comprehensive picture of GE HealthCare Algeria's supply chain environment and operational vulnerabilities. The W2W process mapping, disruption taxonomy, and quantitative analyses collectively reveal a chain exposed to structural fragilities: high lead time variability, extreme geographic sourcing concentration, asymmetric customs behavior, and a persistent visibility gap. These vulnerabilities do not operate in isolation, they interact and amplify one another, giving the problem a systemic character that cannot be addressed through localized fixes alone. This diagnosis serves as the empirical basis for the methodological and analytical work developed in the following chapters, which aim to model disruption propagation, quantify resilience, and translate findings into operational decision-support.

Chapter 2

State of The Art

Introduction

The previous chapter established the industrial and operational context of this study by showing that GE HealthCare Algeria operates a complex, import-dependent, and project-based medical device supply chain. The Wing-to-Wing process is structured around the coordination of several interdependent activities: order validation, factory preparation, international transport, customs clearance, warehouse transit, local delivery, and installation, exposing the chain to multiple forms of uncertainty whose downstream consequences are neither isolated nor proportional to their origin.

This chapter develops the theoretical foundation required to analyse these issues. It does not aim to provide a generic review of supply chain management, nor does it replicate the full technical apparatus of the methods employed later in the dissertation. Instead, it organises the existing literature around a single analytical thread: the ripple effect in supply chains, its nature, the structural conditions that generate it, the chain configurations that amplify it, the resilience framework through which it is managed, and the methodological apparatus available to prioritise, simulate, and ultimately mitigate it.

The chapter is organised progressively. Section 2.1 introduces the foundational concepts of supply chain management and the specific structural characteristics of medical device supply chains. Section 2.2 develops the core theoretical lens of the study, the ripple effect, covering its conceptual properties, structural causes, and methodological treatment. Section 2.3 examines supply chain resilience as the response framework to ripple propagation. Section 2.4 discusses visibility, digital decision support, and agentic AI as emerging enablers of resilience. Section 2.5 positions Multi-Criteria Decision-Making and Discrete Event Simulation as the two complementary methodological pillars of this work. Finally, Section 2.6 identifies the research gap and proposes the conceptual framework of the study.

2.1 Supply Chain Management and Medical Device Supply Chains

2.1.1 Foundations of Supply Chain Management

The supply chain can be defined as a hierarchical, dynamic network of autonomous companies from primary supplier to end customer, interconnected through physical, informational, financial, and

knowledge flows, with the objective of satisfying stakeholders through coordination, integration, and responsiveness [30]. Far from a simple linear sequence, it is an open, complex system whose performance depends as much on inter-firm governance as on individual operational efficiency [31]. The literature approaches this concept through five complementary perspectives: functional (sequence of procurement, production, and distribution activities), strategic (source of competitive advantage in speed, cost, and flexibility), systemic (complex dynamic system of interdependent firms), structural (network of nodes and links), and relational (nexus of trust- and governance-based inter-firm relationships) [32–36].

Supply chain operations are sustained by four interdependent flow categories, each constituting a distinct ripple propagation channel. Physical flows carry goods upstream to downstream and are the primary medium through which disruptions materialise as delivery delays. Informational flows, encompassing orders, forecasts, and delivery statuses, coordinate decisions across nodes; poor visibility amplifies disruption spread. Financial and administrative flows, encompassing payments, contracts, and customs documentation, represent a critical bottleneck in import-dependent chains. Knowledge flows capture the technical expertise that determines recovery speed at a disrupted node [30]. A shock affecting any one of these flows inevitably cascades into the others, extending the scope and duration of the original disruption, a dynamic that is especially consequential in low-visibility, import-dependent configurations such as GE HealthCare Algeria [32].

Supply Chain Management (SCM) refers to the strategic and operational coordination of actors, flows, and processes to maximise overall network performance. Simchi-Levi et al. [37] define it as the efficient integration of suppliers, manufacturers, and distributors to deliver the right product at the right place, time, and cost. Chopra and Meindl [38] extend this to the joint management of product, information, and financial flows aimed at maximising total chain surplus, while Mentzer et al. [39] characterise SCM simultaneously as a management philosophy, its practical implementation, and a set of coordinated processes. From a governance standpoint, SCM reduces transaction costs across firm boundaries [40], making it a framework for managing the very complexity identified as a structural root cause of the ripple effect.

2.1.2 Specificity and Structure of Medical Device Supply Chains

Medical device supply chains differ from conventional industrial supply chains because their performance directly affects healthcare delivery. Delays in critical equipment may disrupt hospital operations, diagnostic capacity, and patient access to care, making these supply chains strategically important for health-system resilience [41].

The medical device sector is highly heterogeneous, encompassing simple consumables, implantable devices, diagnostic instruments, imaging systems, patient monitoring technologies, and complex hospital platforms. High-value medical equipment, such as MRI, CT, radiology, or interventional systems, is particularly demanding from a supply chain perspective. These products are often configured according to customer requirements, manufactured in specialised plants, transported internationally, subject to technical documentation and regulatory controls, delivered to prepared hospital sites, and installed by qualified engineers. Their delivery is therefore not a simple product-flow problem, but a coordinated project involving commercial, technical, regulatory, and logistics actors. As illustrated in Figure 2.1, this project-based delivery structure spans multiple interdependent stages, each introducing specific sources of delay and disruption.




Figure 2.1: Medical devices supply chain. [8]

This project-based structure has direct implications for resilience. In many inventory-based supply chains, demand variability or supply delay can be partly absorbed through safety stock. In the case of high-value medical equipment, local inventory buffers are often limited because products are expensive, bulky, technologically sensitive, and sometimes configured for specific customers. Resilience therefore depends less on stockholding and more on anticipation, visibility, coordination, process reliability, and the ability to recover from disruptions before they affect the committed delivery date [42–44].

A defining structural feature of this domain is the capital equipment dimension. High-value medical devices are not replenished through continuous inventory cycles; they are procured on a make-to-order basis, driven by hospital budget cycles and equipment replacement schedules, with demand that is both infrequent and difficult to forecast under shock conditions [8]. This make-to-order logic eliminates the inventory buffers that standard disruption models rely upon as absorption mechanisms. When demand surges, as occurred dramatically during the COVID-19 pandemic where ventilator orders increased by a factor of twenty, the chain has no slack to absorb the signal, and amplification propagates upward through every tier simultaneously. This is a direct instantiation of the lean paradigm root cause identified in Section 2.2: the structural absence of buffers transforms any demand shock into an unimpeded ripple.

2.1.3 Import Dependence, Component Complexity, and Regulatory Constraints

Medical device supply chains are highly internationalised, relying on globally distributed manufacturing sites, specialised suppliers, logistics platforms, and transport corridors. While this structure enables technological specialisation and efficiency, it also increases exposure to disruptions such as

component shortages, transport congestion, regulatory changes, or documentation errors [41].

In import-dependent markets, local subsidiaries have limited control over upstream production and international logistics. They depend on foreign manufacturing schedules, freight availability, customs procedures, import authorisations, and regulatory approvals. Consequently, vulnerabilities arise not only from physical disruptions but also from administrative, informational, and financial delays. This dependence remains significant even in advanced economies. Around 90% of N95 respirators used in the United States were imported, while Malaysia supplied nearly 60% of global medical glove exports [8]. These vulnerabilities are often more severe in developing countries, where medical device supply chains may rely almost entirely on imports and operate with limited inventory buffers. In such contexts, customs clearance, banking procedures, and currency fluctuations can become critical constraints on supply chain performance [45]. This situation is particularly relevant in Algeria, where advanced medical devices are largely imported and local production remains limited.

Medical devices also integrate mechanical, electronic, digital, and software-based components, creating dependencies on broader industrial ecosystems such as semiconductors and precision manufacturing. As a result, disruptions affecting these sectors can propagate to medical device production [41, 46]. Regulation is another defining characteristic: requirements related to product registration, import authorisation, traceability, and compliance are essential for safety and quality but may also increase lead-time uncertainty. Delays in documentation or administrative procedures can interrupt product flows even when manufacturing and transport capacity are available.

2.1.4 Project-Based Delivery and Installation Synchronisation

High-value medical devices are typically delivered through project-based processes requiring coordination among logistics, installation teams, hospitals, and regulatory actors. Consequently, disruptions may affect not only product availability but also installation schedules and operational readiness.

A disruption may have different consequences depending on its timing and interaction with other stages. A factory delay may be absorbed if a sufficient planning buffer exists before the committed on-site date. A missing item may create a severe delay if it blocks installation, even when the main equipment has already arrived. A customs delay may force the rescheduling of delivery and technical teams. A visibility gap may prevent proactive communication with the customer and reduce the time available for recovery.

For the present dissertation, this specificity justifies considering the customer order as the central unit of analysis. The problem is not only whether products are available somewhere in the global network, but whether each order can move through the Wing-to-Wing process within the expected time window. Delivery reliability, delay propagation, and installation readiness therefore become the central indicators of supply chain resilience [47, 48].

2.2 Supply Chain Disruptions and the Ripple Effect

2.2.1 Nature, Definition, and Conceptual Properties of the Ripple Effect

Before examining how disruptions propagate through supply chains, it is necessary to understand what a disruption is and why it is categorically different from the routine uncertainties that supply chain managers deal with on a daily basis.

Chopra and Sodhi [49] identify nine broad categories of supply chain risk, ranging from procurement failures and forecast inaccuracies to capacity shortfalls and outright disruptions. Among these, disruptions occupy the highest tier of severity. They are characterised by low frequency, high impact, and the absence of reliable historical data from which probability distributions can be inferred with confidence. Unlike operational risks, which are recurrent, predictable, and manageable through standard inventory and capacity buffers, disruption risks belong to the category of deep uncertainty: events whose likelihood cannot be meaningfully estimated from past occurrences alone [2]. This distinction has direct methodological consequences. Standard inventory optimisation models, designed to buffer against demand variability, are ill-suited to handle the structural shock of a disruption. The supply chain management literature has therefore progressively developed a dedicated stream of research centred on disruption dynamics, propagation mechanisms, and recovery strategies, a stream anchored by the concept of the ripple effect.

As shown in Figure 2.2, disruption propagation follows a cascading pattern across supply chain echelons, with performance degradation accumulating at each downstream stage.




Figure 2.2: Disruption propagation in the supply chain [3]

The ripple effect in supply chains was formally defined and systematised by Ivanov, Sokolov, and Dolgui across a series of foundational contributions beginning in 2014. In its most complete form-

lation, Dolgui et al. [3] define the ripple effect as the impact of disruption propagation on supply chain performance and the disruption-based scope of changes in supply chain structural design and planning parameters. In simpler terms, the ripple effect occurs when a disruption, rather than remaining localised at its point of origin, cascades downstream and degrades the performance of the entire chain.

A more operational formulation is offered by Ivanov et al. [2]: the ripple effect results from the propagation of an initial disruption towards other supply chain stages across supply, production, and distribution networks. The scope and economic impact of this rippling depends on two factors: the robustness reserves available to absorb the initial shock, such as inventory buffers and capacity redundancy, and the speed and scale of recovery measures deployed in response.

This definition introduces five conceptual properties, robustness, flexibility, resilience, stability, and viability, that structure the entire theoretical framework and must be understood in relation to one another. A critical distinction concerns robustness and flexibility: both rely on redundancy, but robustness uses it directly (a safety stock absorbs a three-day supplier delay without any change to the execution plan), whereas flexibility uses it indirectly (when a primary supplier fails entirely, inventory and capacity are reallocated to activate a backup sourcing route, fundamentally changing the chain's behaviour). This distinction matters for investment decisions: a supply chain that invests in robustness alone will withstand small shocks but fail catastrophically when those reserves are depleted. It is also important to distinguish risk mitigation inventory from classical safety stock: the latter buffers against demand variability, whereas the former is held specifically against disruptive risk events and managed with a different triggering and replenishment logic [3]. Detailed definitions and mechanisms for each of the five properties are provided in Appendix B.1.

2.2.2 Root Causes, Structural Dynamics, and Control Layers

The ripple effect is not an inevitable consequence of all disruptions. Its occurrence and severity depend on structural features of the supply chain that, when present, transform a localised event into a system-wide cascade. Drawing on Ivanov et al. [2], Dolgui et al. [3], and Ivanov and Dolgui [50], six root causes are identified, spanning both strategic-structural and operational-process dimensions.

At the strategic level, complexity and globalisation create extended, geographically dispersed networks where a single disrupted node can destabilise the entire chain without a coordinated contingency response. The lean paradigm and limited redundancy eliminate the buffers that would otherwise absorb shocks: JIT delivery, single sourcing, and zero safety stock are rational under stable conditions but ensure that any capacity disruption propagates immediately downstream. Geographical specialisation amplifies this by concentrating production in a small number of nodes, making individual facilities critically fragile. IT dependency and limited visibility add an informational dimension: when upstream disruption signals are detected too late, mitigation cannot be activated before downstream effects materialise.

At the operational level, two causes are particularly relevant to import-dependent, make-to-order pipelines. Long lead times extend the window during which upstream disruptions compound undetected before local teams can intervene [3, 41]. Process interdependence ensures that a delay at any single stage, factory preparation, consolidation, customs, or field installation, shifts all downstream milestones, since each stage conditions the next with no decoupling buffer between them.

A detailed description of each cause, its supply chain mechanism, ripple propagation impact, and associated countermeasures is provided in Appendix B.2.

All six root causes are present in the GE HealthCare Algeria supply chain studied in this thesis.

The chain operates across multiple continents with no local manufacturing base, relies on JIT import patterns under working capital constraints, sources specialised medical imaging equipment predominantly from European and American suppliers, depends on IT-mediated documentation systems for customs clearance, faces international lead times of several weeks, and operates a tightly sequenced Wing-to-Wing process where each milestone conditions the next, a configuration examined in depth in Chapter 3.

Ivanov et al. [2] model the supply chain as a dynamic, controllable system subject to disruption inputs, and propose a four-layer hierarchical control framework that describes how the system responds as disruption severity escalates. See Appendix B.4 for the full framework. A foundational observation is that, unlike mechanical systems where a pendulum returns to equilibrium by natural law, supply chain stability is human-driven. There is no automatic restoring force. Managers must actively decide to intervene, diagnose the appropriate control layer, and coordinate a response; the transition from one layer to the next is not a physical threshold but a managerial judgement [2].

This framework holds particular structural significance for the GE HealthCare Algeria case. The zero-inventory, make-to-order architecture of the chain means that Layer 2, the robustness layer built on inventory buffers and redundancy reserves, is structurally absent. Any disruption that exceeds the planning buffer collapses directly into Layer 3 response: replanning and contingency execution. There is no intermediate absorption stage. This architecture therefore front-loads all resilience investment into Layer 1 prevention and Layer 3 agility, compressing the detection-to-replanning window in a way that is invisible to frameworks developed for inventory-holding chains. This observation will be developed as a central structural finding in Chapter 4.

2.2.3 The Bullwhip Effect, the Efficiency Trade-off, and Boundaries of the Ripple Effect

The ripple effect should be distinguished from the bullwhip effect, the well-established phenomenon whereby small demand fluctuations become amplified as orders move upstream in a supply chain's ordering behaviour [51]. Ivanov and Dolgui [50] provide the clearest articulation of why these are fundamentally distinct phenomena operating on different registers of supply chain dynamics. See Appendix B.3 for a full comparative framework.

As shown in Figure 2.3, the two effects differ in their direction of propagation, their triggering mechanism, and their appropriate mitigation strategies.




Figure 2.3: Bullwhip and ripple effects [9]

The distinction has direct implications for which analytical tools are appropriate. Bullwhip-mitigation strategies, vendor-managed inventory, demand signal sharing, order smoothing, are powerless against a disruption that eliminates a supplier entirely or blocks a customs corridor for weeks. The ripple effect, by contrast, concerns disruption propagation: how a disturbance at one point generates cascading operational consequences including delivery delays, service-level deterioration, capacity losses, recovery costs, and customer dissatisfaction [3, 50]. In the GE HealthCare Algeria case, the dominant issue is not demand amplification along a traditional upstream chain; it is whether disruptions affecting upstream preparation, item completeness, logistics, or administrative visibility can propagate through the Wing-to-Wing process and affect delivery reliability. This makes the ripple effect the appropriate theoretical lens for the dissertation.

A central tension running through the ripple effect literature is what Ivanov et al. [2] call the efficiency–flexibility–resilience trade-off. An efficient supply chain, one operating lean with single sourcing, just-in-time delivery, and full capacity utilisation, performs optimally under stable conditions, minimising operating costs and maximising throughput. But efficiency eliminates the very buffers that would absorb a disruption; when a shock occurs, it propagates immediately, with nothing to slow it down. Efficiency thus amplifies the ripple effect. Flexibility dampens the ripple effect by providing alternative execution paths, but comes at a cost: maintaining multiple suppliers, coordinating redundant options, and holding flexibility-enabling inventory all increase operating expenses under normal conditions. Resilience, as the combination of robustness and flexibility, represents the fullest control over ripple propagation, but also requires upfront investment in redundancy that generates no return in the absence of a disruption [44]. Recent literature increasingly highlights the need to balance efficiency and resilience rather than treating them as opposing objectives [44]. The optimal strategy is therefore not to maximise any single dimension but to achieve the best possible service level at the minimum necessary redundancy cost, accepting some efficiency loss in exchange for bounded exposure to catastrophic performance degradation. This trade-off defines the strategic problem that the resilience literature seeks to solve, and it is precisely this tension that makes resilience building particularly challenging in import-dependent

chains operating under working capital constraints, as is the case at GE HealthCare Algeria.

2.2.4 Ripple Effect in Healthcare and Medical Device Supply Chains

Healthcare supply chains provide a particularly relevant context for studying disruption propagation because supply failures may affect service continuity and patient care. Recent crises showed that medical supply chains can be exposed simultaneously to demand surges, supply shortages, logistics bottlenecks, export restrictions, and procurement constraints [41, 52]. While much attention has been given to pharmaceuticals, personal protective equipment, and general medical supplies, high-value medical device supply chains remain less developed in the resilience literature.

Medical device supply chains present specific forms of ripple propagation. Component shortages can delay manufacturing. Regulatory changes can affect product availability. Competition for electronic components can create cross-industry exposure. Documentation errors can block customs clearance. Missing parts can prevent installation even when the main equipment has arrived. Import restrictions or delays in international transport can shift the entire delivery and installation plan [41, 46].

The ripple effect in such systems is therefore both network-based and process-based. It is network-based because disruptions may originate in suppliers, factories, logistics corridors, or regulatory environments. It is process-based because their consequences are observed through the order trajectory: waiting time, delayed milestones, deviation from committed dates, and downstream rescheduling [48]. This dual nature supports the methodological choice to model the order as a stochastic flow object in the simulation chapter, and it is precisely this literature gap, the absence of process-level ripple effect studies for high-value medical device import chains, that the present dissertation seeks to address.

2.2.5 Methodological Approaches to Ripple Effect Analysis

The analytical treatment of the ripple effect has generated a diverse methodological corpus. Ivanov et al. [2] and Dolgui et al. [3] organise this corpus into four primary streams, subsequently extended to seven by Dolgui and Ivanov [4]. The defining criterion for evaluating these methods is whether they can capture the time-dependent, dynamic nature of disruption propagation, the key feature that distinguishes the ripple effect from a static, point-in-time risk event. A full comparative overview of the seven streams is provided in Appendix B.5.

The comparative advantage of simulation over optimisation for ripple effect analysis is particularly well-established in the empirical literature. Ivanov [9] demonstrates that simulation can capture three features invisible to optimisation models: the non-linear relationship between disruption duration and cost (a three-day disruption produces a marginal cost increase, while a ten-day disruption produces an increase of several hundred percent); the amplification of upstream disruptions relative to downstream ones; and the consistently partial nature of recovery, where post-disruption performance never fully returns to its pre-disruption baseline within the same planning horizon. These findings directly inform the modelling choices made in Chapter 3 of this thesis, where discrete-event simulation is adopted as the primary method, complemented by graph-theoretic network analysis for the multi-echelon propagation model.

2.3 Supply Chain Resilience

2.3.1 Conceptual Evolution of Supply Chain Resilience

The preceding sections established that supply chain disruptions, once initiated, do not remain localised: they propagate across echelons, amplify through lean structures, and generate performance losses that far exceed the original disturbance. Managing this propagation dynamic requires more than reactive damage control; it demands a structured capacity to absorb, adapt, and recover. This capacity is what the literature terms supply chain resilience (SCRES).

The concept of resilience entered supply chain management from ecology and materials science, where it described the ability of a system to return to its original state after deformation. Christopher and Peck [42] were among the first to formalise this notion in a supply chain context, defining resilience as the ability of a system to return to its original state or move to a new, more desirable state after being disturbed, a definition that embeds both recovery and adaptation as legitimate outcomes. Ponomarov and Holcomb [53] subsequently defined supply chain resilience as the adaptive capability to prepare for unexpected events, respond to disruptions, and recover from them by maintaining continuity of operations at the desired level of connectedness and control over structure and function.

Hohenstein et al. [54], in a systematic review of 67 peer-reviewed articles spanning 2003 to 2013, confirmed that no single universal definition had emerged, but identified four phases consistently present across the literature: readiness, response, recovery, and growth. These phases reflect a temporal arc, from pre-disruption preparedness through active response and restoration, to the possibility of emerging from a disruption in a stronger competitive position than before. Recent resilience research has further moved beyond a narrow focus on preparedness, increasingly emphasising adaptation, viability, and the ability to maintain critical functions under long-lasting uncertainty [47, 55].

This four-phase structure complements the control layer framework introduced in Section 2.2. Where the control layers describe the escalating intensity of managerial intervention during a disruption event, the resilience phases describe the broader organisational capability that must be built before, during, and after disruption. Resilience, in this sense, is not a single property but a portfolio of capabilities distributed across the disruption lifecycle. The Ivanov corpus formalises this as four dimensions: preparedness and mitigation (pre-disruption, proactive), and stabilisation and recovery (post-disruption, reactive), a taxonomy directly aligned with the proactive and reactive strategy distinction identified by Hohenstein et al. [54].

2.3.2 Resilience Capabilities and Performance Indicators

The literature identifies several capabilities that contribute to supply chain resilience. Visibility allows the organisation to detect disruption signals and understand the status of orders, suppliers, shipments, and resources. Flexibility allows changes in schedules, routes, priorities, or resource allocation. Redundancy provides alternative sources, additional capacity, or buffers. Collaboration supports information sharing and joint problem solving among partners. Agility enables rapid adjustment when conditions change [42–44, 56, 57].

These capabilities are complementary rather than substitutable. Visibility without flexibility may identify problems without enabling response. Flexibility without visibility may be activated too late. Redundancy without coordination may remain unused. Collaboration without reliable information may not produce effective decisions. Resilience therefore depends on the combined ability

to detect, interpret, decide, and respond, a portfolio logic rather than a single-investment problem.

Two structural enablers emerge consistently across both frameworks. The first is redundancy, maintained as safety stock, backup sourcing capacity, or slack production resources, which provides the material foundation for both robustness and flexibility as defined in Section 2.2. The second is flexibility, understood as the ability to reconfigure processes and sourcing structures in response to disruption. Christopher and Peck [42] extend this to the network level, arguing that resilience must be designed into supply chain structures through deliberate sourcing strategy, supply chain mapping, and agility, the capacity to respond rapidly to unpredictable changes in demand or supply. Critically, neither redundancy nor flexibility operates effectively in isolation from visibility: without upstream and downstream information flows, disruption signals arrive too late for proactive response to be meaningful.

In the GE HealthCare Algeria context, these capabilities take a concrete operational meaning. Visibility concerns factory readiness, item completeness, shipment status, document availability, customs progress, warehouse reception, delivery scheduling, and installation readiness. Flexibility concerns the ability to reschedule deliveries, adjust customer commitments, prioritise critical orders, or activate recovery shipments. Collaboration concerns coordination among global factories, logistics providers, customs brokers, local supply chain teams, project managers, field engineers, and customers.

A resilience analysis requires indicators that capture both performance degradation and propagation behaviour. Traditional indicators such as cost, lead time, and service level are useful but insufficient to analyse the ripple effect. Average lead time may hide extreme delays, stage-level bottlenecks, and differences between disruption scenarios. To understand disruption propagation, it is necessary to measure where delays occur, how they move across stages, and whether they are absorbed or amplified [48, 50, 58]. In an order-based medical device supply chain, relevant indicators include on-time delivery rate, mean delay, maximum delay, delay variability, stage-level sojourn time, queue time, resource utilisation, and customer-date deviation. For ripple-effect analysis, additional indicators include the ripple delta, attenuation factor, amplification factor, and Super-Additivity Ratio. These indicators make it possible to compare a baseline situation with disrupted scenarios and to evaluate the severity of propagation [3, 48, 50].

2.3.3 From Resilience Capabilities to Managerial Action

A frequent limitation of resilience studies is that they identify general capabilities without translating them into actionable managerial decisions. For an engineering dissertation, resilience should not remain an abstract concept: it should lead to operational levers that managers can understand and implement.

In the studied supply chain, possible resilience levers include earlier confirmation of item completeness, pre-shipment readiness checks, improved documentation control, customs process monitoring, modality-specific planning buffers, exception alerts, recovery shipment procedures, supplier escalation rules, and scenario-based planning. These actions are directly connected to the resilience capabilities discussed above: visibility, flexibility, collaboration, redundancy, and agility [43, 44, 56, 57]. The purpose of the study is therefore not only to show that disruptions exist, but to identify which disruptions deserve attention, how they propagate, and what type of managerial response can reduce their impact. The simulation results developed in Chapter 3 and the mitigation strategy framework in Chapter 4 operationalise this connection.

2.4 Visibility, Digital Decision Support, and Agentic AI for Resilience

2.4.1 Visibility as a Condition for Resilience

Visibility is one of the most important enablers of supply chain resilience. It refers to the ability to observe, access, and interpret information about supply chain status, events, and risks. In disruption-prone systems, visibility is not simply a reporting function: it is a condition for anticipation, coordination, and timely response [41, 57, 59].

In an import-dependent medical device supply chain, visibility must cover several layers. Upstream visibility concerns factory production status, item completeness, and shipment readiness. Logistics visibility concerns transport status, estimated arrival dates, and freight documentation. Administrative visibility concerns regulatory files, banking procedures, and customs clearance. Downstream visibility concerns warehouse reception, delivery planning, site readiness, and installation resources.

