Hybrid Distribution Strategies for Improved Supply Chain Efficiency

A Case Study at Procter & Gamble

TIL5060 - TIL Thesis

T.P. Van der Hulst

Delft University of Technology

TU Delft

P&G

Hybrid Distribution Strategies for Improved Supply Chain Efficiency

A Case Study at Procter & Gamble

by

T.P. Van der Hulst

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended on Wednesday July 16, 2025.

Student number:	4963601	
Project duration:	February, 2025 - July, 2025	
Thesis committee:	Dr. J.M. (Jaap) Vleugel	Chair Supervisor TU Delft (CEG)
	Ir. M.B. (Mark) Duinkerken	Supervisor TU Delft (ME)
	Dr.ir. A.J. (Arjan) van Binsbergen	Supervisor TU Delft (CEG)
	K. (Kasper) Vermaat	Supervisor Procter & Gamble
Faculty:	Civil Engineering and Geosciences, Delft University of Technology	

This is an anonymized version of the original thesis. Certain confidential or sensitive information has been redacted or removed in accordance with company regulations and data protection guidelines.

Preface

More than a year ago, I decided to join P&G for writing my master's thesis, because of my curiosity to the world of fast-moving consumer goods (FMCG). Over the past months, I have had the opportunity to explore an interesting topic at the intersection of supply chain optimization, sustainability, and operational decision-making. The experience of working on a real-world distribution challenge within a multinational company like P&G has not only deepened my technical and analytical skills, but also taught me what it is like to be part of a fast-moving environment with my hands on the operation.

This thesis marks the final step in completing my studies at TU Delft, and with that, the end of my time in Delft. Over the past years, I have gained not only academic knowledge, but also valuable insights about myself and the world around me. Looking back, I feel well prepared to take the next steps, carrying with me the lessons, friendships, and experiences that have shaped me along the way.

I would like to express a big thank you to Kasper, my supervisor and colleague at P&G, for his support, our valuable feedback meetings on Tuesdays, the quick lessons on things I didn't know about Rotterdam, and for his trust throughout the process. I am also very grateful to my academic supervisors, Jaap and Mark, for their structured guidance, thorough reviews of my hand-ins, and helpful discussions that helped shape this thesis into its final form.

Finally, I want to thank my family and friends, with special thanks to my roommates at the Kade, for their steady support and trust in me over the past months. Your reassurance that things will always work out in the end helped me keep perspective and stay motivated.

I hope this thesis contributes to ongoing discussions on making supply chains more responsive, efficient, and sustainable, encourages further exploration of hybrid distribution strategies, and has offered a meaningful step forward for P&G.

T.P. Van der Hulst
Delft, July 2025

Executive Summary

This report presents a thesis research into outbound distribution strategies in the downstream supply chain of a large fast-moving consumer goods (FMCG) company. The focus is on evaluating the integration of direct distributions from production plants alongside the existing conventional two-echelon distribution structure with usage of centralized distribution centers. The research was initiated in collaboration with Procter & Gamble (P&G), aiming to identify cost-efficient and sustainable solutions within their European distribution network. The case study centers on shipments of Category X and Category Y products from two manufacturing plants to customers in the Benelux market. A mathematical optimization model was developed and implemented in Python using P&G's historical shipment data, to assess the impact of allowing direct plant shipments (DPS) in parallel to the regular routing via distribution centers (DCs).

The research follows a Research and Design methodology. After reviewing literature on supply chain management, hybrid distribution strategies and modeling approaches, a current state analysis of P&G's supply chain was conducted to identify modeling requirements and contextual constraints. These insights informed the formulation of a flow-based two-echelon vehicle routing problem (2E-VRP) that jointly optimizes logistics cost and emissions. The model incorporates practical constraints such as loading thresholds, product eligibility, and customer-specific demand profiles. A structured set of scenarios and configurations was then used to evaluate the results on several key performance indicators, including logistics cost, emissions, service level performance, and stock allocation efficiency.

The results demonstrate that hybrid distribution can offer substantial logistics cost savings, especially for high-volume, long-distance flows. In experiments where DPS was allowed as a distribution method next to the two-echelon DC-based shipments, logistics costs decreased by up to 43.1% for long-haul routes to customers in The Benelux with Category Y products from Plant 2, and by 15.8% for shorter-distance, lower-volume routes like Category X products from the nearer located Plant 1. Service level performance also improved consistently, with average gains of +1.0 to +1.5 percentage points across all scenarios, seeming small but impacting greatly the commercial profit. DPS integration can additionally relieve pressure on the DC by reallocating DPS volume to the plant, thereby improving stock allocation efficiency and reducing the total weekly logistics costs by up to 4.8%. A break-even analysis confirmed that the viability of shipping demand as DPS is highly route-specific, with required minimum volumes ranging from 5 to 28 floor positions (FP) as most cost-efficient solution, depending on distance and demand characteristics. These findings initially highlight that uniform Full Truckload (FTL) thresholds are suboptimal, and that route-level differentiation enhances cost efficiency.

However, the environmental performance of DPS is significantly more nuanced. While average emissions decreased by up to 11.7% in high-volume, long-distance cases, particularly for Category Y with high vehicle utilization, these environmental gains did not generalize across all configurations. For short transport distances on Category X shipments and scenarios with low FTL thresholds, emissions sometimes increased due to reduced vehicle fill rates (VFR) with increased usage of DPS. Notably, lowering the FTL threshold led to greater DPS adoption but often worsened environmental outcomes. The model consistently identified two-echelon flows as a vital aspect of an emission-optimal strategy, due to their superior vehicle utilization and consolidation. This counterintuitive yet robust result, confirmed through scalarization and ϵ-constraint sensitivity analyses, demonstrates that maximizing DPS share as a cost-efficient solution does not inherently support sustainability goals under current system conditions.

A central insight from the analysis is that cost-optimal and sustainability-optimal routing structures may diverge. While customer-specific DPS thresholds and selective decentralization reduce logistics costs, they often come at the expense of emissions performance. Without additional policy levers, such as route-level emissions caps, minimum VFR constraints, or decentralized incentives, carbon pricing alone does not sufficiently promote environmentally beneficial DPS use in the model. This underscores a fundamental trade-off in hybrid distribution systems: DPS improves cost and service performance under the right conditions, but its environmental benefits are conditional and require more than economic optimization to be realized.

These findings give direct implications for practice. To support a sustainable hybrid distribution strategy, companies like P&G must adopt a cross-functional approach that balances economic and environmental objectives. This includes using differentiated DPS thresholds, aligning forecasting and planning of stock allocation strategies with DPS implementation and customers' demand profiles, and integrating explicit emissions constraints into operational planning. A nuanced and policy-aware application of DPS is essential to achieve long-term gains across the multiple key performance indicators cost, service, sustainability, and inventory performance dimensions.

Beyond its practical relevance, this research contributes to academic literature on hybrid distribution strategies by addressing the integration of direct plant shipments in a regional, cross-border FMCG context. It introduces a flow-based adaptation of the two-echelon vehicle routing problem (2E-VRP) that jointly optimizes logistics cost and emissions, while accounting for real-world constraints such as SKU eligibility, FTL thresholds, and demand variability. This extends on existing similar problems that solely focus on cost optimization. By demonstrating the trade-offs between cost and sustainability within a hybrid distribution network, the study highlights the need for integrated environmental constraints, offering a bridge between theoretical modeling and practical decision-making in complex supply chains.

Nomenclature

Frequently Used Abbreviations

Abbreviation	Definition
2E	Two-Echelon
2E-VRP	Two-Echelon Vehicle Routing Problem
3PL	Third-Party Logistics
AOV	Advanced Order Visibility
ATW	Access Time Windows
CD	Cross-Docking
CF	Customer Freight
CLP	Customer Load Preparation
CRM	Customer Relationship Manager
DC	Distribution Center
DPS	Direct Plant Shipments
EOQ	Economic Order Quantity
FE	First-Echelon
FEFO	First Expired, First Out
FP	Floor Positions
FTL	Full Truckload
Cat X&Y	Category X & Category Y (categories in scope)
ISF	Inter Site Freight
KPI	Key Performance Indicators
LTL	Less-Than-Truckload
MILP	Mixed-Integer Linear Programming
OSA	On-Shelf Availability
SCM	Supply Chain Management
SE	Second-Echelon
SKU	Stock Keeping Unit
SNO	Supply Network Operations
SS	Safety Stock
VFR	Vehicle Fill Rate
VRP	Vehicle Routing Problem

Contents

Preface i

Executive Summary ii

Nomenclature iv

1 Introduction 1
- 1.1 Background Procter & Gamble . . . 1
  - 1.1.1 Supply Network Operations . . . 1
- 1.2 Problem Statement and Trigger for Research . . . 2
- 1.3 Background and Rationale from the Literature . . . 2
- 1.4 Research Scope: A Case Study at P&G . . . 3
- 1.5 Objective and Research Questions . . . 5
- 1.6 Relevance and Academic Contribution . . . 6

2 Methodology 7
- 2.1 Full Research Framework . . . 7
- 2.2 Research Approach . . . 8
- 2.3 Case Study Approach . . . 8
- 2.4 Data Collection and Analysis . . . 8
- 2.5 Modeling and Experimental Approach . . . 9
  - 2.5.1 Applicable Modeling Approach . . . 9
  - 2.5.2 Experimental Design . . . 9

3 Literature Review 10
- 3.1 Purpose and Scope . . . 10
  - 3.1.1 Literature Search Strategy . . . 10
- 3.2 Supply Chain Management . . . 11
  - 3.2.1 Definition and Relevance of SCM . . . 11
  - 3.2.2 Demand Management . . . 12
  - 3.2.3 Operations Management . . . 13
  - 3.2.4 Procurement Management . . . 13
  - 3.2.5 Logistics Management . . . 13
  - 3.2.6 Stakeholders in SCM . . . 14
- 3.3 Hybrid Distribution Methods . . . 15
  - 3.3.1 Key Performance Indicators . . . 15
- 3.4 Modeling Approach in Distribution Problems . . . 16
  - 3.4.1 The Vehicle Routing Problem . . . 16
  - 3.4.2 MILP Models in VRP . . . 19
- 3.5 Discussion . . . 19
- 3.6 Conclusion . . . 19

4 Current State Analysis P&G 21
- 4.1 Stakeholder Analysis . . . 21
- 4.2 Demand Management at P&G . . . 23
  - 4.2.1 Demand Forecasting . . . 23
  - 4.2.2 Demand Planning . . . 25
  - 4.2.3 Customer Order Management . . . 25
- 4.3 Logistics Management at P&G . . . 25
  - 4.3.1 Inventory Management at the DC . . . 25
  - 4.3.2 Stock Allocation in a Hybrid Distribution Setup . . . 26
  - 4.3.3 Transport . . . 27
  - 4.3.4 Plants in Scope for DPS Integration: Plant 1 - Category X . . . 28
  - 4.3.5 Plants in Scope for DPS Integration: Plant 2 - Category Y . . . 28
- 4.4 Service Level Metric . . . 28

4.5 Conclusion . . . 29
4.5.1 Current Distribution Practice at P&G . . . 29
4.5.2 Opportunity for Hybrid Distribution in FBNL Scope . . . 30
4.5.3 Key Design Requirements . . . 30
4.5.4 Practical Conditions for Customer Eligibility . . . 31

5 Hybrid Distribution Model: a Flow-Based Adaptation of the 2E-VRP 32
5.1 Model Formulation . . . 32
5.1.1 Assumptions . . . 33
5.1.2 Mathematical Notation . . . 34
5.1.3 Objective Function: Logistics Cost Minimization . . . 35
5.1.4 Constraints . . . 36
5.1.5 Incorporating Emissions in the Objective Function . . . 37
5.1.6 Service Level Performance . . . 38
5.2 Data Preparation . . . 39
5.2.1 Data Overview . . . 39
5.2.2 Region-based Transport Costs . . . 39
5.2.3 Distance matrix . . . 39
5.2.4 Demand Data: Narrowing the Customer Set . . . 40
5.3 Model Implementation . . . 42
5.4 Model Verification . . . 43
5.4.1 Model Consistency . . . 43
5.4.2 Behavioral Testing with Parameters . . . 43
5.4.3 Behavioral Testing with Decision Variables . . . 45
5.4.4 Solver and Performance Checks . . . 46
5.4.5 Conclusion . . . 46
5.5 Model Data Validation . . . 47
5.5.1 Validation of Two-Echelon Input Data (C0) . . . 47
5.5.2 Validation of Direct Shipment Cost Inputs (C1) . . . 47
5.5.3 Conclusion . . . 48
5.6 Conclusion . . . 48

6 Experimental Design 49
6.1 Demand Scenarios: DPS-Eligible Streams . . . 49
6.1.1 Impact of Including Stored SKUs on Weekly Demand Patterns . . . 50
6.2 Model Configurations: Policy Settings . . . 51
6.3 Customer and Time Selection for Experimental Analysis . . . 51
6.4 Summary of Experimental Design . . . 53
6.5 Post-Optimization and Sensitivity Analyses . . . 53
6.5.1 Stock Allocation Impact Analysis . . . 53
6.5.2 Break-Even Analysis of FTL Thresholds . . . 53
6.5.3 Sensitivity of Emission Pricing on Share of DPS (λ Variation) . . . 53
6.5.4 ε-constraint Analysis: Emissions Minimization Under Cost Caps . . . 53
6.6 Key Performance Indicators (KPI's) . . . 54
6.6.1 KPI's Optimized in the Model . . . 54
6.6.2 KPI's Evaluated Post-Optimization . . . 54
6.7 Experimental Execution and KPI Collection . . . 54
6.8 Conclusion . . . 54

7 Computational Results 56
7.1 Model Feasibility and Performance . . . 56
7.2 Cost and Environmental Impact . . . 57
7.2.1 Scenario S1 and S2 - Category X from Plant 1 . . . 57
7.2.2 Scenario S3 and S4 - Category Y from Plant 2 . . . 58
7.3 Service Level and Network Effects . . . 59
7.3.1 Scenario S1 and S2 - Category X from Plant 1 . . . 60
7.3.2 Scenario S3 and S4 - Category Y from Plant 2 . . . 61
7.3.3 Estimated Commercial Impact of Service Level Improvements . . . 62
7.4 Key Scenario Comparison . . . 62
7.5 Pairwise Comparison of Key Performance Indicators . . . 63
7.6 Stock Allocation Impact (Cash) . . . 64

7.6.1 Observed Inventory Shift and Cost Impact . . . 64
7.6.2 Interpretation and Broader Implications . . . 64
7.7 Customer-Specific DPS Threshold Analysis . . . 65
7.7.1 Demand-Based Break-Even Analysis . . . 65
7.8 Sensitivity Analysis: Emissions-Oriented Routing Behavior . . . 67
7.8.1 Impact of Emission Pricing on Share of DPS . . . 67
7.8.2 ε-constraint: Emissions Minimization Under Cost Caps . . . 68
7.8.3 Synthesis of Findings on Emissions-Oriented Routing Behavior . . . 69
7.9 Conclusion . . . 69

8 Conclusion 71
8.1 Sub-questions . . . 71
8.2 Main Findings . . . 72

9 Discussion and Recommendations 74
9.1 Scientific Contribution . . . 74
9.2 Limitations of the Current Model and Assumptions . . . 75
9.3 Directions for Future Research and Model Development . . . 75
9.4 Case-Specific Strategic Recommendations for P&G . . . 76

References 78

A Scientific Paper 81

B Transport Data 91

C Statistics Demand Data 93

D Model Output for Data Validation 95

E Computational Results 96
E.1 Results Cost and Environmental Impact per Scenario . . . 96
E.2 Pairwise Comparison of Logistics Cost and Emissions Results . . . 97
E.3 Results Service Level and Network Effects per Scenario . . . 99
E.4 Pairwise Comparison of Service Level Results . . . 99
E.5 Detailed Results Demand Scenario 1 . . . 101
E.6 Detailed Results Demand Scenario 2 . . . 102
E.7 Detailed Results Demand Scenario 3 . . . 103
E.8 Detailed Results Demand Scenario 4 . . . 105
E.9 Numerical Results ε-constraint Analysis . . . 106
E.10 Data Demand-Based Break-Even Analysis . . . 106

F Python Model 107

1 Introduction

In today's fast-moving and competitive market, companies are constantly looking for ways to make their supply chains more efficient, flexible and sustainable. One of the key challenges in doing so is figuring out the most efficient way of distributing their products through the supply chain. Traditionally, many supply chains rely on multi-echelon networks with distribution centers (DCs), but this is not always the most efficient or sustainable option, especially when customer demand or shipment volumes vary. When searching for efficient distribution, interest has grown in utilizing hybrid distribution strategies, which combine different delivery methods like distributions directly from the plant with distributions via multi-echelon networks. Hybrid strategies could offer companies a better balance between economic and environmental costs and delivery service to the customer.

This research will explore the potential and the effect of integrating such hybrid approaches through an in-depth case study within the supply chain operations of Procter & Gamble in France and The Benelux. By assessing both the practical feasibility and strategic impact of integrating direct shipments alongside multi-echelon distributions, this research aims to contribute to the existing knowledge on hybrid distributions in the literature, and provide actionable insights for building more efficient and sustainable supply networks. This introductory section will outline the background and context of the research and the case study at Procter & Gamble, and state the research objective, research questions and content of the report.

1.1. Background Procter & Gamble

Procter & Gamble (P&G) is a leading global consumer goods company with a commitment to innovation, high quality and sustainability (Procter & Gamble, 2023). Founded in 1837, the company established a strong foundation in more than 180 countries, offering products and new initiatives that improve daily life of their consumers. P&G offers a wide range of categories including beauty, grooming, health care, fabric care, home care, baby and family care, with a mission to deliver superior value to consumers while fostering long-term growth and development. As the company continues to evolve in response to changing market dynamics, its robust supply chain practices play a critical role in ensuring efficiency, sustainability, and resilience in delivering high-quality products to millions of households worldwide.

1.1.1. Supply Network Operations

Supply Network Operations (SNO) is the department that is responsible for all critical processes of aligning demand and supply through collaborative planning, and ensuring shipments from P&G's warehouses, plants and distribution centers, by means of customer orders and transportation methods. The entire supply chain integrates the various functions, including sales, marketing, finance and production, to create a unified plan to optimize inventory levels, enhance service levels, and improve overall operational efficiency. The SNO process enables the company to respond swiftly to market changes, accurately meet customer demands, and effectively manage resources. This holistic approach not only drives profitability but also supports the company's commitment to sustainability by maximizing use of resources throughout the supply chain (Procter & Gamble, 2023).

Procter & Gamble distributes their products across the entire world. In Europe, SNO is divided in clusters to manage the market operations efficiently. This research focuses on one specific cluster: France & The

Benelux (FBNL). Next to the geographical division, the selling brands are operated through five industry-based sector business units: Baby, Feminine and Family Care; Beauty; Health Care; Grooming; and Fabric and Home Care (Procter & Gamble, 2025). The products are produced in different plants all over Europe, often sector-specific. From the plants, the products are distributed to P&G's DC's, from where orders are shipped Business-to-Business (B2B) to DC's of P&G's customers. The scope of this research will be within the FBNL cluster for products only of two categories, from now on named as Category X & Category Y, specifically shipped to customers in the Benelux (BNL).

1.2. Problem Statement and Trigger for Research

Within P&G's supply chain, the need to continuously improve operational efficiency has become increasingly pressing. In the current market context of rising costs, sustainability and evolving customer expectations, traditional distribution strategies are under strain. The multi-echelon model, where goods are shipped from production plants to distribution centers (DCs), and subsequently dispatched to customers, has long been the backbone of P&G's supply network. However, several operational challenges now underscore the urgency to reassess this model.

Firstly, there is growing internal pressure to reduce logistics costs and increase cost-effectiveness across the European network. With fluctuating demand levels and shipment volumes across markets, the current centralized model does not always allow for the most economically optimal routing of goods. Potentially, direct shipments from plants could offer a more flexible and cost-efficient alternative for certain routes or customer segments, yet these have not been structurally integrated in all existing planning frameworks.

Secondly, the current reliance on DCs creates operational bottlenecks, like pressure on space and outbound transport scheduling. During peak seasons or in times of supply disruptions, this pressure can translate into delays and suboptimal service levels. By allowing a more dynamic use of direct shipments alongside conventional flows, hybrid distribution strategies could alleviate this burden and release valuable capacity within the DCs.

Finally, service level performance remains a top priority for P&G to deliver their superior quality. On-time delivery, responsiveness to changing customer requirements, and differentiated service for strategic customers are essential metrics for customer satisfaction. A hybrid distribution strategy could provide the opportunity to design more agile and responsive delivery setups that can better meet specific customer needs.

In light of these challenges, P&G is seeking to explore hybrid distribution strategies that combine the benefits of direct shipments from plants with the robustness of the multi-echelon network. While it sounds promising in theory, the actual implementation of such strategies involves complex trade-offs between cost, service, and operational feasibility. This research aims to investigate these trade-offs and assess the potential for a structurally integrated hybrid distribution model within P&G's supply chain.

1.3. Background and Rationale from the Literature

Scientific research into the facets of supply chain management and the utilization of distribution approaches is done in chapter 3. An analysis of the existing literature gives valuable insights and moreover leads to knowledge gaps where this research jumps in. This section contains a concise rationale from the literature review to give background to the proposed research objective and questions.

Supply chain management (SCM) has undergone significant transformations driven by evolving market demands, technological advancements, and the increasing emphasis on sustainability. The fundamental elements of SCM, such as creating customer value, enhancing competitive advantage, and maintaining inter-organizational collaboration, remain relevant. In the realm of logistics management, various distribution network strategies, including direct shipments, multi-echelon distribution, and cross-docking, offer unique advantages and face specific challenges (Dondo et al., 2011; Potoczki et al., 2024). The integration of these strategies into hybrid distribution models has been highlighted as a potential strategy for optimizing overall supply chain efficiency (Azizi & Hu, 2020; Musa et al., 2010). The Vehicle Routing Problem (VRP) and its more complex variants, such as the two-echelon vehicle routing problem (2E-VRP), provide robust frameworks for analyzing and optimizing distribution networks (Sluijk et al., 2023; Zhou et al., 2024). Including practical constraints and sustainability metrics into these models can further enhance their applicability to real-world scenarios, ensuring more efficient, resilient, and sustainable supply chain networks.

Considering the background of literature, this research aims to evaluate the impact of integrating direct plant shipments alongside two-echelon distribution methods as a hybrid distribution strategy to improve supply chain efficiency. Academically, this research contributes to the existing literature by addressing several gaps. It explores the integration of hybrid distribution strategies, specifically integrating direct shipments within a multi-echelon system, an area that has not been extensively studied. Additionally, while prior studies often focus on urban logistics, this research extends the analysis to regional cross-border distribution networks, providing insights in a different operational context. Furthermore, this study will also assess the trade-off between logistics cost optimization and sustainability metrics such as carbon emissions, alongside the impact on customer service level, network and storage impacts that hybrid distribution strategies cause. While service level is frequently mentioned as a conceptual performance dimension, few empirical studies have quantified how hybrid distribution structures affect delivery reliability, fulfillment rates, or responsiveness, especially in larger, cross-border supply chains. By considering these factors together, this research provides a more in-depth assessment of the effects of integrating direct shipments within an existing distribution framework.

1.4. Research Scope: A Case Study at P&G

In the FBNL cluster, P&G has established two major DCs in Belgium and Northern France to serve the Benelux and French markets, respectively.

producing Category X and Category Y products for FBNL, other production plants are given in Table 1.1.

Table 1.1: Cluster FBNL Overview

Main Production Plants	
Other Production Plants	
Main Distribution Centers (DCs)	

The map in Figure 1.1 gives a schematic overview of the product flow from the given plants. The orange lines give the flow of Category X products that are shipped from plants to the DCs. The green lines give the flow of Category Y products,

This research focuses solely of the production plants Plant 1 and Plant 2, since they account for the largest segment of Category X and Category Y (Cat X&Y) products. Both plant locations have a slightly different physical structure, which will be elaborated on here.

[Figure anonymized for confidentiality reasons]

Figure 1.1: Plants and Flows for Category X and Category Y to DCs

The hub in Plant 1 consists of more than a plant, because it has an automated warehouse located directly next to the plant. Category X Segment A products that are produced in this plant, are automatically packed on pallets and stored in the DC next to the plant. These pallets are shuttled to the DC 1 before they are shipped to BNL customers. Next to Category X Segment A products produced in Plant 1, the Plant 1 hub has other import flows of Category X products from external plants, and export flows to external markets.

[Figure anonymized for confidentiality reasons]

Figure 1.2: Product Flows at the Plant 1 Hub

[Figure anonymized for confidentiality reasons]

Figure 1.3: Product Flows at the Plant 2 Hub

Figure 1.2 gives an overview of the Plant 1 hub. The inflow of products in Plant 1 consists of 3 streams, which are given in the pink arrows:

1. Category X Segment A SKUs products from Plant 1;
2. Category X Segment SKUs from plants outside Europe, entering Plant 1 via transshipment at the import DC;
3. Category X Segment SKUs from European plants that are shipped to non-European markets.

The dotted line shows that external plants ship products from the Category X category and other categories to DC 1 as well, which further serves the Benelux customers.

The outflow of products from the hub can be given in 2 main streams, given in blue arrows:

1. Shipments to customer DCs in the Benelux via DC 1;
2. Shipments to other P&G sites, including Plant 2, to serve French customers and markets outside FBNL.

Figure 1.3 gives an overview of the hub Plant 2. The inflow of products in Plant 2 DC consists of:

1. Category Y Segment A SKUs products from Plant 2;
2. Category Y Segment SKUs from external plants to distribute to FBNL market.

It is again visible that external Category Y plants also ship to DC 1 directly to stock inventory there for shipments to Benelux customers.

The outflow of products is given with purple lines:

1. Shipments to customer DCs in the Benelux via DC 1;
2. Shipments to customer DCs in France and other P&G sites to serve markets outside FBNL.

The goal of this study is to optimize the supply chain by investigating the opportunity to integrate Direct Plant Shipments (DPS) from the two production hubs directly to customer DCs in The Benelux (BNL). These direct flows, schematically represented by the green arrows in Figure 1.2 and Figure 1.3, would bypass the intermediate DC 1. Multiple stakeholders have emphasized the strategic importance and operational potential of this integration.

To systematically assess this opportunity, four distinct DPS demand scenarios are defined. These scenarios differ based on the product category (Category X or Category Y), the origin hub (Plant 1 or Plant 2), and whether only locally produced products or also stored imports are included. The configurations are as follows:

- **C1 – Plant 1 Production Only:** Direct shipments to Benelux customers of Category X Segment A SKUs produced at Plant 1 (inflow 1 in Figure 1.2).
- **C2 – Plant 1 Production + Storage:** Direct shipments from Plant 1 of both produced Category X Segment A SKUs and imported stored Category X SKUs (inflow 1, 2 & 3 in Figure 1.2).

- **C3 – Plant 2 Production Only:** Direct shipments to Benelux customers of Category Y Segment A SKUs produced at Plant 2 (inflow 1 in Figure 1.3).
- **C4 – Plant 2 Production + Storage:** Direct shipments from Plant 2 of both produced and stored/imported Category Y SKUs (inflow 1 & 2 in Figure 1.3).

These scenarios are designed to reflect realistic internal logistics options for enabling or restricting DPS flows. They serve as the foundation for the experiments described in chapter 6, where each scenario is evaluated under external conditions to determine its operational and strategic value.

An important requirement for a direct plant shipment is the order volume. The specifics of FTL are measured in the number of pallets and floor positions, which are elaborated on in chapter 4. Assessment will be required of which specific customers order specific products in the right order volumes. For these potential direct plant shipments, three key performance indicators (KPI's) will be assessed: cost (including logistics costs and emissions), service level, and cash (stock allocation efficiency) (Desmet, 2018). Since not all products are ordered in sufficient volumes, implementation of DPS will lead to hybrid distribution methods to customers. Customers can receive direct shipments from Plant 1 and Plant 2, next to the conventional shipments with consolidated orders from DC 1. This research therefore focuses not only on potential economic and environmental benefits, but also on how DPS affects inventory distribution (e.g., pressure on the DC) and customer service performance. The choice for these KPI's is obtained form previous studies and the literature (see subsection 3.3.1).

1.5. Objective and Research Questions

The objective of this research is to evaluate the impact of integrating direct plant shipments (DPS) into the existing two-echelon distribution network as a hybrid strategy, with the aim of optimizing performance across the three interconnected pillars of the supply chain triangle: cost (including logistics costs and emissions), service level, and cash (stock allocation efficiency) (Desmet, 2018).

This study builds upon both academic literature and industry practices. It starts with a comprehensive literature review on current distribution methods, particularly focusing on the roles and challenges of using distribution centers versus direct plant shipments in large FMCG networks. It will then explore the factors influencing the feasibility and effectiveness of integrating new distribution methods, identifying the key performance indicators and requirements through desk research and a stakeholder analysis. Simultaneously, the research will analyze Procter & Gamble's current supply chain setup in the France and Benelux (FBNL) region to identify practical constraints and opportunities. Subsequently, a modeling framework will be developed to evaluate hybrid configurations with an experimental design. The model aims to simulate and compare network performance under several DPS configurations, focusing on the three KPI's.

This research objective has led to a main research question for this thesis and sub-questions to structure the research and to ensure the investigation of both the theoretical and practical aspects of a hybrid distribution strategy. The first four questions will entail the extensive literature review and analysis of P&G's current practices, the fifth question leads to the development and implementation of a suitable modeling approach and the final sub-question will assess the potential benefits and impacts of a hybrid strategy. In-depth research of all sub-questions leads to answering the main research question:

What is the impact of integrating direct plant shipments in an existing two-echelon distribution network - as a hybrid strategy - on logistics and environmental costs, service level performance, and stock allocation efficiency in a cross-border supply chain?

To answer the main research question (MRQ), the following sub-questions (SQ) are formulated:

1. What are the current logistical processes and challenges associated with using distribution centers versus direct plant shipments, and how can their efficiency be measured?
2. What is a suitable modeling approach to evaluate the added value of direct shipping in hybrid distribution strategies, and what theoretical requirements should be considered when modeling a supply chain?
3. What is the current practice of P&G's distribution model in the FBNL region, and what are the opportunities and limitations for integrating a hybrid distribution approach that includes direct plant shipments?
4. What are the key design requirements for modeling the integration of direct shipments into P&G's hybrid distribution network, considering both theoretical foundations and the practical conditions

for customer eligibility?

5. How should the proposed modeling approach be structured to evaluate the impact of direct shipping on supply chain efficiency?
6. How do different configurations of integrating direct plant shipments affect logistics and environmental costs, service level, and stock allocation in the modeled hybrid distribution scenarios?

The first two SQ's are answered in section 3.6. SQ3 and SQ4 are answered in section 4.5. SQ5 is answered in section 5.6, and final SQ6 is answered in section 7.9. The MRQ will be answered in the final conclusion, in chapter 8. In chapter 2, the methodology and full framework for answering the research questions is further explained. It will furthermore explain why a case study approach is suitable for this research, how data is collected and evaluated, and it gives a schematic overview of how the model solving approach will look like.

1.6. Relevance and Academic Contribution

This thesis contributes to the academic field of supply chain management by addressing three key gaps identified in the literature. First, while hybrid distribution strategies that combine direct plant shipments (DPS) with two-echelon networks have been conceptually explored, few studies have developed and evaluated such models across multiple performance dimensions. This research presents a structured framework to assess hybrid distribution through a flow-based adaptation of the two-echelon vehicle routing problem (2E-VRP), explicitly considering cost, environmental, and service objectives. Second, the study broadens the scope of performance evaluation by incorporating emissions and service reliability alongside logistics costs. These dimensions are often discussed in theory but rarely integrated into operational models of hybrid networks. This approach allows for a more comprehensive assessment of supply chain performance. Third, whereas most existing work focuses on last-mile or urban logistics, this thesis shifts the focus to regional, cross-border supply chains within a multinational FMCG context. The model incorporates practical constraints such as SKU eligibility, full-truckload (FTL) thresholds, and customer-specific policies, demonstrating how theoretical models can be adapted to real-world complexities.

By applying the lens of the supply chain triangle, balancing cost, service, and cash, this research introduces a structured and holistic evaluation of hybrid distribution strategies. Using P&G's cross-border supply chain in the FBNL region as a real-world case study, the results aim to provide not only operational recommendations but also strategic insights into how hybrid logistics models are efficient in multinational settings. As such, this work bridges theoretical modeling, stakeholder impact, and practical implementation, offering valuable contributions to both academic research and industry application.

2 Methodology

This section will outline the full methodology of the research, including the research framework and case study approach, data collection and data analysis and the modeling and experimental approach.

2.1. Full Research Framework

To start with a summarizing overview of the research, Figure 2.1 presents the complete research framework. This figure integrates the methodological components, connecting the problem definition, research questions, and two-phased research approach to the modeling and experimental design. It shows how the Research phase lays the groundwork through literature review, stakeholder analysis, and case study exploration, while the Design phase translates these findings into a flow-based 2E-VRP optimization model. The framework also highlights the feedback loops between the analysis and model design to the MRQ and further outcomes.

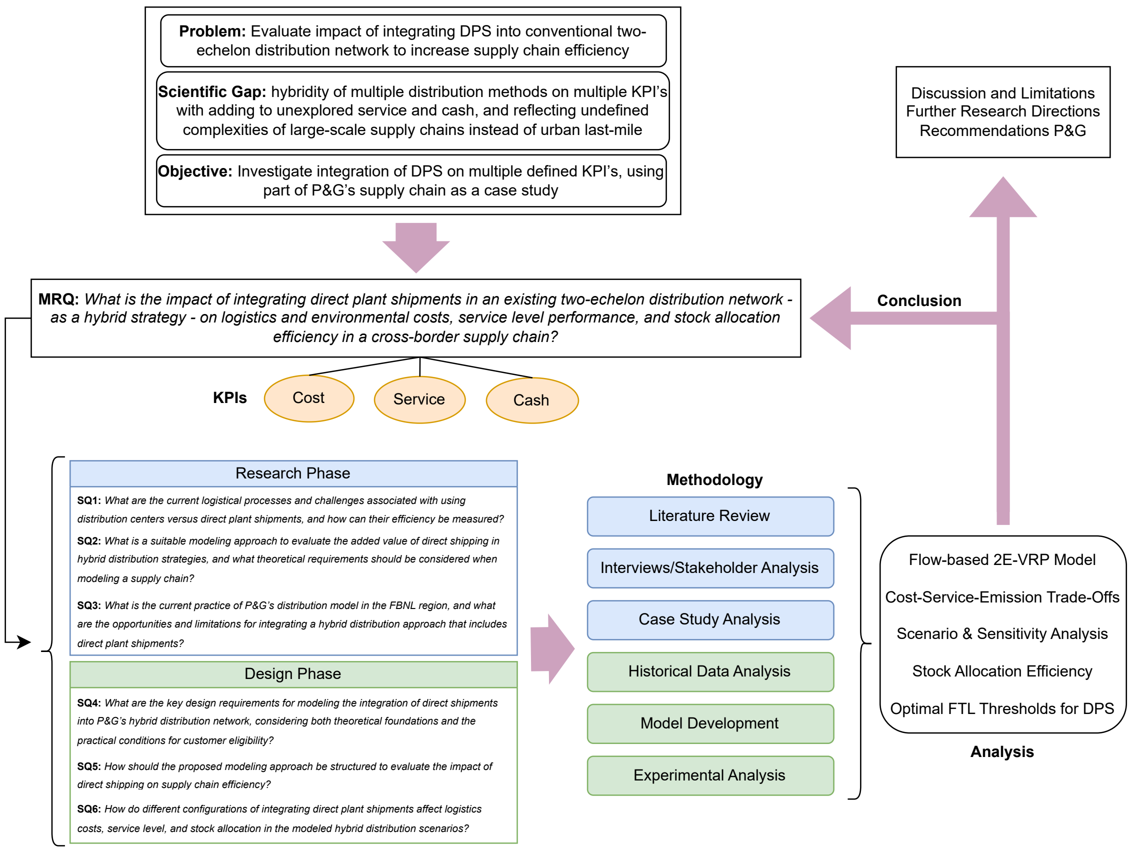

Figure 2.1: Full Research Framework

2.2. Research Approach

This research will consider two phases; the Research phase and the Design phase. The Research phase focuses on understanding the current situation. This includes evaluating P&G's existing logistical processes in the France and Benelux (FBNL) region, identifying key challenges, mapping influencing factors, selecting a suitable modeling approach, and analyzing current practices and potential customers. These topics are addressed with the first three research questions listed in Table 2.1. Once a clear understanding is established, the Design phase follows. In this phase, a modeling approach will be developed and tailored to the findings from the Research phase. Furthermore, the experiments will be performed and computational results are obtained with the model.

The overall research uses a mixed-methods approach, combining both qualitative and quantitative techniques. Qualitative insights from internal stakeholder interviews and operational observations will complement quantitative analysis using shipment data and model-based simulations. This combination is intended to ensure that the final recommendations are both data-driven and practically relevant. An overview of how the research questions are distributed across both phases is given in Table 2.1.

Table 2.1: Sub-questions and Corresponding Methods

Sub-question	Research Methods	Data Collection
Research Phase		
SQ1	Literature review, Supply chain process mapping	Scientific literature, internal supply chain documentation
SQ2	Conceptual modeling, Theoretical framework development	Scientific literature
SQ3	Case study analysis, Stakeholder consultation	Semi-structured interviews, internal reports, operational datasets
Design Phase		
SQ4	Requirement elicitation, Literature review	Semi-structured interviews, internal reports, scientific modeling principles
SQ5	Mathematical modeling, Model formulation	Derived requirements, historical demand and shipment data
SQ6	Scenario-based simulation, Sensitivity analysis	Model outputs, quantitative performance metrics

2.3. Case Study Approach

A case study is a robust research tool to explore the practical application of theoretical insights in a real-world context. This thesis research uses a case study approach focused on Procter & Gamble's supply chain operations in the FBNL region. It provides valuable insights and practical solutions that could be beneficial not only for P&G, but also for other companies in similar industries. Real-world applications of hybrid distribution strategies are still underrepresented in the growing literature on multi-echelon distribution and hybrid delivery models. This case study aims to contribute to closing that gap by analyzing how such strategies could be applied in P&G's operations. By collecting and analyzing both quantitative data (e.g. shipment flows, order patterns, storage volumes) and qualitative data (e.g. stakeholder experiences, operational limitations), the research aims to evaluate the benefits and trade-offs of hybrid distribution strategies.

2.4. Data Collection and Analysis

The qualitative data collection and analysis includes an extensive literature review to establish a theoretical foundation and semi-structured interviews with key stakeholders at P&G. Quantitative data collection consists of historical shipment and inventory data to analyze the demand, identify trends and assess feasibility of customers. Additionally, internal reports and documents will provide context on current distributions and provide material for data validation. The quantitative analysis will formulate a Mixed-Integer Linear Programming (MILP) model within an adapted Two-Echelon Vehicle Routing Problem (2E-VRP) framework, inspired on existing models in literature, to directly evaluate the optimization of logistics and environmental costs, and analyze the impact on service levels and stock allocation efficiency under different configurations. The experimental data analysis will further assess trade-offs and model robustness, which can provide insights into the benefits of hybrid distribution strategies.

2.5. Modeling and Experimental Approach

To quantitatively assess the impact of integrating DPS into P&G's distribution network, this research develops a shipment-assignment optimization model adapted on the Two-Echelon Vehicle Routing Problem (2E-VRP), as reviewed in section 3.4. The modeling process begins by translating the case-specific logistics context into a flow-based hybrid distribution model, further detailed in chapter 5. The overall modeling framework follows a structured sequence of steps, as outlined in Table 2.2, ensuring both theoretical rigor and practical applicability.

Table 2.2: Quantitative Modeling Research Framework

Model Formulation	Define the full mathematical formulation as a Two-Echelon Vehicle Routing Problem (2E-VRP).
Data Preparation	Collect, pre-process and analyze historical shipment and inventory data to implement in the model.
Model Implementation	Encode the MILP model using Python and the FICO Xpress Python API as optimization tool.
Model Verification	Test the model on a controlled dataset to verify logical outcomes when adjusting decision and parameter values.
Model Data Validation	Align model output with theoretical expectations to assess the real-world accuracy.
Experimental Design	Define experimental conditions by combining configurations and demand scenarios. Structure these experiments to evaluate various hybrid distribution strategies.
Solution and Analysis	Execute the optimization model across all defined experiments and perform sensitivity analyses to assess impacts on key performance indicators.

This methodological structure ensures that the research results provide a comprehensive assessment of the hybrid distribution model's impact on cost-efficiency, service level performance, and strategic inventory placement, the three key dimensions of distribution network effectiveness in this case.

2.5.1. Applicable Modeling Approach

The model draws conceptual inspiration from several recent studies in two-echelon distribution modeling. In particular, the 2E-VRPDDATW model developed by Zhou et al. (2024), which incorporates dynamic demands, time window constraints, and direct deliveries, offers valuable insights into how the complexities in this case can be embedded into an optimization model. However, in contrast to traditional VRP formulations that emphasize route sequencing and vehicle scheduling, the model developed in this research simplifies the structure into a flow-based framework. This means that vehicle tours and return routing are not explicitly modeled. Instead, shipments are treated as point-to-point flows within capacity and cost constraints. This modeling choice is grounded in the operational reality of P&G, which is elaborated on in section 5.1. where outbound logistics are executed by a third-party logistics provider (3PL).

2.5.2. Experimental Design

The optimization model is formulated as a Mixed-Integer Linear Programming (MILP) that integrates real-world logistics constraints such as vehicle capacities, regional cost structures, and VFR thresholds. It minimizes logistics costs and monetized emissions for a given demand scenario, while also supporting configuration-based evaluations of service level and stock allocation efficiency.

To evaluate the impact of integrating DPS, the model is tested across a structured set of 48 experiments. These combine four demand scenarios (reflecting different eligibility settings for DPS from Plant 1 and Plant 2) with four internal model configurations that vary policy rules such as DPS activation and shipment size thresholds, over three representative weeks. This scenario-based approach enables a comprehensive comparison of trade-offs across cost, emissions, and service KPIs under varying conditions of operational flexibility.

3
Literature Review

3.1. Purpose and Scope

The purpose of this literature review is to develop a solid understanding of supply chain management literature, with a focus on distribution methods in large-scale supply chains, and to identify opportunities for improved efficiency and cost optimization. This foundation supports the development of the proposed case study at P&G.

The review is structured into three main sections: (1) Supply Chain Management, (2) Hybrid Distribution Methods, and (3) Modeling Approaches in Distribution Research. The first section is organized around four key pillars: demand management, operations management, procurement management, and logistics management. The current state analysis in chapter 4 will later focus specifically on demand and logistics management. After establishing a general understanding of supply chain management, the second section reviews literature on the integration of hybrid distribution methods. The third section discusses modeling approaches commonly used in distribution research to inform the methodology for this study. Each section begins with an overview of key sources and summarizes the main findings per study.

3.1.1. Literature Search Strategy

An effective search strategy is crucial for conducting a structured and relevant literature review (Wohlin et al., 2022). This review follows a comprehensive approach, beginning with the formulation of preliminary research questions, followed by keyword selection and database searches using Google Scholar and Scopus. The key terms and concepts used are listed in Table 3.1. To broaden the scope, a hybrid strategy was applied, combining direct database queries with backward and forward snowballing (Wohlin, 2014). Snowballing is an iterative process of identifying additional studies by examining reference lists and citations. Inclusion and exclusion criteria were defined to ensure quality and relevance, with a focus on studies published within the last fifteen years. This structured approach supports the development of a solid knowledge base and helps identify research gaps relevant to the case study at Procter & Gamble.

Table 3.1: Concepts and keywords for literature review

Concept groups	**Supply Chain Management; Logistics and Transportation; Vehicle Routing; Mathematical modeling**
Keywords	Supply Chain Management: supply chain design, distribution networks, facility location planning
	Logistics and Transportation: direct shipments, direct deliveries, multi-echelon networks, two-echelon systems, hybrid distribution strategies
	Vehicle Routing: vehicle routing problem, two-echelon vehicle routing problem, Integer optimization
	Mathematical modeling: mixed-integer linear programming, routing optimization
Truncation	(Supply Chain Management) OR (Logistics Transportation) AND (Vehicle Routing) OR (Mathematical Modeling)

3.2. Supply Chain Management

This first section of the literature will review general literature on supply chain management (SCM). After introducing SCM, the review is structured in the pillars Demand Management, Operations Management, Procurement Management and Logistics Management. Based on the structure of Altekar (2023), these four aspects are overarching for all relevant aspects in SCM. This is schematically given in Figure 3.1.

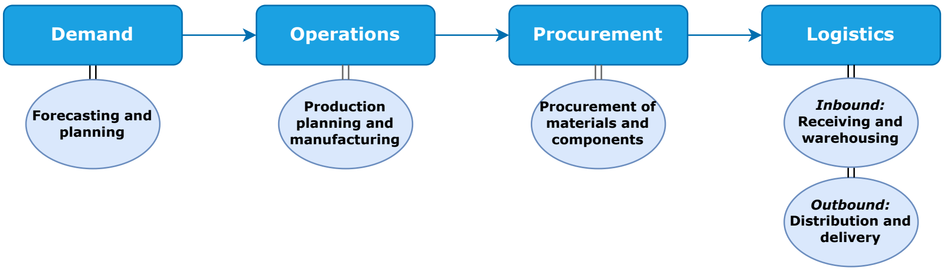

Figure 3.1: Overarching Pillars in SCM

Demand Management in subsection 3.2.2 will explain the significance of demand forecasting, planning and management in the supply chain. Operations Management in subsection 3.2.3 will dive in the coordination of production processes and ensuring operational efficiency. Procurement Management in subsection 3.2.4 discusses the importance of sourcing materials and managing supplier relationships. The section on Logistics Management in subsection 3.2.5 will include transportation, warehousing, and distribution strategies.

Relevant Studies

Table 3.2: Key Studies Literature Supply Chain Management

Source	Key findings
Altekar, R.V. (2023)	Defining the supply chain and key elements of effective supply chain management.
Helo, P., & Hao, Y. (2022)	Technology enhances operations in SCM such as planning, scheduling, optimization and transportation.
Guo, Y., Liu, F., Song, J. S., & Wang, S. (2024)	Effective inventory management strategies significantly mitigate the impact of demand surges and supply disruptions.
Marchuk, V. Y., Sergiy, G., Karpun, O., & Garmash, O. (2020)	Warehousing logistics is essential for maintaining the flow of goods, accounting for 40% of logistics costs.
Rebs, T., Thiel, D., Brandenburg, M., & Seuring, S. (2019)	Stakeholder pressures significantly impact a company's SCM performance.
Azizi, V. & Hu, G. (2020)	Direct shipments eliminate delays associated with intermediate distribution centers, enhancing responsiveness to customer demand.
Dondo, R., Méndez, C.A., & Cerdá, J. (2011)	Multi-echelon distribution strategies involve multiple layers of facilities to efficiently manage the flow of goods from production to customers.
Potoczki, T., Holzapfel, A., Kuhn, H., & Sternbeck, M. (2024)	Cross-docking consolidates products at a central facility, quickly transferring them to delivery vehicles to reduce storage time and improve efficiency.

3.2.1. Definition and Relevance of SCM

SCM can be defined as the systematic, strategic coordination of traditional business functions within a company and across various businesses within the supply chain, aimed at enhancing long-term performance (Min et al., 2019). It involves managing product flows and integrating functions such as procurement, production, and distribution to improve efficiency and responsiveness. The supply chain is also simply defined as the bridge between demand and supply (Altekar, 2023), beginning with the procurement of raw materials and ending in product consumption. According to Altekar (2023), SCM is based on five pillars: Make to Order, Non-core Outsourcing, Multi-tier Supplier Partnership, Multiparty Net-Logistics, and Highly Sophisticated IT Systems. These reflect the shift toward customer-driven production, external partnerships, multi-party logistics collaboration, and technology-enabled optimization.

A conceptual model for modern SCM is illustrated in Figure 3.2, referred to as the Collaborative Planning Forecasting Replenishment (CPFR) approach (Altekar, 2023), in which all supply chain partners collaborate to meet defined customer demands. The ultimate goal is to deliver value by integrating internal and external operations.

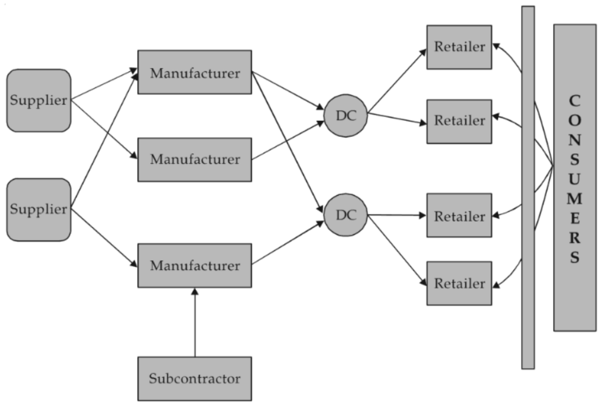

Figure 3.2: CPFR Approach as SCM

SCM has evolved in response to changing market conditions and technological advancements, which require firms to maintain competitive advantage while addressing increasingly complex customer demands. Although the core principles, creating customer value, competitive advantage, and collaboration, remain intact, the tools and technologies have shifted significantly (Min et al., 2019). In today's digital economy, supply chains are more customer-centric, emphasizing omnichannel strategies and real-time data to enhance responsiveness (Zinn & Goldsby, 2017).

Sustainability has also gained importance, with companies adopting responsible sourcing and environmentally friendly practices (Prashar, 2023). Firms are now accountable for the entire product lifecycle, making sustainability a central part of SCM strategies. Overall, modern SCM focuses on leveraging technology and collaboration to manage growing complexity effectively.

SCM in retail focuses on optimizing goods flow from suppliers to customers, balancing efficiency and service quality (Potoczki et al., 2024). Retailers often manage complex logistics networks with warehouses and distribution centers. Efficient SCM integrates cross-docking (CD) to consolidate smaller shipments into more efficient loads (Vogt, 2010). This reduces logistics costs, improves truck utilization, and minimizes inventory. Flexibility is crucial to respond to market fluctuations while ensuring on-time delivery. In short, successful retail SCM relies on supplier collaboration, demand forecasting, and technology to improve visibility and efficiency (Potoczki et al., 2024).

3.2.2. Demand Management

Demand management is a key element of SCM, focusing on forecasting, planning, and managing customer demand to ensure efficient supply chain operations (Altekar, 2023). Accurate forecasting helps minimize inventory costs while meeting customer needs, using historical sales data, market trends, and economic indicators. A common method is collaborative planning, forecasting, and replenishment (CPFR), as shown earlier in Figure 3.2. Through collaboration with suppliers and retailers, businesses can improve forecast accuracy and responsiveness. Segmentation is another useful strategy, allowing supply chains to tailor responses to different customer groups or product lines. Access to demand information across the supply chain enables rapid responses to market changes (Mahmood & Kess, 2014). Close collaboration enhances forecast accuracy and overall responsiveness.

SCM typically applies two strategies: demand-pull and supply-push (Chiang & Huang, 2021). Demand-pull focuses on unmet customer needs, emphasizing downstream collaboration and customer integration. Supply-push, by contrast, involves technology development and upstream integration with suppliers to drive innovation. While demand-pull supports responsiveness, supply-push facilitates proactive market development. Effective supply chains strike a balance between the two, using customer insights and supplier capabilities to improve performance.

3.2.3. Operations Management

Operations management is essential to SCM, coordinating production processes and improving efficiency across the supply chain network (Altekar, 2023; Helo & Hao, 2022). By digitalizing processes and integrating stakeholders, it enhances responsiveness and supports competitive advantage. Best practices include using advanced technologies, such as artificial intelligence, to increase visibility and enable proactive decision-making, streamlining logistics and improving customer satisfaction (Helo & Hao, 2022).

Lean manufacturing (LM), rooted in the Toyota Production System, focuses on efficiency and waste reduction through just-in-time (JIT) production (Vanichchinchai, 2019). It aligns with operations management goals to deliver the right products at the right time and cost. Recent studies emphasize extending lean principles beyond internal operations to strengthen collaboration with supply chain partners. Building trust enhances both operational and supply chain performance. Still, balancing cost-focused, transactional approaches with collaboration-based strategies remains a challenge. Integrating LM offers a framework for creating sustainable, efficient, and responsive supply chains.

Inventory Management

Inventory management is a cornerstone of both operations and logistics management, essential for balancing customer demand with cost efficiency. It involves determining optimal stock levels using methods such as Economic Order Quantity (EOQ), safety stock calculations, and Just-in-Time (JIT) strategies, which align inventory with production schedules and demand forecasts while reducing holding costs and waste (Guo et al., 2024).

From a resilience perspective, inventory strategies such as pre-positioning safety stock and employing multiple sourcing channels help mitigate disruptions and demand surges (Guo et al., 2024). Techniques like Vendor-Managed Inventory (VMI) and collaborative planning enhance coordination and responsiveness by sharing inventory data among supply chain partners.

Inventory management is closely linked to warehousing, which serves as the backbone of goods storage and distribution. Warehousing can represent up to 40% of logistics costs and plays a direct role in service level performance (Marchuk et al., 2020). Modern warehouses increasingly rely on automation and robotics, while techniques like cross-docking reduce storage time by moving goods directly from inbound to outbound transport, especially valuable for perishable or high-turnover items. By integrating lean principles and inventory technologies such as real-time tracking systems, organizations can improve operational efficiency, enhance visibility, and support a more agile and resilient supply chain.

3.2.4. Procurement Management

According to Ross (2015), procurement management is a core component of SCM, involving the strategic sourcing, acquisition, and oversight of goods, services, and resources essential for operations. In FMCG supply chains, characterized by high volume, fast turnover, and product variety, procurement must ensure a steady flow of raw materials, packaging, and finished goods to support continuous production and wide-scale distribution.

Effective procurement extends beyond cost reduction to include supplier collaboration, risk mitigation, and demand-driven sourcing that aligns with shifting consumer preferences. Since procurement costs often represent a large share of total revenue, optimized sourcing directly influences profitability, resilience, and responsiveness. Digital tools such as e-sourcing, predictive analytics, and real-time supplier monitoring increasingly streamline these processes (Ivanov & Dolgui, 2020).

In hybrid distribution contexts, procurement must be closely integrated with multi-channel logistics to ensure timely and efficient material and product flows (Ross, 2015). As FMCG supply chains grow more global and complex, procurement management plays a strategic role in driving efficiency, cost savings, and long-term sustainability.

3.2.5. Logistics Management

Distribution Network Strategies in Supply Chain

In SCM, the design of distribution networks is essential for efficiently delivering products to customers (Dondo et al., 2011). Suppliers adopt various strategies based on product types and operational goals. In a single-echelon strategy, manufacturers ship directly to customers. Distributor storage, by contrast, uses distribution centers (DCs) or warehouses to maintain inventory closer to demand points, particularly effective for fast-moving goods. Two-echelon networks combine factories and warehouses, enabling both centralized shipments and local fulfillment.

Cross-docking offers an alternative, consolidating products at a central facility for immediate transfer to outbound vehicles, avoiding storage time and enhancing efficiency (Potoczki et al., 2024). Unlike DCs, cross-dock (CD) facilities do not store or manipulate goods but focus on rapid throughput. In summary, three core distribution strategies can be identified: direct shipping, multi-echelon distribution, and cross-docking. Although cross-docking is often embedded in multi-echelon systems, it functions as a distinct operational method. Each approach is selected based on product characteristics, cost efficiency, and service level requirements, requiring a tailored design of distribution networks.

Direct Shipments

Direct shipments, or single-echelon distribution, involve sending products from factories directly to customers or retailers without intermediate storage (Azizi & Hu, 2020). This strategy enables faster delivery and reduces handling costs, especially for Full Truck Load (FTL) shipments. It is highly responsive to customer demand but can be less efficient when product volumes are low, resulting in Less-than-Truckload (LTL) shipments and underutilized capacity (Potoczki et al., 2024). Coordinating delivery schedules and managing transportation costs becomes more complex when direct shipping is integrated with other methods like cross-docking or multi-echelon networks (Azizi & Hu, 2020). As a result, this strategy is less commonly chosen, since multi-echelon systems often offer greater optimization in routing, timing, and vehicle use (Sitek & Wikarek, 2015).

Multi-echelon Distribution

Multi-echelon distribution incorporates multiple facility layers, such as suppliers, warehouses, and DCs, to manage goods flow from production to end customers (Dondo et al., 2011). This setup improves inventory balancing and enhances responsiveness by enabling stock to be positioned closer to various demand points. It allows consolidation, optimized routing, and lower transport costs. However, managing these layers increases complexity, requiring close coordination of inventory levels and delivery schedules across locations. This can make service level maintenance and stockout prevention more difficult in volatile markets (Potoczki et al., 2024).

Cross-docking in Multi-echelon Distributions

Cross-docking is a strategy in which products are received at a central facility and immediately transferred to outbound vehicles without long-term storage or manipulation (Potoczki et al., 2024). It supports transportation efficiency and reduces holding costs by consolidating inbound shipments and enabling quick outbound dispatch. Unlike direct shipping or traditional warehousing, cross-docking facilitates the rapid flow of goods, enhances truck utilization, and supports just-in-time (JIT) logistics (Hosseini-Nasab et al., 2023). However, it introduces challenges such as the need for precise coordination of inbound and outbound flows, longer lead times compared to direct shipping, and infrastructure investments.

3.2.6. Stakeholders in SCM

Across all segments of the supply chain, various stakeholders play a critical role in shaping supply chain performance. Their influence significantly affects the efficiency and effectiveness of SCM practices. According to Rebs et al. (2019), stakeholder pressures, such as those from customers, suppliers, governments, and shareholders, can strongly impact a company's supply chain outcomes. In demand management, customers influence product design and delivery expectations, requiring companies to adapt accordingly. In operations management, employees and internal management are key to implementing improvements and maintaining quality. Procurement is shaped by suppliers, who influence sourcing strategies, costs, and material availability. In logistics, regulatory bodies enforce compliance standards that directly affect transportation and distribution.

The interplay between stakeholder demands and a company's internal capabilities determines how effectively supply chain operations are managed. As stakeholder influence varies in intensity, organizations must strategically navigate these pressures to optimize SCM practices and performance (Rebs et al., 2019).

3.3. Hybrid Distribution Methods

Relevant Studies

Table 3.3: Key Studies Literature Hybrid Distribution Methods

Source	Key findings
Azizi, V. & Hu, G. (2020)	Emphasize the significance of developing efficient distribution systems by integrating various decision-making problems in a framework for hybrid distribution strategies, such as the capacitated location of distribution centers, vehicle routing, and direct shipments.
Musa, R., Arnaout, J. P., & Jung, H. (2010)	Address transportation in a cross-docking network that combines direct and indirect shipments from suppliers to customers, using integer programming and ant colony optimization to minimize transportation costs.
Ma, H., Wang, Q., & Xu, X. (2011)	Propose a two-stage method for shipment consolidation and cross-docking with direct delivery options, balancing transportation, inventory, and scheduling costs.
Mohammadi, Z., Barzinpour, F., & Teimoury, E. (2020)	A multi-objective model incorporates direct and indirect deliveries in a sustainable processed food supply chain, and it shows optimization of economic and environmental objectives.
Hosseini-Nasab, H., Nasrollahi, S., Bagher Fakhrzad, M. & Honarvar, M. (2023)	A literature gap is identified regarding the combined use of direct and indirect shipments within cross-docking networks for optimizing transportation costs.

The previous section reviewed three established distribution strategies. However, the integration of these methods into hybrid distribution networks remains underexplored in certain contexts. Recent logistics research increasingly highlights the potential of combining direct and indirect shipment methods to improve supply chain performance.

Azizi and Hu (2020) emphasize the importance of integrated distribution systems that simultaneously consider facility location, vehicle routing, and direct shipment decisions. Their study is relevant for hybrid strategies as it frames distribution network design as a unified decision-making problem to optimize overall efficiency.

Other researchers have also addressed combined distribution approaches. Musa et al. (2010) developed a model incorporating both direct and indirect deliveries within a cross-docking network, applying integer programming and ant colony optimization to minimize transportation costs. Similarly, Ma et al. (2011) proposed a two-stage solution method for shipment consolidation problems, balancing cost and scheduling trade-offs in systems using both cross-docking and direct delivery.

Mohammadi et al. (2020) explore hybrid methods in the processed food industry, where product shelf life makes direct delivery crucial. Yet, indirect shipments remain valuable for shipment consolidation. Their multi-objective model compares direct and indirect approaches based on both economic and environmental outcomes—such as transport costs and carbon emissions. Direct shipments can reduce handling stages, potentially improving both cost and sustainability performance.

Hosseini-Nasab et al. (2023) focus on cross-docking networks and highlight the cost implications of direct versus indirect shipment strategies. Contrary to earlier findings, they argue that direct shipments over long distances, especially for perishable goods, can lead to higher costs and inefficiencies. In contrast, indirect shipments via cross-docking can reduce both transport costs and delivery times. Importantly, this study points out a gap in the literature regarding hybrid use of direct and indirect methods within cross-docking systems, an area with potential for performance improvement in complex, multi-node networks.

Together, these studies underline the value and complexity of integrating multiple distribution strategies. Yet, the body of literature remains fragmented. Few studies address the simultaneous design and optimization of hybrid systems in real-world, high-volume environments like in FMCG. This gap presents an opportunity for further research, as pursued in this research.

3.3.1. Key Performance Indicators

Evaluation of distribution methods, particularly the effectiveness of hybrid approaches, relies on key performance indicators (KPIs). Most studies assess hybrid distribution models using economic and environmental measures.

Economic indicators often focus on maximizing profit by minimizing costs such as plant and DC establishment, transportation (for both direct and indirect shipments), raw materials, and inventory holding (Mohammadi et al., 2020). Fishani et al. (2022) similarly use total cost minimization to evaluate their hybrid model's efficiency. Musa et al. (2010) and Ma et al. (2011) also focus on transportation-related costs, including time, truck setup, fleet size, and inventory holding.

On the environmental side, carbon emissions are the primary metric. In the case of perishable products, environmental KPIs can also include waste management costs (Fishani et al., 2022). These indicators are essential for assessing the sustainability of hybrid distribution models.

Hosseini-Nasab et al. (2023) use a combination of KPIs to evaluate cross-docking and hybrid strategies. Their primary economic KPI is total transportation cost, which includes facility, shipping, and truck use costs. They also assess flow efficiency and cross-dock allocation. Environmental performance is measured by minimizing unnecessary transport movements, thus reducing emissions.

Across studies, transportation costs and carbon emissions are the most frequently used KPIs. While cost reductions primarily benefit the operating party by increasing profitability, they often align with sustainability goals by reducing emissions, highlighting the dual value of optimized hybrid distribution strategies.

3.4. Modeling Approach in Distribution Problems

Relevant Studies

The following table presents key studies that provide background on modeling approaches relevant to distribution problems, such as the Vehicle Routing Problem (VRP) and Mixed-Integer Linear Programming (MILP). These works serve as a theoretical foundation for this research. Specific examples of 2E-VRP model variants applied in the literature are summarized separately in Table 3.5.

Table 3.4: Key Studies Literature Modeling Approaches in Distribution Problems

Source	Key findings
Nielsen et al. (2024)	Provides a systematic review of vehicle routing problems (VRPs), offering a comprehensive overview of recent developments and their applicability to multi-echelon distribution networks, particularly in logistics optimization.
Clarke & Wright (1964)	Introduced the foundational savings algorithm for solving the classic VRP, laying the groundwork for many subsequent variants, including capacitated and multi-echelon routing models.
Dondo et al. (2011)	Applied Mixed-Integer Linear Programming (MILP) to optimize total transportation cost in 2E-VRP with cross-docking, demonstrating the effectiveness of MILP in modeling complex distribution systems with intermediate depots.
Azizi & Hu (2020)	Developed a comprehensive MILP-based framework integrating distribution center location, vehicle routing, and direct shipments, offering a versatile modeling approach for hybrid distribution strategies.
Garrido et al. (2025)	Emphasized the robust optimization capabilities of MILP in supply network design, highlighting its ability to handle multi-objective and multi-constraint problems in both strategic and operational planning.
Sluijk et al. (2023)	Identified a significant research gap in the utilization and optimization of satellite facilities in 2E-VRP networks, and emphasized the value of integrating direct deliveries and stochastic variables such as uncertain demand and travel times.

3.4.1. The Vehicle Routing Problem

Nielsen et al. (2024) refers to multiple studies that give a systematic review on vehicle routing problems. By backward snowballing in this article, many relevant studies can be found that give insights into vehicle routing problems that are occurring in multi-echelon networks like in this case study.

The Vehicle Routing Problem (VRP) is an extensively researched topic in the literature. A basic formulation for the VRP is described as discrete quantities of commodities that are to be delivered to n clients who are geographically dispersed around a central depot by m vehicles possessing identical capacity, initially stationed at a central depot (Clarke & Wright, 1964).

In the context of evaluating direct shipments compared to traditional distribution methods, the VRP can serve as a fundamental framework for analyzing the efficiency and cost-effectiveness of various logistics strategies. The evolution of VRPs into more complex variants, particularly multi-echelon distribution networks, is relevant for direct shipment strategies, as these systems include points such as warehouses and depots to enhance distribution efficiency (Mor & Speranza, 2022).

Two-Echelon Vehicle Routing Problem (2E-VRP)

In large freight transportation services with fast-moving consumer goods, two main distribution strategies are applied: direct shipping (single-echelon) and multi-echelon distribution (Sitek & Wikarek, 2015). In most real situations, the multi-echelon, and in particular the two-echelon, distribution is applied to optimize features like the number of vehicles used, transportation costs, loading factors, and lead times.

Optimal routes for the fleet of vehicles to serve customers' orders can be designed using the VRP. The problem of routing vehicles in a two-echelon distribution network is known as the Two-Echelon Vehicle Routing Problem (2E-VRP) (Sluijk et al., 2023). In such a problem, the distinction can be made between first-echelon (FE) and second-echelon (SE) vehicles, which are, respectively, the delivery operations between the supplier and the satellite, and the satellite and the customer (Guastaroba et al., 2016). The satellite refers to intermediate facilities, like DCs, that consolidate the transshipment of goods between vehicles on both echelons.

An illustration of a simplified two-echelon distribution network is given in Figure 3.3.

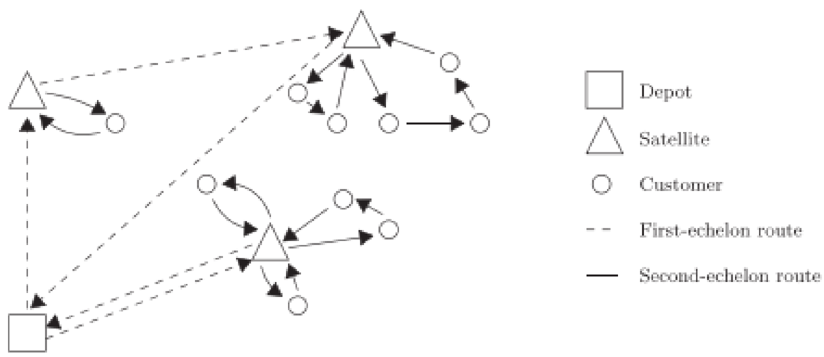

Figure 3.3: Two-echelon Distribution Network (Sluijk et al., 2023)

An extension to the 2E-VRP is the consideration of capacity constraints. In the Two-Echelon Capacitated Vehicle Routing Problem (2E-CVRP), freight delivery, containing FE between the plant and the satellite and SE between the satellite and the customer, is subject to vehicle and intermediate distribution center (satellite) capacity limits (Sitek & Wikarek, 2015). In this problem, timing of the deliveries is still ignored. Over the past few years, other variants of the 2E-CVRP have appeared in the literature (Sluijk et al., 2023). Examples of incorporated aspects that are experienced in real-life vehicle routing applications include time windows, pick-up and delivery operations, multiple commodities, heterogeneous vehicles, and stochastic demand and travel times.

Variations of the 2E-VRP

Sluijk et al. (2023) discuss the existing literature on the 2E-VRP and highlight a significant research gap aiming at the utilization of satellite operations (distribution centers) within distribution networks. It presents the opportunity for future research to explore the implications of optimizing satellite utilization, alongside the potential for direct deliveries from the plant to customers in regions accessible by FE vehicles.