The difficulty is that these data are often distributed across different systems, actors, and formats. Local teams may receive information from global ERP systems, emails, freight forwarder portals, customs platforms, spreadsheets, and customer communications. This fragmentation increases the risk of late detection: a disruption may already be propagating before all actors understand its operational consequences. The latency between the occurrence of a disruption signal and its detection by a human decision-maker defines the central operational problem that modern decision support systems seek to close.

2.4.2 From Passive Monitoring to Active Decision Support

Traditional dashboards provide useful information, but they often remain descriptive: they show what has happened or what is currently happening. Resilience requires a more active form of decision support: the ability to identify at-risk orders, anticipate downstream consequences, compare scenarios, and recommend corrective actions. In the GE HealthCare Algeria case, the same disruption may have very different consequences depending on timing, modality, order status, and remaining buffer. A factory delay does not automatically imply a final delivery delay if the planning buffer is sufficient. A missing item may have limited impact if discovered early, but it may become critical if discovered close to installation. This complexity justifies the use of modelling and simulation: operational teams cannot rely only on intuition when disruptions interact with stochastic lead times, queues, resources, and downstream commitments [52, 60, 61].

Recent research on knowledge graphs and large language models suggests that digital technologies can support multi-tier supply chain visibility by extracting, linking, and interpreting heterogeneous information from multiple sources [59]. This literature is relevant because it shows that visibility is no longer only a dashboard issue: it can become a decision-support capability that connects dispersed information to disruption analysis.

2.4.3 Digital Twins, Agentic AI, and Governance of AI-Based Decision Support

Digital twins have become increasingly important in supply chain management because they allow firms to represent, monitor, and test supply chain systems in a virtual environment. A supply

chain digital twin can combine process data, network structure, resource constraints, and scenario logic to support decision-making under uncertainty [47, 62, 63]. For resilience analysis, digital twins and simulation models are particularly useful because they allow managers to test disruptive events without disturbing real operations, supporting stress testing, scenario analysis, bottleneck identification, and performance comparison.

However, digital twins and simulation models do not automatically generate managerial understanding. Their outputs can be technical, model-dependent, and difficult to interpret for non-specialists. This creates a need for interfaces that can translate simulation results into operational insights. Recent work on LLM-supported supply chain tools suggests that generative AI can help users interact with analytical models, ask what-if questions, and understand simulation outputs more easily [64, 65].

Agentic AI extends traditional decision-support systems by combining information retrieval, reasoning, tool use, memory, and natural-language interaction. In supply chain contexts, such systems can monitor disruption signals, analyse simulation outputs, compare scenarios, and generate explanations for decision-makers [63, 66, 67]. Rather than replacing simulation or managerial judgement, agentic AI complements MCDM and DES by improving accessibility, accelerating analysis, and supporting decision-making through conversational interfaces and retrieval-augmented knowledge systems [64, 68]. As depicted in Figure 2.4, the functional architecture of such systems integrates memory, planning, tool use, and action modules into a coherent decision-support layer.




Figure 2.4: Functional architecture of an LLM-based autonomous agent. [10]

The use of AI in supply chain resilience also raises important limitations. AI systems may generate incorrect explanations, rely on incomplete data, misinterpret operational context, or suggest actions that are not feasible under real constraints. In regulated sectors such as medical devices, decision support must remain transparent, auditable, and controlled by human decision-makers. Human oversight is therefore essential [65]. This is particularly important in healthcare-related supply chains. Decisions affecting medical equipment delivery may influence hospital planning and healthcare service continuity. AI-based support should therefore not be presented as autonomous decision-making: a more rigorous position treats AI as an assistance layer that retrieves information, summarises scenarios, supports interpretation, and accelerates decision-making, while final decisions remain under managerial responsibility. In this dissertation, agentic AI is accordingly positioned as a prospective decision-support extension rather than as the core method. The core analytical methods remain disruption prioritisation and discrete-event simulation.

2.5 Multi-Criteria Decision-Making and Discrete Event Simulation as Methodological Pillars

2.5.1 Positioning and Necessity of the Multi-Criteria Approach

Before the ripple effect can be quantified through simulation, a prior methodological challenge must be addressed: which among the many possible disruption scenarios should be modelled? The supply chain studied in this thesis is subject to a heterogeneous set of disruption sources, supplier failures, customs delays, currency shocks, IT outages, missing items, that cannot all be simulated with equal depth within the scope of a single study. A structured prioritisation mechanism is therefore necessary to direct modelling effort toward the scenarios whose ripple propagation potential is highest. This is the methodological role of Multi-Criteria Decision-Making (MCDM) in the present research: not as an independent analytical contribution, but as a pre-simulation filter that ensures the DES experiments of Chapter 3 are concentrated on the most critical disruption vectors.

The prioritisation of disruptions within a supply chain constitutes a complex decision-making problem, characterised by the multiplicity of evaluation criteria, their often conflicting nature, and the uncertainty inherent in expert judgements. Classical single-criterion approaches, based on a single indicator such as frequency or financial impact, prove insufficient to capture the multidimensional nature of disruption risk [1]. Multi-criteria decision-making constitutes a methodological framework for simultaneously analysing multiple aspects of a complex problem while integrating the preferences of different stakeholders [69]. It is particularly well suited to decision-making contexts characterised by high uncertainty or multidimensional strategic stakes.

2.5.2 Selection of the Fuzzy AHP–TOPSIS Hybrid

Table 2.1: Comparative analysis of MCDM methods (adapted from Hsu et al. 1)

Method	Advantages	Limitations
AHP	Hierarchises and quantifies judgements; easy to implement	Top-down linear structure; assumes independence between criteria
ANP	Models dependencies and feedbacks between criteria	Supermatrix construction is lengthy and complex
TOPSIS	Simultaneous ranking with respect to both positive and negative ideal solutions; intuitive visualisation	Criterion weights are not provided by the method
DEMATEL	Analyses direct and indirect causal interactions between factors	Relative weighting of criteria is not taken into account
ELECTRE	Handles incomparabilities; suited to outranking problems	Complex interpretation; delicate parameterisation

As shown in Table 2.1, no single method fully satisfies all requirements simultaneously [1]. The selection of the Fuzzy AHP–Fuzzy TOPSIS hybrid approach satisfies three requirements: the ability to weight multiple and conflicting criteria, robustness in the face of the uncertainty inherent in expert judgements, and ranking efficiency when the number of alternatives is large. AHP provides a

rigorous framework for criteria weighting but does not offer a mechanism for ranking alternatives. TOPSIS, conversely, excels at ranking but does not generate criteria weights on its own. Their sequential combination exploits the complementarity of their respective strengths: Fuzzy AHP determines criteria weights by managing the imprecision of expert comparisons, while Fuzzy TOPSIS ensures efficient and transparent ranking of disruption scenarios according to those weights [70].

The fuzzy extension of both methods is justified by the very nature of the data collection process. Judgements of domain experts, drawn here from four different functions within GE HealthCare Algeria (logistics, project management, after-sales service, and commercial), are inherently subjective and imprecise. Fuzzy logic enables this imprecision to be modelled in a structured manner, substituting confidence intervals represented by triangular fuzzy numbers (TFNs) for crisp values, thereby improving the reliability and robustness of the results [70]. demonstrate, across two real-world case studies in the manufacturing and transportation sectors, that this combination constitutes a significant contribution for decision-makers facing complex situations, enabling more precise identification of priority risks and more robust selection of mitigation strategies.

It should be noted that the use of Fuzzy TOPSIS in this work differs from its classical application: rather than ranking alternative solutions, it is applied to the ranking of disruptions themselves, the thirteen disruptive scenarios identified within the GE HealthCare Algeria supply chain constituting the alternatives of the model. This adaptation ensures that simulation modelling effort is concentrated on the most critical disruption vectors. The full technical procedure, including triangular fuzzy number arithmetic, Chang's extent analysis for weight extraction, and the Fuzzy TOPSIS closeness coefficient computation, is provided in Appendix B.6.

2.5.3 Discrete Event Simulation: Foundations and Paradigm Selection

Simulation constitutes a fundamental methodological framework for the analysis of complex dynamic systems in which direct experimentation is impossible or prohibitively costly. Shannon [71] defines it as the process of designing a model of a real system and conducting experiments with that model for the purpose of understanding the behaviour of the system and evaluating various strategies for its operation. Law and Kelton [72] emphasise the temporal dimension: a simulation is the imitation of the operation of a real-world process over time, constructed so that inferences about the characteristics of the real system can be drawn from studying the model. Banks et al. [73] identify three operationally distinct objectives: understanding the internal dynamics of the system, evaluating alternative configurations without modifying the physical system, and producing quantitative indicators to support decision-making.

In the context of ripple effect analysis, these properties acquire specific methodological significance. The low-frequency, high-impact nature of major supply chain disruptions renders historical data structurally insufficient for direct statistical inference, and excludes controlled experimentation on the real chain. Simulation therefore constitutes the only rigorous framework for measuring the impact of disruptive scenarios, testing alternative resilience strategies, and quantifying performance indicators, mean delay, service rate, cumulative absorbed delay, under controlled disruption conditions [9].

The simulation literature distinguishes three dynamic paradigms whose application to supply chains constitutes an established corpus [73]. System Dynamics (SD), developed by Forrester [74] and applied to supply chains by Sterman [75], models the system as a set of aggregate stocks and flows governed by differential equations in continuous time. This paradigm is structurally suited to the analysis of inventory management policies, bullwhip dynamics, and feedback loops at the network level. Its fundamental limitation for ripple effect analysis is the absence of entity-level resolution:

individual orders cannot be tracked, disruptions cannot be assigned per order, and the terminal delay at the individual order level, the central metric of the present research, cannot be directly measured [9]. Agent-Based Modelling (ABM) operates at the level of individual actors whose local behavioural rules generate emergent macro-level dynamics. Giannakis and Louis [76] applied ABM to adaptive supply chain networks, demonstrating its capacity to capture heterogeneous actor responses to disruptions. This paradigm is appropriate when behavioural heterogeneity between actors constitutes the primary object of study; its structural overhead is unjustified when the central question concerns queue formation and delay propagation along a processing pipeline [73]. Discrete Event Simulation (DES) represents the system as a set of entities whose state changes only at discrete points in time corresponding to events: entity arrivals, service completions, resource seizures and releases, and disruption triggers. This structure establishes a direct correspondence with the physical architecture of an order processing pipeline, where each node is a resource pool with defined concurrent capacity, queue formation under congestion is an endogenous model outcome, and disruption propagation manifests as a cascade of events at each downstream node. Ivanov [9] positions DES explicitly as the reference paradigm for ripple effect analysis. The summary of the three paradigms is provided in Appendix B.7.

2.5.4 DES Application and Software Platform

The application of DES to ripple effect analysis rests on a structural correspondence between the DES formalism and the physical mechanisms of propagation in a supply chain. Ivanov [9] demonstrates, through DES experiments, three empirical properties of ripple propagation that are structurally invisible to optimisation models. First, the relationship between disruption duration and propagated cost is non-linear: a three-day disruption produces a marginal cost increase, while a ten-day disruption produces an increase of several hundred percent. Second, upstream disruptions amplify more severely than downstream ones, as their delay signal accumulates across all intervening processing stages before reaching the terminal node. Third, post-disruption recovery is structurally partial: performance does not return to its pre-disruption baseline within the same planning horizon. These three properties are emergent outputs of the DES model; they are not postulated as assumptions but arise from the interaction of stochastic service distributions, capacity constraints, and event-driven propagation logic. As illustrated in Figure 2.5, the simulation research methodology for ripple effect analysis integrates scenario design, model execution, and performance measurement into a structured experimental framework.




Figure 2.5: Framework for investigating the ripple effect on the supply chain with the help of simulation research methodology [9]

AnyLogic is selected as the simulation environment for this thesis because it is the only platform providing native support for all three simulation paradigms, DES, ABM, and System Dynamics, within a unified Java-based environment [77]. Its Process Modelling Library provides a native seize–delay–release structure that maps directly onto the order processing pipeline architecture, with queue formation under congestion as an endogenous model outcome rather than an externally imposed assumption. Ivanov [9] used AnyLogic in his foundational study of ripple effect propagation through DES, establishing a methodological precedent that the present research directly extends. A comparative overview of available DES platforms and their suitability for supply chain disruption modelling is provided in Appendix B.9.

2.6 Research Gap and Conceptual Framework

2.6.1 Synthesis of the Literature

The review conducted in this chapter establishes a well-developed theoretical and methodological corpus on supply chain disruptions and the ripple effect. The Ivanov–Dolgui–Sokolov framework provides a rigorous definition of the ripple effect, a structural account of its root causes, and a four-layer control hierarchy for managing its consequences. The resilience literature identifies the capabilities, visibility, flexibility, redundancy, agility, collaboration, that determine whether a supply chain can absorb, adapt, and recover. The methodological literature demonstrates that DES is the appropriate paradigm for capturing time-dependent disruption propagation, and that MCDM provides a structured prioritisation filter for directing simulation effort. Agentic AI emerges as a promising extension for translating quantitative simulation outputs into accessible operational decision support.

2.6.2 Identification of the Research Gap

Despite this rich theoretical foundation, three interlocking gaps remain unaddressed in the existing literature.

Gap 1: Endogenous intra-firm disruption sources. The dominant formulation of the ripple effect in the Ivanov–Dolgui corpus focuses on exogenous shocks, supplier failures, natural disasters, geopolitical disruptions, that propagate forward through the supply network. Endogenous disruption sources, originating within the firm’s own operational processes rather than in the external supply environment, receive comparatively limited treatment. In the GE HealthCare Algeria supply chain, the primary resilience risk is not a factory failure or a logistics catastrophe: it is the missing items disruption, an intra-firm quality-control failure discovered at the final installation node that triggers a backward feedback loop and a recovery shipment, a mechanism that is structurally invisible to forward-propagation models. This endogenous, feedback-generating disruption type represents a contribution to the ripple effect corpus that existing frameworks do not accommodate.

Gap 2: Zero-inventory, import-dependent supply chains. The ripple effect literature, including the four-layer control framework, implicitly assumes inventory-holding, production-based networks as its baseline. Safety stock, robustness reserves, and capacity buffers are central instruments of the Layers 1–2 response. This assumption fails structurally when applied to zero-inventory, fully import-dependent chains operating under developing-country logistics conditions, where Layer 2 is architecturally absent. No existing study applies the ripple effect framework to a supply chain of this type, and none examines the specific compounding dynamics of customs clearance, banking procedures, and regulatory constraints as simultaneous active disruption drivers in a developing-country import context. The literature on medical product supply chains acknowledges such constraints in general terms [45], but their interaction with the ripple effect framework remains unaddressed.

Gap 3: Compound simultaneous disruption in medical device chains. The ripple effect literature has developed substantial evidence on single-disruption scenarios but remains sparse on compound disruption, the simultaneous occurrence of multiple disruption mechanisms whose joint propagation may be super-additive. This gap is particularly consequential in medical device supply chains, where endogenous and exogenous disruptions may co-occur and interact across a tightly coupled installation-dependent delivery pipeline.

These three gaps converge on a single structural problem: the existing ripple effect framework was not designed for zero-inventory, import-dependent, project-based medical device supply chains operating in a developing-country context, and in particular for chains where the primary disruption source is endogenous rather than exogenous.

2.6.3 Conceptual Framework of the Study

This dissertation addresses the identified gap through a three-component analytical framework. The first component is a disruption prioritisation layer, a Fuzzy AHP–Fuzzy TOPSIS model applied to the thirteen disruption scenarios identified within the Wing-to-Wing process, producing a ranked priority order that directs simulation effort. The second component is a ripple effect propagation model, a Discrete Event Simulation implemented in AnyLogic 8, modelling the four-node Wing-to-Wing pipeline under D1 (missing items, endogenous) and D3 (factory retard, exogenous) disruption scenarios with a structured Design of Experiments covering isolated and compound conditions. The third component is a resilience strategy framework, mapping simulation findings onto Ivanov’s four-layer control hierarchy and three-level decision taxonomy, producing actionable mitigation recommendations calibrated to GE HealthCare Algeria’s zero-inventory, import-dependent architecture.

This framework makes three contributions to the ripple effect literature: extending the corpus to endogenous intra-firm disruption sources; adapting the four-layer control hierarchy to a zero-inventory, import-dependent chain where Layer 2 is structurally absent; and quantifying compound disruption super-additivity in a medical device supply chain operating under developing-country logistics conditions. The operationalisation of this framework is developed in Chapter 3.

Conclusion

This chapter has developed the theoretical and methodological foundation of the dissertation through a progressive analytical thread centred on the ripple effect. Beginning with the foundational concepts of supply chain management and the specific structural characteristics of medical device supply chains, it established the architecture through which disruption propagates. The review of the ripple effect literature provided the central theoretical lens, its definition, conceptual properties, structural root causes, control dynamics, and methodological treatment, while the resilience framework identified the capabilities and indicators required to measure and manage propagation consequences. Visibility and agentic AI were positioned as emerging enablers of proactive resilience, while MCDM and DES were established as the two methodological pillars of the study. The research gap identified in Section 2.6 motivates the specific contributions of the present work: extending the ripple effect framework to endogenous disruption sources, zero-inventory import-dependent chains, and compound simultaneous disruption in a developing-country medical device supply chain context.

The following chapter operationalises this framework, the disruption prioritisation model, the simulation architecture and the agent.

Chapter 3

A Multi-Method Modelling Framework for Supply Chain Disruption Analysis

3.1 Introduction

The preceding chapter established the theoretical and methodological foundation of this dissertation, identifying the ripple effect as the central analytical lens, positioning Fuzzy AHP–TOPSIS and Discrete Event Simulation as the two methodological pillars, and motivating the three research gaps that the present work addresses. This chapter operationalises that foundation by constructing the full analytical apparatus applied to the GE HealthCare Algeria case.

The chapter is organised around three interdependent tasks. The first, developed in Section 3.2, is the structural adaptation of the ripple effect framework to GE Algeria’s operational reality: a zero-inventory, Make-to-Order, import-dependent chain in which orders, not inventory, are the fundamental unit of analysis. This reframe is not cosmetic; it determines every subsequent modelling decision, from network construction to performance measurement. The second task, developed in Section 3.3, is the prioritisation of the thirteen identified disruptions through a Fuzzy ANP–Fuzzy TOPSIS model, producing the ranked order that directs simulation effort toward the two highest-criticality scenarios. The third task, developed in Section 3.4, is the mathematical formalisation of the reduced simulation pipeline: the stochastic order-flow model, the JIT scheduling logic, the D3 factory delay operator, and the D1 missing-items recovery loop, that translates the conceptual network into a computable experimental structure.

Together, these three tasks produce the simulation model whose experimental results are presented and interpreted in the following chapter.

3.2 Modelling GE HealthCare Algeria's Supply Chain

3.2.1 Contextualizing the Ripple Effect Framework for GE Health-Care Algeria

The ripple effect literature, and Ivanov's simulation framework in particular, is built around a common structural assumption: supply chains hold inventory, operate production processes, and manage stock replenishment policies. Disruption propagation is modelled as the degradation of these stock levels and production capacities across nodes. GE HealthCare Algeria departs fundamentally from this assumption. It manufactures nothing locally, holds no finished-goods inventory, and operates on a strictly Make-to-Order (MTO), import-dependent basis. Every order is sourced from GE's global manufacturing network, consolidated in France, cleared through Algerian customs, and delivered directly to the client site. There is no buffer stock to absorb disruptions, no reorder point to trigger replenishment, and no bill of materials to sequence local production. Applying the standard Ivanov framework to this chain without adaptation would produce a structurally misaligned model. The first methodological task of this chapter is therefore to establish the adapted modelling foundation appropriate to GE Algeria's operational reality.

3.2.2 The Structural Reframe: Orders as the Unit of Analysis

In a conventional inventory-holding chain, the system state is described by stock levels: how much is held, where, and against what replenishment policy. Disruptions are measured by their effect on these levels, stockouts, fill-rate degradation, order backlogs. None of these state variables exist in GE Algeria's chain.

The adaptation proposed here is a structural reframe: the order replaces inventory as the fundamental unit of analysis. Rather than tracking stock quantities through nodes, the model tracks individual orders as they progress through a pipeline of actors and processes. The system state at any point in time is defined by the set of active orders, their current position in the pipeline, and their elapsed and remaining processing time. Capacity is reframed accordingly, not as production volume, but as processing throughput: the number of orders a node can handle per unit time. Disruptions are modelled as throughput reductions or processing delays at specific nodes, and their propagation is measured by the downstream delay induced on individual orders.

This reframe is not merely a notational convenience. It reflects a genuine structural difference in how value flows through the chain. In GE Algeria's model, an order is simultaneously a commercial commitment, a logistics object, and a financial transaction: it is processed materially, informationally, and financially as it moves from actor to actor. This multi-dimensional nature of the order is preserved in the network model described in Section 3.2.3.

3.2.3 Network Representation of GE HealthCare Algeria's Supply Chain

The supply chain is represented as a directed network $G = (\mathcal{N}, \mathcal{E})$, where $\mathcal{N}$ is the set of actor-nodes and $\mathcal{E}$ is the set of directed edges connecting them. The network was constructed from internal GE Algeria documentation, process interviews, and operational flow mapping conducted in the preliminary phase of this research.

Modelling assumption: nodes represent the actors through which an order passes during its lifecycle-

cle. An order enters the network at the customer and exits at the customer site upon successful installation. Between these two endpoints, it is processed by a sequence of actors each performing a distinct operational function before being handed to the next.

The network comprises 23 nodes and 46 directed edges, structured across four actor categories:

- Internal nodes (GE Algeria operational actors): Commercial/Sales Team, Order Management, Finance & Banking Coordinator, PJM/Project Management, OFS Coordination, Logistics Team, PDR Stock, Service Department, Field Engineers, PMI/Project Manager Installation.
- Internal-Global nodes (GE global entities): VEGA Zone / GE Regional Management, GE Factories (multi-site), France Consolidation Centre (Saint-Witz), Backup Warehouses (Chicago / Singapore PDR only), OFS Coordination.
- External nodes (Algerian market and partners): Customer (Public & Private), Banks / Financial Institutions, Aramex / 3PL Distribution, LSI Suppliers, LSI Stock / LSI Flow, Customer Site / Readiness.
- Regulatory/Hybrid nodes: COSIR / Market Intelligence, Algerian Regulatory Environment, Customs / Regulatory Clearance (Dédouanement).

The 46 edges are classified into three flow types, reflecting the three modes in which an order is processed between nodes:

- Physical flows (material): the movement of equipment, spare parts, and documentation between actors, for example from GE Factories to the France Consolidation Centre, or from Customs to the delivery warehouse.
- Informational flows: order data, planning signals, approval requests, and status updates, for example from Order Management to PJM, or from COSIR to the Commercial Team.
- Financial flows: invoicing, payment validation, letter-of-credit processing, and commercial approvals, for example between Finance & Banking Coordinator and Banks / Financial Institutions.

The full 23-node network is presented in Appendix C.1. Feedback loop edges, representing return flows such as rejection notifications, clarification requests, and rework signals, are structurally present in the network and are relevant to disruption propagation dynamics.

The full 23-node network constitutes the conceptual reference model for GE Algeria's supply chain. Section 3.3 derives from this reference a reduced operative subnetwork retaining only the stages that contribute stochastic content to the simulation.

3.2.4 The Date-Based Performance Framework

GE HealthCare Algeria tracks each order through four key dates that form the empirical backbone of the simulation model. The **Order Date** marks the point at which an order is placed or accepted into the system, and therefore its entry into the order pipeline. The Customer Required Date (CRD) records the date the customer expects delivery. The Confirmed On-Site Delivery Date (COSD) is the date GE commits to deliver, serving as the planned performance target against which execution is measured. Finally, the Actual Transfer Date (ATD) records the date the order physically reaches the customer site.

The primary performance metric derived from these dates is the order-level delay:

$$\Delta_k = \text{ATD}_k - \text{COSD}_k \quad (3.1)$$

where $\Delta_k > 0$ indicates a late delivery and $\Delta_k \leq 0$ indicates on-time or early delivery. A 15-day grace window is applied operationally, such that an order is classified as on-time if and only if $\Delta_k \leq 15$ days. This classification is applied consistently across all modalities and disruption scenarios in the simulation.

The COSD, rather than the CRD, is used as the reference point for delay measurement. This is a deliberate modelling choice: COSD represents GE's committed delivery date, a contractual and operational target that GE controls, whereas CRD is the customer's requested date, which may be commercially aspirational and not operationally binding. Measuring delay against COSD therefore captures GE's internal performance accuracy rather than the gap between customer expectation and reality.

3.3 Prioritization of Supply Chain Disruptions through a Multi-Criteria Decision Making Approach

3.3.1 Framing the Problem and Method Selection

The state-of-the-play chapter identified thirteen disruptions that affect the wing-to-wing (W2W) supply chain process at GE HealthCare Algeria. These disruptions span a wide range of natures: some are financial, others are operational, logistical, or regulatory. A company cannot address all of them with the same intensity and at the same time. Resources, attention, and mitigation capacity are limited, which means that prioritizing interventions becomes essential.

This leads to the first research sub-question driving this section: among the thirteen identified disruptions, which ones have the most significant impact on W2W process performance, and how can they be objectively ranked?

The problem is framed as a multi-criteria decision-making (MCDM) problem. Each disruption is evaluated against a set of criteria that capture its severity from different angles (impact, frequency, propagation, and controllability), and an aggregated score is computed to produce a final ranking. The disruptions ranked highest are those that should be prioritized in the resilience strategy.

MCDM is a branch of decision theory that deals with problems involving multiple, often conflicting, evaluation criteria. In this case, the goal is to rank alternatives, where each alternative is one of the thirteen identified disruptions, in order to determine which deserve priority attention from a resilience perspective. The problem falls into the category of ranking problems ($P.\gamma$), and fits naturally into an MCDM framework for three reasons. First, assessing the criticality of a supply chain disruption cannot be reduced to a single metric: a disruption can be rare but devastating, or frequent but manageable, so multiple dimensions must be considered simultaneously. Second, the evaluation of these dimensions is inherently qualitative and involves expert judgment, which calls for a methodology that can handle linguistic assessments and uncertainty. Third, the final output must be a ranked list, not merely a classification, in order to support actionable decision-making.

A fuzzy logic framework was employed throughout to address the linguistic ambiguity and imprecision inherent in supply chain risk assessment. Expert judgments are often expressed through qualitative terms such as “more important” or “high impact,” which cannot be accurately cap-

tured using precise numerical values. Fuzzy set theory provides an effective means of modelling this uncertainty, allowing linguistic assessments to be represented in a more realistic and reliable manner.

The MCDM model follows the standard structure: a set of alternatives (the 13 disruptions), a set of criteria (detailed below), a group of decision-makers (experts from GE HealthCare Algeria), a weighting method (Fuzzy ANP) to determine the relative importance of each criterion, and a ranking method (Fuzzy TOPSIS) to score and order the alternatives.

The Analytic Hierarchy Process (AHP), developed by Saaty [78], is one of the most widely used methods for criteria weighting in MCDM. It works through pairwise comparisons: the decision-maker compares each pair of criteria and expresses which one is more important and by how much. Classical AHP relies on precise numerical judgments, which is unrealistic in practice. Fuzzy AHP addresses this limitation by replacing the crisp Saaty scale with triangular fuzzy numbers (TFNs), which represent a range of values rather than a single point. Chang's extent analysis [79] was used to derive fuzzy weights from the aggregated comparison matrices. Consistency of all expert judgment matrices was verified via the Consistency Ratio (CR); all values fell below the 0.10 threshold.

TOPSIS (Technique for Order of Preference by Similarity to Ideal Solution), introduced by Hwang and Yoon [80] and extended to the fuzzy domain by Chen [81], ranks alternatives based on their distance to two reference points: a Fuzzy Positive Ideal Solution (FPIS), representing the best possible performance on all criteria, and a Fuzzy Negative Ideal Solution (FNIS), representing the worst. Each alternative is scored using a closeness coefficient (CC), which measures how close it is to the ideal solution. The choice of TOPSIS is motivated by its simplicity, interpretability, and robustness: unlike methods such as VIKOR or PROMETHEE, TOPSIS produces a straightforward numerical score that is easy to communicate to practitioners and handles fuzzy evaluations naturally.

3.3.2 Evaluation Criteria and Expert Panel

The selection of evaluation criteria is one of the most critical steps in an MCDM study. Criteria must be relevant to the problem context, grounded in the literature, and limited in number to avoid redundancy and cognitive overload for the experts.

Four families of criteria emerge consistently from the supply chain risk assessment literature. The first covers probability, frequency, detectability, and velocity of onset [82, 83]. The second captures multi-dimensional consequences across economic, operational, and service channels [82, 84]. The third addresses cascading spread across the network [85, 86]. The fourth reflects the organisation's absorptive and adaptive capacity [87, 88]. The full hierarchical organisation of candidate criteria derived from this review is presented in Figure 3.1, which served as the basis for the contextual filtering process described below.




Figure 3.1: Hierarchical organization of candidate criteria derived from the literature review.

After contextual filtering to remove non-discriminating, correlated, or unreliably evaluable criteria, four evaluation criteria were retained.

C1 (Global Impact) measures the overall severity of the disruption’s consequences on operations, costs, and, where applicable, the availability of critical medical equipment, covering operational, financial, and patient-safety dimensions. C2 (Probability & Frequency) evaluates how frequently the disruption tends to occur in the operational context of GE HealthCare Algeria. C3 (Propagation and Amplification Effect) assesses the disruption’s tendency to spread across the supply chain and generate cascading effects from one department or function to another. C4 (Controllability) measures the organisation’s capacity to anticipate, control, or correct the disruption using available levers, and is scored inversely: a high score indicates strong control capacity, while a low score indicates limited control. C4 is therefore treated as a cost criterion in the TOPSIS model, meaning that a disruption that is difficult to control receives a higher criticality score.

A panel of four experts from GE HealthCare Algeria was assembled, each representing a key function involved in the W2W supply chain process, as shown in Table 3.1. This cross-functional composition ensures that evaluations capture disruption impact from multiple perspectives simultaneously.

Table 3.1: Expert panel composition.

Expert	Function	Profile
E1	Logistics	Manages supply chain operations, order flows, and delivery coordination.
E2	Project Management (PJM)	Oversees the post-order phase until equipment arrives in the country, ensuring that client deadlines are met.
E3	After-Sales Service (SAV)	Responsible for equipment maintenance, spare parts, and field interventions.
E4	Commercial	Handles customer relationships, contracts, and commercial commitments.

Data collection was carried out through a structured questionnaire administered individually during interview sessions. The first section elicited pairwise criteria comparisons mapped to triangular fuzzy numbers via the extended Saaty scale. The second section collected a 13×4 evaluation matrix with scores from 1 (very low) to 5 (very high) per disruption per criterion.

3.3.3 Results: Weights, Ranking, and Simulation Scope

Pairwise comparisons were aggregated using the geometric mean method [89] and consistency was verified for all individual and aggregated matrices ($CR < 0.10$ in all cases). Weights were derived via Chang's extent analysis [79]. The final criteria weights are reported in Table 3.2 and visualized in Figure 3.2, which confirms the dominant influence of C3 (Propagation Effect) and C1 (Global Impact).

Table 3.2: Criteria weights derived from Fuzzy ANP (Chang's extent analysis).

Criterion	Name	Weight (w)	Share (%)
C1	Global Impact	0.2912	29.12%
C2	Probability / Frequency	0.2230	22.30%
C3	Propagation Effect	0.2984	29.84%
C4	Controllability (inverse)	0.1873	18.73%




Figure 3.2: Final criteria weights from Fuzzy ANP.

C3 (Propagation Effect) carries the highest weight, closely followed by C1 (Global Impact), reflecting expert consensus on their dominant influence on supply chain performance.

The final ranking of the thirteen disruptions, with their closeness coefficients, distances to ideal solutions, and criticality levels, is presented in Table 3.3. The corresponding closeness coefficients are visualized in Figure 3.3, and the positioning of each disruption in the d^+/d^- space is shown in Figure 3.4, where disruptions in the upper-left zone (low d^+ , high d^-) are the most critical. Figure 3.5 provides a complementary reading through the heatmap of averaged expert evaluations.

Table 3.3: Final ranking of disruptions by Fuzzy TOPSIS closeness coefficient (CC).

Rank	Code	Disruption	d^+	d^-	CC	Criticality
1	D1	Missing Items	0.3300	0.4941	0.5996	High
2	D11	Spare Parts Network Isolation	0.3388	0.5004	0.5963	High
3	D13	Insufficient Upstream Visibility	0.3535	0.4664	0.5688	High
4	D4	Sales/Delivery Gap	0.3712	0.4561	0.5514	High
5	D3	Factory Delays	0.3775	0.4528	0.5454	High
6	D5	Financial Block (Pending)	0.3969	0.4226	0.5156	High
7	D6	Banking Block	0.4096	0.4098	0.5001	High
8	D8	FE / Site Unavailability	0.4111	0.4035	0.4954	Moderate
9	D12	ANPP Non-Homologation	0.4587	0.3551	0.4363	Moderate
10	D7	Customs Clearance Block	0.4606	0.3534	0.4342	Moderate
11	D2	LSI (Local Supplier Item)	0.5060	0.3143	0.3832	Moderate
12	D9	Post-Signature Cancellation	0.5566	0.2652	0.3227	Low
13	D10	FMI (Field Modification Instr.)	0.5580	0.2551	0.3137	Low




Figure 3.3: Closeness coefficients (CC) for the thirteen disruptions.




Figure 3.4: Distribution of disruptions in the d^+/d^- space. Disruptions in the upper-left zone (low d^+ , high d^-) are the most critical.




Figure 3.5: Heatmap of averaged expert evaluations.

The seven disruptions with a closeness coefficient above 0.50 constitute high-criticality disruptions and should receive priority attention within the organisation's resilience strategy. D1 (Missing Items, $CC = 0.5996$) was ranked first, driven by consistently high scores on C1 (Global Impact) across all four experts. Missing items directly hinder the delivery or installation of medical imaging equipment, creating operational bottlenecks with potential implications for healthcare service continuity. Its high propagation score reflects cascading effects throughout the supply chain, from order management and transportation to installation and after-sales support. D11 (Spare Parts Network Isolation, $CC = 0.5963$) ranks second: regulatory and foreign exchange constraints limit GE HealthCare Algeria's access to the global spare parts network, increasing lead times and reducing maintenance supply chain responsiveness. D13 (Insufficient Upstream Visibility, $CC = 0.5688$) acts as a risk amplifier, reducing the organisation's ability to anticipate and manage other disruptions when information flows from factories and international logistics partners are delayed or incomplete. D4 (Sales/Delivery Gap, $CC = 0.5514$) generates pressure on logistics, procurement, and project execution when commercial commitments do not align with operational realities. D3 (Factory Delays, $CC = 0.5454$) is largely outside GE HealthCare Algeria's direct control; despite being exogenous, such delays are absorbed entirely by the downstream supply chain and significantly affect delivery performance. At the lower end of the ranking, as visible in Figure 3.3, D9 (Post-Signature Cancellation) and D10 (Field Modification Instruction) exhibit relatively low criticality scores: although they may generate significant costs when they occur, their lower frequency and more limited propagation effects reduce their overall criticality.

A notable observation is the convergence of rankings across experts. Despite the four experts coming from different functions and holding different mental models of the supply chain, their evaluations led to a consistent final ordering, suggesting that the perception of disruption criticality is broadly shared across functions at GE HealthCare Algeria.

Two disruptions were selected for detailed simulation: D1 (Missing Items) and D3 (Factory Delays). D1 was retained because it obtained the highest closeness coefficient ($CC = 0.5996$), making it the most critical disruption according to the multi-criteria evaluation. D3, ranked fifth ($CC = 0.5454$), was selected despite its lower position because it directly affects production and order fulfillment processes represented within the simulation model and therefore enables a detailed assessment of disruption propagation and recovery dynamics. Conversely, although D11 (Spare Parts Network Mobility) and D13 (Insufficient Upstream Visibility) were ranked second and third respectively, they were not included in the simulation phase. These disruptions involve broader organisational, regulatory, and network-level mechanisms that extend beyond the boundaries of the discrete-event simulation model developed in this study. The selection therefore combined two criteria: the criticality level identified by the Fuzzy AHP–TOPSIS ranking and the suitability of the disruption for representation within the simulation environment.

3.3.4 Disruption Subnetwork Mapping and Network Reduction

Both disruption scenarios were mapped onto the full network $G = (\mathcal{N}, \mathcal{E})$ introduced in Section 3.2.3. The resulting disruption subnetwork, shown in Figure 3.6, traces the propagation paths activated by each scenario.

For D3, propagation is strictly forward: the disruption originates at the GE Factories node, where production output rate is reduced by a stochastic factor δ . The delay induced at the factory stage propagates downstream through OFS Coordination, PJM, the France Consolidation Centre, Customs / Regulatory Clearance, the Service / Logistics node, and finally Field Engineers, ultimately delaying delivery at the Customer Site.

For D1, the propagation structure is fundamentally different. The disruption is discovered at the

Field Engineers node, when installation reveals that a required component is absent from the delivered equipment. Upon discovery, the Field Engineers directly initiate a re-order command without escalating through PJM or any intermediate coordination actor. This command triggers a complete second import cycle traversing the Factory, all three logistics transfer stages, Consolidation, Customs, and last-mile delivery executed via air freight. D1 is therefore a feedback-loop disruption with direct FE-initiated re-ordering, structurally distinct from D3's forward propagation.

The conceptual subnetwork includes coordination actors, OFS Coordination, PJM, and the Service Department, that relay information and trigger decisions but impose no stochastic processing delay on the order flow. Modelling them explicitly would add structural complexity without adding stochastic content. They are collapsed into the D1 command delay (Triangular(1, 2, 4) days), which absorbs the combined effect of discovery, diagnosis, and re-order authorisation in practice.




Figure 3.6: Disruption subnetwork for D1 and D3 mapped onto the full GE Algeria network.

The simulation retains four operative processing stages connected by three logistics transfer segments, as detailed in Table 3.4. Under D3, orders traverse the 4-stage pipeline once, with a stochastic factory delay modifier applied at the Factory stage. Under D1, orders traverse the full pipeline a second time via air freight after discovery at the Field Engineers stage, with all three transfer segments switching mode and distribution accordingly.

Table 3.4: Operative simulation pipeline: 4 processing stages and 3 transfer segments.

Pipeline Element			Disruption Context	
Element	Type		Position	Role in Disruption
Factory	Processing stage		Origin	D3 injection point (δ modifier); emergency part production in D1 recovery.
Transfer 1: Factory → Consolidation	Logistics	seg-	Post-factory	Sea freight (normal path); air freight (D1 recovery path).
France Consolidation Centre	Processing stage		Mid-chain	Consolidation processing; distinct delay distributions per path.
Transfer 2: Consolidation → Customs	Logistics	seg-	International leg	Mode switches from sea to air under D1.
Customs / Regulatory Clearance	Processing stage		Algeria entry	Stochastic clearance delay; modality-dependent; applies to both paths.
Transfer 3: Customs → Customer Site	Logistics	seg-	Last mile	Present in both paths; reduced duration under D1.
Field Engineers	Processing stage		Terminal	D1 discovery point; directly commands re-order; installation blocked during recovery.

This pipeline with its two disruption activation mechanisms is the basis for the mathematical formalisation developed in Section 3.4.

3.4 Mathematical Formalisation

This section presents the mathematical model of GE Algeria’s supply chain pipeline. The model is formulated generically: distributions are identified by family and parameter roles rather than specific numerical values, so that the framework remains valid across different operational contexts.

All quantities are expressed as rates or durations in a common time unit (days in the simulation implementation). Rates are in orders per day; durations are in days.

3.4.1 The Order as a Stochastic Flow Object

The structural reframe established in Section 3.2.4, orders as the unit of analysis, requires a formal definition of what an order is in mathematical terms before any pipeline dynamics can be specified.

Each order k is defined as a tuple:

$$\Omega_k = (m_k, v_k, \text{CRD}_k, w_k^{\text{start}}, \text{COSD}_k, \text{ATD}_k, \xi_k^{D1}, \xi_k^{D3})$$

where the components are defined as follows:

- $m_k \in \mathcal{M} = \{\text{CT}, \text{RAD}, \text{MR}, \text{MAMMO}, \text{VASC}, \text{NM}, \text{PET}\}$ is the equipment modality, drawn at order creation from the modality probability vector $\mathbf{p} = (p_m)_{m \in \mathcal{M}}$, with $\sum_m p_m = 1$.
- $v_k \in \{\text{urgent}, \text{planned}\}$ is the order urgency class, where $(v_k = \text{urgent}) = \varphi$.
- CRD_k is the Customer Required Date, assigned stochastically at creation (Section 3.4.3).
- w_k^{start} is the planned production release date (Wing-to-Wing start), back-calculated from CRD_k (Section 3.4.3).
- COSD_k is the Committed On-Site Date, GE's delivery promise to the customer (Section 3.4.3).
- ATD_k is the Actual Time of Delivery, recorded at order completion: the only date not known at creation.
- $\xi_k^{D1} \in \{0, 1\}$ is the D1 disruption indicator: $(\xi_k^{D1} = 1) = r_{D1}$, assigned at order creation.
- $\xi_k^{D3} \in \{0, 1\}$ is the D3 disruption indicator: $(\xi_k^{D3} = 1) = r_{D3}$, assigned at order creation.

The two disruption indicators are assigned independently at order creation and are statistically independent of each other and of modality. Orders arrive according to a Poisson process with rate λ orders per day, so that the inter-arrival time between consecutive orders is exponentially distributed with mean $1/\lambda$ days.

3.4.2 Pipeline Stage Formalisation

Once released into production at w_k^{start} , order k traverses a directed pipeline of processing stages and transfer delays. Two complete paths exist: the normal sea-freight path (all orders) and the air-freight recovery path (D1 orders only, activated post-discovery). Both paths share the same stage structure but differ in their stochastic distributions.

Each processing stage i is characterised by a triple $(C_i, S_i(m), \mathcal{F}_i)$, where C_i is the concurrent capacity (number of orders processable simultaneously), $S_i(m)$ is the random service duration for modality m , and $\mathcal{F}_i$ is the distribution family. The stage definitions are presented in Table 3.5.

Table 3.5: Processing stage definitions: capacity and service distributions.

Stage	Capacity C_i	Distribution family $\mathcal{F}_i$	Path
Factory	C_F	Log-normal, modality-specific	Normal
France Consolidation	C_K	Log-normal, modality-specific	Normal
Customs	C_X	Log-normal, modality-specific	Normal
Field Engineers	C_E	Log-normal, modality-specific	Normal
Factory (recovery)	C_F	Triangular, modality-specific	D1 recovery
Consolidation (recov.)	C_K	Triangular	D1 recovery
Customs (recovery)	C_X	Log-normal, modality-specific	D1 recovery

Recovery stages share resource pools with their normal-path counterparts. A D1 recovery order therefore competes for the same Factory, Consolidation, and Customs capacity slots as orders on the normal path. This is the mechanism through which D1 disruptions propagate capacity pressure systemically across the pipeline.

Between processing stages, orders traverse transfer delays representing physical logistics movement. These are not resource-constrained: they impose duration only, with no concurrency limit. Three transfer segments exist on the normal path; all three reappear on the D1 recovery path with modified distributions reflecting the air-freight mode switch, as shown in Table 3.6.

Table 3.6: Transfer segment distributions: normal path vs. D1 recovery path.

Segment	Normal path	D1 recovery path
τ^1 : Factory $\rightarrow$ Consolidation	Triangular(a_1, b_1, c_1)	Triangular(a_1^r, b_1^r, c_1^r), $b_1^r \ll b_1$
τ^2 : Consolidation $\rightarrow$ Customs	Triangular(a_2, b_2, c_2)	Triangular(a_2^r, b_2^r, c_2^r), $b_2^r < b_2$
τ^3 : Last mile $\rightarrow$ Customer Site	Triangular(a_3, b_3, c_3)	Triangular(a_3^r, b_3^r, c_3^r), $b_3^r < b_3$

The reduction in transfer durations under D1 recovery, particularly τ^1 , formalises the air-freight premium: faster transit at higher logistics cost, activated only under disruption.

The queue at stage i at time t evolves according to the work-conserving equation:

$$Q_i(t+1) = \max(0, Q_i(t) + \lambda_i(t) - P_i(t)) \quad (3.2)$$

where the effective throughput $P_i(t)$ is:

$$P_i(t) = \min\left(Q_i(t) + \lambda_i(t), \frac{C_i}{[S_i(m)]}\right) \quad (3.3)$$

Here $\lambda_i(t)$ is the arrival rate at stage i (orders/day), C_i is the concurrent capacity (orders), and $[S_i(m)]$ is the expected service duration (days/order), so that $C_i/[S_i(m)]$ is the maximum throughput rate (orders/day). When $Q_i(t) + \lambda_i(t) \geq C_i/[S_i(m)]$, the stage is capacity-limited and a queue begins to build. This transition from arrival-limited to capacity-limited behaviour is the precise mechanism through which upstream disruptions propagate downstream.

The arrival rate at stage j is determined by the output of the preceding stage i , attenuated by α_{ij} and lagged by the intervening transfer delay τ_{ij} :

$$\lambda_j(t) = \alpha_{ij} \cdot P_i(t - \tau_{ij}) \quad (3.4)$$

The transfer delay τ_{ij} is a random variable drawn from the triangular distributions specified in Table 3.6, adding a second source of variance to the propagation signal beyond the processing variability at stage i itself.

The propagation equation introduces the inter-stage attenuation factor α_{ij} , which quantifies how much of stage i 's throughput change is transmitted to stage j . In the general form, $\alpha_{ij} \in \mathbb{R}$: when $\alpha_{ij} \in (0, 1]$, the intervening transfer buffer absorbs part of the signal (attenuation); when $\alpha_{ij} > 1$, stage j degrades more severely than stage i 's output change alone would predict, which occurs under batch-arrival dynamics when recovery-path orders complete simultaneously (amplification); when $\alpha_{ij} < 0$, a reduction in stage i 's throughput relieves congestion at stage j (demand starvation). α_{ij} is operating-regime dependent, not a fixed structural constant. Its empirical estimation from simulation output is presented in Section 4.4.3.

3.4.3 JIT Scheduling Logic

GE Algeria's MTO, zero-inventory reality means that releasing an order too early generates work-in-progress with no productive outlet. The model therefore implements a just-in-time release discipline.

For each order k with modality m , the planned production release date is back-calculated from the CRD:

$$w_k^{\text{start}} = \text{CRD}_k - L_m - \sigma_m \quad (3.5)$$

where L_m is the nominal pipeline lead time for modality m and σ_m is a modality-specific safety margin. The order waits in a pre-production holding state until simulation time reaches w_k^{start} :

$$\text{wait}_k = \max(0, w_k^{\text{start}} - t_k^{\text{creation}}) \quad (3.6)$$

The Customer Required Date is drawn from a bimodal distribution reflecting the two demand regimes present in GE Algeria's order portfolio:

$$\text{Urgent } (v_k = \text{urgent}) : \quad \text{CRD}_k = t_k^{\text{creation}} + L_m + U_k, \quad U_k \sim \text{Uniform}(0, u_{\max}) \quad (3.7)$$

$$\text{Planned } (v_k = \text{planned}) : \quad \text{CRD}_k = \max(t_k^{\text{creation}} + L_m + \eta, t_k^{\text{creation}} + H_k), \quad H_k \sim \mathcal{N}(\mu_H, \sigma_H^2) \quad (3.8)$$

The floor constraint $L_m + \eta$ ensures the CRD is never set below what the pipeline can physically achieve, preventing structurally infeasible commitments.

The COSD is GE's internal delivery commitment, calibrated against operational load:

$$\text{COSD}_k = w_k^{\text{start}} + L_m + B_k \quad (3.9)$$

where B_k is a queue buffer term that grows with system load at order creation:

$$B_k = \begin{cases} 0 & \text{if } v_k = \text{urgent} \\ \min(\beta_{\max}, |A(t)| \cdot \beta_0) & \text{if } v_k = \text{planned} \end{cases} \quad (3.10)$$

with $|A(t)|$ denoting the number of active orders at time t , β_0 the per-order buffer increment, and $\beta_{\max}$ the buffer cap. This load-aware formulation means that under high congestion, GE commits to later dates, absorbing capacity pressure into the promise rather than into the delay measurement.

To prevent factory overload from clustering of same-type orders, a minimum spacing constraint is enforced on planned orders. Let $R_m(t)$ denote the w^{start} of the most recently released planned order of modality m prior to time t . Then for any new planned order k of modality m :

$$w_k^{\text{start}} \geq R_m(t) + \varepsilon \quad (3.11)$$

where $\varepsilon = \varepsilon_{\min}$ is the minimum spacing parameter. Urgent orders are exempt. This rule distributes factory load across time without requiring explicit production scheduling.

3.4.4 D3 Formalisation: Factory Delay

D3 represents a sustained reduction in factory production speed. It acts at a single stage (Factory), affects only the forward path, and carries no feedback component.

For a D3-affected order ($\xi_k^{D3} = 1$), the factory service time is modified by a multiplicative disruption operator:

$$S_F^{D3}(m, k) = \frac{S_F(m)}{\Gamma_k}, \quad \Gamma_k \sim \text{Lognormal}(\log(\delta), \sigma_\Gamma^2) \quad (3.12)$$

where $\delta \in (0, 1]$ is the factory speed factor. Since $\delta < 1$, $[\Gamma_k] < 1$ and the division produces a longer factory duration. The log-normal wrapper introduces stochastic variation around the severity level δ . The expected factory service time under D3 is:

$$[S_F^{D3}(m)] = \frac{[S_F(m)]}{\delta} \quad (3.13)$$

D3 acts exclusively on the normal-path Factory stage and does not affect the D1 recovery Factory block. D3's propagation is strictly forward and additive. The delay injected at the Factory stage, $\Delta_k^F = S_F^{D3}(m, k) - S_F(m, k)$, travels downstream through all subsequent stages and transfer segments. Under congestion at downstream stages the delay may be amplified; D3 therefore has both a direct effect (factory duration extension) and a potential indirect effect (downstream queue amplification), the relative magnitude of which is an empirical output of the simulation.

3.4.5 D1 Formalisation: Missing Parts and Recovery Loop

D1 is structurally more complex than D3. It is not an injection at the origin of the chain: it is a discovery event at the terminal stage, triggering a complete second pipeline traversal via an expedited path.

The D1 event is discovered at the Field Engineers stage, after the order has completed the full normal pipeline. At discovery time t_k^{disc} :

$$\text{blocked}_k = \text{true}, \quad t_k^{\text{flag}} = t_k^{\text{disc}} \quad (3.14)$$

The order is frozen at the customer site, consuming a slot in the active order count $|A(t)|$, which feeds into the COSD buffer calculation (Equation 3.10) for subsequently created orders.

The recovery sub-process is a complete second pipeline traversal, whose sequence is formalised in Equation 3.15:

$$\text{Recovery}(k) : \underbrace{D^{\text{cmd}}}_{\text{command}} \rightarrow \underbrace{S^{Fr}(m)}_{\text{Factory}_r} \rightarrow \underbrace{\tau_{\text{air}}^1}_{\text{air}} \rightarrow \underbrace{S^{Kr}}_{\text{Consol.}_r} \rightarrow \underbrace{\tau_{\text{air}}^2}_{\text{air}} \rightarrow \underbrace{S^{Xr}(m)}_{\text{Customs}_r} \rightarrow \underbrace{\tau_{\text{air}}^3}_{\text{air}} \rightarrow \text{FE install} \quad (3.15)$$

where each stage draws from its own distribution family, distinct from the normal path, reflecting the expedited air-freight mode. The total recovery duration is:

$$R_k = D^{\text{cmd}} + S^{Fr}(m) + \tau_{\text{air}}^1 + S^{Kr} + \tau_{\text{air}}^2 + S^{Xr}(m) + \tau_{\text{air}}^3 \quad (3.16)$$

All terms in Equation 3.16 are independent random variables. The Customs recovery stage is the dominant source of variance.

For a D1-affected order, ATD_k is recorded only after the full recovery cycle completes and the missing component is installed:

$$\Delta_k = \text{ATD}_k - \text{COSD}_k \quad (3.17)$$

The dominant contribution to Δ_k under D1 is R_k : the recovery loop duration added on top of the normal pipeline sojourn time, against a COSD that was set before the disruption was known to exist.

3.5 Conclusion

This chapter has constructed the complete modelling and methodological apparatus of the dissertation in three successive steps. The structural reframe, replacing inventory with orders as the unit of analysis, established the appropriate foundation for GE Algeria's zero-inventory, import-dependent chain and determined the network representation, the date-based performance framework, and the delay metric Δ_k governing all subsequent measurement. The Fuzzy ANP–Fuzzy TOPSIS prioritisation model produced a convergent, cross-functionally validated ranking of the thirteen disruption scenarios, directing simulation effort toward D1 (Missing Items, CC = 0.5996) and D3 (Factory Delays, CC = 0.5454). The mathematical formalisation then translated the reduced four-stage pipeline into a computable stochastic model, specifying the JIT scheduling discipline, the D3 multiplicative delay operator, and the D1 feedback recovery loop with its full second-cycle air-freight traversal.