Research into this dual or hybrid approach should quantify the benefits of direct shipments, assess their impact on overall routing costs, and develop distribution strategies that foster a more balanced and efficient use of satellite resources. This could eventually lead to a more streamlined and responsive distribution system. The P&G case study can be a practical example of testing such a hybrid approach using single- and two-echelon distributions, with a quantified impact on the solution cost as a result. Moreover, stochastic demand and travel times are often not accounted for in 2E-CVRP models. A case study like this could offer the opportunity to take these variables into account by using available historical data.

Table 3.5 gives an overview of several variations on the VRP model found in literature.

Table 3.5: Cross-table of 2E-VRP Variants and their Modeled Features in Literature

Study	2E	A2E	C	TW	CMC	DD	CD
Sluijk et al. (2023)	✗						
Anderluh et al. (2017, 2021)	✗						
Oliveira et al. (2022)	✗		✗				
Song et al. (2017)		✗	✗				
Marques et al. (2022)	✗		✗	✗			
Grosso et al. (2018)	✗			✗			
Wang et al. (2018)	✗				✗		
Zhou et al. (2024)	✗			✗		✗	
Azizi et al. (2020)	✗						✗

Abbreviations:
2E = Two-Echelon distribution network **A2E** = Adaptive Two-Echelon **C** = Capacitated **TW** = Time Windows **ATW** = Access Time Windows **CMC** = Collaborative Multi-Center **DD** = Direct Deliveries **CD** = Cross-Docking

The multiple studies discuss different variations and end with results and several open gaps. Relevant highlights from the studies are provided in this section.

Both studies of Anderluh et al. (Anderluh et al., 2021, 2017) develop two-echelon city distribution schemes and highlight the importance of considering the city layout in freight distribution strategies. They integrate 'grey zone' customers into an optimization model for the 2E-VRP which balances economic, environmental, and social objectives (Anderluh et al., 2021). The results of both papers can mostly give companies decision recommendations in planning a sustainable city distribution concept.

Oliveira et al. (2022) also focuses on addressing a variant of the 2E-CVRP within the context of city logistics. The study explores the establishment of two-echelon distribution systems where freight from the outskirts of the city is transferred at intermediate locations (satellites) from large urban freighters to smaller, more environmentally friendly city freighters. It addresses the proposed Variable Neighborhood Search (VNS) heuristic as a flexible and capable method for real two-echelon distribution problems with varying operational characteristics.

An extension of the 2E-CVRP is discussed by Song et al. (2017), the Adaptive Two-Echelon Capacitated Vehicle Routing Problem. This model allows multiple depots and enables direct deliveries from depots to customers, bypassing the satellites if that is unnecessary. Unlike 2E-CVRP, where all shipments must go through satellites, A2E-CVRP adapts to real-world logistics by optimizing costs through flexible routing. The multiple shipments this model allows can be seen as shipments possible from multiple plants, but which will always be directed via a DC, it does not inherently model direct shipments from the manufacturer.

Additionally, Marques et al. (2022) presents the first exact algorithm for the multi-trip variant of the 2E-CVRP with time windows (TW). The study demonstrates that allowing for storage and consolidation of freight at satellites can significantly decrease total transportation costs compared to exact synchronization. Their branch-cut-and-price algorithm solves instances notably faster than previous approaches. Moreover, a study by Grosso et al. (2018) explores the effect of access time windows (ATW), solely for classic VRPs in urban areas. Their mathematical model demonstrates higher energy consumption and operational costs for carriers and provides a framework for evaluating and optimizing delivery routes under ATW restrictions.

Another extension of the model, presented by Wang et al. (2018), accounts for collaboration between logistics centers (LCs) and distribution centers (DCs). Unlike the 2E-VRP where depots and satellites operate independently, this 2E-CMCVRP (Two-Echelon Collaborative Multiple Centers Vehicle Routing Problem) allows synergies between different facility types to optimize costs and reduce emissions. The study does not give an indication of the use of direct shipments from the manufacturer or plant; instead, it focuses on collaboration between existing logistics centers within the two-echelon network.

The most extended variant found in the literature that is recently studied is the 2E-VRPDDATW by Zhou et al. (2024). This model involves time windows, access time windows, direct deliveries, and synchronization in a 2E-VRP and will form a good reference for modeling the situation in this specific research.

3.4.2. MILP Models in VRP

Mixed-Integer Linear Programming (MILP) is highly relevant for formulating vehicle routing problems due to its ability to handle complex decision-making scenarios characterized by multiple variables and constraints (Azizi & Hu, 2020). MILP models basically involve problems in which some variables are constrained to be integers while others can be non-integer (real) values.

In the context of SCM, MILP allows for the integration of various components such as the location of distribution centers, vehicle routing, and direct shipment into a cohesive optimization model. It is a robust mathematical technique to derive optimal or near-optimal solutions efficiently. MILP has been widely adopted in supply network design due to its robust optimization capabilities (Garrido et al., 2025). The two-echelon vehicle routing problem with cross-docking (2E-VRPCD) is a well-known example where MILP models have been applied to optimize the total transportation cost by managing freight delivery from a single depot to customers via intermediate depots (Dondo et al., 2011).

Zhou et al. (2024) use MILP to integrate direct deliveries, time window constraints, and access time window constraints into the VRP. This approach is necessary to accurately model the practical constraints and synchronization requirements between the first and second echelons. The ability of MILP models to integrate multiple components and constraints into a unified optimization model makes it a powerful tool for solving complex vehicle routing and supply network problems.

3.5. Discussion

Despite significant advances in distribution modeling, several important research gaps remain. One key area is the limited integration of hybrid distribution strategies. While direct shipments, multi-echelon distribution, and cross-docking have been studied independently, few works address their simultaneous application. Given the increasing complexity of modern supply chains, understanding how these strategies can be combined to balance cost efficiency, environmental performance, service levels, and flexibility is crucial. This study addresses this gap by developing and evaluating an integrated hybrid distribution model using a real-world case at P&G.

Beyond structural integration, the trade-offs of the impact of hybrid distribution on cost-optimality and sustainability-optimality is also underexplored. Most studies focus on logisitics cost minimization, whereas for example Zhou et al. (2024) and Sluijk et al. (2023), frequently used as reference for this research, do not consider the environmental performance alongside cost-efficiency. This study contributes by evaluating how hybrid strategies create certain trade-offs in balancing both these relevant KPI's in the model.

Another gap is the predominant focus on urban logistics. Much of the existing literature addresses intra-city delivery challenges, such as traffic congestion and narrow time windows. By contrast, cross-country supply chains face different complexities, including long-haul coordination, inventory placement, and balancing flexibility with cost efficiency. The integration of direct plant shipments in such contexts remains largely unexamined. The P&G case study offers a valuable opportunity to assess hybrid distribution in a cross-country FMCG network. Unlike last-mile city logistics, this context involves larger shipment volumes, fewer delivery constraints, and a broader network structure. Understanding the trade-offs between direct and multi-echelon deliveries in such settings can yield practical insights for optimizing large-scale distribution strategies.

Together, these observations form the basis for answering SQ1 and SQ2 in the following conclusion. By addressing the identified gaps, this research aims to contribute to the development of more integrated, resilient, and service-oriented distribution models, with actionable implications for both academia and industry.

3.6. Conclusion

This conclusion synthesizes the main insights from the literature review and addresses SQ1 and SQ2, laying the foundation for the case study.

Supply Chain Management (SCM) involves the strategic coordination of procurement, operations, and logistics to improve efficiency and responsiveness. Key pillars include demand forecasting, lean operations, strategic sourcing, and effective outbound logistics. Within distribution, three strategies stand out: direct shipping, multi-echelon networks, and cross-docking. Each offers specific trade-offs in terms of cost, complexity, and responsiveness. To optimize such systems, researchers often use modeling approaches like the Vehicle Routing Problem (VRP), with Mixed-Integer Linear Programming (MILP) commonly applied to handle complex logistical constraints.

1. What are the current logistical processes and challenges associated with using distribution centers versus direct plant shipments, and how can their efficiency be measured?

Distribution centers (DCs) support cost-efficient consolidation, improved routing, and inventory management, but add complexity, coordination burdens, and lead times. Direct plant shipments reduce handling and may offer faster delivery, particularly for high-volume or geographically close customers. However, they can be less efficient for low-demand items due to limited vehicle utilization. Measuring the efficiency of both strategies requires performance indicators related to cost, emissions, and customer service levels, such as transportation cost, delivery time, and order fulfillment rates. Optimization tools like the VRP enable comparative evaluation of these strategies under varying demand and capacity scenarios.

2. What is a suitable modeling approach to evaluate the added value of direct shipping in hybrid distribution strategies, and what theoretical requirements should be considered when modeling a supply chain?

To evaluate hybrid strategies combining direct and DC-based deliveries, advanced VRP variants such as the Two-Echelon VRP (2E-VRP) and its extensions (e.g., 2E-CVRP, 2E-VRPDDATW) are effective. These models can incorporate real-world constraints like capacity, time windows, and stochastic demand. MILP provides a robust framework for modeling such systems, enabling scenario comparison and optimal configuration selection. Theoretical modeling should include dynamic variables such as demand, distances, vehicle constraints, and customer service levels. Hybrid modeling reflects modern logistics practices more accurately and supports better decision-making for balancing cost, service, and flexibility. This study builds on that by applying such models to a cross-country FMCG supply chain at P&G.

4
Current State Analysis P&G

Based on the literature review, stakeholder interviews, and additional data analysis, the current practices at P&G, and specifically in the case study, will be given in this section. The most relevant pillars of the supply chain for this research have followed from the literature review (section 3.2). The focus in this current state section will be on Demand and Logistics Management. A brief overview is given in Figure 4.1. Before evaluating the current practices of demand and logistics at P&G, an analysis of relevant stakeholders that are involved in these processes will be given.

All information in this chapter is based on verified internal sources at P&G (including interviews, internal documentation, and validated data analyses), and is therefore not referenced explicitly in the text.

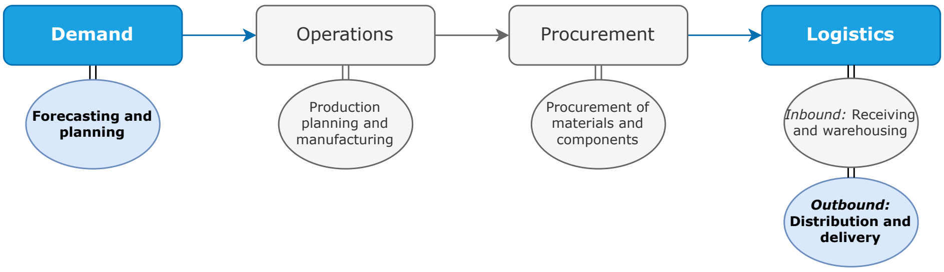

Figure 4.1: SCM Overview (Focus on Demand and Logistics)

The entire supply chain process starts with demand, which serves as the foundation for all subsequent activities. Accurate forecasting is crucial for predicting customer needs across the markets and different product categories. This information leads to global production and supply planning, ensuring efficient allocation of resources to meet the forecasts. P&G performs worldwide supply chain operations, with production plants and customers across all continents. While many of P&G's production facilities operate as self-sustaining units within their respective regions, certain products are produced and shared globally, needing robust overseas logistics capabilities. The global share of products adds complexity to the supply chain, as it requires coordination between various manufacturing plants and warehouses. Additionally, plants rely on raw materials from non-P&G suppliers, which further complicates the operational and procurement landscape.

After completing production and procurement, the focus shifts to logistics management, which entails both inbound and outbound logistics. The inbound logistics involve transportation and storage of raw materials and components necessary for the final production. Outbound logistics refer to the distribution of finished products to distribution centers and ultimately to customers. This research will focus on distribution methods as part of the outbound logistics. Moreover, it is relevant to analyze the impact on demand forecasting and planning if hybrid distribution approaches are applied.

4.1. Stakeholder Analysis

This section outlines the key stakeholders involved in the demand and logistics management processes at P&G. The analysis is based on interviews and practice observations conducted during the Research phase. Each stakeholder plays a distinct role in the planning, execution, and optimization of the supply chain.

Demand Planning and Forecasting - The Demand Planning and Forecasting teams are responsible for generating accurate sales forecasts and translating them into actionable production and distribution plans. This involves both statistical modeling and close collaboration with commercial teams to align forecasts with promotions, seasonal trends, and market dynamics. Their forecasts serve as the basis for supply planning, inventory management, and distribution decisions, namely the forecast on product level is the eventual trigger for production.

Supply Chain Leaders - Supply Chain Leaders bridge demand and supply by ensuring that production aligns with the forecast. They coordinate with manufacturing sites and DCs to ensure supply continuity and efficiency. This means that all products that are needed at the DC because of the demand for that market,need shuttles from plant locations to the DC. The Supply Chain Leaders at Plant 1 and Plant 2 are responsible for the production and distribution planning with an horizon of 12 weeks.

Demand Requirement Planners - The transportation process of the products must ensure timely delivery of goods from production plants or warehouses to customers. At P&G, these stakeholders are called the Demand Requirement Planners (DRP). They are responsible for the deployment of products from the plant to the DC, so, in this scope, from Plant 1 and Plant 2 to the DC 1. A DRP is focused on family level of a plant and has several DCs in his/her scope. For example, the produced Category X Segment A is deployed to the DCs by one DRP, based on the forecast of these DCs of the products on this specific family level. The task of a DRP is limited to authorizing and planning the volumes that have to be shuttled, while the transportation planning for the actual execution of these volumes is done by Transport Planners.

Inventory Managers (SIP) - Inventory managers monitor stock levels at warehousing locations to ensure product availability while minimizing holding costs. Within P&G, these stakeholders are called Site Integrated Planning (SIP) Leaders.

Warehouse Operators - Warehouse Operators are responsible for the day-to-day management of stock within the warehouse. Their tasks include inventory handling, order picking, and ensuring timely dispatch of goods to customers or distribution hubs. At the hub location in Plant 1, you could split the warehouse operations into three teams; the first team is responsible for the product flow and stocking in the CIMAT (automated warehouse), the second team is responsible for the management of the shipments of orders, and the third team is responsible for the executing operations (picking, loading, transportation). The hub location in Plant 2 has a similar set-up, further elaborated in subsection 4.3.5.

3PL - An external stakeholder, , manages warehousing operations under contract (third-party logistics). They are owner of the DC and the transportation business from and to the DC. P&G hires their services, which entails inventory holding costs and handling costs for utilizing the DC for consolidation of the products.

Transport Operations Managers - Transport Operations Managers (T-OPS) are responsible for executing physical shipments across the supply chain, ensuring that the products move efficiently from production sites to DCs or directly to customers. They manage carrier relationships, monitor real-time transport flows, and ensure compliance and documentation for the transportation are in place. A key responsibility lies in coordinating both full truckloads and consolidated shipments. Full loads incidentally ship directly to customers, while partial loads or mixed-product shipments are routed and consolidated through DC 1. T-OPS could track the potential DPS that could happen from plants to customer DCs. This makes that they have a central role in optimizing transport strategies and shaping the future structure of the direct or two-echelon distribution network.

Customers - P&G has a big scope of customers in Belgium and The Netherlands, which are large retailers and smaller distribution companies of P&G's products. Customers are served to their own DC locations, there is no execution to small retail locations or to the end consumer. The retailers and distributors are the end-consumer of P&G's supply chain. Their ordering patterns, expectations, and feedback influence upstream planning and logistics operations.

Order Management - The Order Management (or Availability Management) team is responsible for processing and tracking customer orders from entry to delivery. They ensure that orders are accurately captured and executed within the agreed lead times. Unlike aligning strictly with available inventory, the team often 'cuts' orders rather than delaying them to manage inventory constraints. This management involves coordination with customer service, inventory management, and transport execution to guarantee a seamless end-to-end order fulfillment process. The order management is a coordinating role done from the Supply Network Operations (SNO) office and is not physically located at the plants or DCs.

Customer Relationship Managers (CRM) - Customer Relationship Managers (CRMs) act as a bridge between P&G and its customers. In addition to the responsibilities of Order Managers, CRMs handle order processing on the customer side, manage delivery expectations, and resolve service issues to ensure a smooth customer experience. Each CRM operates at customer level, managing orders across all product categories. This means they are not limited to specific Cat X&Y plants, but these products are part of the customers' orders they manage. Furthermore, CRMs are responsible for generating customer-specific forecasts (Business Intelligence) and for sharing forecasts for all events in the upcoming 2-3 months in collaboration with the customer, including any changes or updates.

Suppliers (Out of Scope) - Although not in the scope of this research, suppliers play a critical upstream role in ensuring the availability of raw materials. Their lead times and reliability influence the downstream planning and logistics processes. Many above mentioned stakeholders are in daily contact with suppliers to maintain control over the production and distribution planning for in- and outbound logistics of the Cat X&Y products.

Figure 4.2 provides a schematic overview of the key stakeholders involved in the demand and logistics management processes at P&G, and illustrates the relationships and flows between them. The swimlane structure categorizes stakeholders into three functional domains: Planning, Execution, and Customer. The arrows indicate the direction of information flow, operational handover, or decision-making influence, offering a visual representation of how responsibilities are distributed and interconnected across the end-to-end network.

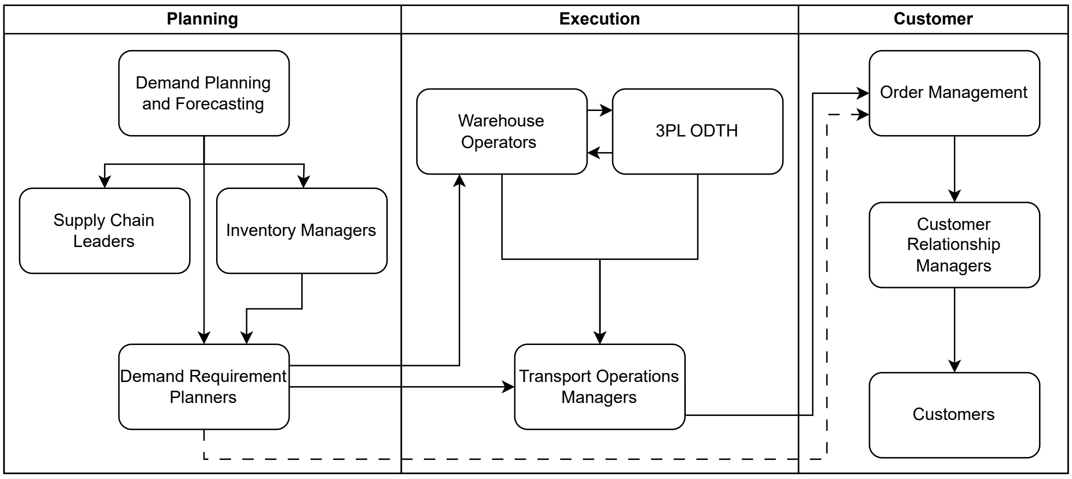

Figure 4.2: Stakeholder Swimlane Diagram

4.2. Demand Management at P&G

This section outlines how demand is managed within the P&G supply chain, starting from forecasting and planning to customer order handling. The aim is to provide an integrated view of how demand triggers supply and logistics activities across the European network.

4.2.1. Demand Forecasting

The forecasting process is essential to initiate the supply chain of goods. As mentioned in the literature review, forecasts are critical for minimizing inventory costs while meeting customer expectations (Altekar, 2023). Forecasts are made on both customer and product category levels. Within P&G, each product category is split into multiple value streams or families, each containing several product codes that correspond to the final consumer products.

Every product progresses through five stages before reaching the market: Discover, Design, Qualify, Ready, and Launch. The Discover phase involves researching consumer and business interest to identify promising ideas. The Design phase focuses on early learnings and feasibility assessments. In the Qualify phase, packaging, marketing, and customer plans are finalized. The Ready phase ensures all systems are prepared for expansion and local market entry. Forecasting becomes essential here to trigger production and plan shipments to the correct distribution centers. Forecasts are made on family level to drive production, while commercial forecasts are created with customer teams and managed centrally from the European Supply Network Hub (SNH). This hub coordinates production across internal and external European plants, ensuring they receive the necessary raw materials. Forecasts thus guide planning and production to meet projected demand. Once demand inputs are used for supply execution, the product can enter the Launch phase. It is essential to adjust forecasts post-launch to ensure resource efficiency.

P&G's demand forecasting process includes three complementary cycles:

Together, these cycles form a layered structure that balances short-term responsiveness with long-term alignment.

[Figure anonymized for confidentiality reasons]

Figure 4.3: Demand Forecasting Cycles

Final demand forecasts are built from three components: a statistical baseline, market intelligence, and manual adjustments (Figure 4.4). The statistical forecast uses historical sales and predictive tools. Market intelligence from customer teams, category leads, and external trends (e.g., promotions or macroeconomic shifts) is then integrated. Lastly, business insights and tools like cannibalization adjustments are applied. Once finalized, forecasts are shared across Supply Planning, Finance & Accounting, and Business Planning functions to enable end-to-end decision-making.

[Figure anonymized for confidentiality reasons]

Figure 4.4: Forecast Building Process (P&G, 2025)

Currently, demand forecasts are created and applied at DC level, with the Belgium DC serving BNL customers and the Northern France DC serving FR customers. This setup triggers replenishment from production plants to DCs. However, in a Direct Plant Shipment (DPS) setup, where shipments bypass DCs, forecasts must be allocated to the respective plant instead.

This shift would require a hybrid forecast ownership model where both DC- and plant-level forecasts coexist. In a hybrid distribution system, forecasts must specify not only the required volume but also the fulfillment source. Forecasting and planning systems must incorporate routing logic and customer-specific requirements. Proper allocation helps ensure accurate inventory holding, minimize handling and storage costs, and maintain flexibility to evaluate the cost and service trade-offs in hybrid models. Over- or understocking at either DCs or plants should be avoided.

4.2.2. Demand Planning

Plant 1 and Plant 2 produce Cat X&Y, respectively. Plant 1 produces one large Category X family locally, while others are sourced externally (see Table 1.1). Similarly, three major Category Y families are produced in Plant 2, and two others are imported. All products sold in the FBNL market are transported to DCs in Belgium and Northern France based on demand planning. Monthly forecasts trigger production and inventory levels at the DCs, further elaborated in section 4.3.

4.2.3. Customer Order Management

Order management is a daily process.

[Figure anonymized for confidentiality reasons]

Figure 4.5: Life of an Order (Day 1 to Day 3) (P&G, 2025)

On customer side, order volumes are typically generated automatically. Promotions or other activities may lead to manual adjustments. Customers send their mixed-product orders to P&G systems, which are verified and routed to the appropriate DCs for picking and shipping. Sometimes, customer orders are unstructured, using wrong product codes or placing multiple small orders instead of one. This requires manual corrections by P&G employees. To incentivize efficiency, P&G offers discounts for Full Truckload (FTL) orders and for timely, correctly submitted orders.

AOV

With Advanced Order Visibility (AOV), customers place orders at least 17 days before Day 3. This improves supply-side planning and reduces the risk of order cuts or delays. AOV enhances efficiency through fewer stockouts, better truck utilization, and less DC waiting time.

4.3. Logistics Management at P&G

Logistics management is a relevant pillar within SCM for this study. In this section, the current practice of specifically inventory management and transport are outlined.

4.3.1. Inventory Management at the DC

Based on accurate demand forecasting and planning, the right amount of inventory is available to meet customer needs without over- or understocking. Though demand management is critical for holding the correct inventory at the correct locations, the focus for this case analysis shifts to the logistics aspect. Effective inventory management within logistics ensures that inventory is tracked, stored and distributed efficiently while minimizing the costs. An important consideration in the inventory management strategy of P&G is the maintenance of safety stocks at the different locations. Safety stock, or buffer stock, is the inventory level that protects against unforeseen events such as inaccuracies in demand forecasts or unexpected delays in supply (ABC Supply Chain, 2023).

Demand uncertainty arises from fluctuations in customer needs, which can be triggered by the flexibility P&G delivers to customers. Ordering patterns of customers are not always consistent and the forecasting process is critical to give the best estimates. Lead time uncertainty is caused by the unpredictable nature of supply chain processes, including order delays, production lead times, transit times and the customization process of products. Figure 4.6 gives an illustration of how inventory levels are managed using safety stock (SS) in a supply chain like Procter & Gamble's. The green line shows how stock decreases over time as products are sold, and new stock is deployed to locations to avoid hitting the safety stocks. The red dots show the point of reordering or reproducing in advance of reaching the SS volume. Combining the Economic Order Quantity (EOQ) with SS, ensures optimization of ordering quantities and protection against uncertainty.

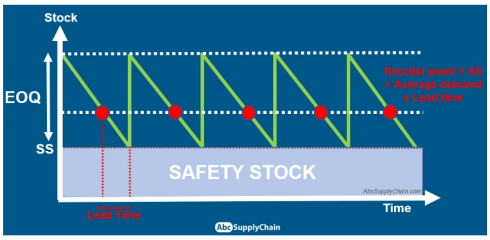

Figure 4.6: Safety Stock Illustration (ABC Supply Chain, 2023)

P&G applies the First Expired, First Out (FEFO) principle to prioritize dispatch of products nearing expiration. Because incoming products often start as safety stock, this principle ensures that SS is not treated separately but remains part of the dynamic inventory flow, helping to manage shelf life efficiently.

In the current setup, demand forecasts and inventory flows are centered around distribution centers (DCs), with customer orders fulfilled from stock held at these sites. This two-echelon structure allows consolidation of products from multiple plants before shipment. However, this centralized approach also creates dependency on DC operations and stock levels. As a result, integrating DPS would significantly impact inventory positioning and flow decisions across the network, which is explored in the next section.

4.3.2. Stock Allocation in a Hybrid Distribution Setup

Integrating DPS into the distribution network introduces important changes to how and where inventory is held. This directly affects the third key performance indicator of this study: Stock Allocation, referring to how inventory is strategically positioned across the supply chain. With DPS, some demand would be fulfilled directly from plants, reducing the need for inventory at DCs, particularly at externally managed sites like DC 1, where P&G incurs storage and handling costs. Avoiding these costs by shifting inventory upstream can result in significant financial benefits. Additionally, direct shipments reduce internal handling, sorting, and transfer activities at the DC. A crucial note for this research is that the Plant 1 and Plant 2 offer sufficient physical space to accommodate this upstream stock reallocation, making such a shift operationally feasible within the new strategy.

However, this upstream shift introduces operational challenges. Direct fulfillment requires sufficient stock availability at the plant at the moment of dispatch, placing more pressure on production planning and forecast accuracy. Moreover, plants like Plant 1 lack the infrastructure for customer-specific services such as case-picking or Customer Load Preparation (CLP), which includes pallet customization, box opening, and labeling. As a result, only certain products and customers are suitable for DPS.

While DC 1 is sometimes compared to a cross-dock facility, it does not fully align with cross-docking principles. Rather than acting as a transient hub, DC 1 serves as a consolidation center with strategic stockholding. It facilitates planned consolidation of products from different plants to complete customer-specific shipments, making it an essential buffer in the current network. One of the advantages of introducing DPS from Plant 1 and Plant 2 is the potential cost reduction at the 3PL-operated DC 1, by lowering storage and handling needs for volumes that can bypass the DC altogether.

4.3.3. Transport

Transportation plays a crucial role in P&G's supply chain design and is a central component in evaluating the potential for hybrid distribution strategies. The current setup follows a two-echelon structure in which finished products are first transported from the production Plant 1 and Plant 2 to the DC 1 (first echelon, or FE), and subsequently from DC 1 to customer locations in Belgium and the Netherlands (second echelon, or SE).

Transport costs between the plants and the DC are referred to as Inter Site Freight (ISF) costs. This is a fixed shuttle price per vehicle. Customer deliveries from DC 1 are associated with Customer Freight (CF) costs. These costs are fixed per destination region and based on haulier agreements. While they do not directly scale with exact distance, costs differ per region based on the distance. Every customer location is part of one region and has the respective CF transport cost. The same hauliers and vehicle types are often used across both FE and SE movements, with similar truck capacities and vehicle fill rate (VFR) considerations.

While routing decisions are partly automated based on forecast logic and shipment volumes, manual interventions by transport planners are common to prevent inefficiencies. For example, FTL orders from Plant 1 may be routed through DC 1 if additional products need to be added to the order from other plants. Since part of the order fills a full truck with origin from the plant, orders are directed through DC 1 unnecessarily, either due to system defaults or no manual correction. This shows the network is not fully equipped to support DPS for Cat X&Y within the region. An example where the direct shipment setup is currently being used more structural for another category, is discussed in subsection 4.5.2.

Transport benefits of DPS

Several transportation-related benefits are associated with a hybrid distribution strategy that includes DPS. First, direct shipments bypass intermediate handling at the DC, avoiding the associated unloading, storage, and reloading costs. This also helps alleviate storage pressure at the DC, which has been reported as a current problem. Secondly, DPS reduces the number of transport steps and trucks needed across the network. By shipping full truckloads directly from the plant, Vehicle Fill Rate (VFR) is optimized, potentially reducing the total number of trips, fuel consumption, and emissions, an outcome that aligns with P&G's environmental goals. Thirdly, DPS reduces reliance on shuttle capacity between plants and the DC. These shuttle movements (ISF) are not unlimited, so a decrease in needed shuttles is beneficial. When large customer orders are delivered directly from the plant, the freed-up shuttle capacity can be redeployed to handle smaller or urgent orders requiring DC consolidation. This indirectly increases responsiveness from the DC in cases where flexibility is needed.

VFR and FTL/LTL

The distribution model distinguishes between Full Truckload (FTL) and Less-Than-Truckload (LTL) transport.

This is an initial internal policy decision for considering a shipment as eligible for dispatch as direct flow. Because a DPS flow would always involve full pallets, because no case-picking is possible at plants, they allow for better double-stacking and inherently achieve a high VFR.

An overview of key cost and operational parameters is presented in Table B.1 in Appendix B. For all shipments originating from the DC 1, fixed Customer Freight (CF) costs are available for each destination region based on historical haulier contract data. These rates serve as a reliable reference for modeling transportation costs within the two-echelon distribution setup, where the DC acts as the consolidation and dispatch point. For DPS from Plant 1 and Plant 2, historical CF cost data is only partially available from incidental vehicle trips on this direct flows. Consequently, several regions lack observed freight rates from these origins. To enable consistent modeling across the whole network with reliable data input, missing values have been estimated. These assumptions are elaborated in section 5.2.

It is observed that the CF rates from Plant 1 to customer locations are generally higher than those from Plant 2. This discrepancy is due to the fact that the available CF values from Plant 1 are based on a small number of historical DPS operations that actually occurred, which turned out to be relatively expensive. These specific cases, although limited in number, have been contractually executed and therefore reflect real, but possibly non-representative, costs. This could lead to relatively higher DPS CF cost in the model with the estimated rates, which would only support the eventual cost reductions with DPS more, if the actual rates would be lower.

4.3.4. Plants in Scope for DPS Integration: Plant 1 - Category X

Plant 1 is the main production site for Segment A of the Category X category and also serves as an import hub for other Category X segments produced overseas. These imported products arrive via an import DC, where they are palletized and then shipped to .
There is no direct flow from the import DC to DC 1; if DC 1 is the final destination, shipments must first pass through Plant 1 (see Figure 4.7 and Figure 1.2).

[Figure anonymized for confidentiality reasons]

Figure 4.7: Setup Import DC - Plant 1 - DC 1 (P&G, 2025)

4.3.5. Plants in Scope for DPS Integration: Plant 2 - Category Y

Plant 2 is the primary production site for the main families within the Category Y category. While its setup is comparable to Plant 1, the local distribution infrastructure differs slightly.

For the purpose of this research, Plant 2 is considered as a single origin point for DPS to BNL customers. The internal split between CIMAT and Big Box is not further distinguished in the analysis, as the focus lies on evaluating the feasibility of DPS at the plant level.

4.4. Service Level Metric

In this research, the impact of DPS on service level will be assessed to determine whether DPS improves or decreases this service level metric. To measure service, historical service levels per shipping location are used to estimate the effect of DPS on service. From P&G data, service levels are given per origin shipping location and for specific destinations, such as customer groups or markets. For this research, the service level of the past 8 months from the three locations, Plant 1, Plant 2, and DC 1, is taken to make an estimation for the change in service level with new routing decisions as with DPS. This is further elaborated in subsection 5.1.6.

4.5. Conclusion

This section answers SQ3 of the Research Phase: *What is the current practice of P&G's distribution model in the FBNL region, and what are the opportunities and limitations for integrating a hybrid distribution approach that includes direct plant shipments?* The analysis of P&G's two-echelon setup, supported by stakeholder input and logistics examples, revealed tangible opportunities to integrate DPS into the current network. However, several operational and planning limitations must be addressed to ensure feasibility. These insights form the foundation for the Design Phase of the research. SQ4 guides this transition and is formulated as follows: *What are the key design requirements for modeling the integration of direct shipments into P&G's hybrid distribution network, considering both theoretical foundations and the practical conditions for customer eligibility?*

SQ3 and SQ4 will be answered in this section by separate discussion of the current state, the opportunities, the model requirements and finally a reference to the section where customer eligibility is treated. The requirements are categorized into hard and soft requirements and linked to the research KPIs and relevant stakeholders, ensuring that the model developed in the next phase is both theoretically and practically grounded.

4.5.1. Current Distribution Practice at P&G

The current distribution model of P&G in the FBNL region is structured as a two-echelon network, in which Cat X&Y goods from the Plant 1 and Plant 2 respectively are transported to the DC1 before reaching customer DCs. This setup enables product consolidation and value-added services at DC 1, but also introduces additional handling costs, inventory holding costs, and routing inefficiencies. Where incidental usage of direct shipments in historical data is seen, the need for exploring full hybrid distribution strategies with direct shipments from Plant 1 and Plant 2, comes from the increasing pressure on the DC, the service level targets to be maintained, and on the other hand the urge of looking for cost optimizing solutions. Moreover, it is observed that many shipments from the DC's are FTL shipments that are entirely shuttled from the plant to the DC before being delivered to the customer, often without requiring further consolidation. This has raised the question of whether DPS could structurally replace certain two-echelon distributions, especially when customer orders are large enough to justify a FTL shipment from the plant.