Three structural features of this model carry forward into the experimental chapter: D1 and D3 differ in propagation topology, not only severity; the zero-inventory architecture collapses Layer 2 of Ivanov's four-layer hierarchy; and the load-aware COSD formulation couples disruption propagation to performance measurement in a way absent from standard models. The following chapter activates this apparatus through a fifteen-run Design of Experiments across four disruption families.

Chapter 4

Simulation-Based Analysis of Disruption Propagation

4.1 Introduction

The preceding chapter constructed the complete modelling apparatus of this dissertation: the order-based network representation of GE HealthCare Algeria’s supply chain, the Fuzzy ANP–Fuzzy TOPSIS disruption prioritisation model that directed experimental effort toward D1 (Missing Items) and D3 (Factory Delays), and the mathematical formalisation of the four-stage stochastic pipeline under both disruption mechanisms. This chapter activates that apparatus.

The chapter is organised in three sections. Section 4.2 describes the AnyLogic discrete-event simulation model: its agent definition, resource pool configuration, stochastic service functions, and the key design decisions that translate GE Algeria’s operational specificities into computable model behaviour. Section 4.3 presents the Design of Experiments: the fifteen-run partial factorial plan spanning four disruption families No Disruption, Missing Items, Factory Retard, and Combined together with the Super-Additivity Ratio as the compound disruption metric. Section 4.4 presents and interprets the simulation results across five analytical dimensions: aggregate KPI behaviour, node-level sojourn inflation and attenuation factors, disruption intensity and super-additivity, modality-level vulnerability, and statistical validation.

The central question driving this chapter is whether, and through what mechanisms, disruption propagates through GE Algeria’s zero-inventory, import-dependent pipeline – and whether that propagation is symmetric between endogenous and exogenous disruption sources. The results answer this question in a way that revises the dominant assumptions of the ripple effect corpus.

4.2 Simulation Model Architecture

The mathematical framework developed in section 3.4 is operationalised as a Discrete-Event Simulation (DES) model built in AnyLogic 8 (Process Modelling Library). DES was selected as the simulation paradigm because the system is fundamentally event-driven: state changes occur at discrete points in time corresponding to order arrivals, service completions, resource seizures, and disruption discoveries. The continuous-time formulation of queue evolution in eq. (3.2) serves as the theoretical substrate; the DES engine resolves it event-by-event with full stochastic fidelity.

4.2.1 Model Overview and Agent Definition

The model is agent-based at the order level. Each order is instantiated as an **Order** agent at the moment of its arrival event, carrying all attributes defined in the order tuple Ω_k (section 3.4) as mutable agent variables. Mutable variables rather than fixed parameters are essential: dates such as COSD_k and ATD_k and flags such as ξ_k^{D1} must be writable at runtime as the simulation evolves. The complete agent variable set is listed in Appendix D.2.

Figure 4.1 shows the AnyLogic model canvas with the full process flow for the normal path and the D1 recovery sub-path, illustrating how the two paths share resource pools at the Consolidation and Customs stages.




Figure 4.1: AnyLogic model canvas: full process flow for the normal path and the D1 recovery sub-path.

4.2.2 Resource Pool Specification

The four processing stages are implemented as resource pools, enforcing the capacity constraints C_i defined in table 3.5. An order must seize a resource unit from the pool before its service delay begins and releases it upon completion. When all resource units are occupied, arriving orders queue and wait implementing the FIFO discipline with no preemption.

Recovery-path stages share resource pools with their normal-path counterparts. This is a deliberate modelling choice: in GE Algeria’s actual operation, the France Consolidation Centre and Algerian Customs do not maintain separate capacity for recovery orders. A D1 recovery order therefore competes for the same slots as all currently active normal orders, making congestion under high D1 rates an endogenous model outcome rather than an assumption.

4.2.3 Stochastic Service Functions

Service durations at each stage are computed by dedicated functions, one per stage per path type, accepting modality as input. This encapsulates all distributional logic in a single callable location. The function architecture ensures that distributional changes for sensitivity analysis require modification in one place only, without touching the process flow blocks. The functions and their distribution mappings are listed in Appendix D.1.

4.2.4 Key Design Decisions

The model incorporates several design decisions that depart from a naïve DES implementation. Each decision reflects a specific operational feature of GE Algeria's chain that a simpler model would fail to capture.

Bimodal Demand: Urgent and Planned Orders: GE Algeria's order portfolio is not homogeneous in its demand structure. Two distinct demand regimes coexist: urgent orders (representing emergency clinical needs or unforeseeable replacement requirements) and planned orders (representing programmed hospital investment cycles, often budgeted a year or more in advance). These two regimes have structurally different CRD distributions, as formalised in eq. (3.7) and eq. (3.8). Treating all orders under a single demand model would systematically underestimate the variance of the CRD distribution and misprice the tail risk of late delivery. The bimodal structure was inferred from GE Algeria order data and confirmed through company interviews.

JIT Release Discipline: Orders are not released into production immediately upon creation. This reflects GE Algeria's MTO discipline: initiating production before the pipeline can physically deliver within the committed window has no operational benefit and consumes factory capacity unnecessarily. The JIT release discipline (eqs. (3.5) and (3.6)) ensures that each order enters the Factory stage at the latest moment consistent with achieving its COSD. This is not merely a scheduling detail it is the mechanism that makes the model's COSD and delay measurements meaningful, since delay can only be defined against a target that reflects a genuine operational commitment.

Load-Aware COSD with Queue Pressure Buffer: The COSD is not a static deadline computed solely from lead time. It incorporates a buffer B_k that grows with the number of active orders in the system at the moment of order creation (eq. (3.10)). This reflects GE Algeria's actual commitment practice: when the pipeline is congested, account managers do not promise the nominal lead time they extend the commitment window to absorb the expected queueing delay. Omitting this buffer would cause the model to systematically over-predict the late-delivery rate under congestion, because orders would be judged against commitments that were never operationally realistic.

D3 as a Multiplicative Service-Time Modifier: The D3 factory Delay is implemented as a division of the nominal factory service time by a log-normally distributed factor Γ_k (eq. (3.12)). This formulation was chosen over a simple additive delay for two reasons. First, it ensures that the Delay effect is modality-proportional: a longer factory process (MR, ~ 48 days) suffers a proportionally larger absolute delay than a shorter one (RAD, ~ 10 days) for the same δ , which is consistent with the operational interpretation of a production speed reduction. Second, the multiplicative structure preserves the positivity of service times under all realisations of Γ_k , whereas an additive model with stochastic variation could theoretically produce zero or negative service times at the tail.

D1 as a Feedback Loop with Resource Competition: The D1 disruption is not implemented as a one-off penalty added to a single order's completion time. It is a structural feedback loop: a complete second pipeline traversal is activated, competing for shared resource pools with all currently active normal orders (eq. (3.15)). This design choice has a direct consequence for

simulation outcomes: under high D1 rates, recovery orders simultaneously flood the Factory, Consolidation, and Customs pools, inducing congestion that degrades the performance of non-disrupted orders. This cross-order congestion effect which would be invisible under a simple penalty model is the primary mechanism behind the super-linear delay scaling observed at high disruption intensities in the results.

Same-Modality Spacing Enforcement: The spacing constraint (eq. (3.11)) is enforced at order creation for planned same- modality orders. Two helper functions, `getLastReleasedAt(m)` and `setLastReleasedAt(m, t)`, maintain the most recent release time per modality. This prevents burst arrivals of same-type orders from saturating the factory’s capacity for that modality, smoothing the load profile in a manner consistent with GE Algeria’s operational planning practice.

4.2.5 Output Collection and Performance Metrics

At order completion, the function `recordOrderCompletion(k)` is called. It computes and records the final performance metrics for order k and updates the system-level accumulators. Two levels of output are collected.

Order-Level Output

Each completed order appends a row to the dataset `orders_tab` containing the full attribute vector: all dates (*Order Date*, *CRD*, *COSD*, *ATD*), the delay Δ_k , the on-time classification, modality, urgency class, disruption flags, and node-level sojourn times (t_{Factory} , $t_{\text{Consolidation}}$, t_{Customs} , t_{FE}). This order-level dataset supports all post-hoc statistical analysis, including modality-stratified delay distributions and disruption-specific breakdowns.

Simulation Output: Order-Level Records

The primary output of the simulation is a structured CSV file in which each row corresponds to a single completed order. The table records the full temporal trace of the order through all pipeline stages, enabling per-order delay computation and node-level sojourn time analysis in post-processing. The columns are grouped by concern in table 6.13.

System-Level Aggregates

Four scalar counters are accumulated during the run and written alongside the order-level records as summary statistics:

- `total_orders_completed`: total number of completed orders.
- `total_orders_late`: count of orders with $\Delta_k > 15$ days.
- `cumulative_delay_days`: $\sum_k \max(0, \Delta_k)$ over all completed orders.
- `d1_occurrences`, `d3_occurrences`: cumulative disruption event counts.

The on-time delivery rate (OTD) and mean delay $\bar{\Delta}$ are derived from these counters as post-processing outputs:

$$\text{OTD} = 1 - \frac{\text{total_orders_late}}{\text{total_orders_completed}} \quad (4.1)$$

$$\bar{\Delta} = \frac{\text{cumulative_delay_days}}{\text{total_orders_completed}} \quad (4.2)$$

These two scalar metrics, OTD and $\bar{\Delta}$, serve as the primary response variables in the Design of Experiments presented in section 4.3. The full order-level records in `orders_tab` additionally support node sojourn time analysis, attenuation factor back-calculation (α_{ij}), and per-scenario trajectory comparisons, as detailed in section 4.4.

4.3 Design of Experiments

4.3.1 Experimental Framework

The simulation model is used to evaluate supply chain performance across a structured set of disruption scenarios. The experimental design follows a **partial factorial** approach, drawing on the methodology of Law & Kelton [72] and Montgomery [90] for stochastic simulation experiments.

The four controllable factors are:

1. **Simulation horizon** (T): the total duration of each simulation run, varied across 800, 1 200, and 2 000 days to assess duration sensitivity.
2. **Missing Items rate** (p_{D1}): the per-order probability that a D1 event occurs, i.e. the fraction of orders discovered to have missing parts at installation. Values tested: 0, 0.15, 0.25, 0.35.
3. **Factory Retard rate** (p_{D3}): the per-order probability that a D3 event occurs, i.e. that the order's factory processing time is subject to a retardation modifier. Values tested: 0, 0.10, 0.20, 0.30.
4. **Factory retard severity** (δ): the speed factor applied to the factory service time under D3, $\delta \in (0, 1]$, where lower values indicate more severe retardation (eq. (3.12)). Values tested: 0.60, 0.70, 0.80.

Both p_{D1} and p_{D3} are per-order Bernoulli probabilities assigned independently at order creation, not annual event rates. A value of $p_{D1} = 0.25$ means that each order has a 25% probability of being affected by a missing-parts event, independently of all other orders.

The full factorial design over these four factors at their respective level counts would require $3 \times 4 \times 4 \times 3 = 144$ experimental runs. Given the computational cost of stochastic simulation and the primary analytical objective — isolating the effects of D1 and D3 individually and in combination — a partial factorial design of **15 runs** was constructed. This reduction is justified by four principles, detailed below.

4.3.2 Design Principles

Principle 1: Baseline isolation.

A no-disruption baseline (E1–E3) is included at three horizon lengths as the reference point against which all disruption effects are measured. This provides the counterfactual $\bar{\Delta}_0$ and OTD_0 and simultaneously establishes whether simulation duration affects baseline performance in the absence of disruption.

Principle 2: Factor isolation.

D1 and D3 are evaluated independently before being combined. Pure D1 runs (E4–E6, $p_{D3} = 0$) and pure D3 runs (E7–E10, $p_{D1} = 0$) isolate the individual effect of each disruption mechanism before joint effects are examined. This structure allows clean attribution of observed performance degradation to its source.

Principle 3: Severity gradient.

Each factor is varied across its operational range in the isolation runs, producing a severity gradient. For D1, p_{D1} increases from 0.15 to 0.35 (E4–E6). For D3, both p_{D3} and δ are jointly varied from mild ($p_{D3} = 0.10$, $\delta = 0.70$) to severe ($p_{D3} = 0.30$, $\delta = 0.60$) (E7–E10). This gradient enables identification of the intensity threshold at which performance degrades non-linearly.

Principle 4: Compound interaction.

Runs E11–E13 combine active D1 and D3 at three severity levels (High/High, Low/Low, Medium/Medium), enabling measurement of the super-additivity ratio (SAR) as defined in eq. (4.3). Runs E14 and E15 replicate the Medium/Medium compound scenario at 800 and 2000 days respectively, isolating the duration effect under disruption.

4.3.3 Super-Additivity Ratio

The Super-Additivity Ratio (SAR) quantifies whether the joint effect of D1 and D3 exceeds or falls short of the sum of their individual effects. For any compound scenario e_c combining D1 at rate p_{D1}^* and D3 at intensity (p_{D3}^*, δ^*) , let $\bar{\Delta}(e_c)$ denote the mean delay under compound disruption, and let $\bar{\Delta}_{D1}$, $\bar{\Delta}_{D3}$, $\bar{\Delta}_0$ denote the mean delays under the corresponding pure D1, pure D3, and baseline scenarios respectively. Then:

$$\text{SAR} = \frac{\bar{\Delta}(e_c) - \bar{\Delta}_0}{(\bar{\Delta}_{D1} - \bar{\Delta}_0) + (\bar{\Delta}_{D3} - \bar{\Delta}_0)} \quad (4.3)$$

An SAR of exactly 1 indicates perfect additivity. $\text{SAR} > 1$ indicates super-additivity (disruptions amplify each other). $\text{SAR} < 1$ indicates sub-additivity (disruptions partially offset each other, e.g. through demand starvation). The SAR is computed for compound scenarios E11–E13 and reported in section 4.4.4.

4.3.4 Run-Length and Sufficiency Justification

The simulation horizon of $T = 1,200$ days was selected to provide a sufficiently long steady-state observation window while respecting the single-run per scenario design constraint. With an expected inter-arrival time of $1/\lambda$ days, the expected number of completed orders over the $T = 1,200$ -day horizon is:

$$[N] = \lambda \cdot T \quad (4.4)$$

This produces a sample large enough for the key performance statistics $\bar{\Delta}$, OTD, stage-level queue lengths to be estimated with acceptable precision. The run length also ensures that rare but high-impact disruption events (D1 at 1 event/year: ~ 3 –4 events per run) occur a statistically meaningful number of times within each scenario.

The present study conducts a single simulation run per scenario, without Monte Carlo replication. This is acknowledged as a methodological limitation: point estimates of $\bar{\Delta}$ and OTD are reported without confidence intervals, and the within-scenario variance contribution cannot be separated from the between-scenario signal. The design was constrained by the AnyLogic Personal Learning Edition environment, which does not support built-in Monte Carlo experiment types. The single-run results are therefore interpreted as indicative of directional effects and relative ordering of scenarios, rather than as precise absolute estimates. Confidence interval computation using a replication loop remains a recommended extension of this work [72, 91].

Warm-up: No explicit warm-up period is applied. Orders begin arriving from $t = 0$ and the system initialises empty. The first ~ 100 days are allocated to pipeline fill-up: given an inter-arrival rate of λ orders per day and a nominal pipeline depth of several weeks across four stages, approximately 100 days are required for a representative order population to accumulate across all stages. This fill-up window isolates the subsequent 1 100-day observation period as a clean three-year equivalent window, consistent with GE Algeria’s annual planning cycle. The warm-up fraction ($\sim 8\%$ of the primary horizon) is small enough that it does not materially distort long-run statistics.

4.3.5 Experimental Plan

The full 15-run experimental plan is summarised in Table 4.1, covering all factor combinations across the four disruption families.

Table 4.1: Full experimental plan: 15 simulation runs

Run	Scenario	T (d)	p_{D1}	p_{D3}	δ	DOE role
E1	No Disruption	1200	0	0	—	Baseline (mid horizon)
E2	No Disruption	800	0	0	—	Baseline (short horizon)
E3	No Disruption	2000	0	0	—	Baseline (long horizon)
E4	Missing Items — Low	1200	0.15	0	—	D1 isolation, low frequency
E5	Missing Items — Medium	1200	0.25	0	—	D1 isolation, medium frequency
E6	Missing Items — High	1200	0.35	0	—	D1 isolation, high frequency
E7	Retard — Low/Medium	1200	0	0.10	0.70	D3 isolation, low rate, medium severity
E8	Retard — Medium/Low	1200	0	0.20	0.80	D3 isolation, medium rate, low severity
E9	Retard — High/High	1200	0	0.30	0.60	D3 isolation, high rate, high severity
E10	Retard — High/Medium	1200	0	0.30	0.70	D3 isolation, high rate, medium severity
E11	Combined — High/High	1200	0.35	0.30	0.70	Compound, stress level
E12	Combined — Low/Low	1200	0.15	0.10	0.80	Compound, low level
E13	Combined — Med/Med	1200	0.25	0.20	0.70	Compound, medium level (mid horizon)
E14	Combined — Med/Med	800	0.25	0.20	0.70	Duration sensitivity (short)
E15	Combined — Med/Med	2000	0.25	0.20	0.70	Duration sensitivity (long)

4.3.6 Response Variables

Two primary and two secondary response variables are recorded for each run.

The first primary metric is the **mean delay** $\bar{\Delta}$, defined as the mean of $\Delta_k = \text{ATD}_k - \text{COSD}_k$ across all completed orders. The second primary metric is the **on-time delivery rate** (OTD), which measures the fraction of orders satisfying $\Delta_k \leq 15$ days (eq. (4.1)). These two indicators are complementary: $\bar{\Delta}$ is continuous and sensitive to the magnitude of delay, while OTD is a binary threshold metric sensitive to the frequency of threshold crossings. Together they provide a fuller characterisation of performance degradation than either metric alone.

Two secondary metrics are also collected. The cumulative delay $\Sigma\Delta$ aggregates total late-delivery days across all orders and is proportional to the aggregate service-level cost. The disruption event counts n_{D1} and n_{D3} record the realised number of disruption occurrences per run and are used to verify injection rate calibration.

4.4 Simulation Results and Analysis

4.4.1 Overview of Experimental Outcomes

Table 4.2 summarises the four primary KPIs across all 15 experiments. All delay values are computed as $\Delta_k = \text{ATD}_k - \text{COSD}_k$, with on-time classification defined as $\Delta_k \leq 15$ days.

Table 4.2: KPI summary across all 15 experiments

Run	Scenario	n orders	OTD	$\bar{\Delta}$ (d)	$\Sigma\Delta$ (d)
E1	No Disruption (1200d)	153	0.72	9.78	1 524
E2	No Disruption (800d)	83	0.84	6.09	528
E3	No Disruption (2000d)	263	0.75	10.49	2 815
E4	Missing Items Low (0.15)	138	0.76	10.35	1 462
E5	Missing Items Medium (0.25)	134	0.57	27.46	3 702
E6	Missing Items High (0.35)	134	0.37	43.78	5 880
E7	delay Low/Medium	157	0.68	12.11	1 936
E8	delay Medium/Low	141	0.70	11.91	1 700
E9	delay High/High	126	0.69	9.90	1 279
E10	delay High/Medium	134	0.74	8.54	1 185
E11	Combined High/High	121	0.45	46.07	5 588
E12	Combined Low/Low	133	0.62	17.93	2 423
E13	Combined Medium/Medium (1200d)	149	0.53	22.82	3 410
E14	Combined Medium/Medium (800d)	84	0.65	16.16	1 369
E15	Combined Medium/Medium (2000d)	246	0.48	27.98	6 933

Baseline OTD note. The baseline OTD of 72–84% under No Disruption (E1–E3) reflects the structural characteristics of GE Algeria’s COSD-setting mechanism, not a modelling deficiency. The COSD incorporates a queue-pressure buffer B_k that grows with system load (eq. (3.10)), and the stochastic pipeline introduces inherent lead-time variance that a minority of orders cannot absorb within the 15-day grace window. This operating-regime OTD is the realistic chain baseline, not an ideal target, and is consistent with GE Algeria’s reported empirical performance.

4.4.2 OTD Rate, Mean Delay, and Duration Sensitivity

The KPI heatmaps in Figure 4.2 provide a cross-scenario overview of OTD and mean delay for all 15 experiments, making the relative severity of each disruption family immediately visible.




Figure 4.2: KPI heatmaps across all 15 experiments: OTD rate (top) and mean delay in days (bottom).

Baseline stability across simulation horizons

Under No Disruption (E1–E3), OTD remains stable between 72 % and 84 % and mean delay between 6.1 and 10.5 days across the 800-, 1 200-, and 2 000-day runs, as shown in Figure 4.2. The marginal drift from 1 200 to 2 000 days (mean delay: 9.8 to 10.5 days) reflects natural queue accumulation as more orders are processed over a longer horizon, not a disruption effect. The model reaches a near-stationary operating regime within 1 200 days under undisrupted conditions, validating the chosen primary simulation horizon.




Figure 4.3: OTD rate versus disruption intensity across scenario families.

Missing Items, monotonic dose-response: The Missing Items experiments (E4–E6) produce the clearest dose-response pattern in the dataset, as shown in Figure 4.3. OTD decreases monotonically: 76 % $\rightarrow$ 57 % $\rightarrow$ 37 % as the disruption rate increases from 0.15 to 0.35. Mean delay rises from 10.4 to 43.8 days. The slope steepens between Medium (0.25) and High (0.35), consistent with queue saturation: once recovery-path orders constitute a sufficient fraction of pipeline load, the residual normal-path capacity becomes insufficient to prevent queue growth, and per-order delay accelerates non-linearly. At High rate (E6), cumulative absorbed delay (5 880 days) marginally exceeds the Combined High/High scenario (E11, 5 588 days), establishing the endogenous missing-items mechanism as at least as systemically damaging as the worst compound scenario tested.

Delay Factory, impact present but intensity-independent: Delay experiments (E7–E10) produce OTD values of 68–74 % and mean delay 8.5–12.1 days, statistically indistinguishable from the No Disruption baseline (Figure 4.3). Factory delay does therefore exert a modest downward pressure on OTD relative to the undisrupted ideal, but its defining characteristic is the absence of any dose-response relationship: increasing disruption rate and severity from E7 (Low/Medium) to E9 (High/High) produces no correlated degradation in OTD or mean delay. E9 ($\bar{\Delta} = 9.9$ d) actually performs better than E7 ($\bar{\Delta} = 12.1$ d). The conclusion is not that factory delay is harmless, but that its intensity is not a meaningful predictor of terminal delivery performance at the tested levels – the chain absorbs additional delay severity within the same performance envelope regardless of how severe the delay is.

Combined disruptions, simulation duration as a performance factor: As shown in Figure 4.4, the Combined Medium/Medium block (E13–E15) reveals that simulation duration becomes a significant performance determinant under disruption in a way that is absent under No Disruption. Under baseline conditions (E1–E3), extending the horizon from 800 to 2,000 days leaves OTD essentially unchanged (72–84 %). Under compound disruption, the same extension drives OTD from 65 % (800 d) down to 48 % (2,000 d) with mean delay rising from 16.2 to 28.0 days. The implication is direct: disruption transforms simulation duration from an irrelevant parameter into a critical one. Longer exposure to compound disruption accumulates damage that shorter runs cannot capture, meaning the 1,200-day primary window structurally underestimates long-run performance degradation under persistent compound scenarios.




Figure 4.4: Sensitivity of OTD and mean delay to simulation duration under baseline and compound disruption conditions.

4.4.3 Node Sojourn Times, Ripple Deltas, and Attenuation Factors

Node-level sojourn inflation

Figure 4.5 plots mean sojourn time per stage across all scenarios, making it possible to trace how delay accumulates or dissipates along the pipeline under each disruption family.




Figure 4.5: Mean sojourn time (days) per stage per scenario.

The ripple delta heatmap in Figure 4.6 complements the sojourn plot by expressing each stage's sojourn as a signed deviation from the No Disruption baseline, isolating the incremental contribution

of each disruption family to node-level delay.




Figure 4.6: Ripple delta heatmap: signed deviation of mean sojourn time from No Disruption baseline, per stage and scenario.

The dominant finding is the Field Engineers terminal accumulation pattern. Under Missing Items High (E6), FE sojourn reaches 37.0 days versus a 9.1-day baseline (+27.9 days), and 32.8 days under Combined High/High (E11, +23.7 days), as visible in Figure 4.5. Delay-only experiments leave FE sojourn unchanged at 8.4–9.3 days. Field Engineers is therefore the primary terminal ripple accumulator under endogenous disruption – not Customs, as theoretical expectations might suggest.

Customs functions as the intermediate amplifier. Its sojourn inflates to 26.5 days under E6 and 30.4 days under E11, versus a 19.1-day baseline (Figure 4.6). The amplification mechanism is cross-path interference: D1 recovery-path orders re-enter the chain through the Customs recovery block, generating congestion that spills onto the normal-path customs queue. This effect is unique to endogenous disruption and absent under pure delay scenarios.

Factory and Consolidation are minimally disruption-sensitive. Factory sojourn is stable at 19.0–24.1 days across all experiments. Consolidation ranges 13.2–21.7 days, with one notable exception: under delay High/Medium (E10), Consolidation sojourn *drops* to 13.2 days, below the 17.3-day baseline – the empirical signature of demand starvation, addressed below.

Inter-stage attenuation factors $\hat{\alpha}_{ij}$

The attenuation factor $\hat{\alpha}_{ij}$ (eq. (3.4)) measures how much of a throughput change at stage i propagates to stage j . Table 4.3 reports estimates across the three interpretable stage pairs and selected scenarios. Cells marked n/a correspond to cases where $|\Delta P_i| < 0.5$ d (unreliable denominator); cells exceeding the display range $[-3, 3]$ are clipped and flagged as artefacts.

Table 4.3: Estimated inter-stage attenuation factors $\hat{\alpha}_{ij}$ by scenario.