The urgency of exploring DPS has also appeared from stakeholder interviews across multiple stakeholder groups. Demand planners and supply chain planners experience pressure to streamline production-to-delivery flows amid rising forecast granularity and demand volatility. DRPs and T-OPS face increasingly complex deployment schedules and rising transport costs, while inventory managers are tasked with minimizing storage footprints with the current capacity constraints at the DC. At the same time, warehouse operators and 3PL providers are confronted with space saturation and costly double handling. For customers and CRMs, ensuring delivery reliability and high service levels remains essential, yet current two-echelon structures can introduce delays or can lead to order cuts, impacting the service measure. By researching and quantifying the potential of DPS integration, this study addresses a critical operational need: enabling P&G to move towards a more responsive and cost-efficient distribution network that helps reduce pressure and inefficiencies across different stakeholder tasks in their supply chain.

4.5.2. Opportunity for Hybrid Distribution in FBNL Scope

This demonstrates that hybrid distribution strategies, combining direct shipments with two-echelon flows, are not entirely new to P&G's network. Building on this precedent, there is an opportunity to investigate whether similar approaches could be beneficial for the Cat X&Y categories within the FBNL scope. Customers with high-volume, homogeneous orders that meet FTL thresholds may be suitable candidates for DPS. Exploring this potential requires careful analysis of demand patterns, routing logic, and plant capabilities, but the fact that real implementations exist elsewhere strengthens the case for evaluating hybrid distribution in this region. Furthermore, this research is needed because of the need for a detailed evaluation of the impact of full integration of DPS on the defined KPIs, to see whether the hybrid strategy is beneficial on the long-term.

4.5.3. Key Design Requirements

To support the integration of DPS into P&G's hybrid distribution model, several operational and planning requirements have been identified based on the current state analysis. These requirements clarify both the technical feasibility and organizational readiness for implementing DPS. They are categorized into two groups: hard requirements, which are essential prerequisites for DPS to function reliably, and soft requirements, which are desirable enhancements that improve overall performance but are not strictly necessary for feasibility.

Each requirement has been linked to one or more of the KPI's: cost (logistics and environmental costs), service (service level performance), and cash (stock allocation efficiency). Additionally, the requirements have been mapped to the primary stakeholders responsible for, or impacted by, that condition. This stakeholder-KPI-requirement matrix ensures alignment between the modeling approach and practical implementation concerns across P&G's supply chain.

Table 4.1: Hard and Soft Requirements for Integrating DPS, Linked to KPIs and Stakeholders

Requirement	Impacted KPI	Primary Impacted Stakeholder
Hard Requirements		
Physical network layout	Cost/Service/Cash	Warehouse Operators, Supply Chain Leaders
Fulfill FTL thresholds	Cost	Transport Operations Managers
No splitted or multi-drop orders	Service	Transport Operations Managers
Plant-Level forecast allocation	Cash	Demand Planning and Forecasting
Loading ability DPS orders	Cost/Cash	Warehouse Operators and Transport Operations Managers
Vehicle availability at all origins	Cost/Service/Cash	Transport Operations, Supply Chain Leaders
Weekly demand must be met	Service	DRP, CRM
Scalable plant storage capacity	Cash	Warehouse Operators, Inventory Managers
Soft Requirements		
Reduction of logistics costs	Cost	Supply Chain Leaders, Warehouse Operators
Minimization of handling	Cost	Warehouse Operators, Transport Operations Managers
Ensured safety stock adaptability	Cash	Warehouse Operators, Inventory Managers
Accommodation of order volume variability	Service	DRP, CRM, Customers
Scalable and flexible DPS framework	Service	Supply Chain Leaders
Contribution to sustainability goals	Cost	Supply Chain Leaders, Transport Operations Managers

Hard Requirements:

- **Physical network layout:** The current network in scope with one DC must be respected in the model as there is no possibility for changing the physical locations.

- **Fulfill FTL thresholds:** DPS must be performed with fulfilling the FTL threshold of . Below this threshold, DC consolidation currently remains required based on internal policy decisions.
- **No splitted or multi-drop orders:** Delivery of multi-drop or splitted orders is not an option for shipments directly from the plant nor the DC, based on internal policy decisions. Every shipment to a customer is originated from one source.
- **Plant-Level forecast allocation:** Forecasting systems must allocate volumes to specific plants rather than all directly to the DC to enable automated identification of DPS opportunities.
- **Loading ability DPS orders:** Plants must support loading of full-pallet, multi-SKU orders as DPS to customers without case-picking or customization.
- **Vehicle availability at all origins:** The 3PL vehicles must be sufficiently available at the plants to be able to perform DPS, next to the ongoing availability of vehicles at the plant.
- **Weekly demand must be met:** In order to integrate DPS successfully maintaining service levels, all weekly demand must be met in baseline.
- **Scalable plant storage capacity:** Plant-side storage should be able to flexibly scale to support additional DPS flows, especially under shifting distribution scenarios.

Soft Requirements:

- **Reduction of logistics costs:** It is wishful that the distribution setup minimizes transport and routing costs across both DPS and two-echelon flows.
- **Minimization of handling:** Reducing intermediate unloading and reloading is desirable to limit operational complexity and avoid additional handling costs at the DC.
- **Ensured safety stock adaptability:** It is preferred that safety stock (SS) levels can dynamically adjust in response to different lead times and demand variability under DPS.
- **Accommodation of order volume variability:** A hybrid system should flexibly serve both large (FTL) and small (LTL) customers without service disruption.
- **Scalable and flexible DPS framework:** The DPS approach should support a configurable structure with a frequent and clear DPS shipment pattern that could also be adjusted to customer needs.
- **Contribution to sustainability goals:** By maximizing VFR and reducing unnecessary shuttle miles, DPS should support P&G's carbon and fuel reduction targets.

Based on the identified requirements and their alignment with the cost–service–cash dimension, the next chapter presents the development of the model to evaluate the performance of hybrid distribution. Requirements can be translated into modeling assumptions in subsection 5.1.1, that will form the basis of modeling and simplifying the distribution system to make trade-offs between two-echelon shipments and DPS.

4.5.4. Practical Conditions for Customer Eligibility

The second part of SQ4 concerns identifying which customers may benefit from integrating DPS into the distribution model. This is evaluated through a detailed statistical analysis of historical demand data, focusing on customer-level demand consistency and volume adequacy on Categories X & Y. The results of this evaluation are given in subsection 5.2.4. These findings serve as the foundation for further modeling and simulation of hybrid distribution strategies in the design phase, giving the model assessment a focus on fewer customers that could actually be eligible for structurally using DPS.

5 Hybrid Distribution Model: a Flow-Based Adaptation of the 2E-VRP

As discussed in section 2.5, the Research Framework for the modeling approach is stated in Table 2.2. This chapter covers the first steps of that framework: model formulation, data preparation, model implementation, verification, and validation. Once the model is set up, the different configurations will be reported in the Experimental Design (chapter 6).

The developed model is designed to support strategic decision-making by simulating trade-offs between the three defined KPIs, based on the supply chain triangle; minimizing logistics costs and emissions (cost), while assessing the implications for service level (service) and stock allocation efficiency (cash). The optimization identifies the most cost- and emission-efficient distribution strategy for a given demand configuration. Based on this optimal solution, the resulting effects on the other two KPIs, service level and stock allocation, will be evaluated as post-analysis after the optimization.

5.1. Model Formulation

This section outlines the mathematical structure of the optimization model developed for the analysis. While inspired by traditional two-echelon vehicle routing problem (2E-VRP) formulations, such as those by Zhou et al. (2024), Sluijk et al. (2023), and others (subsection 3.4.1), the model applied in this research adopts a simplified structure tailored to the real-world context at P&G. Unlike classical 2E-VRP models that explicitly route vehicles across multiple destinations, include return requirements, and enforce time synchronization between stages, this model focuses on shipment assignment rather than route construction. It does not model detailed vehicle routing or sequencing, return trips to the origin, or time-synchronized flows between first-echelon (plant to DC) and second-echelon (DC to customer) shipments.

Instead, the model assumes decoupled inbound and outbound flows at the DCs, enabled by sufficient inventory buffering to avoid temporal coordination constraints. This approach aligns with P&G's operational reality, where transportation is fully outsourced to their 3PL providers. Vehicle routing, reuse, and return logistics are managed externally and not by P&G itself. Thus, the core decision for P&G is whether a given customer shipment should be fulfilled directly from the plant or routed through the DC, with the goal of minimizing cost while meeting demand and capacity constraints.

Transport costs are calculated per shipment leg, and both legs of a two-echelon shipment are charged separately. As such, even though each two-echelon delivery involves two vehicle movements, the model does not explicitly track the number of vehicles used. This simplification is justified by the absence of fixed vehicle start costs and the fact that transportation is fully outsourced; P&G incurs transport charges per trip rather than per vehicle. Consequently, the model evaluates the share of shipments executed via DPS versus 2E, but this does not reflect the actual number of vehicle movements, which would be higher for 2E due to the two legs.

Two distribution methods are modeled:

- **Direct Plant Shipments (DPS):** Products are shipped directly from the plant to the customer;

- **Two-Echelon Distribution (2E):** Products are routed via an intermediate DC before reaching the customer.

The model captures these flows using binary decision variables that indicate whether a shipment is assigned to a specific leg, and continuous variables that represent the volume of goods shipped on each leg. Vehicle capacity constraints are respected, and each shipment is modeled as a full truckload move or fixed-volume allocation. However, vehicle availability is not a limiting factor and is therefore not explicitly modeled. Return trip costs and empty ride penalties are embedded in the transport rates provided by the 3PL and used as input to the model.

This modeling approach reflects a *flow-based 2E-VRP variant*, suitable for evaluating hybrid distribution strategies. It provides strategic insights into how direct and two-echelon shipments can be optimally combined, without the computational complexity of route-level optimization. The formulation retains compatibility with the broader 2E-VRP literature, particularly drawing on the MILP-based hybrid models and cost structure decompositions proposed by Zhou et al. (2024) and Sluijk et al. (2023), while adapting the framework to the decision-making realities faced by P&G. As detailed in subsection 3.4.1, the literature offers various 2E-VRP extensions that address practical logistics challenges. Table 3.5 summarizes the key features of these studies.

5.1.1. Assumptions

The model development is grounded in a structured approach that begins with identifying the requirements from the current state (subsection 4.5.3). These requirements present fundamental needs that can be partly translated to modeling assumptions. The following assumptions provide fundamental details to define the model's constraints and parameters, ensuring that the model accurately reflects the real-world logistics setup of P&G and operational conditions, with simplified model assumptions to be able to perform the analysis.

Network Structure

- The network consists of two plant locations (origins), one DC (satellite), and multiple customer delivery locations (destinations).
- Connections exist from each plant to both the DC and the customers, and the DC is also connected to all customers. There are no connections between plants or between customers themselves.
- Three types of shipment legs are modeled:
  - First echelon (FE): from plant to DC
  - Second echelon (SE): from DC to customer
  - Direct shipment (DPS): from plant to customer
- DPS flows are modeled as part of the FE vehicle set.

Transportation and Vehicle Use

- Two homogeneous vehicle sets are modeled: V for FE shipments (plant origin), W for SE shipments (DC origin).
- Each vehicle has a capacity of , the demand unit in the model.
- The minimum volume required to trigger an FTL shipment is , based on internal P&G policy.
- DPS is only allowed for orders above

  .
- Orders below cannot be split, combined, or shipped to multiple destinations (no multi-drop).
- All transport is executed by a 3PL. Vehicle availability is assumed sufficient. Return routing and repositioning are irrelevant and not modeled.
- Transport costs are fixed per FTL and are determined by origin-destination pairs, regardless of distance.
- While 2E shipments involve two vehicle legs, vehicle counts are not explicitly tracked as transport is charged per shipment leg rather than per vehicle.

Vehicle Fill Rate and Capacity

- For the FE of a 2E distribution, Inter-site Freight (ISF) shipments, the shuttles between the plants and DC, are always stacked to FTL, also mixed with other SKUs to be shuttled. Multiple shuttles are driven daily between the two locations with a fixed cost per vehicle.
- For the SE of a 2E distribution, shipments with Cat X&Y SKUs are always consolidated with other SKUs to the customer, to fulfill the FTL threshold. Transport cost in the objective function are measured per unit to evaluate the cost for specifically Cat X&Y.

Region-Based Transport Costs

- Outbound transport costs (Customer Freight: CF) from plants or DCs to customers are region-based, as is practice at P&G.
- Each customer belongs to a geographic region with a fixed FTL transport rate (see subsection 5.2.2).
- ISF costs are fixed per FTL between Plant 1 and DC 1, and Plant 2 and DC 1.
- Transport cost rates for CF are fixed and given in Table B.2 in Appendix B.

Customer Demand and Service Constraints

- Each customer is served by only one source per shipment (plant or DC).
- Forecasted demand is flexibly allocable to either a plant or the DC, depending on feasibility of DPS and the stock allocation choices made.
- Orders below 1 pallet are not eligible for DPS. Therefore, the analysis uses FP as unit to only include full-pallet orders.
- Unloading costs are uniform across all transport methods and are for the cost of the customer, so not included in the optimization.

Storage and Inventory Allocation

- Safety stock levels are assumed to adjust dynamically with distribution flow shifts.
- Plant storage capacity is assumed scalable to accommodate DPS volume.

5.1.2. Mathematical Notation

Indices and Sets

tabularx

$p \in \mathcal{P}$	Set of origin nodes (Plant 1 and Plant 2)
$d \in \mathcal{D}$	Set of satellite nodes (DC 1)
$c \in \mathcal{C}$	Set of destination nodes (customers)
$N = \mathcal{P} \cup \mathcal{D} \cup \mathcal{C}$	Set of all nodes in the network
$v \in \mathcal{V}$	Set of vehicles assigned to first echelon (plant origin)
$w \in \mathcal{W}$	Set of vehicles assigned to second echelon (DC origin)

Parameters

The parameters represent the known inputs for the model, which can be extracted directly from the data or assumed based on other sources if it is lacking in the dataset.

D_{pc}	Weekly demand from origin p to customer c (in FP)
Q	Vehicle capacity (in FP)
$dist_{ij}$	Travel distance between node i and node j, where $(i, j) \in N \times N$
c_{pd}^{ISF}	Fixed ISF transportation cost from plant $p \in \mathcal{P}$ to DC $d \in \mathcal{D}$ (in €/FTL)
c_{dc}^{CF}	Fixed CF transportation cost from DC $d \in \mathcal{D}$ to customer $c \in \mathcal{C}$ (in €/FTL)
c_{pc}^{CF}	Fixed CF transportation cost from plant $p \in \mathcal{P}$ to customer $c \in \mathcal{C}$ (in €/FTL)
c_i^{load}	Average loading cost at origin node $i \in \mathcal{P} \cup \mathcal{D}$ (in €/FP)
c_j^{unload}	Average unloading cost at destination node $j \in \mathcal{D} \cup \mathcal{C}$ (in €/FP)
c_i^{store}	Average storage cost at origin node $i \in \mathcal{P} \cup \mathcal{D}$ (in €/FP)

All cost parameters are assumed to be non-negative:

$$c_{pd} \geq 0, \quad c_{dc} \geq 0, \quad c_{pc} \geq 0, \quad c_i^{\mathrm{load}} \geq 0, \quad c_j^{\mathrm{unload}} \geq 0,$$

Decision Variables

The decision variables in this model include both binary routing decisions and continuous flow quantities. Binary variables determine whether a specific transport leg is used by a vehicle in the network, while flow variables represent the corresponding transported volumes.

Binary routing variables:

$x_{pdv} \in \{0,1\}$	Binary variable indicating whether vehicle $v \in \mathcal{V}$ is assigned to transport SKUs from plant $p \in \mathcal{P}$ to DC $d \in \mathcal{D}$ (first leg of the two-echelon strategy).
$y_{dcw} \in \{0,1\}$	Binary variable indicating whether vehicle $w \in \mathcal{W}$ is assigned to transport SKUs from DC $d \in \mathcal{D}$ to customer $c \in \mathcal{C}$ (second leg of the two-echelon strategy).
$z_{pcv} \in \{0,1\}$	Binary variable indicating whether vehicle $v \in \mathcal{V}$ is assigned to a direct shipment from plant $p \in \mathcal{P}$ to customer $c \in \mathcal{C}$ (DPS).

Continuous flow variables:

$q_{pcv} \geq 0$	Quantity of FP transported directly from plant p to customer c using vehicle $v \in \mathcal{V}$ (DPS).
$\theta_{pdv} \geq 0$	Quantity of FP transported from plant p to DC d using vehicle $v \in \mathcal{V}$ (first echelon of the two-echelon route).
$\theta_{dcw} \geq 0$	Quantity of FP transported from DC d to customer c using vehicle $w \in \mathcal{W}$ (second echelon of the two-echelon route).

The structure of the network and binary routing arcs is visualized below:

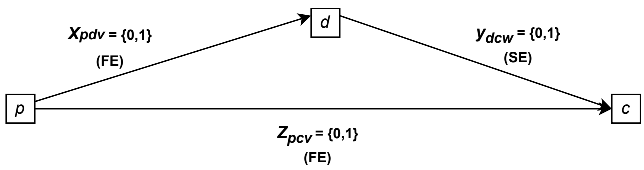

Figure 5.1: Leg Structure of the Modeled Two-echelon Hybrid Network

5.1.3. Objective Function: Logistics Cost Minimization

To identify the conditions under which DPS outperform 2E distributions, the model minimizes the total logistics cost required to satisfy customer demand. These costs include fixed routing costs (per shipment) and variable loading/unloading costs (per FP), enabling evaluation of hybrid distribution configurations.

Inspired by Zhou et al. (2024) and Sluijk et al. (2023), the objective function structure builds on their two-echelon distribution models, but replaces arc-based flow routing with only binary decision variables that indicate whether a shipment is carried out along a given node pair. This formulation allows for flexible modeling of hybrid distribution strategies while preserving a clear cost breakdown across routing and handling components.

$$
\begin{aligned}
\min \bar{Z}_{\text{cost}} = & \sum_{p\in\mathcal{P}} \sum_{d\in\mathcal{D}} \sum_{v\in\mathcal{V}} \Big[\underbrace{\theta_{pdv} \cdot c_p^{\text{load}}}_{\text{loading cost at plant}} + \underbrace{x_{pdv} \cdot \left(\frac{\theta_{pdv}}{Q} \right) \cdot c_{pd}^{\text{ISF}}}_{\text{transport cost plant}\to\text{DC}} + \underbrace{\theta_{pdv} \cdot c_d^{\text{unload}}}_{\text{unloading cost at DC}} \Big] \\
& + \sum_{d\in\mathcal{D}} \sum_{c\in\mathcal{C}} \sum_{w\in\mathcal{W}} \Big[\underbrace{\theta_{dcw} \cdot c_d^{\text{load}}}_{\text{loading cost at DC}} + \underbrace{y_{dcw} \cdot \left(\frac{\theta_{dcw}}{Q} \right) \cdot c_{dc}^{\text{CF}}}_{\text{transport cost DC}\to\text{customer}} \Big] \\
& + \sum_{p\in\mathcal{P}} \sum_{c\in\mathcal{C}} \sum_{v\in\mathcal{V}} \Big[\underbrace{q_{pcv} \cdot c_p^{\text{load}}}_{\text{loading cost at plant}} + z_{pcv} \cdot \underbrace{c_{pc}^{\text{CF}}}_{\text{transport cost plant}\to\text{customer}} \Big]
\end{aligned} \tag{5.1}
$$

The objective function minimizes the total logistics cost associated with two alternative routing strategies:

- **Two-echelon distribution (via DC):**
  - $\theta_{pdv} \cdot c_p^{\text{load}}$ — Variable loading cost per FP at the plant.

  - $\left(\frac{\theta_{pdv}}{Q}\right) \cdot c_{pd}^{\mathrm{ISF}}$ — Transport cost from plant p to DC d, proportional to FTL.
  - $\theta_{pdv} \cdot c_{d}^{\mathrm{unload}}$ — Variable unloading cost per FP at the DC.
  - $\theta_{dcw} \cdot c_{d}^{\mathrm{load}}$ — Variable loading cost per FP at the DC for outbound delivery.
  - $\left(\frac{\theta_{dcw}}{Q}\right) \cdot c_{dc}^{\mathrm{CF}}$ — Transport cost from DC d to customer c, proportional to FTL.
- **Direct Plant Shipment (DPS):**
  - $q_{pcv} \cdot c_{p}^{\mathrm{load}}$ — Variable loading cost per FP at the plant.
  - $z_{pcv} \cdot c_{pc}^{\mathrm{CF}}$ — Transport cost from plant p to customer c per activated direct shipment.

Note that the transport costs for the two-echelon terms of the equation are divided by the vehicle capacity Q, since the corresponding flows θ_{pdv} and θ_{dcw} are assumed to be able to occupy only part of a FTL; the remaining capacity will always be filled by other product flows, allowing transport costs to be shared. In contrast, for direct plant shipments, the transport cost c_{pc}^{CF} is not divided by Q because the flow q_{pcv} is assumed to be the only load on the vehicle and thus bears the full transport cost of the trip.

The binary routing variables $x_{pdv},\ y_{dcw},\ z_{pcv}$ indicate whether a specific vehicle is assigned to a routing leg in the network. The continuous flow variables $\theta_{pdv},\ \theta_{dcw},\ q_{pcv}$ represent the volume of goods transported by the assigned vehicles and determine the variable (FP-unit based) cost contributions. Although the first-echelon flow θ_{pdv} does not directly reference customer indices, its volume is indirectly driven by downstream customer demand fulfilled via the DCs.

Table 5.1: Overview of cost components in the objective function

Cost component	Explanation
c_{pd}^{ISF}	Fixed ISF cost for FTL shipments from plant p to DC d (Table B.1)
c_{dc}^{CF}	Fixed CF cost for FTL shipments from DC d to customer c (Table B.2)
c_{pc}^{CF}	Fixed CF cost for FTL direct shipments from plant p to customer c (DPS) (Table B.2)
c_{p}^{load}	Loading cost per pallet at plant p.
c_{d}^{load}	Loading cost per pallet at DC d.
c_{d}^{unload}	Unloading cost per pallet at DC d.

5.1.4. Constraints

The following constraints ensure the feasibility and integrity of the routing and flow decisions within the hybrid distribution network model.

Demand satisfaction per origin plant

$$\sum_{v \in \mathcal{V}} q_{pcv} + \sum_{d \in \mathcal{D}} \sum_{w \in \mathcal{W}} \theta_{dcw} = D_{pc} \quad \forall p \in \mathcal{P},\ \forall c \in \mathcal{C} \tag{5.2}$$

The full demand originating from each plant for each customer must be satisfied, either directly via DPS or indirectly via a two-echelon flow through the DC.

Flow activation and vehicle capacity constraints

$$q_{pcv} \leq Q \cdot z_{pcv} \qquad \forall p \in \mathcal{P},\ \forall c \in \mathcal{C},\ \forall v \in \mathcal{V} \tag{5.3}$$

$$\theta_{pdv} \leq Q \cdot x_{pdv} \qquad \forall p \in \mathcal{P},\ \forall d \in \mathcal{D},\ \forall v \in \mathcal{V} \tag{5.4}$$

$$\theta_{dcw} \leq Q \cdot y_{dcw} \qquad \forall d \in \mathcal{D},\ \forall c \in \mathcal{C},\ \forall w \in \mathcal{W} \tag{5.5}$$

Flow variables can only be positive if the corresponding routing arc is activated. Each flow is also bounded by the vehicle capacity Q.

FTL eligibility and volume range for DPS

$$\begin{aligned} q_{pcv} &\leq Q \cdot z_{pcv} \\ q_{pcv} &\geq \tau \cdot z_{pcv} \end{aligned} \quad \forall p \in \mathcal{P},\ \forall c \in \mathcal{C},\ \forall v \in \mathcal{V} \tag{5.6}$$

Where represents the minimum threshold for direct shipment eligibility (DPS), and is the maximum vehicle capacity.

DC flow balance

$$\sum_{p\in\mathcal{P}}\sum_{v\in\mathcal{V}}\theta_{pdv} = \sum_{c\in\mathcal{C}}\sum_{w\in\mathcal{W}}\theta_{dcw} \quad \forall d\in\mathcal{D} \tag{5.7}$$

The amount of goods arriving at a distribution center (DC) must equal the amount of goods leaving that DC, aggregated over all vehicles.

Vehicle task exclusivity per echelon

$$\sum_{p\in\mathcal{P}}\sum_{c\in\mathcal{C}} z_{pcv} + \sum_{p\in\mathcal{P}}\sum_{d\in\mathcal{D}} x_{pdv} \leq 1 \quad \forall v\in\mathcal{V} \tag{5.8}$$

Each first-echelon vehicle can be assigned to only one task: either a direct plant-to-customer shipment (DPS) or a plant-to-DC leg of a two-echelon route, but not both simultaneously.

$$\sum_{d\in\mathcal{D}}\sum_{c\in\mathcal{C}} y_{dcw} \leq 1 \quad \forall w\in\mathcal{W} \tag{5.9}$$

Each second-echelon vehicle can serve at most one DC-to-customer connection to maintain simple routing and capacity logic per vehicle.

DPS assignment: one customer per vehicle

$$\sum_{p\in\mathcal{P}}\sum_{c\in\mathcal{C}} z_{pcv} \leq 1 \quad \forall v\in\mathcal{V} \tag{5.10}$$

Each vehicle assigned to direct plant-to-customer shipments (DPS) may serve at most one customer. This constraint reinforces the full-truckload nature of DPS flows and simplifies routing logic. While this condition is logically implied by the vehicle task exclusivity constraint (5.8), it is included explicitly to improve model transparency and reduce computational complexity by narrowing the feasible space.

Variable domains

$$x_{pdv},\ z_{pcv} \in \{0,1\} \qquad \forall p\in\mathcal{P},\ \forall d\in\mathcal{D},\ \forall c\in\mathcal{C},\ \forall v\in\mathcal{V} \tag{5.11}$$

$$y_{dcw} \in \{0,1\} \qquad \forall d\in\mathcal{D},\ \forall c\in\mathcal{C},\ \forall w\in\mathcal{W} \tag{5.12}$$

$$q_{pcv},\ \theta_{pdv},\ \theta_{dcw} \geq 0 \qquad \forall p\in\mathcal{P},\ \forall d\in\mathcal{D},\ \forall c\in\mathcal{C},\ \forall v\in\mathcal{V},\ \forall w\in\mathcal{W} \tag{5.13}$$

Binary routing decisions and non-negative continuous flow variables, where non-negativity is enforced directly in the variable definitions by setting the lower bound to zero.

5.1.5. Incorporating Emissions in the Objective Function

In addition to minimizing the logistics costs with the previous formulation, it is increasingly important to consider the environmental impact of transportation in distribution network design. This research incorporates an environmental objective that quantifies transport-related emissions based on the total distance traveled by vehicles in the network:

$$\bar{Z}_{\text{emissions}} = \gamma\cdot\left(\sum_{p\in\mathcal{P}}\sum_{d\in\mathcal{D}}\sum_{v\in\mathcal{V}}\frac{\theta_{pdv}}{Q}\cdot x_{pdv}\cdot dist_{pd} + \sum_{d\in\mathcal{D}}\sum_{c\in\mathcal{C}}\sum_{w\in\mathcal{W}}\frac{\theta_{dcw}}{Q}\cdot y_{dcw}\cdot dist_{dc} + \sum_{p\in\mathcal{P}}\sum_{c\in\mathcal{C}}\sum_{v\in\mathcal{V}} z_{pcv}\cdot dist_{pc}\right) \tag{5.14}$$

Here, $dist_{ij}$ denotes the travel distance between nodes i and j, and γ = 0.90 kg CO_2/km per vehicle is based on empirical truck emissions estimates (Rodríguez et al., 2020). This hybrid emissions objective combines flow-based and vehicle-based components. For two-echelon (2E) flows, emissions are proportional to the transported volume θ and distance traveled, normalized by vehicle capacity Q, and incurred only if a vehicle is assigned to the route (via x_{pdv} or y_{dcw}). This reflects partial vehicle usage and shared emissions. In contrast, DPS are modeled as dedicated trips; if a direct shipment is assigned ($z_{pcv} = 1$), the entire distance $dist_{pc}$ is counted as emissions, regardless of the transported volume.

To incorporate emissions into the optimization process, the model uses a scalarized objective function that combines logistics cost and emissions, both expressed in monetary terms:

$$\bar{Z}_{\text{total}} = \bar{Z}_{\text{cost}} + \lambda\cdot\bar{Z}_{\text{emissions}} \tag{5.15}$$

Here, λ (€/kg CO_2) represents a carbon pricing factor that can reflect external carbon markets, such as the EU ETS price of €0.065/kg (International Carbon Action Partnership, 2024), or can serve as an internal weight to explore the sensitivity of the distribution choices to emissions. This scalarization method follows the approach proposed by Pereira et al. (2022).

This formulation is used consistently across all experiments, including the main cost optimization scenarios. Changing the value of λ adjusts the relative importance of emissions in the objective function but does not alter the model structure. It remains a single-objective optimization problem, solved for a given value of λ to obtain cost-optimal solutions under different environmental priorities.

To investigate the trade-offs between logistics cost and environmental impact more systematically, two complementary sensitivity analyses were conducted. The first uses the scalarization method to explore routing decisions under varying carbon prices (λ). The second applies an ε-constraint approach, in which emissions are minimized directly as the main optimization function under constrained logistics costs to assess behavior under a sustainability-first objective. The setup of both analyses is detailed in subsection 6.5.3 and subsection 6.5.4, with outcomes discussed in section 7.8.

5.1.6. Service Level Performance

To evaluate customer service performance under different distribution strategies, this research applies a post-solution estimation method based on historical service benchmarks. Rather than modeling real-time product availability, the approach estimates the overall service level as a weighted average of past performance from each shipping origin (plants and DC). This allows for a realistic and interpretable assessment of how changes in shipment allocations, so increasing DPS, influence expected service levels. The method supports both operational evaluation and commercial interpretation, aligning with internal KPIs at P&G.

Post-Solution Estimation Based on Historical Service Levels

Service level rates from Plant 1 and Plant 2 are estimated based on the incidental shipments that have occurred to certain customers in the historical data of service performance. The total shipped volume where this rate is based on is very low, so there might be bias. Nevertheless, as product availability tends to be higher at plants, where production occurs, shipments would typically achieve higher fulfillment rates compared to DCs, which can have stockouts and allocation constraints.

- $SL_{\text{Plant 2}}$ = average historical service level for shipments from Plant 2;
- $SL_{\text{Plant 1}}$ = average historical service level for shipments from Plant 1;
- $SL_{\text{DC 1}}$ = average historical service level for shipments from the DC 1;
- $\text{Vol}_{\text{Plant 2}} = \sum_{c\in\mathcal{C}} \sum_{v\in V} q_{\text{Plant 2},\, cv}$ = total volume shipped directly from Plant 2 to customers;
- $\text{Vol}_{\text{Plant 1}} = \sum_{c\in\mathcal{C}} \sum_{v\in V} q_{\text{Plant 1},\, cv}$ = total volume shipped directly from Plant 1 to customers;
- $\text{Vol}_{\text{DC 1}} = \sum_{c\in\mathcal{C}} \sum_{w\in W} \theta_{\text{DC 1},\, cw}$ = total volume shipped from DC 1 to customers.

Figure 5.2 shows the service level rates of the past months, for both Category X and Category Y, for shipments from the DC and from incidental plant shipments. Plant rates are generally higher, due to the bias of the low number of historical shipments that occurred.

[Figure anonymized for confidentiality reasons]

Figure 5.2: Monthly Service Level Rates per Category and Origin

Considering these historical patterns, averages are taken in the post-solution estimation, given in Table 5.2. For the DC 1, the average of Cat X&Y together is taken for the evaluation.

Table 5.2: Average SL % per Origin Shipment Location

$\text{SL}_{\text{DC 1}}$	X%
$\text{SL}_{\text{Plant 2}}$	X%
$\text{SL}_{\text{Plant 1}}$	X%

The estimated overall service level to BNL customers is calculated as a weighted average across the three origin locations in Equation 5.16. The equation will support the comparison of scenarios with and without DPS, to see if more DPS would enhance the service level KPI.

$$\text{SL}_{\text{estimated}} = \frac{\text{SL}_{\text{Plant 2}} \cdot \text{Vol}_{\text{Plant 2}} + \text{SL}_{\text{Plant 1}} \cdot \text{Vol}_{\text{Plant 1}} + \text{SL}_{\text{DC 1}} \cdot \text{Vol}_{\text{DC 1}}}{\text{Vol}_{\text{Plant 2}} + \text{Vol}_{\text{Plant 1}} + \text{Vol}_{\text{DC 1}}} \tag{5.16}$$

Impact of Service Level Changes on Net Outside Sales

In addition to influencing operational performance, service levels also have a direct financial impact.

This relationship highlights the strategic importance of maintaining high service levels, particularly from origin points with greater product availability as the plants, as even minor improvements in fulfillment performance can lead to measurable gains in sales. With this, post-solution evaluation of service levels not only informs supply reliability but also provides insight into potential commercial outcomes.

5.2. Data Preparation

5.2.1. Data Overview

To support the model, historical shipment data was extracted from P&G's

Model parameters such as customer demand (D_{pc}), vehicle capacity (Q), and transport distances ($dist_{ij}$) were derived directly from this dataset or assumed based on internal knowledge. Fixed transportation costs, including ISF costs (c_{pd}^{ISF}) for plant-to-DC movements, CF costs (c_{dc}^{CF}, c_{pc}^{CF}) for DC-to-customer and direct plant-to-customer deliveries, were sourced from region-based transport rates provided by the Transport Operations (T-OPS) team or estimated where needed. Loading and unloading costs (c_i^{load}, c_j^{unload}) were extracted from haulier contract data. Finally, storage costs were taken from internal DC and plant operations records.

5.2.2. Region-based Transport Costs

Table B.2 in Appendix B presents the transport costs in €/FTL for all destination regions. These costs represent Customer Freight (CF) transport rates, which are fixed per route between origin and destination. Each customer location is assigned to a specific region, and the transport cost for that region is used accordingly.

The green-highlighted values in the table reflect known CF rates, obtained from historical data provided by P&G's T-OPS. All rates for destinations with origin DC 1 are known. In contrast, several values for direct shipments from Plant 1 and Plant 2 were unavailable, primarily because such direct shipments have not historically occurred so no rates from contracted carriers are yet recorded. To address these gaps, estimates were made for the missing rates. The orange-highlighted cells indicate these estimated values. The estimation involved calculating the distance from the plant location to the centroid of the corresponding region and then applying a rate per kilometer derived from the known historical rates of occurred direct shipments. This approach yields a reasonable approximation of the fixed CF transport cost per shipping route.

5.2.3. Distance matrix

Distances between origin nodes (plants and DC) and the selected customer locations were obtained with Python using the *openrouteservice* API. The calculation was based on the precise geographical coordinates (latitude and longitude) of each node. The distances are summarized in Table B.3 in Appendix B, where 14 DPS-eligible customers, coming from the analysis in subsection 5.2.4, are listed alongside the region they belong to and their corresponding distance to each origin node. The resulting distance matrix was stored in a Pandas DataFrame, with rows representing origin nodes and columns representing customer destinations. This matrix is provided in Table B.3 in Appendix B.

Although these distances are not used in the transport cost calculations, because transport costs are region-based and fixed, they are essential for the environmental performance evaluation. Specifically, the model multiplies distance traveled by emission factor γ to estimate the CO_2 emissions per shipment route, as detailed in subsection 5.1.5. This allows for a sensitivity analysis of the environmental impact of DPS versus two-echelon distribution.

5.2.4. Demand Data: Narrowing the Customer Set

The full selected demand dataset contains historical weekly shipment volumes for customers, segmented by product category and origin plant. Each row represents a unique combination of customer, category (Category Y or Category X), origin production location and initial storage location: either one of the main production plants (Plant 1 or Plant 2) or the DC 1. The dataset spans 41 calendar weeks: from week 27 of 2024 (starting July 1st) to week 14 of 2025 (ending April 1st). This selection gives a sufficient amount of data and span over multiple seasons to trace variability. Demand is expressed in Floor Positions (FP), which approximates the required vehicle space per shipment.

To reduce the dataset, a first analysis on the demand patterns is performed. To identify which customers are viable candidates for DPS, a statistical screening was conducted. Key metrics were calculated per customer-category combination: mean weekly demand, standard deviation, coefficient of variation (CV), and the number of weeks in which demand exceeded a specific operational threshold. The threshold was set at , giving the minimum volume at which a vehicle delivery is considered FTL-eligible and would be shipped. However, to be considered a structurally suitable DPS candidate, customers were required to demonstrate an average weekly demand of at least , which is the theoretical FTL threshold.

Considering these statistics, 14 DPS-eligible customers were identified whose average demand exceeded for Category X and Category Y. Their demand profiles, including weekly variability and FTL consistency, are detailed in Table C.1 in Appendix C. By narrowing the analysis to these higher-volume customers, the research maintains a practical focus on those with a realistic potential for DPS integration. Additional details on their demand reliability are provided in Table C.3 in Appendix C, which reports the number and share of weeks in which their demand exceeded the threshold. These findings show that most DPS-eligible customers not only surpass the average volume requirement but also display consistent week-by-week FTL-eligible volumes, strengthening their potential for DPS.

The demand behavior across these 14 DPS-eligible customers altogether, is plotted per product category in Figure 5.3 and Figure 5.4. It shows the weekly average demand (in FP) per customer for Category Y and Category X products, respectively. The dotted blue line indicates the average demand across all weeks, where the dotted red line marks the FTL threshold at . These graphs highlight weekly variation and underscore the structural differences in demand levels and volatility between Cat X&Y. It could already be concluded that Category Y appeals more potential for DPS because of the high average weekly demand.

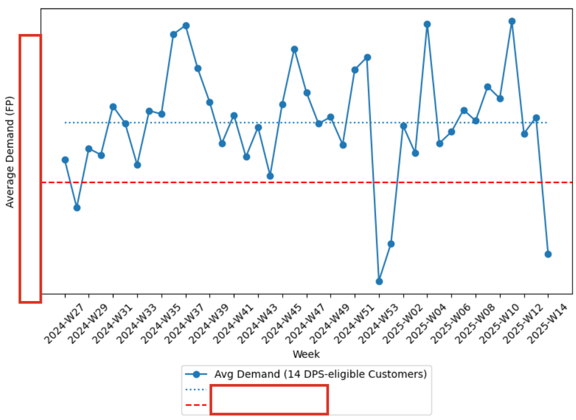

Figure 5.3: Weekly Average Demand per DPS-eligible Customer - Category Y

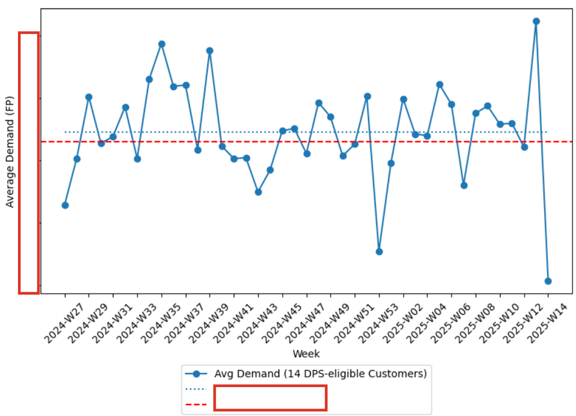

Figure 5.4: Weekly Average Demand per DPS-eligible Customer - Category X

Assessing Other Category Demand to Finalize Customer Set

After identifying the initial selection of 14 DPS-eligible customers with high demand for Cat X&Y products, a further scoping step was conducted to ensure that implementing DPS would not compromise the integrity of existing logistics. In the current distribution system, all of these customers also receive products from other categories, which are shipped via the central DC 1. While DPS is considered for their Cat X&Y flows from Plant 1 and Plant 2, the residual demand for other categories must remain sufficiently high to independently support weekly full truckloads from the DC. This ensures that introducing DPS for Cat X&Y does not erode the volume base needed to sustain efficient DC-based deliveries.

To assess this, demand for non-Cat X&Y was analyzed over the same 41-week period. Figure 5.5 highlights the top five customers with the highest overall demand, showing the number of weeks in which demand exceeded the FTL threshold of . This threshold ensures that at least are filled, accounting for a typical mix of pallet types. The final five selected customers as DPS candidates are: Customer 1, Customer 2, Customer 3, Customer 4 and Customer 5 (geographical locations given in Figure 5.6).

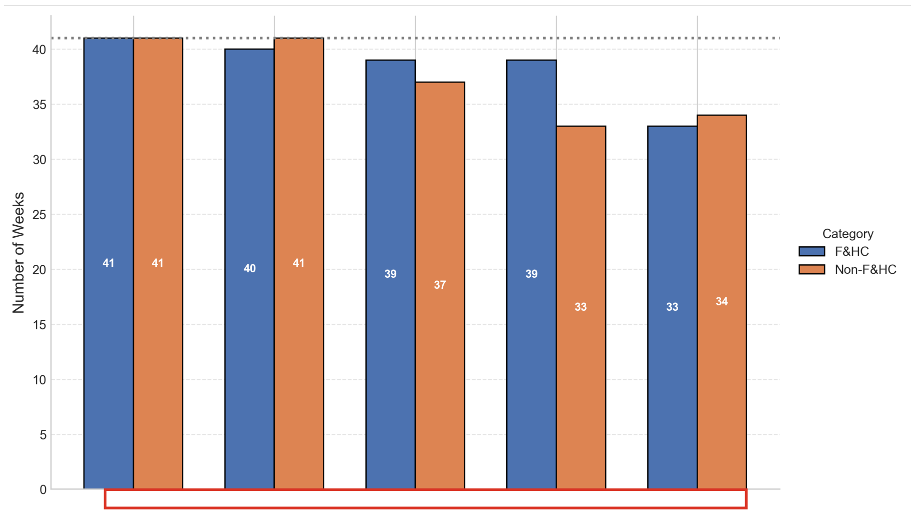

Figure 5.5: Number of Weeks (out of 41) with Demand Above Thresholds for Cat X&Y and Other Categories

Descriptive statistics for these five customers are shown in Table 5.3. The coefficient of variation (CV) is included to indicate the relative demand stability. A low CV suggests more predictable weekly volumes, which supports the feasibility of decoupling Cat X&Y flows (for DPS) from DC-based shipments.

Table 5.3: Demand Statistics for Top 5 Customers

Customer	Weekly Avg Cat X&Y (PAL)	Std Cat X&Y	CV Cat X&Y
Customer 1			
Customer 2			
Customer 3			
Customer 4			
Customer 5			

Customer	Weekly Avg *Non*-Cat X&Y (PAL)	Std *Non*-Cat X&Y	CV *Non*-Cat X&Y
Customer 1			
Customer 2			
Customer 3			
Customer 4			
Customer 5			

[Figure anonymized for confidentiality reasons]

Figure 5.6: Map of Top 5 Customers, Plants and DC

Autocorrelation Test

To evaluate the temporal consistency of weekly demand among the five DPS candidates identified in the previous section, the autocorrelation function (ACF) was computed. ACF measures the correlation between current and past demand values over different time lags. For each customer-category combination, the average ACF over lags 1 to 3 was calculated. This reflects short-term memory in the data, so whether demand in one week is dependent on demand in the preceding weeks.

The results, presented in Table C.4 in Appendix C, show that most average ACF values range between -0.2 and 0.1, suggesting weak short-term autocorrelation. This implies that week-to-week demand is not strongly predictable for and may be subject to operational variability or customer-specific ordering behavior. While this indicates that short-term volatility is present, it does not rule out the value of long-term planning of DPS based on overall averages, especially in combination with volume stability.

Stationarity Analysis: Augmented Dickey-Fuller Test

To complement the ACF analysis, the Augmented Dickey-Fuller (ADF) test was performed to assess the stationarity of the weekly demand series for the same five DPS candidates. Stationarity indicates that key statistical properties (such as mean and variance) remain stable over time, which is essential for forecasting and strategic planning.

As shown in Table C.5 in Appendix C, 7 out of 10 customer-category combinations exhibit stationarity, indicated by p-values below the 0.05 significance threshold. These include both categories for customers Customer 1, Customer 4, and Customer 5, as well as the Cat X&Y for Customer 3. This means that the null hypothesis of non-stationarity can be rejected for these cases.

In contrast, both categories for Customer 2 and the 'Other' category for Customer 3 are found to be non-stationary, indicating higher volatility or underlying trends in those demand series. These cases require more caution in planning and may benefit from further analysis or smoothing techniques. Overall, the predominance of stationary series supports the assumption that historical weekly demand behavior is a reliable basis for DPS planning. The presence of stable demand patterns across most customers reinforces the feasibility of using long-term averages for strategic decision-making.

Conclusion: Statistical Demand Properties of 5 Key DPS Candidates

After reducing the original set of customers to 14 DPS-eligible candidates, a further refinement step resulted in a final selection of 5 customers included as DPS candidates in the analysis. For these 5 customers, in summary, the analysis of short-term autocorrelation and long-term stationarity shows that weekly demand is stable over time, though not highly predictable week to week. These findings support the use of long-term demand averages for DPS planning, while maintaining sufficient reliability in the traditional DC-based flows.

5.3. Model Implementation

The mathematical model was implemented in Python, using FICO Xpress 9.5 as the optimization solver, which is the internal used optimization tool within P&G. The purpose of this implementation is to translate the adapted 2E-VRP with hybrid distribution into a computational framework capable of solving realistic planning scenarios based on historical demand and cost data. The implementation includes routing decisions, shipment quantities, and flow allocations across the direct and two-echelon distribution paths. The Python Model in code format is given in Appendix F.

FICO Xpress 9.5 is a commercial optimization engine specialized in solving large-scale mixed-integer linear programming (MILP) problems. It applies a combination of branch-and-bound, cutting planes, and heuristic algorithms to efficiently search the solution space. The branch-and-bound method systematically explores partial solutions, pruning suboptimal branches based on bounds, while cutting planes are used to eliminate fractional solutions and accelerate convergence. Heuristics help identify good feasible solutions early in the process, improving solver performance. These optimization techniques are handled internally by the solver, allowing the modeler to focus on formulating the problem structure and constraints. The resulting framework allows for flexible scenario testing and analysis of trade-offs between direct shipments and two-echelon deliveries.

5.4. Model Verification

A verification on the model is performed to confirm that the implementation of the hybrid 2E-VRP model behaves as intended and corresponds to the mathematical formulation. Both structural consistency and behavioral response to changes in decision variables and parameters are assessed.

5.4.1. Model Consistency

The model formulation described in subsection 5.1.2 was implemented in Python and solved using FICO Xpress 9.5. Verification was performed on a simplified test instance derived from the real-world dataset, using a small subset of the data (two plants, one DC, and two customers: Customer 1 and Customer 3). This ensures that routing logic, parameter sensitivity, and constraint behavior could be inspected in a controlled environment. Verification started with ensuring alignment between the code and the mathematical model in terms of:

- **Set definitions:** Sets for plants ($\mathcal{P}$), DCs ($\mathcal{D}$), and customers ($\mathcal{C}$) were checked to ensure correct dimensionality and mapping to the historical demand dataset.
- **Variable logic:** All routing variables ($x_{pdv}, y_{dcw}, z_{pcv} \in \{0, 1\}$) and flow variables ($\theta_{pdv}, \theta_{dcw}, q_{pcv} \geq 0$) were confirmed to match their defined purpose in the formulation and produce feasible values under some simple test scenarios.
- **Objective function:** The main objective function for logistics costs was implemented and results of the separate cost components were printed to see if the total cost correctly aligns with the split in cost components. The additional functions for service level and emissions are also tested on consistency and components are split out to ensure the expected output is correctly aggregated.
- **Constraint implementation:** All constraints are cross-checked against their mathematical representation. Specific constraints are manually tested by inspecting solution outputs to ensure logical correctness (e.g., no dual flows to the same customer).

In addition, the demand, distance, and cost datasets were checked for consistency by inspecting printed values for selected customer and location combinations and verifying this with the data files, confirming that data inputs were interpreted and used correctly within the model.

5.4.2. Behavioral Testing with Parameters

To evaluate the robustness and interpretability of the model under extreme parameter settings, a series of behavioral tests of the parameters is conducted. In each test, a single parameter is varied while all other parameters are held constant. This approach allows for the identification of whether the model reacts logically to changes and continues to produce consistent, explainable outputs.

Before performing the parameter variations, the base performance of the model is first evaluated using default parameter settings. This base test scenario considers two customers (Customer 1 and Customer 3) with demand originating from both plants and includes only Cat X&Y products. The corresponding results are shown in Figure 5.4. To assess the model's behavioral robustness, a series of extreme parameter value tests were conducted, summarized in Figure 5.5. These tests examine the impact of varying transport costs, vehicle capacity, and emission factors on model performance.

Table 5.4: Parameter Testing: Base Output

Week Number	46
Customers included	Customer 1, Customer 3
Demand included	Cat X&Y
Total cost	
Total distance traveled	
Total emissions	
Direct shipments (DPS)	
Two-echelon trips (2E)	

Table 5.5: Behavioral Tests with Extreme Parameter Values

Test	Parameter	Test Values
T1	DPS transport cost c_{pc}^{CF}	
T2	SE transport cost c_{dc}^{CF}	
T3	Vehicle capacity Q	
T4	Emission factor γ	

T1: Variation in DPS Transport Cost

Table 5.6 presents the results when varying the DPS cost parameter. When the DPS cost is set to zero, the total logistics cost decreases; however, the number of DPS remains unchanged due to the enforcement of the Full Truckload (FTL) constraint. When this constraint is relaxed and the DPS cost remains zero, the model selects only DPS routes, including for smaller quantities, further decreasing the total cost. In contrast, setting the DPS cost to a prohibitively high value results in all demand being routed through the two-echelon network.

Table 5.6: T1: Parameter Testing: DPS cost c_{pc}^{CF}

DPS cost c_{pc}^{CF} = 0		**DPS cost c_{pc}^{CF} =**	
Total cost		Total cost	
Total distance traveled		Total distance traveled	
Total emissions		Total emissions	
Direct shipments (DPS)		Direct shipments (DPS)	
Two-echelon trips (2E)		Two-echelon trips (2E)	

DPS cost c_{pc}^{CF} = 0 (Relaxed FTL Enforcement constraint)	
Total cost	
Total distance traveled	
Total emissions	
Direct shipments (DPS)	
Two-echelon trips (2E)	

T2: Variation in Second Echelon Transport Cost

In Test T2, the cost of shipping from the DC to the customer (SE) is varied. With the FTL constraint enforced, changing this parameter has limited effect on the shipment strategy. However, once the FTL constraint is relaxed, an increase in SE cost leads to a higher number of direct shipments (DPS), as expected. This confirms that the model responds logically to cost differentials when not constrained by FTL requirements.

Table 5.7: T2: Parameter Testing: SE cost c_{dc}^{CF}

SE cost c_{dc}^{CF} = 0		**SE cost c_{dc}^{CF} = 10000**	
Total cost		Total cost	
Total distance traveled		Total distance traveled	
Total emissions		Total emissions	
Direct shipments (DPS)		Direct shipments (DPS)	
Two-echelon trips (2E)		Two-echelon trips (2E)	
SE cost c_{dc}^{CF} = 0 (Relaxed FTL Enforcement constraint)		**SE cost c_{dc}^{CF} = 10000** (Relaxed FTL Enforcement constraint)	
Total cost		Total cost	
Total distance traveled		Total distance traveled	
Total emissions		Total emissions	
Direct shipments (DPS)		Direct shipments (DPS)	
Two-echelon trips (2E)		Two-echelon trips (2E)	

T3: Variation in Vehicle Capacity

Test T3 examines the impact of vehicle capacity on routing outcomes. Reducing capacity results in significantly more 2E trips and higher total costs due to an increased number of small shipments. On the other hand, increasing vehicle capacity reduces both the number of trips and the total cost. These outcomes are in line with expectations, confirming that the model adjusts routing decisions based on available vehicle capacity.

Table 5.8: T3: Parameter Testing Vehicle Capacity Q

# DPS	
# 2E	
Total cost	

T4: Variation in Emission Factor

Finally, test T4 explores the impact of varying the emission factor γ. As shown in Table 5.9, modifying this parameter directly affects the total emission values, as expected. However, the routing decisions remain unchanged across all values tested. This suggests that the contribution of emissions to the total cost is relatively small, and thus does not drive changes in routing when cost minimization is the primary objective.

Table 5.9: T4: Parameter Testing Emission Factor γ

	$\gamma = 0.90$	$\gamma = 0.10$	$\gamma = 2.00$
# DPS			
# 2E			
Total emissions			

5.4.3. Behavioral Testing with Decision Variables

Another series of behavioral tests is performed by manually fixing key binary and continuous decision variables. These tests simulate extreme routing conditions and isolate specific routing mechanisms, allowing verification of model consistency, feasibility handling, and constraint integration. Each test is performed on the previously mentioned simplified instance. By selectively enabling or disabling shipment options and flow variables, the model's response to controlled scenarios is observed and interpreted in Table 5.10.

Table 5.10: Behavioral Tests on Binary Decision Logic

Test	Description	Result
T1	Force $z_{pc}^{v} = 1$ (DPS activated for all eligible customers); disable DC flows ($x_{pd}^{v} = 0$, $y_{dc}^{v} = 0$)	*Infeasible solution:* Model correctly identifies that not all demand can be routed via DPS due to eligibility, volume, or FTL constraints.
T2	Force $z_{pc}^{v} = 0$ (DPS disabled)	*Feasible solution:* All customer demand is routed via the two-echelon network, confirming that the model correctly switches to 2E routing when DPS is disabled.
T3	Force $y_{dc}^{v} = 0$ for all d, c	*Initially infeasible:* With FTL enforcement active, no feasible solution to route all volume via DPS was possible. After relaxing the FTL constraint, a feasible solution was found using only DPS shipments.
T4	Disable DPS flow: set $q_{pc}^{v} = 0$ for all p, c, v	*Feasible solution:* All flow is routed through the two-echelon network, confirming that the model correctly enforces 2E routing when DPS flow is disabled.
T5	Disable 2E flow: set θ_{pdv} and θ_{dcw} to 0 for all valid indices	*Initially infeasible:* With FTL enforcement active, DPS routing alone could not meet all demand. After relaxing the FTL constraint, a feasible solution was found using only DPS flows.
T6	Fix a single DPS flow: set $q_{pc}^{v} = 33$ for Plant 2 $\to$ Customer 1, disable all other DPS flows to Customer 1	*Feasible solution:* One DPS shipment from Plant 2 to Customer 1 was enforced as intended; the remaining demand was served via two-echelon routes, confirming correct integration of flow and routing logic.

5.4.4. Solver and Performance Checks

To ensure that the model formulation is correct and solvable, several verification and performance checks were performed using the FICO Xpress Solver (v9.5.0). These checks confirm the model's numerical stability, feasibility, and computational efficiency:

- The solver successfully found optimal solutions for small instances, based on a reduced test dataset containing two customers (Customer 1, Customer 3), one DC, and both plants. This subset was chosen to reflect representative routing scenarios while ensuring full transparency of variable assignments. Optimal solutions were found within 5 seconds, with all binary routing decisions correctly assigned.
- The model's integrity was confirmed by toggling constraints such as full truckload (FTL) enforcement and single-task eligibility. These toggles influenced routing decisions and solution structure, as expected.
- For the main test case with 2 customers and 1,000 variables, the presolved problem was significantly reduced to 550 variables and 554 constraints, indicating efficient preprocessing.
- The optimal solution in this test run was obtained with a final objective value of €6,480.25 and zero primal and integer infeasibility.
- The solver achieved this solution in under 0.05 seconds, with only two branch-and-bound nodes explored, demonstrating high solver performance and problem tractability.

These solver and performance checks support the model's suitability for application in more complex test cases and form the basis for the overall behavioral assessment presented in the following conclusion.

5.4.5. Conclusion

The behavioral tests confirm that the model logic is sound, robust, and responsive to variations in key input parameters. For cost-related parameters (T1 and T2), the model appropriately adjusts routing decisions in line with cost incentives, particularly when the full truckload (FTL) enforcement constraint is relaxed. Infeasibilities are correctly triggered when constraints conflict with routing possibilities over the DPS arc, but once relaxed, the model successfully reroutes demand via alternative paths. Changes in vehicle capacity (T3) affect routing complexity and cost efficiency as expected, while adjustments to the emission factor (T4) influence only environmental performance outputs, leaving routing decisions unchanged. These outcomes demonstrate not only the internal consistency of the model but also its transparency and controllability, as routing behavior can be accurately directed through fixed decision variables. Together, these findings confirm the correctness of the formulation and its suitability for analyzing more complex or larger-scale instances.

5.5. Model Data Validation

To ensure that the optimization model is built on realistic and credible inputs, a data validation was performed. This process confirms that the data structures, assumptions, and parameter values used in the hybrid distribution model correctly reflect the operational setup at P&G. Following the principles described in Sargent (2010) for model validation and data credibility, this section focuses on validating the alignment between the model's parameterization and actual logistics data.

The validation is based on data provided by P&G's T-OPS team, including transport cost structures, historical routing volumes, and known rates for both two-echelon and direct shipment flows. While customer-specific invoice data from ERP systems was not available, the reference values used here were sourced directly from T-OPS operational reports and planning documents. As such, they serve as authoritative benchmarks to test the model's data realism and consistency.

5.5.1. Validation of Two-Echelon Input Data (C0)

The standard two-echelon (2E) configuration currently in use at P&G was used as the basis for this validation. In configuration C0, direct shipment options are disabled and all customer demand is routed through the plant-DC-customer chain. The goal is to validate whether the model's inputs and data interpretation correctly reproduce the expected cost and flow behavior for known customers.

Table 5.11 summarizes the model results for two key customers (Customer 1 and Customer 3) under scenario S1 (Category X demand from Plant 1), and compares the modeled transport cost breakdown to the expected historical cost structure used by T-OPS. Full details of the model output are provided in Table D.1 in Appendix D.

Table 5.11: Data Validation of 2E Transport Costs and Volumes (Configuration C0, Week 46)

Customer	Metric	Model Output	Expected Benchmark
Customer 1	Total volume fulfilled		
	Transport cost (Plant 1 → DC 1 → Customer 1)		
Customer 3	Total volume fulfilled		
	Transport cost (Plant 1 → DC 1 → Customer 3)		

The results show that the model cost logic (based on fixed ISF and CF rates, truckload capacity, and proportional volume allocation) corresponds exactly to historical planning logic. This confirms that the model correctly interprets input data, reflects how truckload-based costs are calculated in practice, and maintains fidelity to real-world constraints on volume fulfillment.

5.5.2. Validation of Direct Shipment Cost Inputs (C1)

While direct plant shipments (DPS) are not yet a standard part of P&G's logistics network, historical full-truck DPS shipments have occasionally occurred. The transport rates for these incidental shipments were shared by the Transport Operations (T-OPS) team (see Table B.2) and serve here to validate the cost assumptions applied to DPS arcs in the model.

Table 5.12 presents a comparison of model-derived DPS transport costs with historical FTL rates for two customer destinations in scenario S3 (Category Y from Plant 2). The model's DPS costs were computed using fixed cost parameters and estimated transport distances, assuming FTL usage. A detailed cost breakdown of the model run is given in Table D.2.

Table 5.12: Validation of DPS Cost Input Values Against Historical FTL Rates

Route	Model DPS Cost	FTL Cost Data T-OPS	Deviation
Plant 2 → Customer 1			
Plant 2 → Customer 3			

The results show minimal deviation for the Plant 2-Customer 1 route and a somewhat larger deviation for the Plant 2-Customer 3 route. These variations can likely be attributed to differences in VFR between the modeled shipments and the shipments from the historical data. Despite these differences, the results confirm that the model's DPS cost assumptions are within a reasonable range of operational reality, supporting the credibility of its cost parameterization.

5.5.3. Conclusion

The data validation confirms that the hybrid distribution model is built on consistent, realistic, and credible data inputs. Cost structures for both two-echelon and direct shipment configurations were shown to match operational benchmarks provided by the P&G logistics team. While this validation does not assess predictive performance or model accuracy in new scenarios, it provides confidence that the foundational data and parameter assumptions in the model are aligned with real-world logistics operations.

This type of validation primarily qualifies as data validation, confirming that the model inputs, cost components, and parameter values align with real-world logistics data. While not a full conceptual validation, this step is essential to ensure that the model is grounded in realistic and operationally meaningful data (Sargent, 2010).

5.6. Conclusion

This chapter presented and verified a flow-based optimization model to support strategic decision-making on hybrid distribution at P&G. The model captures both DPS and two-echelon flows, incorporating real-world cost structures and fixed-capacity constraints. Unlike classical VRP models that construct vehicle routes, this formulation abstracts vehicles as capacity units, reflecting P&G's outsourced transport operations where routing is managed by 3PL providers. Transport costs are therefore modeled per shipment leg, and vehicle availability is assumed.

The model was implemented as a Mixed-Integer Linear Program (MILP) using Python and the FICO Xpress solver. Behavioral testing confirmed that the model behaves logically under varying demand scenarios and parameter values. Solver diagnostics further demonstrated its numerical stability and tractability for small to medium instances, with solution times under one second.

To answer SQ5, *How should the proposed modeling approach be structured to evaluate the impact of direct shipping on supply chain efficiency?*; The modeling approach is structured as a flow-based MILP variant of the two-echelon vehicle routing problem. It optimizes shipment assignments from plants or DCs to customers, based on fixed transport costs and vehicle capacities, without explicitly routing vehicles. This abstraction aligns with P&G's operational context and enables scalable analysis of strategic distribution decisions. The structure supports scenario-based evaluation of logistics cost, service level, and emissions outcomes, providing insight into the conditions under which integrating DPS improves performance in a two-echelon network.

6

Experimental Design

This chapter outlines the experimental framework used to evaluate the performance of the hybrid transport network, under different conditions with and without DPS. The goal is to assess the trade-offs between logistics costs, environmental impact, customer service performance and the impact on stock allocation for various structural scenarios and policy configurations of the network.

To support a robust experimental design, a clear distinction is made between exogenous parameters (external, fixed inputs) and endogenous parameters (internally configurable model choices):

- **Exogenous parameters**: Fixed input data that are outside the control of P&G as the decision-maker. These include historical customer order patterns, demand volumes[1], transport costs, handling and storage costs, travel distances, and the baseline carbon emissions factor μ.
- **Endogenous parameters**: Internally configurable model decisions, including the DPS activation status, full truckload (FTL) volume thresholds for DPS eligibility, the structure of the transport network (two-echelon vs. DPS-enabled flows), and the carbon price factor λ used to value emissions.

This classification supports a robust experimental design by clearly separating input assumptions from configurable levers. Based on this, the experiments consist of testing multiple *exogenous scenarios* (demand streams) across different *endogenous configurations* (model settings), as detailed below.

6.1. Demand Scenarios: DPS-Eligible Streams

The four scenarios represent different exogenous demand streams that are eligible for DPS. While the underlying customer demand data is fixed, the eligibility of certain product-origin flows for DPS is an endogenous design decision that reflects internal policy choices. The four demand scenarios define which combinations of product types and origins are allowed to bypass the DC.

- **S1 – Category X ex Plant 1 (Production only)**: Only Category X SKUs produced at Plant 1 are eligible for DPS.
- **S2 – Category X ex Plant 1 (Production + Storage)**: DPS is allowed for both produced and stored Category X SKUs at Plant 1 (in the CIMAT warehouse).
- **S3 – Category Y ex Plant 2 (Production only)**: DPS is allowed for Category Y SKUs produced at the Plant 2 plant.
- **S4 – Category Y ex Plant 2 (Production + Storage)**: DPS is allowed for both produced and stored Category Y SKUs at Plant 2 (FBNL shared codes).

These scenarios reflect different internal strategies regarding which volumes are considered for DPS. Allowing plant-location stored SKUs in addition to produced ones increases potential DPS volume and provides additional flexibility in plant-to-customer deliveries.

To illustrate the structural differences between the demand scenarios, Figure 6.1 provides a schematic representation of the demand streams in a simplified two-echelon network comprising a plant P, a distribution center D, and a customer C. Extended versions of this illustration, tailored to Plant 1 and Plant 2, were previously given in Figure 1.2 and Figure 1.3.

[1] While demand volumes are exogenous, the selection of which demand streams are eligible for DPS is an endogenous choice

The illustration in Figure 6.1 is explained using the example of Category X demand originating from Plant 1. In the baseline configuration (C0), all three Category X Segments of demand are routed via the DC, meaning no direct shipments from the plant to the customer are allowed.

When enabling DPS (in configuration C1-C3) for scenario S1, only those Category X SKUs that are *produced* at Plant 1 (stream 1) are eligible for DPS, enabling direct shipments for a subset of demand. Scenario S2 extends this by also allowing DPS for Category X SKUs that are *stored* at Plant 1 (in the CIMAT warehouse), even if they are not produced there (stream 2). Thus, scenarios S1 and S2 progressively increase the potential DPS volume.

The expected effect of broadening DPS possibilities between S1 and S2, is a shift in volume away from the P-D-C path toward direct P-C deliveries. This would reduce the load through the DC in terms of both throughput volume and storage, and potentially lowering logistics costs and improving service levels.

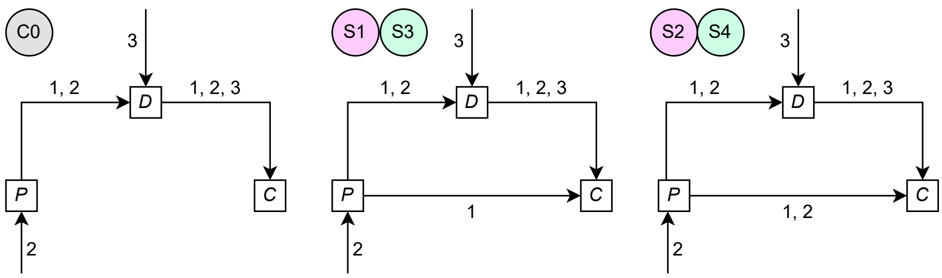

Figure 6.1: Demand Scenarios Illustrative Example

This illustrative example also holds for scenarios S3 and S4, reflecting the demand of Category Y from Plant 2. Scenarios S3 and S4 also progressively increase the potential DPS volume. All scenarios will be run separately through the model to assess the impact of enabling DPS under varying levels of operational flexibility. For each scenario, customer-level demand data will be filtered according to the applicable SKUs and origin constraints.

6.1.1. Impact of Including Stored SKUs on Weekly Demand Patterns

To assess the impact of including stored SKUs not produced at the plant, Figure 6.2 and Figure 6.3 show the weekly demand distributions per customer for Category X and Category Y, respectively. These boxplots compare demand before (S1/S3, blue) and after (S2/S4, orange) adding stored SKUs. The white diamond represents the mean value in each case.

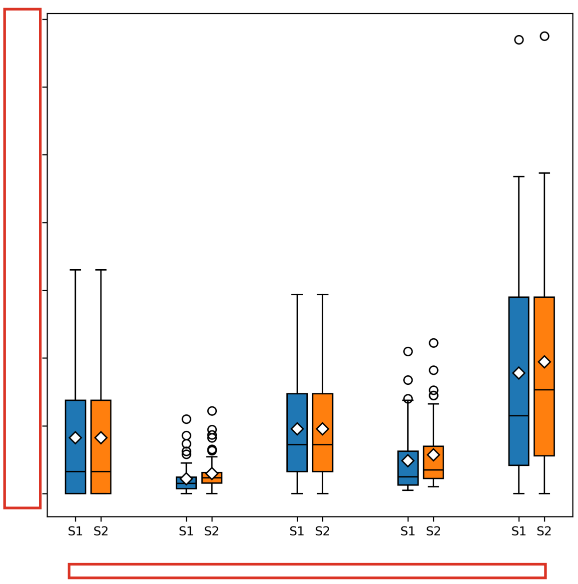

Figure 6.2: Weekly Demand Distribution per Customer for Category X (S1 vs. S2)

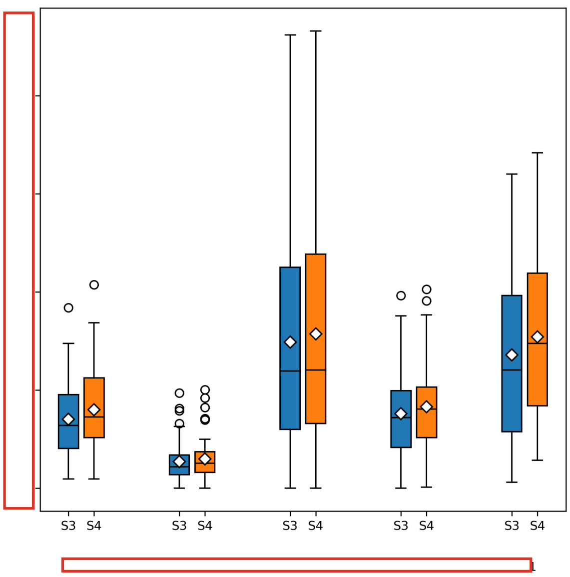

Figure 6.3: Weekly Demand Distribution per Customer for Category Y (S3 vs. S4)

Table 6.1 shows the average percentage increase in weekly demand per customer when stored SKUs are included. The impact varies notably between customers and categories. While Customer 2 and Customer 5 show substantial average increases for Category X, other customers such as Customer 1 and Customer 3 are only affected in the Category Y category, and exhibit no meaningful change in Category X. Overall, the effect of adding storage demand is highly customer- and category-specific. These low additional values can be explained by the fact that the largest segment of Category X and Category Y products within the categories are produced at the plants in this scope. This result suggests that while DPS potential increases slightly for some customers, it does not significantly affect demand volatility.