Stage pair	Comb. H/H	Comb. L/L	MI High	MI Med.	Ret. H/M	Ret. L/M
Factory $\rightarrow$ Consolidation	0.60	0.71	2.21	n/a	−0.20	2.05
Consolidation $\rightarrow$ Customs	(9.18)*	0.03	(5.62)*	0.15	−0.68	−0.71
Customs $\rightarrow$ Field Engineers	−0.01	n/a	0.16	0.05	0.59	0.24

* Artefact: upstream $|\Delta P| < 0.5$ d; values excluded from interpretation.

Three propagation regimes are identifiable from the reliable cells:

Attenuation ($\hat{\alpha} \in (0, 1)$): the dominant regime: Under combined low/low and moderate delay scenarios, Factory $\rightarrow$ Consolidation yields $\hat{\alpha} \in [0.60, 0.71]$, indicating that the transfer buffer absorbs 30–40 % of the upstream delay before it reaches Consolidation. The Customs $\rightarrow$ FE row shows consistently weak propagation ($\hat{\alpha} \leq 0.59$), suggesting that cumulative buffering across three upstream stages effectively decouples FE sojourn from the original disruption signal at this pair.

Amplification ($\hat{\alpha} > 1$): batch-arrival overflow: Factory $\rightarrow$ Consolidation shows $\hat{\alpha} = 2.21$ (Missing Items High) and 2.05 (delay Low/Medium). This superlinear propagation is not a modelling artefact: it arises from batch-arrival dynamics in which recovery-path orders complete simultaneously and arrive at Consolidation in bursts, overwhelming the queue beyond what the linear propagation model in eq. (3.4) anticipates. These cases demonstrate that the static α formulation is insufficient under high-frequency endogenous disruption; a regime-switching formulation is required.

Demand starvation ($\hat{\alpha} < 0$): downstream improvement: Consolidation $\rightarrow$ Customs shows $\hat{\alpha} = -0.68$ and -0.71 under delay scenarios, consistent with the negative sojourn delta at Consolidation observed in Figure 4.6. Reduced factory throughput lowers the arrival rate at Customs, draining its queue below baseline and *improving* downstream performance. This starvation regime constitutes empirical validation of the sub-unity α formulation and demonstrates that $\hat{\alpha}_{ij}$ is operating-regime-dependent, not a fixed structural constant.

4.4.4 Disruption Intensity Effects and SAR

Delay distribution by disruption type

Figure 4.7 presents the per-order delay distributions Δ_k stratified by disruption type, making the structural differences between disruption families directly visible.




Figure 4.7: Delay distributions $\Delta_k = \text{ATD}_k - \text{COSD}_k$ by disruption type.

Four qualitatively distinct delay-generating processes are visible in Figure 4.7:

- **No Disruption:** right-skewed near-zero distribution (mean 9.5 d), consistent with log-normal-like stochastic pipeline delay. The mass to the left of $\Delta_k = 15$ days represents the 72–84% on-time population.
- **Delay Factory:** distributional near-identity with No Disruption (mean 10.7 d). The two histograms are visually indistinguishable, providing the strongest available visualisation evidence for the Mann-Whitney non-significance result ($p = 0.159$).
- **Missing Items:** structural mass shift rightward (mean 27 d), heavy tail to 175 d. A bimodal tendency is visible: one population follows the normal transit timeline; the other follows the extended recovery-path timeline. This is a categorical structural shift, not a variance increase.
- **Combined:** mean 29 d with tail extending to 275 d. The shape reflects two co-existing order populations – those whose transit windows avoid active disruption events, and those caught mid-chain during compound disruption. The 275-day maximum represents simultaneous accumulation across all nodes, the most extreme outcome in the dataset.

Super-additivity ratio (SAR) computation

The SAR quantifies whether compound disruption exceeds, equals, or falls short of the sum of individual disruption effects (eq. (4.3)). Using the baseline E1 ($\bar{\Delta}_0 = 9.78$ d), the D1 component from the comparable pure Missing Items experiment, and the D3 component from the closest delay experiment:

For the **E11 (Combined High/High)** compound scenario:

$$\bar{\Delta}_{D1}^{E6} = 43.78 \text{ d} \quad (\text{Missing Items High, rate 0.35})$$

$$\bar{\Delta}_{D3}^{E9} = 9.90 \text{ d} \quad (\text{delay High/High, closest match})$$

$$\bar{\Delta}^{E11} = 46.07 \text{ d}$$

$$\text{SAR}_{E11} = \frac{46.07 - 9.78}{(43.78 - 9.78) + (9.90 - 9.78)} = \frac{36.29}{34.00 + 0.12} = \frac{36.29}{34.12} \approx \mathbf{1.064} \quad (4.5)$$

The SAR of 1.064 indicates **mild super-additivity**: the compound effect exceeds the sum of individual effects by approximately 6.4 %. This result revises the initial sub-additivity interpretation: when the delay increment is computed correctly from the data ($\Delta_{D3} = 9.90 - 9.78 = 0.12 \text{ d}$, not an assumed $\sim 3 \text{ d}$), the Combined High/High scenario is *not* sub-additive. The mechanism is that factory delay eliminates the buffering capacity that would otherwise attenuate the missing-items signal, amplifying the compound outcome beyond naive additivity.

For the **E12 (Combined Low/Low)** compound scenario:

$$\bar{\Delta}_{D1}^{E4} = 10.35 \text{ d} \quad (\text{Missing Items Low, rate 0.15})$$

$$\bar{\Delta}_{D3}^{E7} = 12.11 \text{ d} \quad (\text{delay Low/Medium, closest match})$$

$$\bar{\Delta}^{E12} = 17.93 \text{ d}$$

$$\text{SAR}_{E12} = \frac{17.93 - 9.78}{(10.35 - 9.78) + (12.11 - 9.78)} = \frac{8.15}{0.57 + 2.33} = \frac{8.15}{2.90} \approx \mathbf{2.81} \quad (4.6)$$

The SAR of 2.81 at low intensity indicates **strong super-additivity**. At low disruption rates, neither D1 nor D3 alone generates significant delay, but their combination produces a disproportionate outcome: the co-occurrence of a missing-item event and a factory delay within the same pipeline window creates a compounded delay that neither mechanism would generate independently. This result is the most operationally significant SAR finding: compound risk at low individual severity is a meaningful resilience threat precisely because it is invisible in single-disruption analysis.

4.4.5 Modality-Level Analysis

Figure 4.8 shows delay distributions by equipment modality and disruption scenario, revealing portfolio-level heterogeneity in ripple exposure that is invisible at the aggregate level.




Figure 4.8: Delay distribution by equipment modality and disruption scenario.

MR: highest tail risk: MRI systems show the largest delay outliers in the dataset (exceeding 250 days under Missing Items and Combined disruptions, as visible in Figure 4.8), driven by MR's 140-day baseline lead time and high unit complexity. A recovery-path MR order must traverse an extended pipeline before delivery, generating the heaviest right tails. Despite this, MR maintains 67–100 % OTD across most experiments: its long lead time provides a wide COSD buffer that absorbs moderate delays without crossing the 15-day grace threshold. MR therefore combines high tail risk with relatively robust median performance.

RAD: highest transit sensitivity: Radiology (42-day lead time) shows OTD degradation across the broadest range of scenarios (29–65 %), including under delay Factory (53 % under E7). RAD's short lead time leaves minimal buffer for any upstream disturbance; it is the modality most sensitive to inherent transit variability even without structural disruption.

VASC, NM and PET: insufficient sample size: Nuclear Medicine and PET exhibit erratic patterns, including 0% OTD in some scenarios, driven by 1–3 orders per run. Results for these modalities are explicitly unreliable and should not be interpreted as indicative of their true disruption exposure. These modalities require targeted simulation with higher arrival rates or longer horizons.

Delay Factory: uniform near-baseline performance across modalities: Delay Factory produces small, near-baseline delay boxes across all modalities (Figure 4.8), including low-lead-time RAD. The modality-level analysis confirms that the delay insensitivity observed in aggregate is not an averaging artefact: it holds consistently at the equipment-type level.

Operational implication: Resilience management cannot be conducted at the aggregate portfolio level alone. Interventions must be modality-specific: pre-shipment parts verification and BOM completeness checking are most critical for CT (high volume) and MR (tail risk); RAD requires transit-time buffer management independent of disruption scenarios.

4.4.6 Statistical Validation

Non-parametric tests: Table 4.4 summarises the Kruskal-Wallis and Mann-Whitney U pairwise tests applied to the Δ_k distributions across disruption type groups, together with rank-biserial correlation effect sizes $r = Z/\sqrt{n}$.

Table 4.4: Statistical test results with effect sizes

Test	Statistic	p -value	Result	Effect r	Interpretation
Kruskal–Wallis (overall)	$H = 80.48$	2.42×10^{-17}	Sig.	—	At least one disruption type shifts Δ_k
Missing Items vs. No Disruption	$U = 76\,558$	2.53×10^{-10}	Sig.	$r = 0.43$	Large effect; endogenous disruption materially shifts Δ_k
Factory Delay vs. No Disruption	$U = 132\,247$	0.159	Non-sig.	$r = 0.06$	Negligible effect; delay indistinguishable from baseline
Combined vs. No Disruption	$U = 72\,856$	1.09×10^{-12}	Sig.	$r = 0.45$	Large effect; compound disruption amplifies endogenous signal

Interpretation of key results

Overall significance: The Kruskal-Wallis test ($H = 80.48$, $p < 10^{-17}$) establishes that disruption type has a statistically significant effect on per-order delay. This confirms that the ripple effect operates and is measurable in GE Algeria’s import-dependent MTO chain.

Missing Items: large and operationally significant effect. The rank-biserial correlation $r = 0.43$ for Missing Items versus No Disruption falls in the *large effect* range by conventional thresholds ($r > 0.3$). This is not merely a consequence of the large sample size: the effect size confirms that the delay shift is operationally meaningful – a randomly selected Missing Items order is more likely to be delayed than a randomly selected No Disruption order with probability ≈ 0.72 . The p -value alone cannot convey this magnitude; the effect size does.

Delay Factory: negligible effect, non-significance confirms the mechanism. The $p = 0.159$ result ($r = 0.06$) is not a limitation of the statistical test – it is the expected outcome of a chain where factory retard operates entirely within the normal processing path without triggering path-switching. The node-level analysis already demonstrated why: Consolidation queue drains under starvation, Customs receives fewer arrivals and processes faster, and the COSD buffer absorbs the modest lead-time increment. The terminal metric Δ_k is insensitive to retard intensity because the chain’s internal mechanisms – starvation decoupling, transit variance envelope, COSD buffering – collectively neutralise the signal before it reaches order completion. The statistical non-significance is therefore a direct validation of these mechanisms, not evidence that retard is unimportant at higher intensities than tested.

Disruption interaction: factory delay as catalyst. The contrast between delay Factory ($r = 0.06$, non-significant) and Combined ($r = 0.45$, $p < 10^{-12}$) provides statistical evidence for a disruption interaction effect. Factory delay alone produces no measurable order-level delay; combined with missing-items disruption, it amplifies the outcome (SAR = 1.06, eq. (4.5)). Delay eliminates the buffering capacity that would otherwise attenuate the missing-items signal, rendering it visible at the terminal metric level. The dominant driver of delay in the combined scenario is the missing-items disruption; factory delay acts as an enabling, not a primary, cause.

4.4.7 Summary of Findings

The fifteen simulation experiments yield eight thesis-level findings that collectively advance ripple effect modelling for import-dependent MTO/JIT medical equipment supply chains.

- F1. Ripple effect is confirmed and statistically significant:** The Kruskal-Wallis test ($H = 80.48$, $p < 10^{-17}$) validates that disruption-driven delay propagation operates in GE Algeria’s chain – a context not previously addressed in the Ivanov/Dolgui/Sokolov corpus. The Missing Items effect size ($r = 0.43$) confirms operational, not merely statistical, significance.
- F2. The recovery path is a damage-control tool, not a resilience mechanism:** Recovery-path orders achieve near-zero OTD across all experiments. Air expediting reduces delay magnitude but cannot restore committed delivery dates; administrative re-processing overhead dominates transit-time savings. True resilience requires upstream prevention – pre-shipment order verification and BOM completeness checking – rather than recovery path investment.
- F3. Asymmetric disruption sensitivity: endogenous $\gg$ exogenous:** The chain is structurally insensitive to exogenous factory delay ($r = 0.06$, $p = 0.159$) but highly sensitive to endogenous missing-items disruption ($r = 0.43$, $p < 10^{-10}$). The asymmetry is qualitative: missing items trigger path-switching (recovery loop), while factory delay operates within the normal pipeline without activating a structurally different processing path. The chain’s transit variability envelope absorbs the delay; no such absorption exists for path-switching events.
- F4. Field Engineers is the terminal ripple accumulator:** Contrary to the prior expectation that Customs would be the dominant bottleneck, Field Engineers carries the largest absolute sojourn inflation: +27.9 days (E6) and +23.7 days (E11), as shown in Figure 4.5. Customs remains the intermediate amplifier (+11.3 days maximum). Terminal accumulation at Field Engineers is driven by scheduling disruption and extended waiting for recovery-path equipment.
- F5. Three distinct $\hat{\alpha}_{ij}$ propagation regimes:** The inter-stage attenuation factor spans three regimes: attenuation ($\hat{\alpha} \in (0, 1]$, dominant under moderate disruption), amplification ($\hat{\alpha} > 1$ at

Factory $\rightarrow$ Consolidation under Missing Items High, driven by batch-arrival overflow), and demand starvation ($\hat{\alpha} < 0$ at Consolidation $\rightarrow$ Customs under delay, driven by upstream throughput suppression). The static α formulation of eq. (3.4) requires regime-switching extension to capture the full range of observed behaviour.

- F6. Compound disruption is mildly super-additive at high intensity:** $\text{SAR}_{E11} = 1.06$: Combined High/High exceeds the sum of individual Missing Items High and delay High/High increments by 6.4 %. At low intensity, $\text{SAR}_{E12} = 2.81$, indicating strong super-additivity when neither disruption alone would produce significant delay. Compound risk at low individual severity is a meaningful resilience threat invisible in single-disruption analysis.
- F7. Modality-differentiated vulnerability requires portfolio-level management:** MR exhibits the highest tail risk (outliers > 250 days, Figure 4.8) but robust median OTD. RAD is most sensitive to transit variability (OTD 29–65 % across disruption types). Delay Factory leaves all modalities near-baseline. Resilience interventions must be modality-specific and cannot be designed at aggregate portfolio level.
- F8. Compound disruption prevents steady-state convergence:** Under Combined Medium/Medium, OTD declines monotonically from 65 % (800d) to 48 % (2 000d) without convergence, as shown in Figure 4.4. The 1 200-day primary window underestimates long-run compound disruption impact by ≈ 5 points OTD and 5 days mean delay. Real-world GE Healthcare Algeria performance assessments over 3–4 year periods may encounter compounding degradation not visible in annual KPI reviews.

Collectively, these findings establish that GE Healthcare Algeria’s primary resilience risk is the endogenous, intra-firm missing-items disruption – not the exogenous factory delays that dominate the existing ripple effect literature. The most effective resilience intervention is upstream order verification, not downstream recovery path investment. Customs and Field Engineers serve as the chain’s intermediate amplifier and terminal accumulator respectively, and should be the focus of capacity and process resilience investment.

4.5 Conclusion

This chapter has presented the simulation model, experimental design, and analytical results that constitute the empirical core of the dissertation. Fifteen simulation runs across four disruption families produced eight thesis-level findings, whose collective implication is unambiguous: GE HealthCare Algeria’s primary resilience risk is the endogenous, intra-firm missing-items disruption, not the exogenous factory delays that dominate the existing ripple effect literature.

Three results deserve particular emphasis as contributions to the corpus. First, the asymmetric disruption sensitivity – endogenous D1 generating a large, statistically significant effect ($r = 0.43$, $p < 10^{-10}$) while exogenous D3 remains statistically indistinguishable from baseline ($r = 0.06$, $p = 0.159$) – establishes that the dominant propagation mechanism in import-dependent MTO chains is path-switching, not lead-time extension. Second, compound disruption is super-additive at both tested intensity levels ($\text{SAR}_{E11} = 1.064$; $\text{SAR}_{E12} = 2.81$), with the low-intensity result being the more operationally consequential: neither D1 nor D3 alone produces significant delay at low rates, but their co-occurrence generates disproportionate damage invisible to single-disruption analysis. Third, compound disruption prevents steady-state convergence within the 1 200-day primary horizon, meaning that long-run performance assessments substantially underestimate persistent compound disruption impact.

These findings directly inform the mitigation strategy framework developed in the following chapter, which maps the identified propagation mechanisms onto Ivanov’s four-layer control hierarchy and translates them into actionable resilience recommendations calibrated to GE HealthCare Algeria’s zero-inventory, import-dependent architecture. However, translating simulation findings into managerial action requires more than a static framework – it requires an interface through which decision-makers can interrogate scenarios, explore trade-offs, and receive grounded recommendations without depending on simulation expertise. The following chapter presents SC Copilot, an agentic AI decision-support layer that operationalises exactly this capability by transforming the AnyLogic simulation model into an accessible natural-language platform for resilience analysis.

Chapter 5

SC Copilot: Design and Implementation of an AI-Powered Decision Support Layer for Supply Chain Simulation

5.1 Introduction

This chapter introduces SC Copilot, an agentic AI decision-support layer developed to transform the AnyLogic simulation model of GE HealthCare Algeria's supply chain into an accessible managerial tool. While the simulation accurately reproduces the complete order fulfilment process and enables disruption testing, its practical use remains limited by the need for manual scenario configuration, simulation execution, data extraction, and result interpretation. Effective use of the model typically requires competencies in simulation modelling, data analysis, and industrial engineering, restricting its accessibility to a limited group of specialists.

SC Copilot addresses these limitations by providing a natural-language interface through which managers can describe disruption scenarios without interacting directly with the simulation model or its technical components. The system automatically translates user requests into simulation parameters, executes the model, analyses the outputs, and generates decision-oriented recommendations. In doing so, it reduces the expertise barrier associated with simulation-based decision making and broadens access to resilience analysis capabilities. The remainder of the chapter presents the conceptual foundations, design objectives, architecture, implementation, and evaluation of SC Copilot as an AI-enhanced decision-support system for supply chain resilience analysis.

5.2 Contextual Positioning of Simulation and Agentic Decision Support

5.2.1 The Gap Between Simulation and Decision

Simulation models have long been recognised as powerful tools for supply chain analysis, particularly through discrete-event simulation, which captures stochastic processes, resource constraints, disruptions, and lead-time dynamics with high realism. However, a persistent challenge remains

the translation of simulation outputs into actionable managerial insights. Simulation results are often complex and require specialised expertise to interpret, while non-expert decision-makers may struggle to link statistical indicators to operational actions. Additionally, exploring alternative scenarios generally requires direct interaction with model parameters, limiting accessibility. To address these challenges, Intelligent Decision Support Systems (IDSS) have been developed to bridge the gap between analytical models and decision-making. More recently, large language models (LLMs) have emerged as a promising solution by providing a natural language interface capable of interpreting simulation outputs and supporting managerial decision processes within a unified framework.

5.2.2 Agentic AI Systems

Our developed solution, SC Copilot, belongs to the emerging class of agentic AI systems. Unlike conventional conversational AI, agentic systems pursue user-defined goals through autonomous execution of multi-step workflows involving planning, tool usage, and adaptive decision-making. Their key characteristics include goal-oriented behaviour, autonomous orchestration of actions, integration with external tools, and the ability to adjust execution based on intermediate outcomes.

SC Copilot exhibits these characteristics by transforming a natural-language request into an end-to-end simulation workflow. Upon receiving a user objective, it automatically extracts simulation parameters, initiates the execution pipeline, runs the AnyLogic model, monitors generated outputs, and performs KPI analysis. Throughout this process, the system adapts to execution states by detecting file readiness, handling errors, and managing follow-up interactions. The final response is therefore not a generic language-model output but a synthesis grounded in the specific numerical results produced by the simulation run.

5.3 Business Objectives and Design Requirements

5.3.1 Operational Context and Stakeholder Needs

As previously discussed, GE HealthCare Algeria operates a fully import-dependent, make-to-order supply chain for medical equipment across multiple modalities and supply chain stages. This configuration exposes operations to several sources of uncertainty, particularly those related to upstream supply performance, customs clearance processes, and local installation capacity. These dimensions constitute the principal sources of uncertainty represented in the simulation model in the previous chapter.

Our solution was designed to support the decision-making needs of logistics and project management stakeholders responsible for maintaining service performance under uncertainty. During the internship, several key requirements were identified. Decision-makers required rapid access to insights, reducing the time needed to evaluate scenarios from hours of manual analysis to a matter of minutes. The solution also needed to be accessible to non-simulation experts by allowing scenarios to be specified in natural language rather than through direct manipulation of model parameters. Furthermore, all recommendations and conclusions had to remain grounded in simulation outputs, ensuring that managerial guidance could be traced to specific performance indicators rather than generic heuristics. The ability to compare alternative scenarios and evaluate the trade-offs associated with different mitigation strategies was also identified as a critical requirement. Finally, results needed to be communicated in a concise, executive-oriented format suitable for management discussions without requiring additional interpretation by simulation specialists.

5.3.2 Formalised Design Requirements

To translate stakeholder needs into engineering requirements, the following formal objectives were established for SC Copilot. Seven formal design requirements govern the development of SC Copilot.

DR-01 – Natural Language Interface. The system shall accept scenario descriptions in free-form text and autonomously extract all required simulation parameters without user interaction with model internals.

DR-02 – End-to-End Automation. The system shall execute the full pipeline (parameter extraction, simulation launch, output collection, KPI computation, and LLM briefing) as a single automated workflow triggered by a single user action.

DR-03 – Executive Output Quality. LLM-generated analysis shall use business vocabulary exclusively, grounded in exact simulation figures, and structured in a fixed five-section format suitable for board-level presentation.

DR-04 – Conversational Follow-up. After a simulation run, the system shall support multi-turn conversational queries about the results, maintaining simulation context across turns.

DR-05 – Scenario Comparison. The system shall maintain a historical record of all simulation runs and support side-by-side comparison of up to four scenarios with delta indicators and a composite recommendation.

DR-06 – Modality-Level Granularity. KPI analysis shall be disaggregated by equipment modality (CT, MR, RAD, MAMMO, VASC, NM, PET) to support targeted operational decisions.

DR-07 – Robustness. The pipeline shall handle AnyLogic execution timing variability, file write latency, and LLM API failures gracefully, with appropriate fallback mechanisms.

5.4 System Architecture

5.4.1 Overview

SC Copilot is built around five main components that work together under the control of a central coordination layer. As shown in Figure 5.1, these components are: (1) the Natural Language Scenario Interpreter, (2) the Pipeline Orchestrator, (3) the AnyLogic Simulation Engine, (4) the Data Processing and KPI Engine, and (5) the LLM Analysis and Conversation Engine. Information is exchanged through files, which allows the Python-based components and the Java-based AnyLogic model to communicate without requiring any changes to the simulation model itself.




Figure 5.1: SC Copilot high-level architecture

5.4.2 Technology Stack

The technology choices reflect the need to integrate three heterogeneous environments: a Java-based simulation platform (AnyLogic), a Python-based data science and web framework ecosystem (Streamlit, pandas), and cloud-hosted large language model APIs. The full technology stack is summarised in Table 5.1.

Table 5.1: SC Copilot technology stack.

Component	Role	Key Technology
AnyLogic 8.x	Discrete-event simulation engine	Java runtime, proprietary model format (.alp)
Python 3.10+	Orchestration, data processing, UI	subprocess, pathlib, threading, queue
Streamlit	Web-based dashboard and chat interface	Component-based reactive UI, session state
Groq API	LLM inference (llama-3.3-70b-versatile)	Low-latency inference, JSON-structured output
pandas	CSV data processing, KPI computation	DataFrame operations, groupby aggregation
pyautogui / pygetwindow	AnyLogic window automation	Win32 API via ctypes, F5 keystroke injection

5.4.3 File-Based Integration Layer

A critical architectural decision was to use the filesystem as the integration medium between the Python orchestration layer and the AnyLogic simulation engine. This choice was primarily driven by the technical limitations of the AnyLogic Personal Learning Edition (PLE) used during the development of this project. Unlike the Professional and Research editions, the PLE version does not provide access to the APIs and integration capabilities required for direct communication between external applications and the simulation model. Consequently, implementing API-based interactions was not feasible within the available environment.

To overcome this constraint, a file-based input/output mechanism was adopted. This approach enables reliable data exchange between the Python components and the simulation model without requiring inter-process communication protocols, AnyLogic API hooks, or custom Java extensions. The integration contract is defined by three files:

- `scenario.csv`: written by the NL Interpreter (`app.py`), read by AnyLogic at simulation startup. Contains one row with six simulation parameter values.
- `orders_tab.csv`: written by AnyLogic at simulation completion. Contains one row per completed order, with columns for modality, delivery dates, delay days, on-time flag, and disruption indicators. This is the primary data artefact consumed by the processing and LLM layers.
- `simulation_history.csv`: written by the KPI Engine (`process_results.py`) after each run. Contains one row per completed simulation, aggregating headline KPIs and modality-level metrics as a JSON-encoded column (`modality_json`). This file drives the scenario comparison panel.

Beyond addressing the limitations of the PLE version, this design also ensures loose coupling

between the simulation model and the agentic system. As a result, SC Copilot remains largely non-invasive with respect to the AnyLogic model, requiring only file read/write operations while avoiding modifications to the model's core simulation logic. This architecture also facilitates future migration toward API-based integration should Professional or Research licenses become available.

5.5 Agent Components

5.5.1 Component 1 – Natural Language Scenario Interpreter

5.5.1.1 Purpose and Approach

The Natural Language Scenario Interpreter (implemented in `app.py`) is the entry point of the SC Copilot pipeline. Its function is to transform an unstructured manager scenario description into a precisely defined set of six numerical simulation parameters, together with a confidence score and a human-readable rationale explaining the mapping.

The interpreter is powered by the Groq-hosted `llama-3.3-70b-versatile` model. The choice of a 70-billion parameter instruction-tuned model for this task reflects the need for reliable structured output extraction in the presence of ambiguous, idiomatic, and domain-specific natural language. The model is prompted with a carefully engineered system instruction that defines the six parameters, their ranges, their default values, and the required output schema.

5.5.1.2 Simulation Parameter Space

The six parameters that define a simulation scenario are presented in Table 5.2.

Table 5.2: Simulation parameter space managed by the NL Scenario Interpreter.