Table 6.1: Mean Percentage Additional Stored SKUs on Weekly Demand per Customer for Category X and Category Y

Customer	Category X (S1 to S2)	Category Y (S3 to S4)
Customer 1		
Customer 2		
Customer 3		
Customer 4		
Customer 5		
Avg		

6.2. Model Configurations: Policy Settings

To test each demand scenario under varying internal logistics policies, four model configurations are defined. These configurations vary key endogenous parameters, such as whether DPS is allowed and what shipment volume thresholds apply. The objective is to observe how performance outcomes change under different internal rules.

- **C0 – Baseline (No DPS)**: Two-echelon network with DC-based delivery only. Used as benchmark in the model.
- **C1 – Full DPS Enabled**: DPS allowed for all eligible orders with . Used to test routing decisions with the most realistic integrated DPS configuration.
- **C2 – DPS Threshold Relaxation**: FTL threshold relaxed to , potentially increasing DPS opportunities. Evaluates the effect of partial relaxation policies on cost-efficiency.
- **C3 – DPS Threshold Removed**: FTL threshold fully removed (τ = 0 FP), allowing DPS for all shipment sizes. Assesses the impact of highly flexible shipment policies.

Each of the four demand scenarios (S1-S4) is tested under the four model configurations (C0-C3), resulting in 16 main experiments.

6.3. Customer and Time Selection for Experimental Analysis

To enable in-depth scenario testing while maintaining computational efficiency, the experimental analysis focuses on a subset of five high-demand customers. These were selected based on a broader demand screening described in subsection 5.2.4, which evaluated weekly shipment volumes for customers over a 41-week period.

The five selected customers (Customer 1, Customer 2, Customer 3, Customer 4, and Customer 5) were chosen due to their consistently high demand in both the Category Y and Category X product category, along with sufficient residual volume in other product categories to sustain DC-based flows. This ensures that DPS integration can be tested without disrupting the viability of existing logistics.

Week Selection for Experimental Input

For each selected customer, the full demand in FP for a set of 3 representative weeks is used as model input. This choice enables the evaluation of different distribution strategies and scenario settings under realistic, disaggregated demand conditions while capturing natural week-to-week volatility. Limiting the analysis to 3 weeks ensures a manageable number of experimental runs, 48 in total, with 16 experiments conducted across 3 weeks, gaining a practical balance between analytical depth and computational feasibility.

The decision to include multiple consecutive weeks is grounded in the statistical demand analysis presented in subsection 5.2.4, particularly the autocorrelation and stationarity results in Figure 5.2.4 and Figure 5.2.4. The autocorrelation analysis revealed weak short-term memory in the demand patterns, which gives sign of operational volatility on a week-to-week basis. At the same time, the stationarity analysis showed that most customer-category combinations are stable in the long term, supporting the validity of using representative historical data for planning purposes. Together, these findings justify the use of three consecutive weeks: they are sufficient to capture short-term volatility while remaining within a stable long-term demand environment, offering a realistic and statistically sound basis for model input without inflating the number of experimental runs beyond feasible limits.

To ensure that the selected weeks reflect the central tendency and variability in weekly demand, a statistical boxplot was created for the total weekly demand summed across the five selected customers, shown in Figure 6.4 and Figure 6.5. These figures depict the interquartile spread, medians, and outliers in demand over the 41 weeks for Category X and Category Y, respectively. The associated descriptive statistics are shown in Table 6.2.

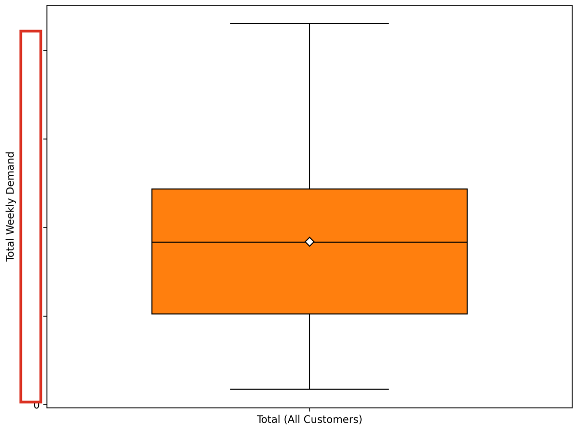

Figure 6.4: Total Weekly Demand Across all 5 Customers in Category X

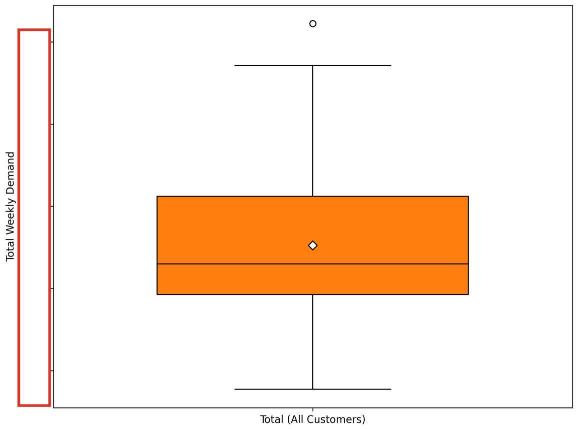

Figure 6.5: Total Weekly Demand Across all 5 customers in Category Y

Table 6.2: Descriptive Statistics of Total Weekly Demand (41 weeks)

Category	Mean	Std	Min	25%	Median	75%	Max	95% CI LL	95% CI UL
Category X									
Category Y									

To ensure representativeness, the three selected weeks are chosen to correspond to statistical markers within the interquartile distribution of total weekly demand. Specifically, one week was selected near the 25th percentile, one near the mean, and one near the 75th percentile of total demand. This approach ensures that the experimental analysis captures a diverse and representative range of operational conditions. The choice for percentiles rather than the 95% confidence interval bounds is deliberate: while confidence intervals quantify statistical uncertainty around the mean, percentiles reflect the actual spread and variation in observed weekly demand. Selecting weeks based on the interquartile range therefore provides a more realistic basis for modeling typical low, average, and high demand scenarios.

For Category X, weeks 36 to 38 from the data align well with this strategy. Week 36 corresponds closely with the 25th percentile, week 38 is near the median, and week 37 approximates the 75th percentile. As shown in Table E.6, the demand volumes of these weeks are respectively 49, aligning with the statistical markers.

However, demand in week 36 to 38 for Category Y falls consistently above the 75th percentile. As such, this selection does not meet the intended criterion of representativeness. Therefore, an alternative set of weeks will be used for Category Y, ensuring that the new selection aligns with the 25th, median, and 75th percentile values as shown in Table 6.2. A revised set of weeks, 30, 31, and 32, was selected to meet the same representativeness criteria. Week 30 approximates the 25th percentile, week 32 lies near the mean, and week 31 falls just above the 75th percentile threshold. Together, these weeks reflect the diversity of demand conditions without including statistical outliers, ensuring realistic and balanced input for scenario-based routing analysis (full demand specifics can be seen in Table E.3 and Table E.8).

6.4. Summary of Experimental Design

The experimental design follows a full-factorial structure, combining different exogenous and endogenous conditions across selected time periods. Summarized, the experiments are defined along the following three dimensions:

- **Demand scenarios (S1-S4)**: These represent different sets of eligible DPS demand, based on product type and origin (as described in section 6.1). Each scenario defines which SKUs can bypass the distribution center.
- **Model configurations (C0-C3)**: These vary internal logistics policy parameters, such as DPS activation and the full truckload (FTL) eligibility threshold, as detailed in section 6.2.
- **Representative weeks (Week 1-3)**: For both category-plant combination, three representative weeks were selected to reflect low, medium, and high demand conditions, based on statistical markers (see section 6.3).

By testing all combinations of these dimensions - 4 demand scenarios × 4 model configurations × 3 representative weeks - a total of 48 experiments is conducted. This setup enables systematic evaluation of internal design decisions (e.g., DPS thresholds) under fixed external conditions (e.g., demand eligibility), while capturing realistic weekly demand volatility. The design also ensures comparability across experiments and supports robust analysis of performance trade-offs.

6.5. Post-Optimization and Sensitivity Analyses

After running the optimization with the flow-based VRP model, and assessing the results for the main KPI's, the following sections describe the post-optimization analyses to be performed, to give further insights into the effects of integrating DPS into a conventional two-echelon network.

6.5.1. Stock Allocation Impact Analysis

To assess how DPS adoption influences inventory positioning, a post-optimization analysis on the KPI Stock Allocation is conducted. Based on the model outputs, this analysis quantifies the shift in volume from DC-based fulfillment to plant-based fulfillment. These shifts are relevant for understanding impacts on storage cost, working capital, and the "cash" dimension of supply chain performance.

6.5.2. Break-Even Analysis of FTL Thresholds

Beyond the predefined configurations (C0–C3), a break-even analysis was performed to explore whether the full truckload (FTL) threshold for DPS eligibility should vary by customer or route. By calculating break-even volumes per plant-customer path, this analysis helps identify more efficient, cost-aligned, customer-specific DPS policies that could improve the robustness of the transport network design.

6.5.3. Sensitivity of Emission Pricing on Share of DPS (λ Variation)

A dedicated sensitivity analysis is conducted to assess how internalizing carbon costs affects routing decisions and emissions outcomes. This was performed under configuration C4 (full DPS enabled) across all four demand scenarios, with the carbon price factor λ varied from €0.00 to €4.00 per kg CO_2. Unlike the main experiments, which vary structural configurations, this analysis keeps the network setup fixed and isolates the effect of emission valuation in the objective function. While the pricing levels are hypothetical, results could offer insight into how increased emphasis on emissions could influence decision-making at P&G.

The sensitivity analysis follows the scalarization approach described by Pereira et al. (2022), where a composite objective function is optimized based on a weighted sum of logistics cost and emissions (see Equation 5.15 in subsection 5.1.5). Varying λ from €0.00 to €4.00 per kg CO_2 allows exploration of trade-offs between cost and emissions under different carbon valuation scenarios.

6.5.4. ε-constraint Analysis: Emissions Minimization Under Cost Caps

To complement the scalarization-based sensitivity analysis, an additional analysis is performed using the ε-constraint method. In this approach, emissions are minimized directly as the sole objective, while the logistics costs are constrained to remain within a specified upper bound. This formulation helps to explore how the model behaves when sustainability is prioritized above cost, and whether alternative routing decisions, such as a shift from DPS to two-echelon flows, emerge under a purely environmental objective.

The use of the ε-constraint method in multi-objective optimization has been widely adopted in emissions-related domains. For instance, Ghorbani et al. (2020) applied this approach in economic emission dispatch, demonstrating its effectiveness in isolating emission-minimizing strategies without introducing scalarization bias. Their study shows that this method supports more transparent exploration of trade-off frontiers, especially when evaluating operational decisions under real-world constraints. Inspired by this rationale, the present study applies a similar approach in a logistics context to evaluate whether emission-optimal routing patterns can be achieved independently of explicit cost minimization, while still respecting economic feasibility constraints.

6.6. Key Performance Indicators (KPI's)

Performance is evaluated using a set of KPI's derived from the research objective (section 1.5), which emphasizes the three pillars of the supply chain triangle: cost, service, and cash (Desmet, 2018). A distinction is made between KPI's that are explicitly optimized within the model and those evaluated post-solution based on model outputs and domain logic.

6.6.1. KPI's Optimized in the Model

- $\bar{Z}_{\text{cost}}$: Total logistics cost (€), including transport, loading, and unloading costs across both two-echelon and DPS routes;
- $\bar{Z}_{\text{emissions}}$: Total kg CO_2 emissions, derived from traveled distances and flow assignments, capturing the environmental impact of transport decisions;
- $\bar{Z}_{\text{total}}$: Combined cost objective (€), incorporating monetized emissions using the carbon price factor λ, thus internalizing environmental costs into the optimization.

6.6.2. KPI's Evaluated Post-Optimization

- Estimated service level ($SL_{\text{estimated}}$): Volume-weighted average fulfillment performance, based on historical delivery reliability per shipping location (see Equation 5.16);
- DPS share and stock allocation: Share of total shipment volume transported via DPS, reflecting the degree of decentralization in the DC and its impact on stock allocation toward plants rather than the DC. This metric captures the number of DPS vs 2E shipments, but does not reflect the total number of vehicle movements, as 2E shipments involve two legs and the model does not explicitly track vehicle usage. Transport costs are modeled per shipment leg, consistent with P&G's outsourced cost structure.

6.7. Experimental Execution and KPI Collection

To evaluate each scenario and configuration, the following standardized workflow is applied:

1. Model Solution: The optimization model is solved using the scalarized objective function, based on the active configuration parameters;
2. Extraction of Decision Variables: Key outputs such as flow allocations, routing decisions, and vehicle assignments are collected post-optimization;
3. KPI Calculation: All relevant performance indicators are computed, including optimized and post-optimization KPI's;
4. Result Aggregation and Visualization: Outcomes are organized into structured summary tables and visualized through comparative plots to support scenario comparison.

This workflow enables systematic assessment of trade-offs across cost, sustainability, and service, in line with the supply chain triangle framework. It also gives understanding of how the expanded DPS eligibility affects the operational efficiency and KPI's.

6.8. Conclusion

This chapter has defined the experimental design used to assess the hybrid distribution model under a range of structural scenarios and internal policy configurations. By distinguishing between exogenous demand scenarios and endogenous configuration levers, the design allows for targeted analysis of how changes in logistics rules affect system performance. Combining customer-level demand inputs, configurable DPS policies, and emissions valuation, the design lays the foundation for the performance comparisons presented in the next chapter, which evaluate the cost-effectiveness and operational impact of integrating direct plant shipments into the current network design.

7 Computational Results

This chapter presents the results of the experiments defined in chapter 6. Each experiment combines one of four exogenous demand scenarios (S1-S4) with one of four endogenous policy configurations (C0-C3), resulting in a structured comparison across varying supply and demand conditions.

In all scenarios, demand is defined at customer-week level and fulfilled through individual shipments per customer, without inter-customer or inter-week consolidation. This modeling choice reflects and simplifies operational realities and ensures that shipment-level decisions (e.g., route choice, vehicle assignment) are evaluated per customer delivery.

The experiments are evaluated using the KPI framework introduced (section 6.6), covering logistics and environmental costs, service level, and stock allocation. The primary objective is to quantify the added value of structurally integrating DPS into the existing two-echelon network in the research scope and to explore how this hybrid strategy affects performance across the cost-service-cash dimensions of the supply chain. Key trade-offs and stakeholder-relevant outcomes are discussed for each scenario–configuration combination.

Full numerical results of all experiments of the model optimization are given per scenario in section E.5, section E.6, section E.7, and section E.8, but the extensive interpretation of the results is given in the following sections with supporting plots, for the Category X scenarios and Category Y scenarios respectively.

Next to the supportive plots to show the results, performance of different distribution configurations is systematically evaluated by pairwise comparison across the defined scenarios, as this approach is widely recommended for simulation-based performance evaluation and ranking (Xiao et al., 2023). For each configuration (C1-C3), results were compared to the baseline configuration (C0) on KPI's logistics costs, emissions, and service level. The comparative analysis is based on the three weekly replications per configuration within each scenario (S1-S4). First, the average and standard deviation of the percentage difference between each configuration and the baseline were calculated per scenario. Then, these scenario-level comparisons were aggregated into final averages across all four scenarios. This two-step averaging approach ensures that each scenario contributes equally to the overall comparison, independent of the absolute values of its KPIs.

7.1. Model Feasibility and Performance

First of all, the model has been checked for feasibility and the appearance of optimal performance. All experiments have found an optimal solution, with solving times beneath 1 second. An example of the performance for the first demands scenario is given in Table 7.1. Further elaboration on solving times is left out of this report, but is controlled and has shown stable results.

Table 7.1: Feasibility and Solve Time Summary (Scenario S1)

Experiment	Status (Optimal/Infeasible)	Solve Time (s)	Gap (%)
S1-C0-W37	Optimal	0.01	0.00000%
S1-C1-W37	Optimal	0.06	0.00001%
S1-C2-W37	Optimal	0.02	0.00000%
S1-C3-W37	Optimal	0.01	0.00000%

7.2. Cost and Environmental Impact

To evaluate the effect of DPS on logistics cost and environmental performance, experiments are performed for demand of the five top customers. This section summarizes the extended results in subsection 7.2.1 and subsection 7.2.2. The impact of the model's routing decisions on the optimization of the logistics costs and emission KPIs are as follows:

- **Logistics costs:** Across all demand scenarios, enabling DPS leads to lower logistics costs. The model tends to prioritize DPS when volumes and travel distances are sufficient, reducing reliance on the DC. The average logistics cost decrease over all experiments is 26.4% under the configuration with a regular FTL threshold.
- **Emissions:** Emission reductions are only observed when DPS volumes support efficient vehicle utilization. For smaller volumes or short travel distances, the environmental benefits are limited or may even reverse under low FTL thresholds due to underutilized vehicle capacity. The average decrease in emissions is 7.0% under regular FTL thresholds.
- **FTL threshold effect:** Lowering the FTL threshold across configurations C1 to C3 results in a marginal additional decrease in logistics costs, driven by the model's increased selection of DPS as the preferred distribution method. On the other hand, lowering the threshold shows a negative impact on emissions, even increasing emissions in certain scenarios.

7.2.1. Scenario S1 and S2 - Category X from Plant 1

To evaluate the cost and environmental impact of integrating DPS for Category X products originating from Plant 1, the results for scenarios S1 and S2 are plotted in Figure 7.1 and are analyzed in this section. Figure 7.1 illustrate the outcomes of each configuration across the selected weeks for both scenarios. Plotted results for the scenarios separately, are found in section E.1 in Appendix E.

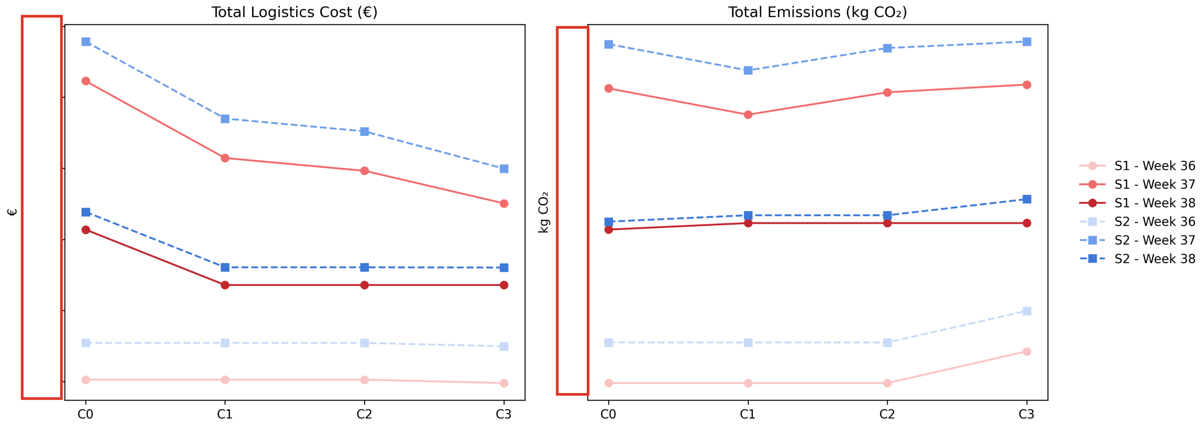

Figure 7.1: Weekly Logistics Cost and CO_2 Emissions for Category X Scenarios S1 and S2

The first observation for both scenarios is that allowing the model to choose DPS as a distribution method consistently results in a reduction in logistics costs. Comparing configurations C1, C2, and C3 to the baseline configuration C0 using a pairwise comparison over all three weeks, yields the percentage differences shown in Table 7.2. These results indicate that enabling DPS always leads to cost savings in logistics. Further relaxing the FTL threshold, moving from C1 to C3, results in modest additional reductions. The primary driver behind these savings is the increased proportion of DPS deliveries, which reduce reliance on DC-based flows. As illustrated in Figure 7.1, certain low-volume demand weeks, such as week 36, do not exhibit logistics cost reductions when DPS is used. This is because the model does not select

DPS for smaller volumes; instead, it opts for two-echelon distribution, which offers better consolidation and is therefore more cost-efficient in such cases.

However, allowing DPS in these scenarios does not lead to a reduction in total emissions. For small shipment volumes, transporting goods via the convential DC yields lower emissions per unit shipped. This is due to the assumption that DC-based vehicles are fully loaded with aggregated demand, thereby distributing emissions over a larger volume of units in the vehicle. In contrast, DPS shipments are modeled as dedicated trips, where the full vehicle emissions are attributed solely to the specific Category X load. Given the short distance between Plant 1 and DC 1, the additional distance savings with DPS are relatively small, limiting their potential environmental benefit. It can be seen that lowering the FTL threshold (in C3 fully to 0), will even lead to an increase of emissions for the demanded volume, if the model performs the scalarized optimization run where the most cost-efficient solution is found.

Table 7.2: Pairwise Comparison of Configurations vs Baseline C0 (Scenarios S1 and S2)

Scenario	Comparison	Logistics Cost Mean % diff	Logistics Cost Std % diff	Emissions Mean % diff	Emissions Std % diff
S1	C1 vs C0	-12.1%	10.5%	-1.2%	4.0%
	C2 vs C0	-13.0%	11.3%	0.4%	1.5%
	C3 vs C0	-16.3%	12.8%	7.4%	10.4%
S2	C1 vs C0	-11.2%	9.8%	-1.1%	3.7%
	C2 vs C0	-12.1%	10.5%	0.4%	1.4%
	C3 vs C0	-15.4%	12.4%	7.7%	7.5%
Avg (S1 & S2)	C1 vs C0	**-11.7%**	10.1%	**-1.2%**	3.9%
	C2 vs C0	**-12.6%**	10.9%	**0.4%**	1.5%
	C3 vs C0	**-15.8%**	12.6%	**7.6%**	9.0%

Scenario S2 expands demand eligibility by including not only plant-produced, but also stored Category X SKUs at Plant 1. Adding the stored volumes to the demand in the model gives an average increase of % over all five top customers (Table 6.1). Despite this higher volume, the relative impact on logistics cost and emissions remains similar to S1 (Table 7.2). Logistics costs continue to decrease significantly in configurations C1 to C3 compared to C0, but the marginal benefit of further lowering the DPS threshold beyond is limited. Emissions results are similar to scenario S1.

Overall, the analysis of scenarios S1 and S2 indicates that enabling DPS for Category X demand from Plant 1 significantly reduces logistics costs, especially when demand volumes are sufficiently high (e.g., week 37). However, for low-volume weeks, the model favors DC-based shipments, which benefit from consolidation with aggregated demand. Emission-wise, allowing DPS shows minimal benefits and lowering the DPS threshold can adversely affect emissions due to underutilized vehicles in DPS shipments.

7.2.2. Scenario S3 and S4 - Category Y from Plant 2

Figure 7.2 present the results of the optimization experiments for scenarios S3 and S4, for Category Y demand from the plant in Plant 2. Separated graphs for both scenarios are seen in section E.1.

As concluded from the demand analysis, compared to the Category X scenarios, Category Y demand of the five top customers is significantly larger, and distances from the plant to both the DC 1 and to customers are substantially longer. As a result, the model is more decisive in optimizing routing through DPS, even under stricter internal FTL thresholds. This indicates that bypassing the DC can yield considerable cost and distance advantages, especially for high-volume, long-haul flows.

The optimization clearly shows a substantial reduction in both logistics costs and emissions upon enabling DPS. Configuration C1 achieves an average logistics cost reduction of over 42.1% relative to the baseline (C0) (Table 7.3). These cost benefits remain relatively stable across configurations C1 through C3, suggesting that the model already prioritizes DPS at higher thresholds due to the compelling cost dynamics. Emissions show a similar improvement, decreasing by up to 11.7% compared to C0. However, just like in Category X, the emission benefits decline with lowering the FTL threshold further. The overall benefits of DPS in Category Y scenarios are notably more pronounced than in Category X, primarily due to the higher-volumes than fill the trucks and the longer distance savings that can be made.

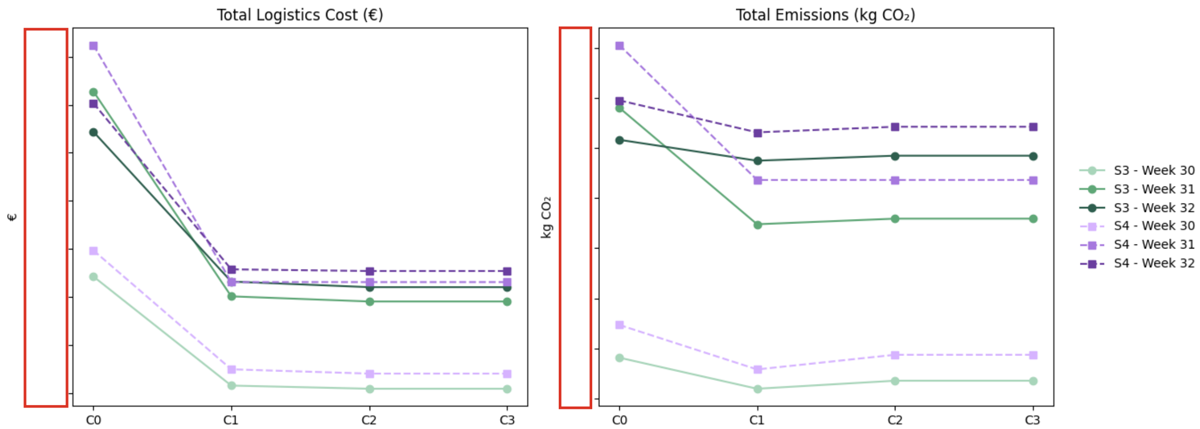

Figure 7.2: Weekly Logistics Cost and CO_2 Emissions for Category Y Scenarios S3 and S4

Scenario S4 expands DPS eligibility to include stored SKUs in addition to those produced at Plant 2 with on average % weekly (Table 6.1). Despite this added volume, the outcomes for configurations C1 to C3 in Figure 7.2 remain nearly identical to those for both scenarios. This suggests that the model had already maximized DPS utilization under the high production volumes of S3 and that a marginal increase in DPS-eligible volume in S4 does not significantly alter routing decisions in finding the optimal solution.

Table 7.3: Pairwise Comparison of Configurations vs Baseline C0 (Scenarios S3 and S4)

		Logistics Cost		**Emissions**	
Scenario	**Comparison**	Mean % diff	Std % diff	Mean % diff	Std % diff
S3	C1 vs C0	-41.6%	4.5%	-10.9%	8.1%
	C2 vs C0	-42.9%	4.4%	-9.5%	8.3%
	C3 vs C0	-42.9%	4.4%	-9.5%	8.3%
S4	C1 vs C0	-42.7%	5.0%	-12.6%	7.7%
	C2 vs C0	-43.3%	4.8%	-11.0%	8.5%
	C3 vs C0	-43.3%	4.8%	-11.0%	8.5%
Avg (S3 & S4)	C1 vs C0	**-42.1%**	4.8%	**-11.7%**	7.9%
	C2 vs C0	**-43.1%**	4.6%	**-10.2%**	8.4%
	C3 vs C0	**-43.1%**	4.6%	**-10.2%**	8.4%

This plateau effect highlights a key insight: in scenarios with high demand and long distances, the cost-optimizing model inherently favors DPS, even at strict FTL thresholds. Consequently, the benefit of lowering the FTL threshold (from to to 0 pallets) becomes less impactful in terms of cost savings, as the model already shifts large volumes directly from the plant whenever possible. However, the FTL threshold does influence VFR and with that emissions. Higher volumes allow for better utilization of DPS vehicles, making direct shipments more environmentally efficient by spreading emissions across larger loads.

Overall, scenarios S3 and S4 underscore the importance of shipment volume and transport distance in determining the effectiveness of DPS. Cost and emissions savings are significant, but lowering the FTL threshold will not further improve this. It is primarily affecting emissions through its influence on VFR efficiency, rather than altering route selection itself.

7.3. Service Level and Network Effects

Service level and network effects resulting from changes in routing decisions are analyzed to assess the operational robustness of each experiment. Key insights across all four scenarios are summarized in this section and further elaborated in subsection 7.3.1 subsection 7.3.2.

- **Service level:** Across all scenarios, enabling DPS consistently improves service levels. Even under restricted configurations, DPS allows for higher fulfillment rates directly from the plant, reducing dependency on DC inventory and increasing responsiveness to customer demand. These improvements translate into commercial value, as discussed in subsection 7.3.3,

.

- **Network effects:** The shift from conventional two-echelon routing to DPS is most pronounced in high-volume and long-distance scenarios (S3 and S4). In these cases, DPS becomes the dominant routing mode as soon as it is enabled. Although extending DPS eligibility to stored SKUs (in S2 and S4) adds some flexibility in vehicle utilization, the incremental effect on routing and service levels is limited. This indicates that the initial demand volume and shipment distance are more decisive factors than extending eligibility in determining routing preferences.
- **FTL threshold effect:** Lowering the internal FTL threshold offers only marginal gains in both DPS share (network effects) and service level. Especially in scenarios with already high DPS adoption, such as S3 and S4, the performance stabilizes quickly after DPS is enabled, suggesting a saturation point in the benefits of further relaxing FTL constraints.

7.3.1. Scenario S1 and S2 - Category X from Plant 1

Figure 7.3 shows the trend of service level performance over the configurations for both demand scenarios in Category X. Next to that, the share or division of routing choices, DPS vs the conventional two-echelon distribution, is given in the bars, for both scenarios separately.

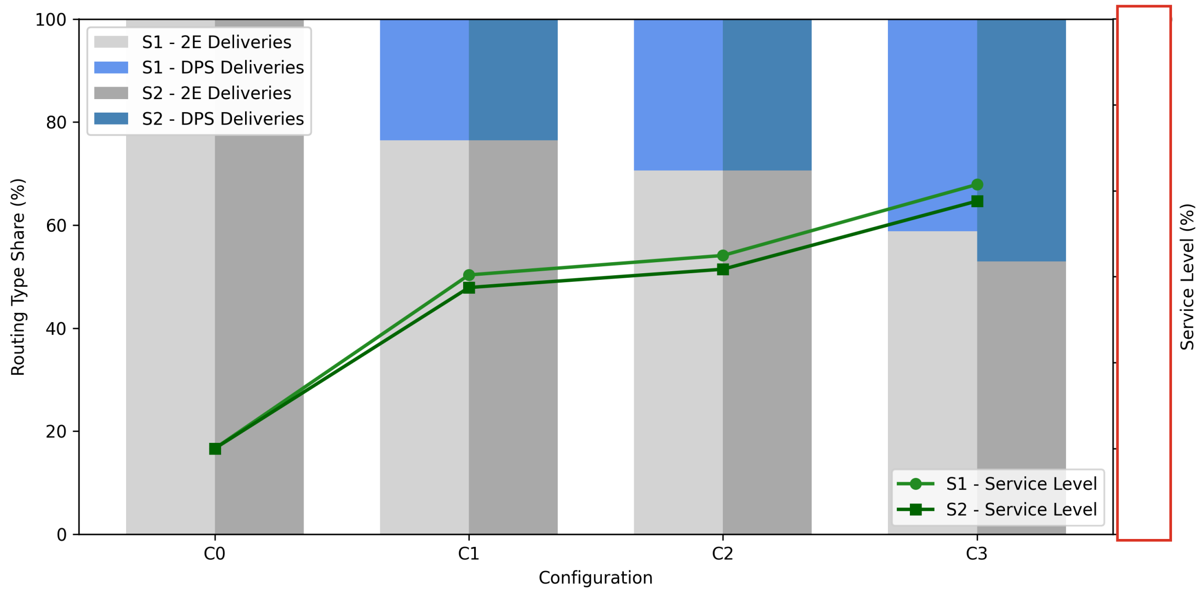

Figure 7.3: Service Level and Routing Choice for Category X in S1 and S2 over the Configurations

Integrating DPS has a clear positive effect on service level performance, as illustrated by the green lines in Figure 7.3. Over the three evaluated weeks, the average service level for the top five customers increases from X% in the baseline configuration (C0) to above X% in configuration C3. This improvement is driven by more flexible and direct fulfillment options enabled by DPS, which reduce dependency on DC inventory availability. As shown in the graph, the share of DPS deliveries increases from 0% in C0 to over 40% in C3, closely tracking the rise in service level.

The pairwise comparison results in Table 7.4 further supports this trend: although C1 already yields notable improvements, the largest relative gain over the baseline is observed in C3. This confirms that while enabling DPS already provides a strong boost to service performance, relaxing the FTL threshold further continues to yield modest additional benefits. These findings highlight that DPS enhances responsiveness in the supply network, especially when volume and delivery frequency allow for more direct routing.

Table 7.4: Pairwise Comparison of Configurations vs Baseline C0 for Service Level (Scenarios S1 and S2)

Scenario	Comparison	Service Level Mean % diff	Service Level Std % diff
S1	C1 - C0	1.0%	0.9%
	C2 - C0	1.2%	1.0%
	C3 - C0	1.5%	0.6%
S2	C1 - C0	1.0%	0.9%
	C2 - C0	1.1%	0.9%
	C3 - C0	1.5%	0.6%
Avg (S1 & S2)	C1 - C0	**1.0%**	0.9%
	C2 - C0	**1.1%**	1.0%
	C3 - C0	**1.5%**	0.6%

The impact of adding the additional stored demand in scenario S2 is minimal; both scenarios exhibit a very similar pattern in service level improvement across configurations C0 to C3. However, service level for S2 is consistently slightly higher than for S1. This difference can be attributed to the additional volume included in S2, which is more frequently assigned to DPS deliveries. As a result, a larger share of demand is fulfilled more efficient with DPS, marginally increasing the overall service level. This suggests that the inclusion of stored demand in scenario S2 has only a limited effect on service level outcomes compared to S1.

7.3.2. Scenario S3 and S4 - Category Y from Plant 2

The results for scenarios S3 and S4 are presented in Figure 7.4 and the pairwise comparison is detailed in Table 7.5. In both scenarios, integrating DPS in configuration C1 leads to a high increase in DPS share, exceeding X% on average, and a corresponding rise in service level from X% to approximately X%. This level remains stable across configurations C1 to C3, indicating that most gains are achieved immediately upon enabling DPS in high-volume scenarios.

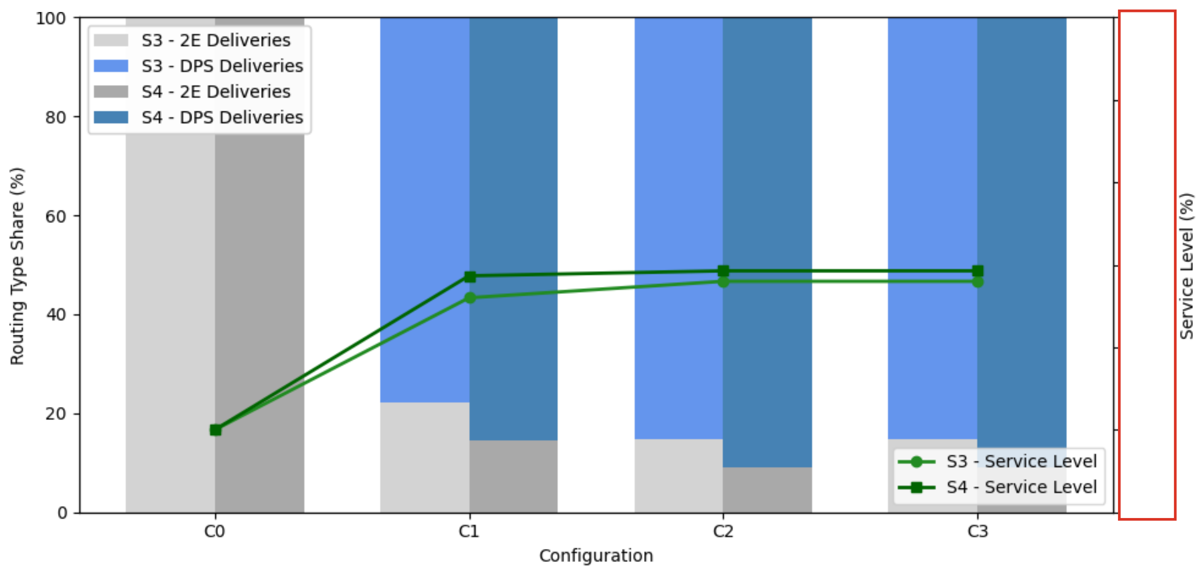

Figure 7.4: Service Level and Routing Choice for Category Y in S3 and S4 over the Configurations

Table 7.5: Pairwise Comparison of Configurations vs Baseline C0 for Service Level (Scenarios S3 and S4)

Scenario	Comparison	Service Level	
		Mean % diff	Std % diff
S3	C1 - C0	1.0%	0.0%
	C2 - C0	1.0%	0.0%
	C3 - C0	1.0%	0.0%
S4	C1 - C0	1.0%	0.0%
	C2 - C0	1.0%	0.0%
	C3 - C0	1.0%	0.0%
Avg (S3 & S4)	C1 - C0	**1.0%**	0.0%
	C2 - C0	**1.0%**	0.0%
	C3 - C0	**1.0%**	0.0%

The consistently high DPS share and service levels suggest that the model quickly converges to an optimal routing structure when demand volumes are high and transport distances are long, as is the case for Category Y shipments from Plant 2. Relaxing the FTL threshold beyond configuration C1 does not yield further improvements, confirming that the system reaches a saturation point in performance once DPS is allowed.

Scenario S4, including stored SKUs in addition to plant-produced SKUs, shows a nearly identical trend to scenario S3. While the added volume does not negatively impact performance, it also does not structurally change the routing outcomes. These results highlight that in this already high-volume, long-distance settings, the model strongly favors DPS from the outset, and further adjustments to eligibility rules or thresholds have limited additional value.

7.3.3. Estimated Commercial Impact of Service Level Improvements

While this study primarily evaluates logistics performance, the service level improvements observed across DPS-enabled configurations also have strategic commercial implications.

, as mentioned in subsection 5.1.6.

To illustrate this relationship, Table 7.6 presents the estimated NOS impact on the shipments of the supply chain in this FBNL scope, derived from the average service level gains achieved in scenarios S1 to S4. These figures are indicative and assume a linear relationship between service level and sales performance, without accounting for category-specific price or margin differences.

Table 7.6: Estimated Relative NOS Impact Based on Service Level Gains

Scenario	Configuration	Service Level Gain (%pt)	Estimated NOS Gain (%)
S1	C0 → C3	$+1.5$	
S2	C0 → C3	$+1.4$	
S3	C0 → C3	$+1.0$	
S4	C0 → C3	$+1.0$	

These results demonstrate that marginal increases in service level, enabled through more flexible routing such as DPS, can lead to meaningful sales improvements for P&G's supply chain on Category Y and Category X in FBNL. The commercial value of enhanced fulfillment performance reinforces the importance of supply chain design decisions not only from a logistics cost perspective, but also in terms of customer satisfaction and potential NOS increase.

7.4. Key Scenario Comparison

This section compares the four demand scenarios to evaluate what form of DPS eligibility yields the greatest benefits under shared policy configurations. Scenarios are shaped by demand volume and the distance between the plant and the DC, which affect both cost and service outcomes. The values in Table 7.7 are calculated averages of the full results in E.1 and Table E.2.

- **Key differences Category X (Plant 1) and Category Y (Plant 2):** The contrasting effects between Category X (Plant 1) and Category Y (Plant 2) scenarios emphasize the role of shipment volume and distance in determining the effectiveness of DPS. From Plant 1, small-volume direct shipments yield minimal benefits, whereas from Plant 2, even moderate volumes justify DPS due to larger distances and associated cost savings. Where the effect on service level performance of integrating DPS is similar, the logistics cost and emission changes are way more impactful for Category Y from Plant 2, visible in Table 7.7.

- **S1 vs. S2 (Category X - Plant 1):** Expanding DPS eligibility to include stored SKUs in scenario S2, which is on average % additional volume (Table 6.1), provides no marginal improvements over production-only flows in scenario S1. Given the relatively small volumes and the short distance between Plant 1 and the DC, direct shipments offer limited added value and can even lead to higher emissions when underutilized vehicles below FTL thresholds would be deployed.
- **S3 vs. S4 (Category Y - Plant 2):** In contrast to S1 and S2, S3 and S4 benefit strongly from DPS, driven by higher volumes and longer transport distances. Scenario S4 includes stored SKUs and increases the total DPS-eligible volume, but performance outcomes remain only slightly higher than S3. This suggests that the model already fully leverages the DPS potential under production-only flows.

Table 7.7: Average Percentage Change from C0 to C1 per Scenario

Metric	**S1**	**S2**	**S3**	**S4**
Average Logistics Cost Change (%)	-12.1	-11.2	-41.6	-42.7
Average Emission Change (%)	-1.2	-1.1	-10.9	-12.6
Average Service Level Change (%)	$+1.5$	$+1.4$	$+1.0$	$+1.0$

7.5. Pairwise Comparison of Key Performance Indicators

As explained in the introduction of this chapter, a pairwise comparison is useful to systematically evaluate the performance of different distribution configurations on the KPI's.

The full pairwise comparison tables are given in Appendix E in Table E.1 and Table E.2, displaying the results for logistics cost and emissions, and service level respectively, including average differences and standard deviations. The aggregated comparison on KPI's costs and emissions across all configurations for Cat X&Y together, is shown in Figure 7.5, while Figure E.6 in section E.2 breaks down the results by product category (Category X vs. Category Y).

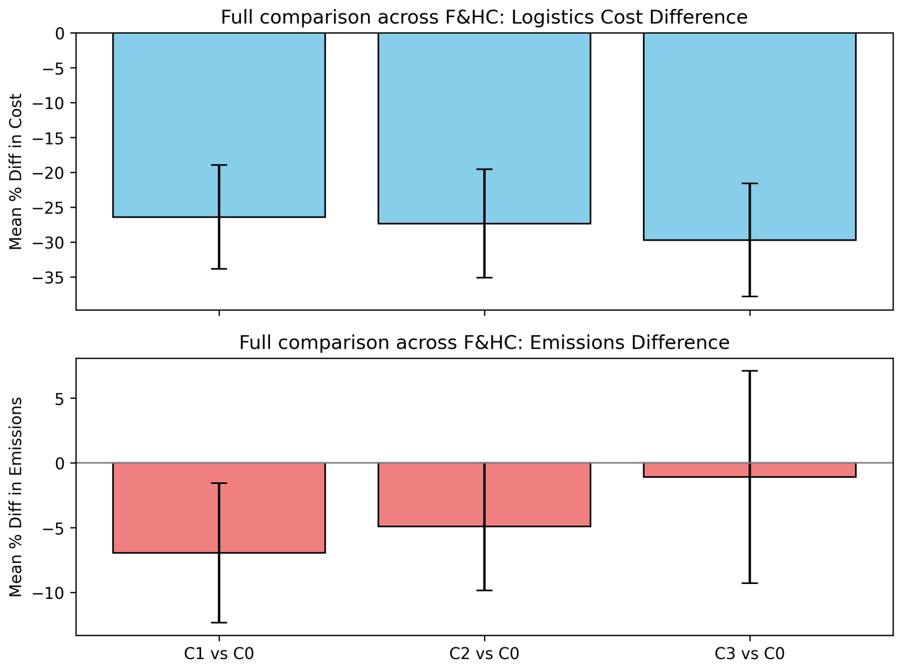

Figure 7.5: Full Comparison Across Both Categories of Logistics Cost and Emissions Savings over the Configurations

Similarly, the numerical results for service level are presented in Table E.2 in section E.4, with an overall view for aggregated Cat X&Y visualized in Figure 7.6. To provide further insight into category-specific performance, separate graphs are provided in the Appendix for Category X and Category Y in Figure E.12.

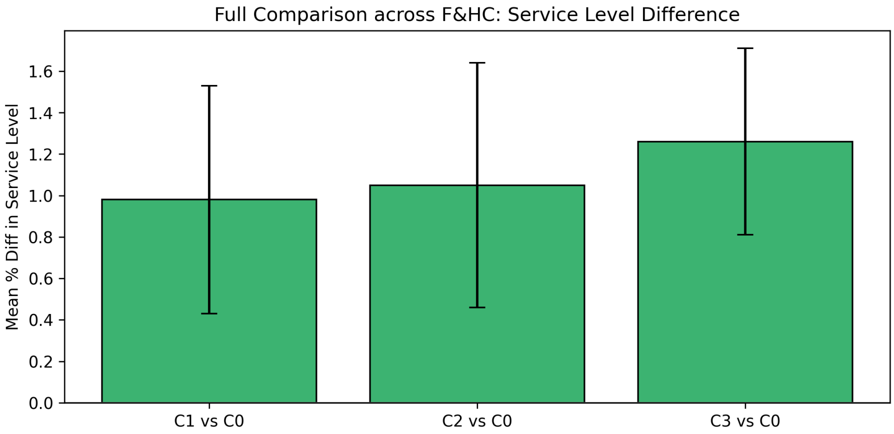

Figure 7.6: Full Comparison Across Both Categories of Service Level Change over the Configurations

The main observations to be done from the pairwise comparison aggregated for Cat X&Y together on the KPIs, is that lowering the FTL threshold of the vehicles for DPS, will have a positive trend on logistics costs, with an ultimate cost benefit of 29.7%. On emissions, a negative trend is seen with lowering the thresholds, where on average a high FTL threshold yields emission savings of 7.0%, but the full relaxed threshold only yields savings of 1.1% compared to C0. The effect on service level shows a positive trend, with an average ultimate service level increase of 1.3% in C3. This can be explained by the increase in DPS share which ultimately leads to service performance increase based on the estimation. These results enable a robust understanding of the relative impact of different configurations across the key performance metrics on an aggregated level.

7.6. Stock Allocation Impact (Cash)

To evaluate how integrating DPS influences inventory positioning and storage costs, a post-optimization stock allocation analysis was conducted. As a representative example, this section focuses on scenario S3 (Category Y - Plant 2), Configuration C1, in week 31, in which DPS is enabled under the threshold setting of FP. The full output of this experiment is given in Table E.10.

7.6.1. Observed Inventory Shift and Cost Impact

The optimization results reveal a substantial reallocation of SKU flows from two-echelon via DC 1 to DPS from Plant 2. In the baseline configuration (C0), the entire weekly Category Y demand of X FP is routed through the DC. In Configuration C1, X FP are fulfilled directly from the plant, while only X FP remain routed via the DC.

This reallocation implies a reduction of X FP in inventory holding requirement at the DC. Since storage costs at the plant are negligible (because P&G owned), financial impact is driven by the cost savings from avoiding DC storage. The estimated storage cost at the DC is per pallet (average B1/B2) per day (Table B.1). A floor position (FP) is assumed to occupy 1.5 pallet positions (2 B1 pallets or 1 B2 pallet), resulting in an average cost of per FP per day. Assuming an average storage duration of X days (based on weekly inbound and outbound flows), the total inventory cost reduction for this single experiment is calculated as follows:

$$\text{Inventory Cost Savings} = \boxed{}$$

This calculation represents the weekly savings, of this particular experiment, in storage costs, due to the shift in stock allocation from the DC to the plant, aggregated over the five top customers.

7.6.2. Interpretation and Broader Implications

The observed experiment illustrates that DPS not only improves service levels and transportation efficiency, but also contributes to inventory cost reduction by relocating stock from higher-cost DC storage in DC 1 to the plants. In this example, over 90% of weekly demand is fulfilled as DPS, effectively minimizing the role of the DC as an intermediate storage point. The margin of savings on logistics costs for this single experiment, which showed a total of €X (Table E.9), is 4.8% on weekly costs. While this percentage may appear modest, a structural weekly cost reduction of 4.8% could translate into substantial financial benefits when scaled across the full distribution network.

Next to logistics cost savings, this stock shift also alleviates pressure on DC storage capacity, which has become an increasingly urgent concern in recent operations. By freeing up space at the DC, DPS offers a potential solution to short-term capacity constraints.

Another important consideration is that reallocation of stock from the DC to the plant implies a shift in demand planning responsibility and forecasting accuracy. While observed cost savings are clearly positive from this case study, they must be weighed against the potential cost and risk of misaligned forecasts. Demand is not always predictable or consistent, particularly for lower volume on for example Category X SKUs. In such cases, maintaining availability at the DC remains critical, as it enables flexible response to small or unexpected orders. Fully relocating such SKUs to the plant could jeopardize service reliability.

However, for high-volume products with regular ordering patterns, such as those in the Category Y category, demand tends to be more stable and predictable. In these cases, keeping a portion of the inventory at the plant would be both operationally feasible and financially advantageous, without compromising responsiveness. This segmentation of SKUs based on volume and variability could help determine which SKUs are suitable for permanent DPS-based fulfillment and which require DC-based buffering.

7.7. Customer-Specific DPS Threshold Analysis

While the experiments reveal the broader impact of DPS thresholds, a more granular analysis can help determine when DPS becomes cost-optimal at different demand levels. One illustrative case is the Sold-To customer linked to Ship-To locations Customer 1 and Customer 5. Although these locations are treated separately in the optimization, their combined demand already triggers a high share of direct shipments (9 out of 10 routes) even under the standard FP threshold. Lowering this threshold to FP or even removing it entirely does not lead to additional DPS selection, indicating that threshold relaxation alone is insufficient if the demand volume per route is too low.

This limitation highlights the need for a different perspective: rather than adjusting thresholds under fixed demand, the following section investigates the minimum demand volume per plant-customer route at which DPS becomes more cost-efficient than two-echelon distributions. This demand-based break-even analysis enables a data-driven approach to determining customer-specific DPS policies.

7.7.1. Demand-Based Break-Even Analysis

To move beyond fixed-threshold testing in the previous experiments, a more insightful approach is to identify the minimum demand volume at which DPS becomes favorable according to the model optimization. Rather than adjusting thresholds under static demand, this break-even analysis incrementally increases weekly demand for specific plant-customer routes, which reveals the break-even point at which the model switches from two-echelon flows to DPS. This method provides a route-level view of where and when DPS is cost-optimal, and lays the groundwork for designing customer-specific thresholds based on underlying cost structures.

Table 7.8 shows the manually identified break-even demand volumes for each relevant origin–destination pair, covering the top five customer locations from both Plant 2 and Plant 1. These values represent the minimum volume for which a direct shipment is selected by the model as cost-optimal solution, thus reflect the optimal threshold P&G could offer each customer to align the incentives with cost-efficiency.

Table 7.8: Minimum Demand Volume Required for DPS Selection by Plant-customer Combination

Origin Plant	Customer Ship To	Break-even Demand Volume (FP)
Plant 2	Customer 1	
Plant 2	Customer 5	
Plant 2	Customer 2	
Plant 2	Customer 3	
Plant 2	Customer 4	
Plant 1	Customer 1	
Plant 1	Customer 5	
Plant 1	Customer 2	
Plant 1	Customer 3	
Plant 1	Customer 4	

To investigate whether these break-even thresholds are systematically related to distance, Figure 7.7 shows the break-even DPS thresholds plotted against the direct transport distance from the origin plant to the customer (see Table E.14 in section E.10 for detailed data for this graph). Each point is color-coded by origin and labeled with the customer name. Thin lines connect the customers per origin to indicate the trend.

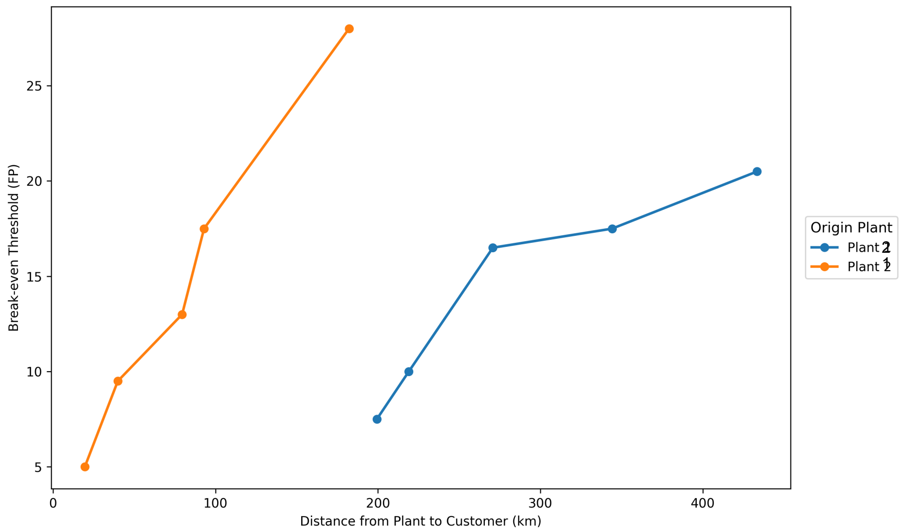

Figure 7.7: Break-even DPS Threshold (FP) vs. Direct Distance from Plant to Customer

This expanded analysis reveals a consistent and intuitive trend. For both origin plants, there is a clear positive relationship between the direct distance to the customer and the cost-efficient break-even DPS threshold. That is, as transport distance increases, the minimum volume needed to justify a direct shipment also increases. This makes sense from a logistics cost perspective: longer direct distances mean higher fixed CF transport costs per shipment, which require more volume to be offset by avoiding handling and transshipment costs at the DC.

For Plant 1 in particular, the relationship is steep. Moving from nearby customers like Customer 2 (X km, threshold X FP) to further ones like Customer 5 (X km, threshold X FP) shows a marked rise in break-even volume. The trend for Plant 2 is more gradual, but still visible, with thresholds rising from X FP (Customer 3, X km) to X FP (Customer 5, X km). This flatter trend at longer distances likely reflects the fact that more of the potential cost savings from bypassing the DC are already realized even at lower volumes. That is, on longer routes, the marginal gain from increasing shipment size decreases because the fixed cost of the direct leg becomes more dominant in the total cost structure.

This analysis demonstrates that the FTL threshold for DPS should not be a one-size-fits-all rule (e.g., X FP), but could instead be customized based on the plant–customer distance and associated cost trade-offs. By aligning the threshold to match the break-even point per route, P&G encourages cost-optimal routing behavior through a more data-driven and tailored DPS policy. However, it is important to note that these break-even thresholds are based solely on minimizing logistics costs and monetized emissions, and do not account for broader operational or sustainability considerations, such as Vehicle Fill Rate (VFR) targets or Scope 3 emission reduction goals.

Model output shows that even when DPS are chosen at the break-even volume, the total emissions (in kg CO_2) are equivalent to those of a fully filled truck. As a result, the emission intensity per unit FP increases when smaller volumes are shipped directly (seen in section 7.2). However, since emissions have limited influence in their monetized form in the current objective function, the model still favors these lower-volume DPS routes under a cost-minimization framework (Equation 5.15).

The following section explores this trade-off explicitly: What happens when environmental objectives are prioritized more strongly in the model? Can carbon pricing or direct emission minimization shift routing behavior away from cost-driven DPS? To explore this trade-off, sensitivity analyses investigate whether stronger environmental objectives, such as carbon pricing or direct emissions minimization, can shift routing behavior away from cost-driven DPS decisions.

7.8. Sensitivity Analysis: Emissions-Oriented Routing Behavior

Building on to the previous analysis, two additional sensitivity analyses are conducted to evaluate how environmental cost considerations influence routing behavior. The first approach applies a scalarization method to assess cost-emissions trade-offs under varying carbon pricing (Pereira et al., 2022). The second analysis applies the ε-constraint method, where emissions are minimized directly under logistics cost caps (Ghorbani et al., 2020). Together, these analyses provide insights into whether environmental incentives encourage shifts between DPS and two-echelon (2E) flows.

7.8.1. Impact of Emission Pricing on Share of DPS

This analysis evaluates how varying levels of carbon pricing affect the share of DPS within the hybrid distribution network. All results correspond to configuration C3, in which the model is free to choose the most volume-efficient shipment sizes without enforcing a FTL threshold. This configuration choice allows for more flexible vehicle utilization and reflects an unconstrained optimization of shipment consolidation.

Using the scalarization method described in subsection 5.1.5, the model optimizes a composite objective function that combines logistics cost and emissions:

$$Z_{\text{total}} = \bar{Z}_{\text{cost}} + \lambda \cdot \bar{Z}_{\text{emissions}} \tag{7.1}$$

The scalarization parameter λ represents a hypothetical carbon price (€/kg CO_2) and is varied from 0.00 to 4.00. For each value of λ, the model produces a new optimal routing configuration, balancing cost and emissions under the given weight.

Figure 7.8 shows how the share of DPS in total weekly demand evolves as a function of the carbon price for all four demand scenarios: Category X (S1 and S2) and Category Y (S3 and S4). The figure clearly illustrates that, across all scenarios, the DPS share remains virtually constant, even as carbon pricing increases significantly.

This suggests that under the current network and parameter configuration, the inclusion of carbon costs in monetized form in the objective function has no meaningful influence on the model's routing decisions. The minimal change in DPS share indicates that the cost savings from 2E flows continue to outweigh any additional emissions penalties, even at higher λ values.

The findings underscore a key insight: carbon pricing in the objective alone is insufficient to drive structural changes in routing or shipment modality. This suggests that complementary policy measures or operational constraints (e.g., stricter emissions caps or sustainability targets) would be needed to shift distribution behavior toward more direct and potentially more sustainable shipment options.

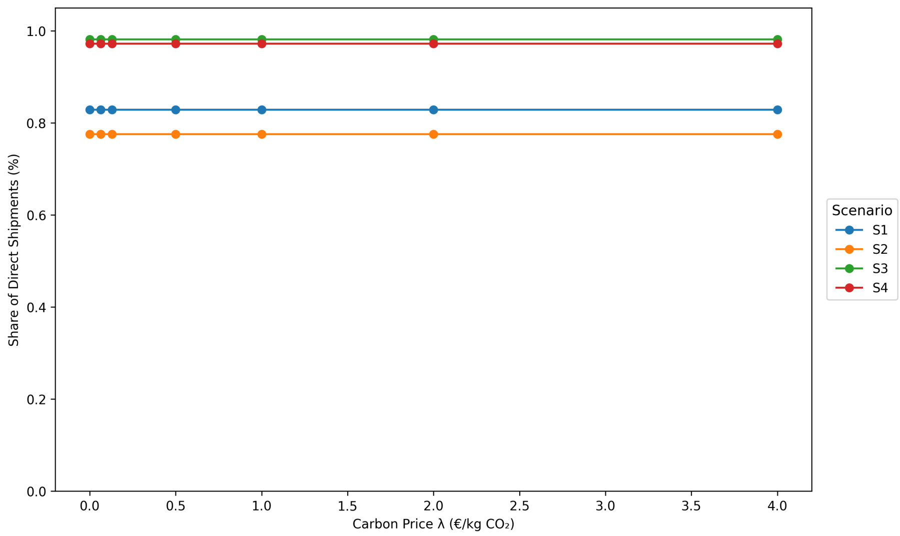

Figure 7.8: Share of DPS Across Carbon Price Levels (λ) for All Scenarios under Configuration C3

7.8.2. ε-constraint: Emissions Minimization Under Cost Caps

To complement the scalarization-based sensitivity analysis, an additional experiment was conducted using the ε-constraint method. Instead of minimizing a weighted combination of logistics cost and emissions in €, this approach directly minimizes total emissions in kg CO_2, subject to an upper bound (ε) on total logistics cost (in €). This framing allows sustainability to take precedence in the objective function, while ensuring economic feasibility remains enforced through explicit cost caps.

The method was applied on two representative demand scenarios, S2 and S4, each evaluated over the three representative weeks. The experiments are run in configuration C3, where DPS is permitted without a FTL threshold, allowing inefficient low-fill routes if chosen. To prevent cost from dominating decisions, several cost caps were set, ranging from the original cost-optimal value from the initial optimization, up to generous allowances of €X.

The resulting routing choices, values of the ε-constraint, and performance metrics are shown in Table E.13 (see section E.9). The relative trade-offs are visualized in Figure 7.9, where all normalized emissions and corresponding DPS shares across the scenarios are aggregated into a single Pareto front. Different colors denote different weeks, and a trend line is fitted across all points.

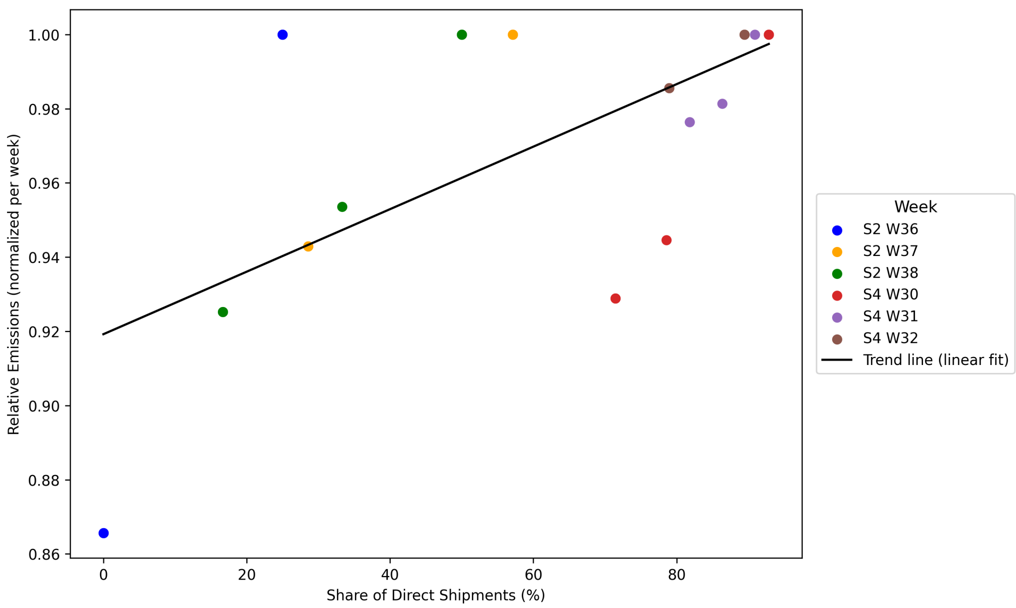

Figure 7.9: Aggregated Relative Emissions vs. DPS Share (Pareto Front across S2 and S4)

A striking outcome is visible: minimizing emissions does not inherently favor higher DPS usage. In fact, across both scenarios, increasing the budget for emissions reduction always led to lower DPS shares. For example, in scenario S4, week 30, allowing emissions to be minimized under a relaxed cost cap of €X led to a reduction in emissions, but also a decrease in DPS share, from x out of X shipments to x out of X shipments. This pattern is visible in the overall trend line, which slopes downward: as relative emissions decrease, DPS share also declines.

This counterintuitive result is explained by structural trade-offs in the model. Two-echelon (2E) routes, although involving intermediate handling and longer absolute trasnport distances, benefit from higher VFR and shorter per-unit delivery distances. The optimization model thus favors consolidated shipments under emission minimization, often reverting back to 2E flows even when the budget allows full cost flexibility. In this context, direct routes can be less efficient on a per-unit emission basis due to lower vehicle utilization.

These findings highlight the need for careful evaluation of sustainability strategies. Simply prioritizing emissions in the objective function does not yield increased use of DPS or more decentralized distribution. Instead, the results demonstrate that emission-optimal decisions often align with high VFR and consolidated routing structures, reinforcing the importance of analyzing emissions per unit delivered rather than just absolute emissions or structural distribution modes.

7.8.3. Synthesis of Findings on Emissions-Oriented Routing Behavior

The sensitivity analyses show that when emissions are prioritized over cost, either through carbon pricing or direct emissions minimization, the model does not seem to favor a bigger share of DPS, or even shifts towards more consolidated 2E flows. This stands in contrast to the demand-based break-even analysis, where DPS emerged as cost-optimal at lower volumes under relaxed thresholds. The results highlight a key trade-off: while cost-efficient solutions favor decentralized routing through DPS, emissions-based optimization favors consolidation and higher VFR, reinforcing the conventional 2E structures.

7.9. Conclusion

To answer SQ6: *How does the integration of direct plant shipments (DPS), under different configuration thresholds, affect logistics and environmental costs, service levels, and stock allocation within hybrid distribution scenarios?*, this concluding section synthesizes the outcomes of the experiments.

Across all scenarios, enabling DPS significantly improves logistics cost performance, particularly in configurations that allow more flexible shipment volumes. The magnitude of improvement is strongly influenced by the volume and transport distance of the flows:

- For high-volume, long-haul shipments (e.g., Category Y from Plant 2), the model consistently prioritizes DPS routes, achieving up to 43% cost reductions in configurations C1-C3. Environmental performance also improved, with average emission reductions of 11.7% compared to the baseline.
- For lower-volume, short-distance flows (e.g., Category X from Plant 1), logistics costs decreased by a smaller margin (12.6% on average), and emissions results were nuanced, sometimes even increasing due to inefficient VFR with DPS usage.

Notably, lowering the FTL threshold from X to 0 pallets increased DPS adoption but often worsened environmental performance. Emission-optimal routing was consistently associated with higher vehicle utilization through two-echelon flows. This counterintuitive but robust finding, confirmed through scalarization and ε-constraint sensitivity analyses, demonstrates that greener routing decisions do not align with higher DPS shares under current system conditions. Without structural changes or additional policy levers, such as route-level emission caps, minimum VFR requirements, or incentives for decentralized delivery, carbon-focused objectives alone may reinforce centralized distribution patterns rather than shift them. As such, cost-optimal and sustainability-optimal routing structures may diverge in hybrid networks without targeted policy support.

From a service perspective (defined as the proportion of fulfilled demand originating from a given shipping location), integrating DPS led to measurable improvements in product availability and delivery reliability across all scenarios. Average service level gains ranged from +1.0 to +1.5 percentage points, with the highest values observed in flexible configurations and high-volume demand weeks. This improved responsiveness translates into commercial benefits on NOS results. The majority of gains occurred upon enabling DPS in the first place (C1), the subsequent threshold reductions delivered diminishing returns.

In terms of inventory and cash, DPS causes a significant reallocation of stock from the DC to the plant. For example, in the Plant 2-Category Y case, over 90% of volume was fulfilled with DPS in configuration C1, reducing storage needs at the DC and savings of 4.8% in weekly logistics costs in a representative case. This shift supports better working capital efficiency, but also implies greater responsibility for forecasting and planning at plant-level. Products with volatile or low demand remain less suitable for such strategies, as the risk of misalignment increases when DC buffering is eliminated.

Finally, the break-even analysis revealed that optimal DPS decisions are route-specific and highly dependent on distance. The volume required to justify cost-efficient DPS ranged from x to X FP, scaling with the plant-customer distance. This suggests that a fixed DPS threshold is suboptimal, and that customized, route-level thresholds would better align operational decisions with cost efficiency. On the other hand, a key insight emerges when comparing these break-even thresholds with the results of the emissions-based sensitivity analyses: even when the model identifies DPS as the most cost-efficient choice under current volume and distance conditions, a stronger valuation of emissions in the objective function causes the model to revert back to 2E flows. This shift underscores a fundamental trade-off in hybrid networks: carbon-optimized routing does not necessarily align with the cost-optimal use of DPS. Instead, emissions-based incentives systematically favor consolidated, high-fill routes, highlighting the need for deliberate coordination between cost and sustainability goals in designing effective DPS policies.

In summary, the integration of DPS in hybrid distribution networks offers substantial cost and service improvements under the right conditions. However, its environmental benefits are nuanced and only realized when shipment volumes allow efficient vehicle utilization. To maximize DPS impact, policies should account for volume, distance, and emission efficiency, possibly through differentiated thresholds and additional constraints. A cross-functional approach is required to balance cost savings, service improvement, environmental goals, and inventory agility, aligning operational decisions with strategic supply chain objectives.