Parameter	Description
D1_RATE_BY_YEAR	Expected missing-parts disruption events per year. Range 0–12. Primary upstream supply risk: a BOM component unavailable at the factory level that prevents shipment.
D3_RATE_BY_YEAR	Expected factory production delay events per year. Range 0–10. Models production overrun events causing late dispatch from the manufacturing site.
FE_CAPACITY	Number of available field engineers in Algeria. Range 1–20. The downstream installation resource that determines last-mile throughput rate.
ORDER_INTERARRIVAL	Mean days between incoming equipment orders. Range 2.0–14.0. Controls the demand intensity arriving at the supply chain.
URGENT_FRACTION	Fraction of orders classified as urgent. Range 0.05–0.80. Urgent orders carry tighter COSD commitments and higher penalty exposure.
SIM_DURATION	Total simulation duration in days. Range 365–5000. Configurable to support short stress-tests or multi-year strategic analyses.

5.5.1.3 LLM Prompting Strategy

The system prompt provided to the LLM specifies the parameter extraction task with a high degree of constraint. Several engineering choices merit discussion:

- **Temperature = 0.2:** a low sampling temperature is applied to maximise determinism in parameter extraction. Scenario interpretation is not a creative task; variability in parameter values across repeated calls for the same input would undermine reproducibility.
- **JSON-only output:** the LLM is instructed to return exclusively a valid JSON object. The parsing logic includes a secondary extraction step using regular expressions to recover valid JSON if the model adds surrounding text despite instructions.
- **Quantitative reasoning:** the prompt explicitly instructs the model to apply arithmetic when quantities are given, rather than treating the description as purely qualitative.
- **Confidence score:** the LLM returns a confidence value between 0 and 1 reflecting the precision of the mapping. This score is surfaced in the UI to alert users when a scenario description was ambiguous and the parameter mapping may not fully reflect their intent.
- **Rule-based fallback:** if the Groq API call fails for any reason (rate limit, network error, quota exhaustion), the interpreter falls back to a keyword-matching rule engine that maps common phrases such as “high disruption”, “reduced capacity”, or “surge demand” to predefined parameter sets. This ensures the pipeline degrades gracefully rather than failing completely.

The full decision logic of the interpreter, including LLM invocation and fallback routing, is represented in Figure 5.2.




Figure 5.2: Decision flowchart for the Natural Language Scenario Interpreter

5.5.1.4 Parameter Validation

All LLM-extracted parameters are passed through a validation function before being written to `scenario.csv`. The validator clamps each value to its defined range, applies type coercion (`FE_CAPACITY` and `SIM_DURATION` are cast to integers), and substitutes defaults for any parameter absent from the LLM response. This validation step is critical: LLM outputs, despite low-temperature settings, may occasionally produce values outside the simulation model’s valid input space, which would cause AnyLogic to fail silently or produce undefined behaviour.

5.5.2 Component 2 – Pipeline Orchestrator

5.5.2.1 Role and Structure

The Pipeline Orchestrator (`run_all.py`) is the central coordination component of SC Copilot. It is responsible for sequencing the five execution phases, managing inter-component timing, handling AnyLogic process automation, monitoring file system state, and emitting structured log events consumed by the UI layer. The orchestrator runs as a subprocess launched by the Streamlit UI in a background daemon thread, enabling the dashboard to remain responsive and display live progress while the pipeline executes.

5.5.2.2 Five-Phase Execution Model

The pipeline is structured into five phases, each corresponding to a distinct operational stage. The phase structure is not merely organisational: phase transition events are emitted as structured log lines that the UI parses to drive the animated progress badge strip displayed on the dashboard. The complete sequence of phases and their inter-dependencies are represented as a UML activity diagram in Figure 5.3.

The SC Copilot pipeline is structured into five sequential phases.

Phase 1 – AI Scenario Interpreter: The orchestrator invokes `app.py` as a subprocess with the scenario text piped via stdin, then validates that `scenario.csv` and `scenario_rationale.json` were successfully produced before allowing the pipeline to proceed.

Phase 2 – Pipeline Initialisation: Stale output files from the previous run, specifically `orders_tab.csv` and `remediation_scenario.csv`, are purged to ensure that the data processing phase does not inadvertently consume residual data.

Phase 3 – AnyLogic Simulation: AnyLogic execution is automated via `pyautogui` through F5 keystroke injection after window focus acquisition. The orchestrator then enters a polling loop that waits for `orders_tab.csv` to appear on disk before advancing.

Phase 4 – Data Stream Monitoring: The size of `orders_tab.csv` is polled at two-second intervals until three consecutive stable readings confirm that AnyLogic has finished writing. This stability check prevents premature data ingestion.

Phase 5 – KPI Diagnostics: `process_results.py` is invoked as a subprocess to compute KPIs, compare them against the baseline, detect remediation needs, and append a summary row to `simulation_history.csv`.




Figure 5.3: UML Activity Diagram of the SC Copilot pipeline orchestrator across three execution domains

5.5.2.3 AnyLogic Automation Strategy

The automation of AnyLogic execution is one of the more technically constrained aspects of the system design. AnyLogic does not expose a command-line execution API for pre-built experiments; triggering a simulation run from an external process requires interaction with the AnyLogic IDE window. SC Copilot addresses this through a two-step approach.

First, `pygetwindow` is used to locate the AnyLogic workspace window by matching its title against known fragments. The Win32 `SetForegroundWindow` API is then called to bring the window to the foreground with OS-level authority, a necessary step because Python-level window focus requests are often blocked by Windows focus-steal prevention mechanisms. Second, any lingering simulation runtime frames from previous runs are closed programmatically before the new run is triggered. Finally, `pyautogui` injects an F5 keystroke into the focused AnyLogic window, which AnyLogic interprets as the “Run” command in its experiment panel.

If the automation sequence fails (AnyLogic window not found, focus acquisition blocked), the orchestrator emits an `ANYLOGIC_WAIT` signal to the UI, which displays a prominent action-required alert instructing the user to switch to AnyLogic and press Run manually. This fallback ensures the pipeline can proceed even when OS-level automation is unavailable.

5.5.2.4 File Stability Monitoring

The file stability monitoring logic in Phase 4 addresses a practical challenge: AnyLogic writes `orders_tab.csv` progressively during simulation execution. Detecting file creation is insufficient, as the file may be partially written when detected. The stability algorithm waits for three consecutive file size measurements, taken at 2-second intervals, to be identical. Only when the file size has not changed across three consecutive polls does the orchestrator conclude that AnyLogic has finished writing, and proceed to Phase 5. A 2-hour absolute timeout protects against scenarios where AnyLogic fails silently without producing output.

5.5.3 Component 3 – Data Processing and KPI Engine

5.5.3.1 Input Data Structure

The primary data artefact produced by the AnyLogic simulation is `orders_tab.csv`. Each row represents one completed order in the simulated supply chain. The columns relevant to SC Copilot’s analysis are presented in Table 5.3.

Table 5.3: Key columns of the `orders_tab.csv` data artefact produced by AnyLogic.

Column	Description
<code>modality</code>	Equipment type: CT, MR, RAD, MAMMO, VASC, NM, PET.
<code>atd</code>	Actual time of delivery (simulation time units, days from epoch).
<code>w2w_start_time</code>	Order creation time. W2W = Wing-to-Wing.
<code>delay_days</code>	Delivery delay relative to COSD: ATD minus COSD. Negative values indicate early delivery.
<code>on_time</code>	Boolean flag: True if <code>delay_days</code> is at most 15 (grace window), aligned with the 15-day OTD definition.
<code>is_missing</code>	Boolean flag: True if the order experienced a missing-parts disruption event.
<code>is_retard</code>	Boolean flag: True if the order experienced a factory production delay event.
<code>cosd</code>	Customer Order Ship Date – the committed delivery date.

5.5.3.2 KPI Computation

The KPI engine (`process_results.py`) computes a two-level KPI hierarchy: headline metrics at the fleet level, and granular metrics disaggregated by modality.

At the headline level, the engine computes: total throughput (number of completed orders), average delay days, on-time delivery rate (OTD), missing-parts incidence rate, factory delay incidence rate, and compound disruption count (orders affected by both disruption types simultaneously). The OTD rate is computed using the 15-day grace window aligned with GE HealthCare’s contractual definition.

At the modality level, for each equipment type group, the engine computes: order count, average delay days, worst-case delay buffer (95th percentile delay), average W2W lead time (computed as ATD minus order creation time), missing-parts count, and per-modality OTD rate. These modality-level metrics are serialised as a JSON string and stored in the `modality_json` column of `simulation_history.csv`, making them available to the LLM without requiring re-reads of the full orders dataset.

5.5.3.3 Baseline Comparison

Each simulation run is automatically compared against a pre-established baseline derived from Scenario E1, a no-disruption, full-capacity, three-year simulation producing 152 completed orders with a documented OTD of 73.7% and a mean delay of 11.4 days. The baseline is stored as `baseline.csv` and read by `process_results.py` at the start of each evaluation cycle.

The comparison logic computes three variance indicators: throughput variance (percentage change relative to baseline), delay variance (percentage change in average delay), and OTD variance (percentage point change in on-time rate). Based on these variances, the engine classifies the current run into one of three operational status categories: STABLE, HIGH BACKLOG, or THROUGHPUT COLLAPSE. If the classification indicates a degraded state, a remediation scenario (`remediation_scenario.csv`)

is automatically generated with conservative parameter recommendations aimed at restoring baseline performance.

5.5.4 Component 4 – LLM Analysis and Conversation Engine

5.5.4.1 Design Philosophy

The LLM analysis engine is the component responsible for transforming structured KPI data into executive-grade narrative intelligence. This is not a generic summarisation task. The engine is designed around three specific engineering constraints: absolute prohibition of statistical jargon in output, strict grounding in simulation-specific numerical values, and maintenance of a fixed output structure compatible with board-level briefing formats.

Three distinct LLM roles are implemented, each with its own system prompt and invocation logic: the Post-Simulation Analysis prompt (for single-run executive briefings), the Scenario Comparison prompt (for multi-run delta analyses), and the Conversational Follow-up prompt (for open-ended queries in the conversational interface).

5.5.4.2 The Post-Simulation Analysis Prompt

The post-simulation analysis prompt is the most elaborately engineered component of the system. It establishes the LLM in the persona of the Director of Supply Chain Operations for GE HealthCare Algeria, briefing the Country Managing Director and CFO after a simulation stress test. Key design elements are as follows:

- **Banned terms list:** a set of statistical terms (P50, P95, Mean, Percentile, Distribution, Standard Deviation, and academic codes F1–F8) are explicitly banned. The prompt provides mandatory business-language replacements: P50 becomes “Typical Delivery Horizon”, P95 becomes “Worst-Case Scenario Buffer”, and Standard Deviation becomes “variability in delivery performance”. The silent translation rule further instructs the LLM to translate banned terms present in the user’s own questions back into executive language before answering.
- **Data grounding mandate:** the LLM is required to extract figures from the `modality_json` payload injected into the context and cite them directly. It is forbidden to produce placeholder ranges or estimate values not present in the data. If a modality is absent, it must state this explicitly.
- **Fixed five-section output format:** Executive Summary, Root Cause & Bottleneck Diagnosis, Modality & Hospital Impact, Estimated Financial Risk, and Strategic Action Plan. This structure ensures consistent output regardless of the specific scenario, enabling systematic comparison across briefings.
- **Length constraint:** 350–500 words, to balance thoroughness with executive attention span.

5.5.4.3 Query Routing Logic

A key architectural feature of the LLM engine is the query routing mechanism implemented in `app_ui.py`. When the user submits a message, the system must determine whether it constitutes a new simulation request (requiring pipeline execution) or a conversational follow-up on existing results (requiring only LLM inference).

The routing function (`_is_simulation_request`) applies a two-stage decision logic. First, it checks for retrospective signal phrases such as “what was”, “explain”, “why did”, “show me”, or “the simulation”, which indicate the user is asking about existing results rather than requesting a new run. If such phrases are detected, the message is routed to the conversational follow-up engine regardless of other content. Second, it scans for simulation-intent keywords such as “simulate”, “what if”, “stress test”, “what happens if”, or “adjust parameters”, which trigger pipeline execution. This routing logic correctly handles the majority of real-world query patterns observed during system testing.

5.5.4.4 Conversational Context Management

The conversational follow-up engine maintains a rich simulation context across turns. When a simulation run completes, the system captures and stores in Streamlit session state: the full structured orders data dictionary, the headline KPIs, the scenario rationale, and the scenario parameters. This context is serialised into a structured text block and injected into the system prompt of every subsequent conversational LLM call, ensuring that follow-up answers are grounded in the specific numerical results of the most recent run rather than in generic domain knowledge.

The conversational history, comprising all prior user messages and assistant responses, is also included in each API call as a sequence of role-tagged messages, enabling the LLM to refer back to prior exchanges naturally. A temperature of 0.4 (slightly higher than the 0.2 used for parameter extraction) is applied to conversational responses, allowing slightly more varied phrasing while maintaining factual grounding.

5.5.5 Component 5 – Executive Dashboard (Streamlit UI)

5.5.5.1 Architecture of the Interface

The SC Copilot dashboard is implemented in Streamlit (`app_ui.py`) and serves as the unified interface through which managers interact with the entire system. The dashboard is designed around a top-down executive information flow: the most critical information (headline KPIs) appears at the top, followed by execution state indicators, then the conversational interface and detailed data panels.

The interface is structured into three functional zones: a left sidebar containing scenario input controls, simulation duration configuration, pipeline status indicators, and system logs; a main content area organised in three tabs (Copilot, Data Ledger, Scenario Comparison); and a real-time auto-refresh mechanism (1.5-second polling interval via `streamlit-autorefresh`) that enables live pipeline progress display without manual page reloads. The overall layout of the dashboard is presented in Figure 5.4.

5.5.5.2 KPI Display and Data Ledger

Three headline KPI cards are rendered at the top of the main content area immediately upon simulation completion: Throughput (total orders processed), Typical Delay Horizon (average delay days against COSD commitment), and On-Time Delivery Rate (OTD percentage against the 80% target). The OTD card applies colour coding: green for OTD at or above 80%, amber for 60–79%, and red below 60%. These thresholds correspond directly to the operational classification logic in `process_results.py`.

The Data Ledger tab renders the full orders dataset as a formatted financial ledger table, disaggregated by modality. For each modality, the table displays order count, average delay days, worst-case delivery buffer, average W2W lead time, missing-parts count, and per-modality OTD rate. A separate disruption impact table shows the split between orders affected by missing-parts events and clean orders, with their respective average delay profiles. This modality-level view is the primary tool for identifying which equipment types require targeted intervention.




Figure 5.4: Layout of the SC Copilot executive dashboard

5.5.5.3 State Management and Concurrency

Streamlit’s execution model, which re-runs the entire script on each user interaction, creates non-trivial state management challenges for a multi-phase asynchronous pipeline. SC Copilot addresses this through a combination of session state persistence and background threading.

All pipeline state variables (`pipeline_running`, `pipeline_done`, `phase`, `last_kpis`, `last_orders_data`, `log_lines`, `scenario_history`) are stored in Streamlit’s `session_state` dictionary, which persists across script re-runs within a user session. The pipeline itself executes in a daemon background thread, communicating with the main thread via a thread-safe queue (Python’s `queue.Queue`). The main script’s polling loop drains this queue on each auto-refresh cycle, updating session state variables that drive the UI display. This architecture ensures that the pipeline’s progress is reflected in the UI in near-real time without blocking the Streamlit event loop.

5.6 Implementation Considerations and Engineering Decisions

5.6.1 Prompt Engineering Constraints

The prompt engineering effort for SC Copilot went through multiple iteration cycles before converging on the final system prompts. The most persistent challenge was preventing the LLM from reverting to statistical vocabulary, particularly the terms P50, P95, and “distribution”, when discussing delay profiles. Initial prompt versions achieved compliance in most cases but failed on

adversarial inputs where the user’s question contained these terms. The final solution was the silent translation rule, which instructs the LLM to detect banned terms in the user’s input and answer using the executive replacement vocabulary without acknowledging the translation. This approach achieved consistent compliance across all tested input patterns.

A secondary prompt engineering challenge was preventing hallucination of modality-level figures. Without an explicit data grounding mandate, the LLM would occasionally produce plausible but fabricated lead time estimates when a modality had limited representation in the orders dataset. The solution was a two-part constraint: (i) an explicit instruction to read modality figures only from the injected `modality_json` payload, and (ii) a prescribed fallback phrase “No delivery records for [modality] were captured in this simulation run” to be used when a modality is genuinely absent, preventing the LLM from filling the gap with estimates.

5.6.2 API Rate Limiting and Fallback

The Groq API applies rate limits that can be encountered during intensive scenario exploration sessions involving multiple rapid-fire simulation runs. The system handles API failures gracefully: if the post-simulation analysis call fails, the assistant message records the error and displays the raw KPI values from session state, ensuring the user retains access to the quantitative results even without the narrative briefing. The rule-based fallback in the NL interpreter similarly ensures that parameter extraction remains available even during API outages.

5.6.3 Environment Configuration

SC Copilot uses a `.env` file co-located with the main scripts for API key management (`GROQ_API_KEY`, `OPENAI_API_KEY`). The `.env` file is parsed at import time using a custom line-by-line parser rather than the `python-dotenv` library, reducing the dependency footprint. Parsed keys are injected into `os.environ` and propagated to subprocesses via the `env` dictionary passed to `subprocess.Popen`, ensuring that the LLM API keys are available throughout the pipeline without requiring system-level environment variable configuration.

5.6.4 Cross-Platform Considerations

The AnyLogic window automation component (`focus_anylogic_and_run`) relies on Win32 API calls via `ctypes` and `pyautogui`, making it Windows-specific. The remainder of the system (LLM calls, data processing, Streamlit UI) is fully cross-platform. The orchestrator handles the absence of `CREATE_NO_WINDOW` on non-Windows systems via a `try/except` block, allowing the pipeline to execute on macOS or Linux without the AnyLogic automation step, falling through to the manual action-required prompt.

5.7 Evaluation Against Design Requirements

5.7.1 Functional Assessment

Table 5.4 evaluates SC Copilot against the design requirements established in Section 5.3.2, based on observed system behaviour during the simulation campaign conducted as part of this dissertation.

Table 5.4: Evaluation of SC Copilot against design requirements.

ID	Requirement	Assessment
DR-01	Natural Language Interface	Fully met. The LLM interpreter correctly extracted parameters from all tested scenario descriptions, including quantitative formulations (“30% fewer engineers”) and qualitative framings (“severe disruption”). Confidence scores below 0.6 correctly flagged ambiguous inputs.
DR-02	End-to-End Automation	Substantially met. The full pipeline executes autonomously in typical operating conditions. The AnyLogic window automation succeeds reliably when the workspace is open and the IDE is responsive. The manual fallback activates in edge cases (window not found, focus blocked).
DR-03	Executive Output Quality	Fully met. Compliance with the banned terms list was achieved across all tested runs after the silent translation rule was implemented. All five output sections are consistently produced with figures grounded in the specific run data.
DR-04	Conversational Follow-up	Fully met. The routing logic correctly distinguishes new simulation requests from follow-up queries in all tested cases. Conversational context is maintained across multi-turn sessions.
DR-05	Scenario Comparison	Fully met. The comparison panel correctly renders delta indicators and recommendations for up to four historical scenarios. The composite scoring logic produces actionable recommendations.
DR-06	Modality Granularity	Fully met. All seven modalities (CT, MR, RAD, MAMMO, VASC, NM, PET) are disaggregated in the KPI engine and surfaced in both the data ledger and the LLM narrative.
DR-07	Robustness	Substantially met. API failure fallbacks, file stability monitoring, and AnyLogic automation fallbacks were all tested and function as designed. The 2-hour AnyLogic wait timeout was not reached in any test run.

5.7.2 Value Contribution

Beyond functional compliance, SC Copilot demonstrates three forms of value contribution relevant to the broader research question of this dissertation – how to operationalise supply chain resilience analysis in a practitioner context.

First, it dramatically reduces the time-to-insight cycle. A scenario analysis that previously required 30–90 minutes of analyst time (parameter encoding, simulation execution, output extraction, nar-

rative synthesis) is compressed to 5–15 minutes of elapsed time with minimal user effort.

Second, it democratises access to the simulation model. The Logistics Manager, who is not a simulation practitioner, can independently explore scenarios without depending on an analyst to translate their questions into model parameters and their results into interpretable findings.

Third, it enforces rigorous communication standards. The banned-terms enforcement and fixed output structure ensure that every briefing meets the same quality standard, regardless of the complexity of the underlying scenario. This consistency is valuable in an institutional context where simulation results must be communicated to decision-makers who are not simulation experts.

5.8 Conclusion

This chapter presented SC Copilot, an AI-driven decision support layer that transforms the AnyLogic simulation model of GE HealthCare Algeria's supply chain into an interactive natural-language platform. By combining LLM-based scenario interpretation, simulation orchestration, automated KPI processing, executive-level analysis generation, and an intuitive dashboard interface, the system bridges the gap between simulation outputs and managerial decision-making.

The proposed architecture enables non-expert users to conduct resilience analyses through natural language interactions while preserving the analytical rigor of the underlying simulation model. Its modular design allows individual components to be upgraded or replaced independently without affecting the overall system, ensuring scalability and future extensibility. As such, SC Copilot represents a validated framework for AI-assisted simulation-based decision support in supply chain resilience management.

General Conclusion

This dissertation set out to address a practical and theoretically underexplored challenge: how to analyse, quantify, and support the management of supply chain disruption propagation in an import-dependent, zero-inventory, make-to-order medical equipment supply chain. Applied to GE HealthCare Algeria, the research developed a three-pillar apparatus comprising a Fuzzy ANP–Fuzzy TOPSIS prioritisation model, a discrete-event simulation of the W2W pipeline, and an agentic AI decision-support layer (SC Copilot), producing empirical findings that revise several assumptions of the standard ripple effect corpus.

Summary

The first chapter established the operational context at GE HealthCare Algeria. Field interviews and internal data revealed a supply chain characterised by total import dependence, seven equipment modalities with heterogeneous lead-time profiles, and a W2W process crossing four institutional actors: GE’s global manufacturing network, the France Consolidation Centre, Algerian Customs, and local Field Engineers. A quantitative diagnosis identified thirteen disruption sources, a high-variance customs process, and a baseline on-time delivery rate below the 80% operational target.

The second chapter situated the research within the supply chain resilience and ripple effect literature, identifying the absence of ripple effect simulation studies for zero-inventory, import-dependent MTO chains as the primary gap justifying the present contribution.

The third chapter operationalised the analytical framework. A structural reframe replaced inventory with orders as the unit of analysis, producing a 23-node directed network centred on the COSD delay metric Δ_k . The Fuzzy ANP–Fuzzy TOPSIS model ranked thirteen disruptions across four criteria, directing simulation effort toward D1 (Missing Items, $CC = 0.5996$) and D3 (Factory Delays, $CC = 0.5454$). The chapter concluded with the mathematical formalisation of the pipeline, JIT scheduling, the D3 delay operator, and the D1 feedback recovery loop.

The fourth chapter ran a fifteen-run Design of Experiments across four disruption families. The key finding is a structural asymmetry: the endogenous missing-items mechanism produces a large, significant delay effect ($r = 0.43$, $p < 10^{-10}$), while the exogenous factory delay is statistically indistinguishable from baseline ($r = 0.06$, $p = 0.159$). D1 activates a complete second pipeline traversal, whereas D3 is absorbed by the chain’s starvation-decoupling and COSD-buffering mechanisms. Field Engineers emerged as the dominant terminal accumulator (+27.9 days under Missing Items High). Compound disruption was super-additive at both intensity levels tested ($SAR_{E11} = 1.064$; $SAR_{E12} = 2.81$), and prevents steady-state convergence within the 1 200-day horizon.

The fifth chapter presented SC Copilot, which combines an LLM-based scenario interpreter, a five-phase pipeline orchestrator, an automated KPI engine, and a Streamlit dashboard. The system compresses a scenario analysis cycle from 30–90 minutes to 5–10 minutes while grounding every recommendation in specific simulation output figures.

Main Contributions

Contribution 1: Structural reframe for zero-inventory MTO chains. Replacing inventory-state variables with order-tracking variables determines the network representation, the performance metric, the capacity formalisation, and the propagation logic. The resulting model constitutes a reusable analytical framework for any import-dependent chain operating without buffer stock.

Contribution 2: Empirical characterisation of asymmetric disruption sensitivity. The statistical demonstration that D3 has no measurable terminal impact while D1 drives large-effect delay shifts ($r = 0.43$) revises the dominant assumption of the Ivanov/Dolgui corpus, which primarily models supply-side shocks propagating forward. In GE Algeria's architecture, the most damaging disruption propagates backward from the terminal stage as a capacity demand on shared upstream resources. Upstream prevention, namely pre-shipment BOM verification, dominates downstream recovery investment as a resilience lever.

Contribution 3: Three inter-stage propagation regimes and demand starvation. The empirical estimation of $\hat{\alpha}_{ij}$ revealed three qualitatively distinct regimes: standard attenuation, batch-arrival amplification under high-frequency D1, and demand starvation under D3 (negative $\hat{\alpha}$ at the Consolidation $\rightarrow$ Customs pair). The starvation regime, rarely empirically demonstrated in the simulation literature, validates the operating-regime dependence of α_{ij} and motivates a regime-switching extension of the propagation equation.

Contribution 4: SC Copilot as a validated AI-assisted decision-support framework. The five-component architecture (NL interpreter, pipeline orchestrator, KPI engine, LLM analysis engine, executive dashboard) constitutes a replicable pattern for embedding agentic AI into industrial simulation workflows, bridging the gap between simulation output and managerial decision-making without requiring model modification.

Limitations

Four limitations bound the generalisability of the reported results. First, the experimental design conducts a single run per scenario without Monte Carlo replication, due to the constraints of the AnyLogic Personal Learning Edition. Point estimates of $\bar{\Delta}$ and OTD are therefore reported without confidence intervals. Second, the simulation model is calibrated on observed distributions and expert judgments but not formally validated against a holdout dataset; goodness-of-fit testing against out-of-sample orders remains an outstanding step. Third, the static formulation of α_{ij} is insufficient under high-frequency D1, where batch-arrival dynamics produce amplification values ($\hat{\alpha} > 2$) the linear model cannot anticipate. Fourth, the SC Copilot automation component relies on Win32-specific window injection, limiting full pipeline automation to Windows environments.

Future Perspectives

Monte Carlo replication. Implementing a replication loop executing $n = 30$ independent runs per scenario and computing 95% confidence intervals for OTD and $\bar{\Delta}$ would convert the

current directional findings into statistically certified benchmarks. This is the most immediately impactful improvement to the experimental infrastructure.

Regime-switching propagation model. A threshold-based formulation in which α_{ij} transitions between attenuation, amplification, and starvation regimes as a function of system load would provide a more faithful analytical representation of the observed dynamics and could serve as the substrate for an online resilience monitoring system.

Stochastic and dynamic extensions. Reformulating D1 and D3 occurrence rates as time-varying Markov processes, or embedding the pipeline within a rolling-horizon optimisation framework, would extend the model's ability to support dynamic resilience planning under evolving geopolitical, regulatory, and operational conditions.

Extension of the SC Copilot agentic layer. Three extensions would expand SC Copilot's operational value: migration to API-based AnyLogic integration enabling fully programmatic simulation control; multi-agent orchestration for simultaneous scenario management; and a retrieval-augmented generation layer grounding LLM responses in GE Algeria's historical order data rather than simulation outputs alone.

The central finding of this dissertation is both analytically specific and managerially direct: in GE HealthCare Algeria's supply chain, the most damaging disruption is not the one that originates furthest upstream. It is the one discovered furthest downstream. Missing parts, detected only at installation, trigger a complete second import cycle and accumulate terminal delay at the Field Engineers stage in a way no amount of expediting can offset. The chain's most effective resilience lever is therefore not a faster recovery path but systematic pre-shipment BOM verification before any equipment leaves the factory.

SC Copilot operationalises this insight as a deployable decision-support tool, enabling managers to stress-test scenarios and receive simulation-grounded recommendations in natural language. The broader implication reaches beyond GE HealthCare Algeria: any supply chain operating without buffer inventory and delivering project-based equipment through international channels faces the same structural exposure, and this class of supply chain is, by the constraints of medical equipment procurement worldwide, far from rare.

Bibliography

- [1] Chia-Wei Hsu, RJ Kuo, and CY Chiou. A multi-criteria decision-making approach for evaluating carbon performance of suppliers in the electronics industry. *International Journal of Environmental Science and Technology*, 11(3):775–784, 2014.
- [2] Dmitry Ivanov, Boris Sokolov, and Alexandre Dolgui. The ripple effect in supply chains: Trade-off ‘efficiency–flexibility–resilience’ in disruption management. *International Journal of Production Research*, 52(7):2154–2172, 2014. doi: 10.1080/00207543.2013.858836.
- [3] Alexandre Dolgui, Dmitry Ivanov, and Boris Sokolov. Ripple effect in the supply chain: an analysis and recent literature. *International Journal of Production Research*, 56(1–2):414–430, 2018. doi: 10.1080/00207543.2017.1387680.
- [4] Alexandre Dolgui and Dmitry Ivanov. Ripple effect and supply chain disruption management: new trends and research directions. *International Journal of Production Research*, 59(1):102–119, 2021. doi: 10.1080/00207543.2020.1840368.
- [5] Precedence Research. Medical devices market size worth USD 1,146.95 billion by 2034. <https://finance.yahoo.com/news/medical-devices-market-size-exceed-131800842.html>, 2025.
- [6] Medical Design & Outsourcing. Medtech big 100: The world’s largest medical device companies, 2025. URL <https://www.medicaldesignandoutsourcing.com/2025-medtech-big-100-worlds-largest-medical-device-companies/>. Accessed: June 2025.
- [7] Fortune Business Insights. Medical devices market size, share & industry analysis, by type (orthopedic devices, cardiovascular devices, diagnostic imaging, ivd, and others), regional forecast, 2025–2032, 2025. URL <https://www.fortunebusinessinsights.com/industry-reports/medical-devices-market-100085>. Accessed: June 2025.
- [8] Peggy G Chen, Edward W Chan, Nabeel Qureshi, Shoshana Shelton, and Andrew W Mulcahy. Medical device supply chains, 2021.
- [9] Dmitry Ivanov. Simulation-based ripple effect modelling in the supply chain. *International Journal of Production Research*, 55(7):2083–2101, 2017. doi: 10.1080/00207543.2016.1275873.
- [10] Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei, and Ji-Rong Wen. A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18:186345, 2024. doi: 10.1007/s11704-024-40231-1.
- [11] Comparing Gamma and Log-Normal Distributions. Wolfram Demonstrations Project, <https://demonstrations.wolfram.com/ComparingGammaAndLogNormalDistributions/>, Year. Accessed: June 2, 2026.
- [12] Romain Wuilbercq, Alessandro Mogavero, and Richard E Brown. Robust multi-disciplinary modelling of future re-usable aerospace planes. In *20th AIAA International Space Planes and Hypersonic Systems and Technologies Conference*, page 3583, 2015.

- [13] DataBridge Market Research. Global medical devices market. <https://www.databridgemarketresearch.com/reports/global-medical-devices-market>, 2024.
- [14] World Integrated Trade Solution (WITS). Medical instruments exports by country, 2023, 2023. URL <https://wits.worldbank.org>. World Bank and UN Comtrade database, accessed June 2026.
- [15] Allied Market Research. Medical devices market size, share, growth report 2033. <https://www.alliedmarketresearch.com/medical-devices-market-A74110>, 2023.
- [16] Statista. Medical devices worldwide market forecast. <https://www.statista.com/outlook/hmo/medical-technology/medical-devices/worldwide>, 2024.
- [17] HERMS Global. Algeria health sector market overview. <https://hermsglobal.com/algeria-health-sector-market>.
- [18] Délégués commerciaux Canada. Profil de marché du secteur de la santé en algérie. <https://www.deleguescommerciaux.gc.ca/algeria-algerie/market-reports-etudes-de-marches/0006431.aspx?lang=fra>.
- [19] PMC / PubMed Central. Incidence de la maladie rénale chronique stade 5 traitée dans le Sud-Est algérien. <https://pmc.ncbi.nlm.nih.gov/articles/PMC11358792/>, 2024.
- [20] Statista Market Forecast. Medical devices Algeria. <https://www.statista.com/outlook/hmo/medical-technology/medical-devices/algeria>, 2024.
- [21] Agence Ecofin. Algérie : plus de 7,5 milliards \$ alloués à la santé dans le budget 2025. <https://www.agenceecofin.com/actualites-services/3010-123007>, 2024.
- [22] Algérie360. Hôpitaux d'urgence et infrastructures modernes : l'algérie investit dans la santé. <https://www.algerie360.com>, 2025.
- [23] 37 Degrés. Industrie pharmaceutique : l'algérie couvre 79% des besoins du marché en médicaments. <https://37degres.dz/index.php/2025/04/28/industrie-pharmaceutique-lalgerie-couvre-79-des-besoins-du-marche-en-medicaments/>, 2025.
- [24] Pharmchoices. Full list of medical device companies in Algeria. <https://pharmchoices.com/list-of-medical-device-companies-in-algeria-2/>, 2024.
- [25] Medica Magazine. Quelle réglementation pour les dispositifs médicaux en algérie ? <https://medica-magazine.com/quelle-reglementation-pour-les-dispositifs-medicaux-en-algerie/>, 2024.
- [26] La Voie d'Algérie. Les dispositifs médicaux importés soumis au contrôle de l'agence. <https://lavoiedalgerie.dz>, 2025.
- [27] TSA Algérie. Importation des dispositifs médicaux en algérie : nouveau tour de vis. <https://www.tsa-algerie.com>, 2025.
- [28] SoftGroup. Serialization now mandatory in Algeria: What the new 2025 decree means for pharma importers. <https://www.softgroup.eu/2025/11/03/blog-algeria-serialization-decree-2025/>, 2025.
- [29] Legal Doctrine. Nouveau cadre technique pour l'importation des produits pharmaceutiques et dispositifs médicaux en algérie. <https://legal-doctrine.com>, 2025.
- [30] Amar Laghouag and Iskander Zouaghi. Supply chain relational dynamics in informal ecosystems: exploring the case of a developing country. *Supply Chain Forum: An International Journal*, 2013.

- [31] Tarik Saikouk, Iskander Zouaghi, and Alain Spalanzani. The bullwhip effect: concretization of entropic information dissipation in supply chain systems. *Journal of Systemics, Cybernetics, and Informatics*, 10(4):75–79, 2012.
- [32] James B. Ayers. *Handbook of Supply Chain Management*. Auerbach Publications, 2 edition, 2006.
- [33] Michael E. Porter. *Competitive Advantage: Creating and Sustaining Superior Performance*. Free Press, 2008. Originally published 1985; 2008 reprint.
- [34] Graham C. Stevens. Integrating the supply chain. *International Journal of Physical Distribution & Materials Management*, 19(8):3–8, 1989. doi: 10.1108/EUM00000000000329.
- [35] Nada R. Sanders. *Supply Chain Management: A Global Perspective*. Wiley, 2012.
- [36] Benita M. Beamon. Supply chain design and analysis: models and methods. *International Journal of Production Economics*, 55(3):281–294, 1998. doi: 10.1016/S0925-5273(98)00079-6.
- [37] David Simchi-Levi, Philip Kaminsky, and Edith Simchi-Levi. *Designing and Managing the Supply Chain: Concepts, Strategies, and Case Studies*. McGraw-Hill, 2 edition, 2003.
- [38] Sunil Chopra and Peter Meindl. *Supply Chain Management: Strategy, Planning, and Operation*. Pearson, 6 edition, 2016.
- [39] John T. Mentzer, William DeWitt, James S. Keebler, Soonhong Min, Nancy W. Nix, Carlo D. Smith, and Zach G. Zacharia. Defining supply chain management. *Journal of Business Logistics*, 22(2):1–25, 2001. doi: 10.1002/j.2158-1592.2001.tb00001.x.
- [40] Iskander Zouaghi and Alain Spalanzani. Supply chains: Ago-antagonistic systems through co-opetition game theory lens. 2009.
- [41] OECD. *Securing Medical Supply Chains in a Post-Pandemic World*. OECD Publishing, 2024. doi: 10.1787/119c59d9-en.
- [42] Martin Christopher and Helen Peck. Building the resilient supply chain. *International Journal of Logistics Management*, 15(2):1–13, 2004. doi: 10.1108/09574090410700275.
- [43] T. J. Pettit, J. Fiksel, and K. L. Croxton. Ensuring supply chain resilience: Development of a conceptual framework. *Journal of Business Logistics*, 31(1):1–21, 2010.
- [44] F. Lücker, A. Timonina-Farkas, and R. W. Seifert. Balancing resilience and efficiency: A literature review on overcoming supply chain disruptions. *Production and Operations Management*, 34(1), 2025. doi: 10.1177/10591478241302735.
- [45] Mary J. Meixell and Vidyaranya B. Gargeya. Global supply chain design: a literature review and critique. *Transportation Research Part E: Logistics and Transportation Review*, 41(6): 531–550, 2005. doi: 10.1016/j.tre.2005.06.003.
- [46] F. Fox, J. Hayes, B. Whelan, D. Casey, and M. Connolly. Key factors impacting a medical ventilator supply chain during the covid-19 pandemic: Lessons for pandemic preparedness. *Disaster Medicine and Public Health Preparedness*, 18:e65, 2024. doi: 10.1017/dmp.2024.55.
- [47] D. Ivanov. Transformation of supply chain resilience research through the covid-19 pandemic. *International Journal of Production Research*, 62(23):8217–8238, 2024. doi: 10.1080/00207543.2024.2334420.

- [48] G. Marcucci, F. E. Ciarapica, G. Mazzuto, and M. Bevilacqua. Analysis of ripple effect and its impact on supply chain resilience: A general framework and a case study on agri-food supply chain during the covid-19 pandemic. *Operations Management Research*, 17(1):175–200, 2024. doi: 10.1007/s12063-023-00415-7.
- [49] Sunil Chopra and ManMohan S. Sodhi. Managing risk to avoid supply-chain breakdown. *MIT Sloan Management Review*, 46(1):53–61, 2004.
- [50] Dmitry Ivanov and Alexandre Dolgui. New disruption risk management perspectives in supply chains: digital twins, the ripple effect, and resilience. *IFAC-PapersOnLine*, 52(13):337–342, 2019. doi: 10.1016/j.ifacol.2019.11.138.
- [51] H. L. Lee, V. Padmanabhan, and S. Whang. The bullwhip effect in supply chains. *Sloan Management Review*, 38(3):93–102, 1997.
- [52] P. Götz, W. L. Auping, and S. Hinrichs-Krapels. Contributing to health system resilience during pandemics via purchasing and supply strategies: An exploratory system dynamics approach. *BMC Health Services Research*, 24:130, 2024. doi: 10.1186/s12913-023-10487-7.
- [53] S. Y. Ponomarov and M. C. Holcomb. Understanding the concept of supply chain resilience. *The International Journal of Logistics Management*, 20(1):124–143, 2009.
- [54] Nils-Ole Hohenstein, Edda Feisel, Evi Hartmann, and Larry Giunipero. Research on the phenomenon of supply chain resilience: a systematic review and paths for further investigation. *International Journal of Physical Distribution & Logistics Management*, 45(1/2):90–117, 2015. doi: 10.1108/IJPDLM-05-2013-0128.
- [55] D. Ivanov and A. Dolgui. Viability of intertwined supply networks: Extending the supply chain resilience angles towards survivability. *International Journal of Production Research*, 58(10):2904–2915, 2020.
- [56] A. Wieland and C. M. Wallenburg. The influence of relational competencies on supply chain resilience: A relational view. *International Journal of Physical Distribution & Logistics Management*, 43(4):300–320, 2013.
- [57] E. Brandon-Jones, B. Squire, C. W. Autry, and K. J. Petersen. A contingent resource-based perspective of supply chain resilience and robustness. *Journal of Supply Chain Management*, 50(3):55–73, 2014.
- [58] S. Hosseini, D. Ivanov, and A. Dolgui. Review of quantitative methods for supply chain resilience analysis. *Transportation Research Part E: Logistics and Transportation Review*, 125: 285–307, 2019.
- [59] S. AlMahri, L. Xu, and A. Brintrup. Enhancing supply chain visibility with knowledge graphs and large language models. *International Journal of Production Research*, 2025. doi: 10.1080/00207543.2025.2575841.
- [60] S. Robinson. *Simulation: The Practice of Model Development and Use*. Palgrave Macmillan, 2 edition, 2014.
- [61] A. M. Law. *Simulation Modeling and Analysis*. McGraw-Hill, 5 edition, 2015.
- [62] D. Ivanov and A. Dolgui. A digital supply chain twin for managing disruption risks and resilience in the era of industry 4.0. *Production Planning & Control*, 32(9):775–788, 2021.
- [63] D. Ivanov. Agentic digital twins: Bridging model-based and ai-driven decision-making support for a new era of supply chain and operations management. *International Journal of Production Research*, 2026. doi: 10.1080/00207543.2026.2630277.

- [64] B. Li, K. Mellou, B. Zhang, J. Pathuri, and I. Menache. Large language models for supply chain optimization. *arXiv preprint arXiv:2307.03875*, 2023.
- [65] Z. Song, Y. Xie, L. Yang, and Y. Zhao. Large language models in supply chain management: A systematic literature review and application framework. *International Journal of Production Research*, 2026. doi: 10.1080/00207543.2026.2641103.
- [66] V. Jannelli, S. Schoepf, M. Bickel, T. Netland, and A. Brintrup. Agentic llms in the supply chain: Towards autonomous multi-agent consensus-seeking. *arXiv preprint arXiv:2411.10184*, 2024.
- [67] S. AlMahri, L. Xu, and A. Brintrup. Automating supply chain disruption monitoring via an agentic ai approach. *arXiv preprint arXiv:2601.09680*, 2026.
- [68] Dmitry Ivanov. Intelligent digital twin (iDT) for supply chain stress-testing, resilience, and viability. *International Journal of Production Economics*, 263:108938, 2023. doi: 10.1016/j.ijpe.2023.108938.
- [69] Efraim Turban and Jay E. Aronson. *Decision Support Systems and Intelligent Systems*. Prentice Hall, 5 edition, 1998.
- [70] Murad Samhoury, Majdoleen Abualeenein, and Farah Al-Atrash. Enhancing supply chain resilience through a fuzzy AHP and TOPSIS to mitigate transportation disruption. *Sustainability*, 17(16):7375, 2025. doi: 10.3390/su17167375.
- [71] Robert E. Shannon. *Systems Simulation: The Art and Science*. Prentice-Hall, 1975.
- [72] Averil M Law and W David Kelton. *Simulation and analysis*. McGraw-hill, 1991.
- [73] Jerry Banks, John S. Carson, Barry L. Nelson, and David M. Nicol. *Discrete-Event System Simulation*. Prentice Hall, 4 edition, 2005.
- [74] Jay W. Forrester. *Industrial Dynamics*. MIT Press, 1961.
- [75] John D. Sterman. *Business Dynamics: Systems Thinking and Modeling for a Complex World*. Irwin/McGraw-Hill, 2000.
- [76] Mihalis Giannakis and Michalis Louis. A multi-agent based framework for supply chain risk management. *Journal of Purchasing and Supply Management*, 17(1):23–31, 2011. doi: 10.1016/j.pursup.2010.05.001.
- [77] Andrei Borshchev. *The Big Book of Simulation Modeling: Multimethod Modeling with AnyLogic 6*. AnyLogic North America, 2013.
- [78] Thomas L. Saaty. *The Analytic Hierarchy Process*. McGraw-Hill, 1980.
- [79] Da-Yong Chang. Applications of the extent analysis method on fuzzy AHP. *European Journal of Operational Research*, 95(3):649–655, 1996. doi: 10.1016/0377-2217(95)00300-2.
- [80] Ching-Lai Hwang and Kwangsun Yoon. *Multiple Attribute Decision Making: Methods and Applications*. Springer, Berlin, 1981. doi: 10.1007/978-3-642-48318-9.
- [81] Ching-Torng Chen. Extensions of the TOPSIS for group decision-making under fuzzy environment. *Fuzzy Sets and Systems*, 114(1):1–9, 2000. doi: 10.1016/S0165-0114(97)00377-1.
- [82] Sunil Chopra and ManMohan S. Sodhi. Managing risk to avoid supply-chain breakdown. *MIT Sloan Management Review*, 46(1):53–61, 2004.

- [83] Paul R. Kleindorfer and Germaine H. Saad. Managing disruption risks in supply chains. *Production and Operations Management*, 14(1):53–68, 2005.
- [84] Christopher S. Tang. Perspectives in supply chain risk management. *International Journal of Production Economics*, 103(2):451–488, 2006.
- [85] Alexandre Dolgui, Dmitry Ivanov, and Boris Sokolov. Ripple effect in the supply chain: An analysis and recent literature. *International Journal of Production Research*, 56(1–2):414–430, 2018.
- [86] Christopher W. Craighead, Jennifer Blackhurst, M. Johnny Rungtusanatham, and Robert B. Handfield. The severity of supply chain disruptions: Design characteristics and mitigation capabilities. *Decision Sciences*, 38(1):131–156, 2007.
- [87] Martin Christopher and Helen Peck. *Building the Resilient Supply Chain*, volume 15. 2004.
- [88] Serhiy Y. Ponomarov and Mary C. Holcomb. Understanding the concept of supply chain resilience. *International Journal of Logistics Management*, 20(1):124–143, 2009.
- [89] James J. Buckley. Fuzzy hierarchical analysis. *Fuzzy Sets and Systems*, 17(3):233–247, 1985. doi: 10.1016/0165-0114(85)90090-9.
- [90] Douglas C. Montgomery. *Design and Analysis of Experiments*. Wiley, 9 edition, 2017.
- [91] Susan M. Sanchez. Work smarter, not harder: Guidelines for designing simulation experiments. *Proceedings of the Winter Simulation Conference*, pages 84–94, 2007.
- [92] Aliakbar Hasani and Amirhossein Khosrojerdi. Robust global supply chain network design under disruption and uncertainty considering resilience strategies: a parallel memetic algorithm for a real-life case study. *Transportation Research Part E: Logistics and Transportation Review*, 87:20–52, 2016. doi: 10.1016/j.tre.2015.12.009.
- [93] Lotfi A. Zadeh. Fuzzy sets. *Information and Control*, 8(3):338–353, 1965. doi: 10.1016/S0019-9958(65)90241-X.
- [94] Arnold Kaufmann and Madan M. Gupta. *Introduction to Fuzzy Arithmetic: Theory and Applications*. Van Nostrand Reinhold, 1991.
- [95] A. Ravi Ravindran. *Operations Research Methodologies*. CRC Press, 2009.
- [96] Michael Wooldridge and Nicholas R. Jennings. Intelligent agents: Theory and practice. *The Knowledge Engineering Review*, 10(2):115–152, 1995.
- [97] Patrick Lewis et al. Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474, 2020.

Chapter 6

Appendices

Appendix A — State of Play

A.1 GE HealthCare’s Main Milestones

Table 6.1: GE HealthCare, Key Historical Milestones

Year	Key Milestone
1896	GE’s Elihu Thomson builds electrical equipment for X-ray production and demonstrates stereoscopic pictures for diagnosing bone fractures.
1913	William Coolidge develops the hot-cathode, high-vacuum X-ray tube. The basic Coolidge tube design is still in use today.
1974	First non-ionic X-ray contrast media introduced. Lower osmolality significantly reduces the risk of adverse drug reactions during diagnostic scans.
1978	First routine total-body CT scanner (7800 CT) made widely available, the first CT system designed by GE HealthCare.
1983	First high-field 1.5T MRI scanner introduced, based on a superconducting magnet 30,000× stronger than Earth’s gravity. Produces cross-sectional images of soft tissues with unprecedented clarity.
2000	First 4D high-resolution ultrasound for women’s health enables visualization of fetal movement <i>in utero</i> . GE also launches the first integrated hybrid infant incubator and warmer.
2015	First NASA-style hospital command center: a digital solution using machine learning and predictive analytics to help hospitals reduce waste and deliver safer care.
2019	Deep Learning Image Reconstruction (DLIR) gains FDA clearance, the first deep learning-based CT image reconstruction technology.
2023	January 4: GE HealthCare spins off from GE and begins trading on Nasdaq (GEHC) with presence in 160+ countries, ~51,000 colleagues, and 4M+ installed units across Imaging, Ultrasound, PCS and PDx.

Appendix B — State of the Art

B.1 Core Conceptual Properties of Supply Chain Disruption Theory

Table 6.2: Core conceptual properties of supply chain disruption theory

Property	Definition	Mechanism	Role in Disruption Response
Robustness	The capacity of a supply chain to absorb perturbations without altering planned execution	Direct use of redundancy reserves: safety stock, backup capacity, pre-contracted alternatives	Pre-disruption buffer — prevents deviation from plan
Flexibility	The ability to adjust and change processes and structures in response to internal and external disturbances	Indirect use of redundancy — reallocation of inventories, capacities, and sourcing facilities	Adaptive response — enables alternative execution paths
Resilience	The ability to maintain, execute, and recover planned performance after a disruption	Combines robustness and flexibility; encompasses four dimensions (see Table 5.8)	Overarching capability — limits ripple scope and duration
Stability	Ensuring continuity of supply chain behaviour despite perturbations	Formal systems-theoretic notion: Lyapunov stability and BIBO (bounded output for bounded input)	Baseline condition — supply chain remains in controlled state
Viability	The ability of the supply chain to survive and adapt when resilience has been exhausted	Triggered when no correcting programme can recover planned performance; requires fundamental structural redesign	Last-resort property — invoked only when all other layers fail

B.2 Root Causes of the Ripple Effect

Table 6.3: Root causes of the ripple effect: structural and operational dimensions

Root Cause	Supply Chain Mechanism	Ripple Effect Impact	Countermeasures
Complexity and globalisation	Geographically dispersed chains with concentrated industrial districts and heavy transport infrastructure dependency; supply concentration amplifies exposure when a single sourcing node is disrupted	Without coordinated contingency policy, recovery is long-lasting and economically severe	Sourcing diversification; global contingency plans; supplier segmentation by disruption risk
Lean paradigm and limited redundancy	Single sourcing, JIT delivery, zero safety stock, 100% capacity utilisation rational under stable conditions but eliminates all buffers and alternative capacity	Any capacity disruption ripples immediately through the chain with nothing to absorb it	Multiple/dual sourcing; backup suppliers; risk mitigation inventory; capacity buffers
Geographical specialisation	Manufacturing concentration in specific nodes or regions; disruption at one node affects the majority of downstream flows simultaneously	Disruptions in one or several nodes affect almost all others; individual nodes become critically fragile	Geographic sourcing diversity; contingency plans by region
IT dependency and limited visibility	Supply chain coordination increasingly dependent on information systems for execution; when upstream signals are detected too late, mitigation cannot be activated before downstream effects materialise	IT disruptions directly cause material flow disruptions; delayed visibility prevents anticipatory action an underestimated and growing propagation channel	Decentralised IT architecture; redundant information channels; cloud services; real-time upstream monitoring
Long lead time	Extended international pipelines increase exposure to uncertainty and reduce the time window available for corrective action [3, 41, 50]	Disruptions upstream have more time to compound before they can be detected and mitigated locally	Lead time reduction strategies; safety time buffers; early-warning monitoring of upstream milestones

Continued on next page

Root Cause	Supply Chain Mechanism	Ripple Effect Impact	Countermeasures
Process inter-dependence	Each stage depends tightly on the timely completion of preceding stages; local delays accumulate and propagate rather than being absorbed at the point of origin	Delay at any single node (factory, consolidation, customs, field engineer) shifts all downstream milestones	Process decoupling where possible; milestone-based visibility; parallel processing of administrative and logistical tasks

B.3 Ripple Effect vs. Bullwhip Effect: Comparative Framework

Table 6.4: Ripple effect vs. bullwhip effect

Feature	Ripple Effect	Bullwhip Effect
Type of uncertainty	Deep uncertainty / hazard	Random, operational uncertainty
Type of risk	Exceptional, disruption risk	Recurrent, operational risk
What is disturbed	SC structures and critical performance	Operational parameters (lead times, inventory levels)
Propagation prevention	Proactive redundancy and flexibility	Information coordination and demand sharing
Post-disturbance response	Short-term stabilisation + medium/long-term recovery	Short-term demand-supply rebalancing
Performance impact	Output performance (revenue, profit)	Current performance (stock-out costs, service level)
SC dynamics type	Structural dynamics — network reconfigures	Inventory dynamics — demand signal distorted
Disruption frequency	Low	High
Impact severity	High — can lead to market share loss or firm closure	Low to moderate

B.4 The Four-Layer Disruption Control Framework

Table 6.5: The four-layer disruption control framework [2]

Layer	Property	Condition	Management Response
Layer 1	Stability	Perturbation does not affect planned execution; the system naturally absorbs the shock	Monitor and continue — no action required
Layer 2	Robustness	Perturbation occurs but is fully absorbed by redundancy reserves without replanning	Use pre-built reserves; no schedule change
Layer 3	Resilience	Perturbation exceeds redundancy; planned performance deviates, but a correcting programme can recover the desired output via rescheduling or rerouting	Execute contingency plan; partial replanning; higher cost; some performance loss likely
Layer 4	Viability	No correcting programme can achieve planned performance; fundamental redesign of the supply chain network is required	Redesign SC; change suppliers; add facilities; revise targets; highest cost and longest recovery time

B.5 Methodological Streams for Ripple Effect Analysis

Table 6.6: Methodological streams for ripple effect analysis [2–4]

Stream	Core Logic	Key Capability	Key Limitation	Relevance to This Thesis
Mathematical optimisation (MIP, stochastic programming)	Models SC design and planning as constrained optimisation problems	Strong for strategic design, facility location, supplier selection	Treats disruptions as static events; cannot capture time-dependent propagation or gradual capacity degradation	Limited — used in adjacent literature [92] but not adopted here

Continued on next page

Table 6.6 – continued from previous page

Stream	Core Logic	Key Capability	Key Limitation	Relevance to This Thesis
Simulation (DES, ABM, System Dynamics)	Represents the SC as a dynamic system evolving over discrete or continuous time; disruptions injected as time-stamped events	Handles time-dependent dynamics, gradual recovery, multiple simultaneous KPIs, and real problem complexity	Requires careful model validation; results are scenario-specific rather than globally optimal	Primary method — DES adopted for this thesis (see Chapter 3)
Control theory (OPC, MPC)	Views the SC as a dynamic machine with state variables, inputs, outputs, and feedback loops	Provides formal mathematical structure for analysing SC stability and resilience	Controllers in SCs are humans — decisions take hours or days; difficult to operationalise with real data	Conceptual reference — provides formal grounding for the four-layer control framework
Complexity and reliability theory	Represents the SC as a network; analyses node criticality, density, and disruption propagation probability	Quantifies structural vulnerability; supports identification of critical nodes and links	Abstracts away operational detail; requires network topology data	Secondary method — adopted in the CoDIT 2026 paper for multi-echelon propagation modelling
Game theory (2021)	Models competitive and cooperative decisions under supply disruption	Captures strategic interactions between SC actors during recovery	High analytical complexity; limited empirical validation	Not adopted
Data-driven analytics (2021)	Machine learning and big data for demand pattern identification and proactive risk detection	Enables real-time risk signals from operational data	Requires large, labelled historical disruption datasets	Not adopted — insufficient disruption event history at GE HealthCare Algeria
Epidemic / SIS models (2021)	SIS models combined with complex network theory for transnational risk transmission	Captures systemic, cross-border propagation dynamics	Primarily applicable to pandemic-scale, multi-country disruptions	Not adopted

B.6 Full Technical Procedure: Fuzzy AHP–Fuzzy TOPSIS

B.6.1 Fuzzy Set Theory and Triangular Fuzzy Numbers

Fuzzy set theory, introduced by Zadeh [93], extends classical binary logic by allowing an element to partially belong to a set. Formally, a fuzzy set $\tilde{A}$ defined on a universe X is characterised by a membership function $\mu_{\tilde{A}} : X \rightarrow [0, 1]$, where $\mu_{\tilde{A}}(x)$ expresses the degree of membership of element x in $\tilde{A}$. A *fuzzy number* is a fuzzy subset of $\mathbb{R}$, convex and normal, meaning that $\sup_{x \in \mathbb{R}} \mu_{\tilde{A}}(x) = 1$ [94].

Among fuzzy number representations, the *triangular fuzzy number* (TFN) is the most widely used in MCDM applications due to its computational simplicity and intuitive interpretability [81]. A TFN $\tilde{A}$ is defined by a triplet (l, m, u) with $l \leq m \leq u$, and its membership function is given by:

$$\mu_{\tilde{A}}(x) = \begin{cases} 0 & \text{if } x < l \\ \frac{x-l}{m-l} & \text{if } l \leq x \leq m \\ \frac{m-l}{u-l} & \text{if } m \leq x \leq u \\ 0 & \text{if } x > u \end{cases} \quad (6.1)$$

where l is the lower bound (pessimistic value), m the modal value (most likely value), and u the upper bound (optimistic value). Let $\tilde{A} = (a_1, a_2, a_3)$ and $\tilde{B} = (b_1, b_2, b_3)$ be two TFNs. The main operational laws are [94]:

$$\tilde{A} \oplus \tilde{B} = (a_1 + b_1, a_2 + b_2, a_3 + b_3) \quad (6.2)$$

$$\tilde{A} \ominus \tilde{B} = (a_1 - b_3, a_2 - b_2, a_3 - b_1) \quad (6.3)$$

$$\tilde{A} \otimes \tilde{B} = (a_1 b_1, a_2 b_2, a_3 b_3) \quad (6.4)$$

$$\tilde{A} \oslash \tilde{B} = \left(\frac{a_1}{b_3}, \frac{a_2}{b_2}, \frac{a_3}{b_1} \right) \quad \text{for } b_1 > 0 \quad (6.5)$$

The inverse of a TFN is defined by:

$$\tilde{A}^{-1} = \left(\frac{1}{a_3}, \frac{1}{a_2}, \frac{1}{a_1} \right) \quad \text{for } a_1 > 0 \quad (6.6)$$

Within the Fuzzy AHP framework, pairwise comparison judgements are collected as verbal terms and converted into TFNs via the extended fuzzy Saaty scale, presented in Table 6.7.

Table 6.7: Extended fuzzy Saaty scale — conversion of verbal judgements into TFNs

Verbal Judgement	TFN (l, m, u)	Inverse TFN $(1/u, 1/m, 1/l)$
Equal importance	$(1, 1, 1)$	$(1, 1, 1)$
Slightly more important	$(1, 2, 3)$	$(1/3, 1/2, 1)$
Moderately more important	$(2, 3, 4)$	$(1/4, 1/3, 1/2)$
Strongly more important	$(3, 4, 5)$	$(1/5, 1/4, 1/3)$
Very strongly more important	$(4, 5, 6)$	$(1/6, 1/5, 1/4)$
Absolutely more important	$(5, 6, 7)$	$(1/7, 1/6, 1/5)$

B.6.2 Fuzzy AHP: Criteria Weighting

The classical AHP, proposed by Saaty [78], rests on four main steps: hierarchical structuring of the problem; pairwise comparison of criteria via the Saaty scale (1 to 9); computation of the priority vector by eigenvalue; and verification of judgement consistency via the consistency ratio CR. In its classical form, AHP assumes that decision-makers' judgements are precise and crisp, which proves unrealistic in contexts where assessments are inherently vague or expressed in linguistic terms [?].

Fuzzy AHP substitutes TFNs for the crisp values, enabling the modelling of imprecision inherent in expert judgements [79]. The fuzzy pairwise comparison matrix $\tilde{A}$ is constructed as:

$$\tilde{A} = \begin{pmatrix} (1, 1, 1) & \tilde{a}_{12} & \cdots & \tilde{a}_{1n} \\ \tilde{a}_{12}^{-1} & (1, 1, 1) & \cdots & \tilde{a}_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ \tilde{a}_{1n}^{-1} & \tilde{a}_{2n}^{-1} & \cdots & (1, 1, 1) \end{pmatrix} \quad (6.7)$$

When several experts participate, their individual matrices are aggregated using the geometric mean method of Buckley [89]:

$$\tilde{a}_{ij}^{\text{agg}} = \left(\tilde{a}_{ij}^{(1)} \otimes \tilde{a}_{ij}^{(2)} \otimes \cdots \otimes \tilde{a}_{ij}^{(K)} \right)^{1/K} \quad (6.8)$$

Weight extraction follows Chang's (1996) extent analysis in four steps.

Step 1: Fuzzy extent value $\tilde{S}_i$

$$\tilde{S}_i = \sum_{j=1}^n \tilde{a}_{ij} \otimes \left(\sum_{i=1}^n \sum_{j=1}^n \tilde{a}_{ij} \right)^{-1} \quad (6.9)$$

Step 2: Degree of possibility $V(\tilde{S}_i \geq \tilde{S}_j)$

$$V(\tilde{S}_1 \geq \tilde{S}_2) = \begin{cases} 1 & \text{if } m_1 \geq m_2 \\ 0 & \text{if } l_2 \geq u_1 \\ \frac{l_2 - u_1}{(m_1 - u_1) - (m_2 - l_2)} & \text{otherwise} \end{cases} \quad (6.10)$$

Step 3: Raw weight $d(C_i)$

$$d(C_i) = \min_{j \neq i} V(\tilde{S}_i \geq \tilde{S}_j) \quad (6.11)$$

Step 4: Normalised weight vector

$$w_i = \frac{d(C_i)}{\sum_{k=1}^n d(C_k)}, \quad \sum_{i=1}^n w_i = 1 \quad (6.12)$$

Consistency is verified on the defuzzified matrices using Saaty's consistency ratio $CR = CI/RI$, where $CI = (\lambda_{\max} - n)/(n - 1)$. A $CR \leq 0.10$ is considered acceptable [78].

B.6.3 Fuzzy TOPSIS: Ranking of Disruptions

The TOPSIS method [80] rests on the principle that the best alternative is simultaneously closest to the fuzzy positive ideal solution (FPIS) and farthest from the fuzzy negative ideal solution (FNIS). The fuzzy extension of Chen [81] proceeds in five steps.

Step 1: Fuzzy decision matrix Each alternative A_i is evaluated with respect to each criterion C_j by a TFN $\tilde{x}_{ij} = (l_{ij}, m_{ij}, u_{ij})$ obtained by converting expert scores via $\tilde{x}_{ij} = (\max(1, k - 1), k, \min(5, k + 1))$ where k is the aggregated mean expert score.

Step 2: Fuzzy normalisation For a benefit criterion: $\tilde{r}_{ij} = (l_{ij}/u_j^*, m_{ij}/u_j^*, u_{ij}/u_j^*)$, $u_j^* = \max_i u_{ij}$. For a cost criterion: $\tilde{r}_{ij} = (l_j^-/u_{ij}, l_j^-/m_{ij}, l_j^-/l_{ij})$, $l_j^- = \min_i l_{ij}$.

Step 3: Weighted fuzzy matrix $\tilde{v}_{ij} = \tilde{r}_{ij} \otimes w_j$

Step 4: Ideal solutions and distances

$$A^+ = \left\{ \tilde{v}_j^+ \right\}_{j=1}^n, \quad \tilde{v}_j^+ = (u_j^{\max}, u_j^{\max}, u_j^{\max}) \quad (6.13)$$

$$A^- = \left\{ \tilde{v}_j^- \right\}_{j=1}^n, \quad \tilde{v}_j^- = (l_j^{\min}, l_j^{\min}, l_j^{\min}) \quad (6.14)$$

Distance between two TFNs is computed by the vertex method:

$$d(\tilde{a}, \tilde{b}) = \sqrt{\frac{1}{3} [(l_1 - l_2)^2 + (m_1 - m_2)^2 + (u_1 - u_2)^2]} \quad (6.15)$$

Step 5: Closeness coefficient and final ranking

$$CC_i = \frac{d_i^-}{d_i^+ + d_i^-}, \quad CC_i \in [0, 1] \quad (6.16)$$

A value $CC_i \rightarrow 1$ indicates a highly critical disruption; $CC_i \rightarrow 0$ indicates a low-criticality disruption. Disruptions are ranked in decreasing order of CC_i .

B.7 Stochastic Input Distributions for DES

The credibility of a DES model depends directly on the structural justification of its input distributions. The simulation literature [72] identifies three distribution families as appropriate for supply chain processing durations.

B.7.1 Log-Normal Distribution

The log-normal distribution is the reference model for production lead times and logistics consolidation durations [72]. A random variable X follows a log-normal distribution, $X \sim \mathcal{LN}(\mu, \sigma^2)$, if $\ln X \sim \mathcal{N}(\mu, \sigma^2)$, with probability density function:

$$f(x; \mu, \sigma) = \frac{1}{x \sigma \sqrt{2\pi}} \exp\left(-\frac{(\ln x - \mu)^2}{2\sigma^2}\right), \quad x > 0 \quad (6.17)$$

B.7.2 Gamma Distribution

The Gamma distribution is the standard model for administrative and customs clearance delays [95]. A random variable $X \sim \Gamma(\alpha, \beta)$ has probability density function:

$$f(x; \alpha, \beta) = \frac{x^{\alpha-1} e^{-x/\beta}}{\beta^\alpha \Gamma(\alpha)}, \quad x > 0 \quad (6.18)$$

B.7.3 Triangular Distribution

The triangular distribution is used for logistics transfer durations and expedited processing times when empirical data is limited [72]. A random variable $X \sim \text{Tri}(a, b, c)$ has probability density function:

$$f(x; a, b, c) = \begin{cases} \frac{2(x-a)}{(c-a)(b-a)} & a \leq x \leq b \\ \frac{2(c-x)}{(c-a)(c-b)} & b < x \leq c \\ 0 & \text{otherwise} \end{cases} \quad (6.19)$$




Figure 6.1: Stochastic input distributions for DES supply chain modelling [11]




Figure 6.2: An example of triangular probability density function [12]

B.7.4 Simulation Paradigm Comparison

Table 6.8: Comparative properties of simulation paradigms for supply chain disruption analysis

Criterion	System Dynamics	Agent-Based	DES
Unit of analysis	Aggregates	Individual actors	Entities (orders)
Temporal resolution	Continuous	Discrete (step)	Event-driven
Individual entity tracking	No	Yes	Yes
Native queue modelling	No	Partial	Yes
Per-order disruption assignment	No	Partial	Yes
Terminal delay measurement	No	Partial	Yes
Implementation complexity	Low	High	Moderate
Key SC reference	Sterman [75]	Giannakis and Louis [76]	Ivanov [9]

B.8 Experimental Design Principles

Montgomery [90] and Law and Kelton [72] define four principles of rigorous experimental design for stochastic simulation of disruptions. **Baseline isolation** requires the inclusion of a no-disruption scenario as the reference point. **Factor isolation** requires each disruption mechanism to be evaluated independently before compound scenarios are examined. **Severity gradient** requires systematic variation of disruption intensity across its operational range. **Interaction testing** requires compound scenarios combining multiple active disruption mechanisms simultaneously, enabling measurement of whether their joint effect is additive, super-additive, or sub-additive [90].

B.9 Software Comparison

Table 6.9: Comparative analysis of DES software platforms for supply chain disruption modelling

Software	Paradigms	Entity tracking	Multi-method	Integrated DoE	Language	SC relevance
Arena	DES	Yes	No	Process Analyser	SIMAN/VBA	Moderate
Flexsim	DES	Yes	No	Flexsim Experimenter	C++/FlexScript	Moderate
AnyLogistix	Optim./Sim.	No	Partial	Integrated scenarios		Strategic SC
SimPy	DES	Yes	No	Manual (Python)	Python	Research
AnyLogic	DES+ABM+SD	Yes	Yes	Monte Carlo, optimisation, sensitivity	Java	High

Appendix C — Modeling

C.1 Full Supply Chain Network Map

Figure 6.3 presents the complete 23-node, 46-edge directed network $G = (\mathcal{N}, \mathcal{E})$ of GE HealthCare Algeria's supply chain, introduced in Section 3.2.3. Nodes are colour-coded by actor category: internal GE Algeria nodes, internal-global GE nodes, external market and partner nodes, and regulatory/hybrid nodes. Edge types distinguish physical, informational, and financial flows.




Figure 6.3: Full supply chain network of GE HealthCare Algeria.

Appendix D — Simulation Parameters

This appendix contains the full numerical specification of the simulation model. All parameters are referenced in section 3.4 of the main chapter.

This annex provides all distributional assumptions and parameter values required to replicate the discrete-event simulation model implemented in AnyLogic 8 (Process Modelling Library, Personal Learning Edition). Parameters are organised into five groups: (A.1) processing-stage distributions, (A.2) transfer-segment distributions, (A.3) order-arrival and modality parameters, (A.4) JIT scheduling parameters, and (A.5) disruption parameters.

D.1 Service Duration Functions

Table 6.10: Service duration functions and distribution families

Function	Stage	Path	Distribution family
<code>factoryDelay(m)</code>	Factory	Normal	Log-normal, modality-indexed
<code>factoryDelayUrgent(m)</code>	Factory	Normal	Log-normal, reduced-severity urgent variant
<code>consolidationDelay(m)</code>	Consolidation	Normal	Log-normal, modality-indexed
<code>customsDelay(m)</code>	Customs	Normal	Log-normal, modality-indexed
<code>feDelay(m)</code>	Field Eng.	Normal	Log-normal, modality-indexed
<code>factoryDelay_r(m)</code>	Factory	Recovery	Triangular, modality-indexed
<code>consolidationDelay_r</code>	Consolidation	Recovery	Triangular (mode-switch)
<code>customsDelay_r(m)</code>	Customs	Recovery	Log-normal, modality-indexed
<code>feWait(k)</code>	Field Eng.	All	Deterministic JIT wait (eq. (3.6))

D.2 Simulation Variables

Table 6.11: Order agent variable definitions

Variable	Type	Unit	Role
<code>orderId</code>	int	—	Unique order identifier
<code>modality</code>	int	—	Modality index ($m \in \mathcal{M}$)
<code>Order_Date</code>	double	days	Simulation time at order creation
<code>crd</code>	double	days	Customer Required Date (eqs. (3.7) and (3.8))
<code>cosd</code>	double	days	Committed On-Site Date (eq. (3.9))

Continued on next page

(Table 6.11 continued)

Variable	Type	Unit	Role
w2wStart	double	days	Planned production release date (eq. (3.5))
atd	double	days	Actual Time of Delivery
delayDays	double	days	$\Delta_k = \text{ATD}_k - \text{COSD}_k$
onTime	bool	—	True iff $\Delta_k \leq 15$
isUrgentOrder	bool	—	True iff $v_k = \text{urgent}$
needsMissingPart	bool	—	D1 indicator ξ_k^{D1}
isD3Disrupted	bool	—	D3 indicator ξ_k^{D3}
hasSlack	bool	—	True iff $w_k^{\text{start}} > t_k^{\text{creation}}$
waitingForRecovery	bool	—	True while D1 recovery path is active
t_Factory	double	days	Sojourn time at Factory stage
t_Consolidation	double	days	Sojourn time at Consolidation stage
t_Customs	double	days	Sojourn time at Customs stage
t_FE	double	days	Sojourn time at Field Engineers stage
t_flagged_out	double	days	Simulation time at D1 discovery

D.3 Resource Pools

Table 6.12: Resource pool configuration

Pool name	Stage	Capacity C_i	Shared with
res_Factory	Factory	C_F	D1 recovery Factory block
res_Consolidation	Consolidation	C_K	D1 recovery Consolidation block
res_Customs	Customs	C_X	D1 recovery Customs block
res_FieldEngineers	Field Engineers	C_E	— (terminal stage)

D.4 Simulation Output

Table 6.13: Structure of the `orders_tab` output CSV

Column(s)	Semantics
<i>Order identity and classification</i>	
<code>orderId</code>	Unique order identifier
<code>modality, needsMissingPart</code>	Equipment modality; D1 susceptibility flag
<code>is_urgent, is_retard</code>	Urgency flag; D3 (factory delay) flag
<i>Order lifecycle timestamps</i>	
<code>created_t, agent_CRD, CPD</code>	Creation time; Customer Required Date; Confirmed Promise Date
<code>w2W_Start</code>	Work-to-work start time
<i>Nominal pipeline stages</i>	
<code>seize_factory_t, at_factory_t, end_factory_t</code>	Factory stage: queue entry, service start, completion
<code>in_transfer_to_dfr_t</code>	Departure from factory to consolidation hub
<code>seize_consol_t, in_consol_t, end_consol_t</code>	Consolidation stage timestamps
<code>to_port_t, at_sea_t</code>	Port departure; sea transit start
<code>seize_customs_t, at_customs_t, end_customs_t</code>	Customs clearance timestamps
<code>last_mile_t</code>	Last-mile delivery start
<i>D1 recovery sub-process (air-freight path)</i>	
<code>check_missing_items_t</code>	Entry into missing-part check
<code>missing_item_blocked_t,</code> <code>missing_item_ordered_t</code>	Block detection; re-order emission
<code>seize_factory_recovery_t,</code> <code>at_factory_recovery_t,</code> <code>end_factory_recovery_t</code>	Recovery factory stage
<code>in_transfer_to_dfr_recovery_t</code>	Transfer to consolidation (recovery path)
<code>seize_consol_recovery_t,</code> <code>at_consol_recovery_t, end_consol_recovery_t</code>	Recovery consolidation timestamps
<code>air_transfer_t</code>	Air-freight departure
<code>seize_customs_recovery_t,</code> <code>at_customs_recovery_t,</code> <code>end_customs_recovery_t</code>	Recovery customs timestamps
<code>last_mile_recovery_t</code>	Recovery last-mile start
<i>Field Engineer dispatch</i>	
<code>seize_fe_t, at_FE_t, end_FE_t</code>	FE queue entry, assignment, task completion
<i>Derived completion fields</i>	
<code>completed_t</code>	Final completion timestamp
<code>completed, onTime</code>	Completion flag; on-time flag ($\Delta_k \leq 15$ days)
<code>delayDays</code>	$\Delta_k = \text{completed_t} - \text{COSD}$

Appendix E — Fuzzy ANP + Fuzzy TOPSIS: Python Model Documentation

The quantitative prioritization model developed in this study was implemented as a self-contained Python script. The following describes its structure, dependencies, and outputs.

E.1 Programming Environment and Libraries

The script was developed in Python and relies on three core libraries: **NumPy** for fuzzy arithmetic and matrix computations, **Pandas** for data structuring and tabular output, and **Matplotlib** for result visualization. All fuzzy operations were coded from scratch without relying on any external MCDM library, ensuring full methodological transparency and reproducibility.

E.2 Inputs

The script takes as input the raw questionnaire responses collected from four GE HealthCare Algeria domain experts representing Logistics, Project Management, After-Sales, and Commercial functions. Two types of input are encoded:

- Verbal pairwise comparisons between the four evaluation criteria (Overall Impact, Probability/Frequency, Propagation Effect, and Controllability).
- Numerical scores from 1 to 5 rating each of the 13 disruptions against each criterion.

All inputs are immediately converted into Triangular Fuzzy Numbers of the form (l, m, u) following the extended Saaty scale.

E.3 Computational Pipeline

The script executes the full MCDM pipeline across 15 sequential blocs, summarized in Table 6.14.

Table 6.14: Structure of the Fuzzy ANP + Fuzzy TOPSIS Python Script

Bloc	Step	Description
1	Raw Expert Data	Encoding of pairwise comparison responses and disruption evaluation scores from the four domain experts.
2	TFN Conversion	Verbal and numerical inputs converted into Triangular Fuzzy Numbers (l, m, u) using the extended Saaty scale.
3	Fuzzy Matrices	Individual 4×4 fuzzy comparison matrices constructed per expert.
4	Geometric Aggregation	Expert matrices aggregated into a single consensus matrix via Buckley's geometric mean method (1985).
5	Fuzzy Weights	Criteria weights derived via Chang's Extent Analysis (1996) through fuzzy synthetic extent values and pairwise possibility degrees.
6	Defuzzification	Centre of Gravity method applied to produce crisp weight estimates for cross-validation.
7	Consistency Check	Saaty's Consistency Ratio computed at both the individual and aggregate levels. Threshold: $CR \leq 0.10$.
8	Decision Matrix	Aggregated crisp decision matrix constructed from the mean scores of the four experts across 13 disruptions and 4 criteria.
9	Fuzzy Decision Matrix	Mean scores converted into TFNs to form the full fuzzy decision matrix $\tilde{X} = [\tilde{x}_{ij}]$.
10	Fuzzy Normalization	Chen's normalization method (2000) applied, distinguishing benefit criteria (C1, C2, C3) from the cost criterion (C4).
11	Weighted Matrix	Normalized fuzzy matrix weighted using the Fuzzy ANP criteria weights to produce $\tilde{V} = \tilde{R} \otimes W$.
12	Ideal Solutions	Fuzzy Positive Ideal Solution (FPIS) and Fuzzy Negative Ideal Solution (FNIS) identified for each criterion.
13	Fuzzy Distances	Distance of each disruption to FPIS (d^+) and FNIS (d^-) computed via the vertex method.
14	Closeness Coefficient	$CC_i = d_i^- / (d_i^+ + d_i^-)$ computed for each disruption and used to produce the final criticality ranking.
15	Visualization	Four diagnostic charts generated: criteria weight bar chart, CC ranking chart, d^+ vs d^- scatter plot, and expert evaluation heatmap.

E.4 Outputs

The script produces two categories of output:

- A ranked summary table classifying all 13 disruptions into four criticality tiers: *Low*, *Moderate*, *High*, and *Very High*.
- Four diagnostic visualizations: a criteria weight bar chart, a Closeness Coefficient ranking chart, a d^+ versus d^- scatter plot, and an expert evaluation heatmap.

Appendix F — Agentic AI: Technical Architecture

An *AI agent* in the operational sense is a software system that perceives its environment through data inputs, maintains an internal representation of that environment, and takes actions in pursuit of a defined objective, without requiring step-by-step human instruction for each action [96].

F.1 Functional Layers

The **data ingestion layer** connects the agent to its information environment through APIs [67]. The **knowledge and retrieval layer** provides access to contextualised institutional memory through Retrieval-Augmented Generation (RAG) [97]. The **reasoning and orchestration layer** constitutes the agent’s cognitive core [68]. The **action and output layer** translates agent reasoning into operational outputs [37, 66].

Table 6.15: AI agent contribution to supply chain visibility dimensions

Visibility mension	Di-	Limitation of Conven- tional Tools	AI Agent Contribution
Temporal		Periodic reporting intro- duces 24–72h detection la- tency	API-driven continuous monitoring reduces latency to near-real time
Relational		Single-source reports can- not identify cross-node causal patterns	LLM reasoning across heteroge- neous sources surfaces non-obvious propagation linkages
Scenario		Simulation requires spe- cialist configuration; inac- cessible under time pres- sure	Natural language interface enables on-demand scenario testing by non- specialist staff
Institutional		Historical disruption knowledge resides in individual memory or dispersed documents	RAG retrieval grounds agent re- sponses in organisation-specific precedent