8 Conclusion

This research examined the effect of integrating Direct Plant Shipments (DPS) alongside conventional two-echelon distribution as a hybrid distribution strategy, focused on a case study of the supply chain of Procter & Gamble in France and The Benelux. The central aim was to evaluate how such integration impacts three main KPIs, cost (including logistics and environmental cost), service (service level), and cash (stock allocation efficiency), in a cross-country distribution network, with a structured Research & Design approach. Using a flow-based adaptation of the Two-echelon Vehicle Routing Problem (2E-VRP), a hybrid distribution model was developed and applied on a selection of high-demand customers in the case study, using real historical data. The model jointly optimizes logistics cost and emissions, with post-evaluation of the impact on service level performance and stock allocation to different storage locations.

This chapter will discuss the conclusion of the sub-questions and the final main findings to formulate an answer to the main research question:

What is the impact of integrating direct plant shipments in an existing two-echelon distribution network - as a hybrid strategy - on logistics and environmental costs, service level performance, and stock allocation efficiency in a cross-border supply chain?.

The key discussion points, scientific contributions, and future research directions are elaborated in chapter 9. Based on these findings, case-specific recommendations for P&G's hybrid distribution strategy are also included in this chapter.

8.1. Sub-questions

Answering the sub-questions throughout the report provided a structured foundation for addressing the main research question. Each sub-question contributed a specific insight that, together, supported the development and evaluation of the hybrid distribution model.

Two-echelon distribution systems offer advantages in terms of shipment consolidation and network efficiency, but they often come at the cost of additional logistics handling, flexibility and pressure on the supply chain. Single-echelon shipments, on the other hand, enhance responsiveness and reduce lead times but are only cost-effective when customer demand volumes are high enough to justify full loaded transport.

A suitable modeling approach identified from the literature is the arc-based formulation of the two-echelon vehicle routing problem (2E-VRP), which supports both routing and flow-based decision-making in multi-layered supply chains. This structure was adapted to evaluate hybrid distribution networks by allowing both direct plant-to-customer shipments (DPS) and two-echelon flows via a distribution center. The model design draws on theoretical requirements such as vehicle capacity constraints, cost minimization, and customer eligibility conditions, while enabling flexibility in representing real-world configurations and policy scenarios. This approach proved well-suited for capturing the complexities of hybrid distribution and comparing alternative strategies in a structured, data-driven way.

P&G's distribution in the FBNL region currently relies on its two-echelon flows via the DC 1. The analysis shows that high-volume, stable-demand customers, especially those farther from the DC like Customer 1 and Customer 5, are strong candidates for DPS, offering cost and service benefits. However, the main limitation remains the loss of consolidation, which can reduce VFR and increase emissions for smaller shipments.

To ensure the model's applicability, several key design requirements were embedded in its formulation. These include hard constraints such as shipment volume thresholds, single-origin shipment rules, and fulfillment of all weekly demand. In addition, soft requirements were addressed where feasible, such as minimizing handling and supporting a scalable DPS framework. The formulation also incorporates real-world factors such as transport cost variability by distance, product allocation logic, and service level measures, ensuring that the model reflects the operational realities of P&G's hybrid distribution network.

The resulting model is a flow-based adaptation of the 2E-VRP, scenario-based and highly configurable, allowing structured experimentation with DPS settings and threshold rules. It supports sensitivity analyses by allowing changes in cost-emission trade-offs, pricing for emissions, and customer allocation strategies, making it as a useful tool for evaluating the impact of internal policies under different demand configurations. Its flexible structure facilitates the assessment of logistics cost, service level, and emissions under a variety of realistic supply chain setups.

Implementing a hybrid distribution strategy that integrates DPS with conventional two-echelon flows can significantly enhance logistics performance and service reliability, particularly for high-volume, long-distance customers. However, the benefits of DPS are conditional: while logistics costs decrease substantially and service levels improve when shipment volumes justify full truckloads, environmental performance often deteriorates due to lower vehicle fill rates in smaller or shorter-distance flows. Stock reallocation from the DC to the plant improves working capital efficiency but increases the demand for accurate forecasting at the source. Additionally, the analysis shows that a fixed DPS threshold is suboptimal. Route-specific characteristics, especially volume and distance, should guide DPS eligibility. But even in this cost-optimal solutions with DPS, the model demonstrates that environmental goals are better served by two-echelon flows in the absence of policy levers such as route-level emission caps or minimum vehicle fill requirements. These findings highlight a fundamental trade-off in hybrid networks: emission-optimal routing does not necessarily align with cost-optimal DPS decisions. To maximize the impact of DPS, future strategies should incorporate differentiated thresholds and environmental constraints, balancing cost, service, and sustainability priorities. A cross-functional approach is essential to align operational policies with strategic supply chain objectives.

8.2. Main Findings

The model results show that integrating DPS leads to a notable reduction in logistics cost, particularly for high-demand customers with longer transport distances from the plant to the customer. Within a hybrid network structure, this selective use of DPS complements the efficiency of two-echelon flows by offering a more responsive alternative where conditions allow. In configurations C1-C3, logistics costs were reduced by up to 43% for high-volume, long-haul flows, such as Category Y from Plant 2. For lower-volume, short-distance flows such as Category X from Plant 1, logistics cost reductions were more modest, averaging around 12.6%. These findings confirm that the cost efficiency of DPS strongly depends on shipment volume and transport distance.

The environmental performance of DPS is more nuanced. While average emissions decreased by up to 11.7% in high-volume, long-distance cases, particularly for Category Y, emissions did not consistently improve across all configurations. In fact, for Category X, emissions sometimes increased, especially in configurations with lower full-truckload (FTL) thresholds, due to reduced VFR. Notably, lowering the FTL threshold from to 0 pallets led to greater DPS adoption but often worsened environmental outcomes. Emission-optimal routing was consistently achieved through two-echelon flows, which offer higher vehicle utilization and lower emissions per unit delivered. This counterintuitive but robust finding, confirmed through scalarization and ε-constraint sensitivity analyses, demonstrates that greener routing decisions do not align with higher DPS shares under current system conditions. Without additional policy levers, such as route-level emission caps, minimum fill rate requirements, or incentives for decentralized delivery, minimizing emissions alone is unlikely to promote more decentralized logistics. As a result, cost-optimal and sustainability-optimal routing structures may diverge, and environmental goals cannot be achieved through carbon pricing or cost minimization alone.

From a service perspective, integrating DPS consistently improves fulfillment performance across all scenarios. Average service level gains ranged from +1.0 to +1.5 percentage points, with most of the improvement realized in the shift from the baseline configuration (C0) to the first DPS-enabled setup (C1). Further relaxing the FTL threshold beyond C1 led to only marginal additional gains, indicating diminishing returns. This pattern suggests that enabling DPS enhances responsiveness and reduces dependency on DC inventory, particularly in the initial phase of implementation, while further adjustments offer limited incremental benefit.

In terms of inventory and working capital, DPS offers a structural opportunity to reduce storage costs by shifting stock from DC's to plants. This shift alleviates pressure on DC capacity and avoids storage costs, particularly because plant storage is available and P&G owned. In cases where DPS is fully allowed and integrated, weekly logistics cost reductions of around 4.8% can be achieved, as illustrated in an example scenario of Category Y from Plant 2. However, this reallocation also shifts forecasting responsibility upstream and may reduce responsiveness for SKUs with volatile or low demand. To fully capture the benefits of DPS without compromising service levels, it is essential to segment SKUs based on demand stability and tailor stock allocation strategies accordingly.

A break-even analysis confirmed that the cost-optimal viability of DPS is highly route-specific. The volume required to justify DPS in a cost-optimal solution varied between floor positions (FP), depending on the plant-to-customer distance. This suggests that a uniform DPS threshold for cost minimization is suboptimal; instead, targeted, route-level thresholds offer greater cost efficiency. However, such thresholds, while economically justified, could undermine vehicle utilization and increase emissions intensity when shipment volumes are too low to maintain high fill rates.

In summary, the integration of DPS improves logistics cost and service level under the right conditions, particularly for high-volume, long-distance customers. However, its environmental benefits are conditional and not guaranteed without additional constraints to safeguard vehicle efficiency. These findings underscore the value of a hybrid distribution strategy that selectively combines DPS and two-echelon flows based on route- and customer-specific trade-offs, rather than fully shifting to either mode. A nuanced, policy-aware application of DPS is necessary to align operational decisions with broader cost-efficient, sustainability, service performance and inventory management goals.

9

Discussion and Recommendations

This chapter firstly reflects on the scientific contribution and key limitations of this research. Next, opportunities for model improvement are discussed and areas for case-specific recommendations to the problem owner P&G are outlined. While the results offer valuable insights on the integration of Direct Plant Shipments (DPS) into P&G's hybrid distribution strategy, several assumptions and simplifications made in the modeling process ask for further consideration.

9.1. Scientific Contribution

This research contributes to the academic field of supply chain management by addressing several key gaps identified in the literature. First, while hybrid distribution strategies have been conceptually proposed, few studies have rigorously modeled and empirically evaluated the integration of direct plant shipments in a two-echelon network across multiple key performance indicators. Where cost minimization is often the sole objective in such models, this study explicitly includes the performance on emissions and service level, thus expanding the scope of performance evaluation. By developing a novel flow-based adaptation of the 2E-VRP that jointly optimizes economic and environmental costs, evaluates the associated change in service level performance, and extends the experiments with post-optimization analyses, this study complements traditional modeling approaches beyond cost-focused frameworks. Using a case study of real-world operations in the FMCG sector, the model demonstrates how hybrid distribution structures affect these multiple objectives under realistic demand and routing conditions.

A central insight emerging from this multi-objective modeling is the persistent trade-off between cost efficiency and environmental performance in hybrid routing decisions. Although DPS often reduces logistics costs, it can increase the total emissions when vehicle fill rates are low. Sensitivity analyses reveal that even when emissions are highly valuated in monetized form or explicitly minimized as the objective, the model tends to still prioritize centralized two-echelon flows due to their superior vehicle utilization and consolidation. This finding indicates that the emissions inefficiency of underutilized vehicles can outweigh the logistics benefits of direct flows. This challenges the assumption that expanding direct shipping inherently supports sustainability goals and highlights a saturation point beyond which further adoption may become counterproductive. The research therefore adds to the theoretical understanding of how hybrid routing interacts with environmental policy goals, suggesting that more explicit integration of emissions, via hard constraints or adaptive policies, may be needed to ensure alignment between economic and ecological objectives.

Secondly, while most prior research has focused on urban or last-mile logistics, this thesis shifts the lens to regional, cross-border distribution within a multinational FMCG network. By evaluating the effects of transport distance and volume thresholds on routing decisions, it enriches theoretical perspectives on hybrid logistics at scale. The model also embeds practical constraints such as SKU eligibility, full-truckload (FTL) thresholds, and customer-specific policies, demonstrating how abstract VRP-based formulations can be adapted for a scenario analysis in real business contexts.

Finally, the case study at P&G provides empirical evidence that hybrid configurations can outperform conventional strategies under specific conditions. It offers strategic insights into when and how direct shipping adds value, and identifies operational trade-offs related to inventory allocation, emissions intensity, and service responsiveness. In doing so, this work bridges the gap between theoretical optimization and real-world supply chain design, contributing to both academic advancement and managerial decision-making.

9.2. Limitations of the Current Model and Assumptions

The model developed in this research is a flow-based adaptation of the two-echelon vehicle routing problem (2E-VRP), intentionally abstracting away from vehicle-level routing, synchronization, and vehicle counting. It assumes unlimited vehicle availability at both the plant and the distribution center (DC), a simplification justified by the involvement of third-party logistics (3PL) providers who manage P&G's transport execution and charge per shipment leg rather than per vehicle. Although two-echelon shipments require two vehicle trips, this is not reflected as a separate KPI, as transport costs are modeled per leg and fixed vehicle start costs are not included. However, omitting specific vehicle assignment, returns, and time synchronization limits the model's operational realism, particularly in contexts with constrained vehicle availability, loading schedules, or driver hours. Incorporating these factors could improve accuracy in high-utilization settings and help identify potential bottlenecks in vehicle coordination, especially during peak periods or in geographies with limited 3PL flexibility.

The current model does not include delivery time windows or lead time as a factor in the optimization or as a KPI. This means that while logistics costs and emissions are optimized, the time responsiveness of the supply chain is not explicitly assessed. DPS can offer shorter lead times due to fewer handling steps, which can improve service performance from a customer perspective. On the other hand, if DPS results in less frequent deliveries to customers, this could lead to increased holding time or uncertainty at customer sites.

The optimization was performed over three consecutive weeks that were statistically selected to represent typical operational variability. While this approach ensures a solid baseline for evaluating model performance, it does not fully capture long-term dynamics such as seasonal demand shifts, promotional effects, or supply chain disruptions. In the FMCG sector, such factors can influence the feasibility and attractiveness of DPS strategies. Expanding the time horizon could help assess how well the model generalizes under varying demand conditions.

While this study incorporates emissions into the objective function through monetized carbon pricing, this approach aligns with P&G's current incentive structure, where total cost minimization remains the dominant driver. However, results show that even at elevated carbon prices, routing decisions remain largely unchanged, revealing that pricing alone does not meaningfully influence route selection under the current network configuration. To address this, the ε-constraint method was used to isolate the impact of pure emissions minimization. These results suggest that stronger integration of emissions into the decision logic, such as hard constraints on emissions per route or minimum vehicle fill requirements, may be required to achieve meaningful sustainability improvements.

Furthermore, although the model assumes unconstrained shipment volumes and no enforced FTL threshold in some configurations, environmental performance was still strongly driven by vehicle utilization. This indicates that emissions inefficiency is structurally embedded in certain direct shipping patterns, particularly for small or fragmented flows. The current model does not prevent low-fill shipments when they are cost-efficient, even if they are environmentally suboptimal. As a result, threshold-based policies should be refined to reflect both cost and environmental trade-offs, ideally through adaptive or route-specific criteria.

The analysis shows that DPS can reduce inventory at the DC and lower storage costs by shifting fulfillment upstream to the plant. In P&G's setting, where plant storage capacity typically exceeds that of the DC, this shift is operationally feasible from a capacity standpoint. However, it increases reliance on accurate decentralized forecasting and tighter planning at the plant level. For low-volume or highly variable customer demand, the removal of the DC buffer may reduce service reliability or increase the risk of stockouts. The current model does not capture these operational trade-offs, nor the potential implications for planning complexity, upstream workload, or overall system resilience.

9.3. Directions for Future Research and Model Development

The limitations outlined above offer several pathways for scientific advancement. This section proposes how future research could build on the current model to enhance realism, applicability, and decision support capacity.

First, future research could explore the integration of time-based routing constraints and vehicle synchronization into hybrid distribution models to assess their impact on feasibility and system performance under more realistic operational conditions. This includes modeling vehicle availability limits, return trips, and multi-stop tours, as well as constraints related to driver working hours, loading windows, and coordination across echelons. Such extensions would enable a more granular analysis of temporal bottlenecks and resource allocation trade-offs, especially during peak demand periods or in regions with limited 3PL flexibility.

In addition to time-based routing, subsequent work should incorporate lead times and delivery frequency into the optimization framework to better evaluate the service-related trade-offs of hybrid distribution strategies. These temporal dimensions affect not only customer responsiveness but also the regularity of deliveries and inventory dynamics at customer sites. Capturing them would allow for more balanced planning between cost, emissions, and service performance.

Moreover, expanding the temporal scope of the analysis to cover a broader set of representative weeks or full-year demand cycles would help capture long-term variability, including seasonal effects and promotional campaigns. While the current selection of statistically representative weeks offers a robust snapshot, testing the model under more extreme or volatile conditions would provide deeper insights into the consistency and robustness of the proposed strategies.

To move beyond the current cost-driven framing of emissions, future research could further explore multi-objective optimization frameworks that treat cost and environmental impact as distinct but equally important objectives. This would support more balanced trade-off analyses and help align model outputs with long-term sustainability goals. Additionally, the exploration of alternative policy mechanisms, such as absolute emission caps or efficiency thresholds, could enrich the understanding of how regulatory or corporate targets shape optimal logistics configurations.

From an inventory management perspective, the decentralization of stock from DCs to plants raises questions about forecasting accuracy, planning responsibilities, and upstream operational dynamics. Future research could focus on developing simulation-based or integrated inventory-transport models to evaluate how such structural shifts affect system resilience and workload distribution under uncertainty, beyond just physical storage concerns.

Finally, future work should consider how to refine DPS threshold policies to align more closely with both cost efficiency and environmental goals. Rather than applying a single static threshold, adaptive strategies that incorporate unit-based emissions constraints and customer-specific thresholds based on distance and shipment profiles could lead to more sustainable and nuanced decisions. This would help avoid promoting low-fill but cost-efficient shipments that are misaligned with broader environmental objectives.

9.4. Case-Specific Strategic Recommendations for P&G

This section presents specific recommendations obtained from the findings of this research, to provide the problem owner P&G with valuable insights for implementing DPS within their hybrid distribution strategy.

Pursue structural DPS implementation with high-potential customers: This research has provided a broad eligibility assessment across the full BNL customer set, identifying customers with favorable demand volumes, routing conditions, and cost-emission profiles for DPS. As a next step, P&G should focus on strengthening collaboration with selected high-potential customers, such as Customer 1 and Customer 5 for Category Y demand from Plant 2, to explore opportunities for more structural and consistent use of DPS. Engaging these customers in tailored agreements or service-level discussions can support more regular DPS flows and help establish best practices for wider rollout across the network.

Limit DPS scope to plant-produced SKUs for operational simplicity: Restricting DPS to SKUs that are both produced and stored at the plant can streamline operations and reduce complexity. Including stored SKUs that require internal transfers or cross-warehouse coordination introduces additional logistical challenges, often for limited additional benefit. A focused DPS scope ensures smoother integration with production schedules, simplifies outbound handling, and still captures a substantial share of the cost and emissions savings potential.

Stimulate AOV-based ordering to support DPS opportunities: To enable better planning and execution of DPS, efforts could be made to encourage customers to place more structural Advance Order Volume (AOV) orders specifically for SKUs produced at the plant. These orders, typically used for mixed shipments from the DC, could be adapted to fit DPS planning needs when aligned with plant-produced product flows. Tailored incentives, such as differentiated discounts for solely DPS-eligible SKUs, may help increase customer adoption. This approach leverages DPS-related cost savings to create a mutually beneficial arrangement that supports operational efficiency.

Differentiate DPS policies based on customer-specific characteristics and environmental thresholds: Rather than applying uniform DPS thresholds, P&G could consider differentiated policies based on break-even volumes, routing efficiency, and environmental impact. Customers closer to the plant or with stable mid-to-high demand may qualify for lower DPS thresholds, provided that vehicle fill rates remain sufficient to avoid emissions inefficiencies. To ensure alignment with sustainability goals, threshold differentiation should be guided by combined economic and ecological criteria, potentially incorporating minimum fill rate constraints or emissions caps into DPS decision rules.

Assess and plan plant storage requirements for DPS volumes: As DPS shifts fulfillment upstream, sufficient storage and handling capacity at the plant becomes a key enabler of success. Although capacity may not be a structural constraint in P&G's current situation, planning for temporary buffers, dynamic volume shifts, and order consolidation is critical to prevent congestion or delays. Ensuring that physical space, labor planning, and outbound processes are adapted to accommodate DPS flows will help maintain reliability and avoid undermining the expected efficiency gains.

Integrate DPS logic into forecasting and planning processes: Decentralized fulfillment models such as DPS increase dependency on accurate local forecasts and coordinated planning. Without the buffering effect of a central DC, mismatches between demand and supply can have a greater impact on service level performance. Embedding DPS considerations into forecasting systems and strengthening alignment between demand planning and transport decisions can mitigate this risk and ensure more reliable order fulfillment.

Further assess DPS feasibility for selected customers from Plant 1: While the analysis shows that DPS is most cost-attractive for Category Y shipments from Plant 2, opportunities also exist for selected customers served by Plant 1. For the customers with relatively high and stable volumes and with favorable location and routing characteristics, like Customer 5 and Customer 4, DPS may offer a viable alternative to conventional distribution. Specific case-by-case evaluation of demand patterns and ordering behavior could inform whether structural DPS agreements would be feasible and beneficial in these cases.

References

ABC Supply Chain. (2023). *Safety stock formula: How to calculate it? (excel included)* [Accessed: 2025-04-12]. https://abcsupplychain.com/safety-stock-formula-calculation/

Altekar, R. V. (2023). *Supply chain management: Concepts and cases* (2nd). PHI Learning Private Limited.

Anderluh, A., Hemmelmayr, V. C., & Nolz, P. C. (2017). Synchronizing vans and cargo bikes in a city distribution network. *Central European Journal of Operations Research*, *25*, 345–376. https://doi.org/10.1007/s10100-016-0441-z

Anderluh, A., Nolz, P. C., Hemmelmayr, V. C., & Crainic, T. G. (2021). Multi-objective optimization of a two-echelon vehicle routing problem with vehicle synchronization and 'grey zone' customers arising in urban logistics. *European Journal of Operational Research*, *289*, 940–958. https://doi.org/10.1016/j.ejor.2019.07.049

Azizi, V., & Hu, G. (2020). Multi-product pickup and delivery supply chain design with location-routing and direct shipment. *International Journal of Production Economics*, *226*, 107648. https://doi.org/10.1016/J.IJPE.2020.107648

Chiang, A. H., & Huang, M. Y. (2021). Demand-pull vs supply-push strategy: The effects of organizational structure on supply chain integration and response capabilities. *Journal of Manufacturing Technology Management*, *32*, 1493–1514. https://doi.org/10.1108/JMTM-08-2020-0324

Clarke, G., & Wright, J. (1964). Scheduling of vehicles from a central depot to a number of delivery points. *Operations Research*, 568–581.

Desmet, B. (2018). *The supply chain triangle of service, cost and cash – essentials* [Accessed May 27, 2025]. Solventure. https://www.bramdesmet.com/wp-content/uploads/2018/03/supplychaintriangle-essentials.pdf

Dondo, R., Méndez, C. A., & Cerdá, J. (2011). The multi-echelon vehicle routing problem with cross docking in supply chain management. *Computers & Chemical Engineering*, *35*, 3002–3024. https://doi.org/10.1016/J.COMPCHEMENG.2011.03.028

Fishani, B. A., Mahmoodirad, A., Niroomand, S., & Fallah, M. (2022). Multi-objective location-allocation-routing problem of perishable multi-product supply chain with direct shipment and open routing possibilities under sustainability. *Concurrency and Computation: Practice and Experience*, *34*. https://doi.org/10.1002/cpe.6860

Garrido, A., Cárdenas-Barrón, L. E., Buitrago, O. Y., & Álvarez-Pomar, L. (2025). A multi-echelon globalized agro-industrial supply chain under conditions of uncertainty: A two-stage fuzzy-possibilistic mixed-integer linear programming model. *Expert Systems with Applications*, *270*. https://doi.org/10.1016/j.eswa.2025.126569

Ghorbani, N., Niknam, T., Ghavidel, S., Anane, M., & Aghaei, J. (2020). Multiobjective multi-verse optimization algorithm to solve combined economic emission dispatch. *Applied Soft Computing*, *87*, 105998. https://doi.org/10.1016/j.asoc.2019.105998

Grosso, R., Munuzuri, J., Escudero-Santana, A., & Barbadilla-Martin, E. (2018). Mathematical formulation and comparison of solution approaches for the vehicle routing problem with access time windows. *Complexity*, 1–10.

Guastaroba, G., Speranza, M., & Vigo, D. (2016). Intermediate facilities in freight transportation planning: A survey. *Transportation Science 50(3):763-789*.

Guo, Y., Liu, F., Song, J. S., & Wang, S. (2024). Supply chain resilience: A review from the inventory management perspective. https://doi.org/10.1016/j.fmre.2024.08.002

Helo, P., & Hao, Y. (2022). Artificial intelligence in operations management and supply chain management: An exploratory case study. *Production Planning and Control*, *33*, 1573–1590. https://doi.org/10.1080/09537287.2021.1882690

Hosseini-Nasab, H., Nasrollahi, S., Fakhrzad, M. B., & Honarvar, M. (2023). Transportation cost reduction using cross-docks linking [Case study]. *Journal of Engineering Research (Kuwait)*, *11*. https://doi.org/10.1016/j.jer.2023.100015

International Carbon Action Partnership. (2024). Eu emissions trading system (eu ets) [Accessed: 2025-05-21].

Ivanov, D., & Dolgui, A. (2020). A digital supply chain twin for managing the disruption risks and resilience in the era of industry 4.0 [Received 05 Apr 2019, Accepted 09 May 2020, Published online: 21 May 2020]. *International Journal of Production Research*, *58*(16), 775–788. https://doi.org/10.1080/09537287.2020.1768450

Ma, H., Miao, Z., Lim, A., & Rodrigues, B. (2011). Crossdocking distribution networks with setup cost and time window constraint. *Omega*, *39*, 64–72. https://doi.org/10.1016/j.omega.2010.03.001

Mahmood, S., & Kess, P. (2014). An overview of demand management through demand supply chain in fashion industry. *INTERNATIONAL JOURNAL OF MANAGEMENT SCIENCE AND BUSINESS ADMINISTRATION*, *2*, 7–19. https://doi.org/10.18775/ijmsba.1849-5664-5419.2014.212.1001

Marchuk, V. Y., Sergiy, G., Karpun, O., & Garmash, O. (2020). Intellectualization of logistics and supply chain management 2 2020. *Article in Electronic Scientific Journal Intellectualization of Logistics and Supply Chain Management*, *1*. https://doi.org/10.46783/smart-scm/2020-2

Marques, G., Sadykov, R., Dupas, R., & Deschamps, J.-C. (2022). A branch-cut-and-price approach for the single-trip and multi-trip two-echelon vehicle routing problem with time windows. *Transportation Science 56(6):1598-1617*. https://doi.org/10.1287/trsc.2022.1136

Min, S., Zacharia, Z. G., & Smith, C. D. (2019). Defining supply chain management: In the past, present, and future. *Journal of Business Logistics*, *40*, 44–55. https://doi.org/10.1111/jbl.12201

Mohammadi, Z., Barzinpour, F., & Teimoury, E. (2020). Designing sustainable supply chain network by considering direct and indirect shipment: Evidence from food industry. *Decision Science Letters*, *9*, 323–336. https://doi.org/10.5267/j.dsl.2020.5.003

Mor, A., & Speranza, M. G. (2022, July). Vehicle routing problems over time: A survey. https://doi.org/10.1007/s10479-021-04488-0

Musa, R., Arnaout, J. P., & Jung, H. (2010). Ant colony optimization algorithm to solve for the transportation problem of cross-docking network. *Computers and Industrial Engineering*, *59*, 85–92. https://doi.org/10.1016/j.cie.2010.03.002

Nielsen, P., Dahanayaka, M., Perera, H. N., Thibbotuwawa, A., & Kilic, D. K. (2024, September). A systematic review of vehicle routing problems and models in multi-echelon distribution networks [More recent paper+ research gaps]. https://doi.org/10.1016/j.sca.2024.100072

Oliveira, B., Ramos, A. G., & Sousa, J. P. D. (2022). A heuristic for two-echelon urban distribution systems. *Transportation Research Procedia*, *62*, 533–540. https://doi.org/10.1016/j.trpro.2022.02.066

Pereira, J. L. J., Oliver, G. A., Francisco, M. B., Cunha, S. S., & Gomes, G. F. (2022, June). A review of multi-objective optimization: Methods and algorithms in mechanical engineering problems. https://doi.org/10.1007/s11831-021-09663-x

Potoczki, T., Holzapfel, A., Kuhn, H., & Sternbeck, M. (2024). Integrated cross-dock location and supply mode planning in retail networks. *International Journal of Production Economics*, *276*. https://doi.org/10.1016/j.ijpe.2024.109349

Prashar, A. (2023). Supply chain sustainability drivers for fast-moving consumer goods (fmcg) sector: An indian perspective. *International Journal of Productivity and Performance Management*, *72*, 2397–2419. https://doi.org/10.1108/IJPPM-10-2021-0582

Procter & Gamble. (2025). Structure & governance: Corporate structure [Accessed: 2025-02-24]. https://us.pg.com/structure-and-governance/corporate-structure/

Procter & Gamble. (2023). Who we are [Accessed: 2025-02-13]. https://us.pg.com/who-we-are/

Rebs, T., Thiel, D., Brandenburg, M., & Seuring, S. (2019). Impacts of stakeholder influences and dynamic capabilities on the sustainability performance of supply chains: A system dynamics model. *Journal of Business Economics*, *89*, 893–926. https://doi.org/10.1007/s11573-019-00940-7

Rodríguez, F., Delgado, O., & Baldino, C. (2020). Co_2 emissions from trucks in the eu: An analysis of the heavy-duty co_2 standards baseline data [Accessed: 2025-05-21].

Ross, D. F. (2015). *Distribution planning and control: Managing in the era of supply chain management*. Springer Science+Business Media New York. https://doi.org/10.1007/978-1-4899-7578-2_11

Sargent, R. G. (2010). Verification and validation of simulation models. In B. Johansson, S. Jain, J. Montoya-Torres, J. Hugan, & E. Yücesan (Eds.), *Proceedings of the 2010 winter simulation conference* (pp. 166–183). IEEE.

Sitek, P., & Wikarek, J. (2015). A hybrid approach to the optimization of multiechelon systems. *Mathematical Problems in Engineering*, *2015*. https://doi.org/10.1155/2015/925675

Sluijk, N., Florio, A. M., Kinable, J., Dellaert, N., & Woensel, T. V. (2023, February). Two-echelon vehicle routing problems: A literature review. https://doi.org/10.1016/j.ejor.2022.02.022

Song, L., Gu, H., & Huang, H. (2017). A lower bound for the adaptive two-echelon capacitated vehicle routing problem. *Journal of Combinatorial Optimization*, *33*, 1145–1167. https://doi.org/10.1007/s10878-016-0028-6

Vanichchinchai, A. (2019). The effect of lean manufacturing on a supply chain relationship and performance. *Sustainability (Switzerland)*, *11*. https://doi.org/10.3390/su11205751

Vogt, J. J. (2010). The successful cross-dock based supply chain. *Journal of Business Logistics*, *31*, 99–119. https://doi.org/10.1002/j.2158-1592.2010.tb00130.x

Wang, Y., Zhang, S., Assogba, K., Fan, J., Xu, M., & Wang, Y. (2018). Economic and environmental evaluations in the two-echelon collaborative multiple centers vehicle routing optimization. *Journal of Cleaner Production*, *197*, 443–461. https://doi.org/10.1016/j.jclepro.2018.06.208

Wohlin, C. (2014). Guidelines for snowballing in systematic literature studies and a replication in software engineering. https://doi.org/10.1145/2601248.2601268

Wohlin, C., Kalinowski, M., Felizardo, K. R., & Mendes, E. (2022). Successful combination of database search and snowballing for identification of primary studies in systematic literature studies. *Information and Software Technology*, *147*, 106908. https://doi.org/10.1016/J.INFSOF.2022.106908

Xiao, H., Zhang, Y., Kou, G., Zhang, S., & Branke, J. (2023). Ranking and selection for pairwise comparison [First published: 10 January 2023]. *Naval Research Logistics*. https://doi.org/10.1002/nav.22093

Zhou, S., Zhang, D., Ji, B., & Li, S. (2024). Two-echelon vehicle routing problem with direct deliveries and access time windows. *Expert Systems with Applications*, *244*. https://doi.org/10.1016/j.eswa.2023.121150

Zinn, W., & Goldsby, T. J. (2017, March). In search of research ideas? call a professional. https://doi.org/10.1111/jbl.12160

A

Scientific Paper

Navigating the Cost-Emission Trade-Off in Hybrid Distribution Networks: Insights from a Flow-Based Vehicle Routing Problem Integrating Direct Plant Shipments

T.P. Van der Hulst, Dr. J.M. Vleugel, Ir. M.B. Duinkerken

Faculty of Civil Engineering and Geosciences, Delft University of Technology

Abstract

Amid increasing pressure to decarbonize logistics operations, hybrid distribution strategies that combine direct plant shipments with conventional two-echelon networks are gaining attention in the fast-moving consumer goods (FMCG) sector. This study explores the trade-offs between logistics costs and environmental performance in such hybrid distribution networks. Using a novel flow-based adaptation of the two-echelon vehicle routing problem (2E-VRP), a multi-objective optimization framework is developed and applied to a real-world FMCG case study. Through scalarization and ε-constraint methods, the results reveal that although direct shipments reduce logistics costs, environmental performance gains are not guaranteed: when emissions are directly minimized, direct shipments are often deprioritized in favor of two-echelon flows due to higher vehicle utilization in the latter. Sensitivity analyses show that carbon pricing alone does not substantially alter routing decisions, while direct emission optimization with cost constraints leads to more sustainable but centralized logistics patterns. This reveals a structural misalignment between purely cost-efficient and environmentally optimal distribution strategies. The findings highlight the need for route-specific eligibility criteria and additional operational constraints, such as minimum vehicle fill rates, to align logistics decisions with sustainability objectives.

***Keywords*:** direct shipments; distribution network design; emissions optimization; hybrid logistics; carbon pricing

1 INTRODUCTION

With increasing attention to sustainability, many companies are re-evaluating traditional logistics structures. A hybrid distribution strategy, where direct shipments from manufacturing plants are combined with conventional two-echelon distribution networks, offers potential benefits such as shorter lead times and increased flexibility. However, the impact of such strategies on both logistics cost and environmental performance remains insufficiently understood in practical applications. This study investigates the trade-offs between logistics cost and carbon emissions in hybrid distribution systems. A flow-based formulation of the two-echelon vehicle routing problem (2E-VRP) is applied to assess how routing decisions are affected when environmental considerations, such as carbon pricing or emissions minimization, are incorporated alongside economic objectives. The analysis is based on a real-world case from a major fast-moving consumer goods (FMCG) company, ensuring practical relevance. The central research question addressed is: *How do environmental incentives, such as carbon pricing or emissions minimization, influence routing behavior in hybrid distribution networks?* The findings provide insight into the extent to which sustainability goals align, or conflict, with cost-driven logistics decisions.

2 LITERATURE REVIEW

2.1 Hybrid Distribution and the 2E-VRP

Hybrid distribution strategies integrate multiple delivery models, such as direct deliveries, multi-echelon flows via distribution centers (DCs), and cross-docking, to optimize logistics networks. While classical distribution networks often separate direct and indirect deliveries, hybrid approaches aim to combine their strengths, offering cost-efficiency and increased flexibility. Recent studies have started to incorporate both direct and indirect flows within distribution models, particularly in controlled or urban logistics settings. For instance, Azizi and Hu (2020), formulates a hybrid approach combining capacitated location, routing and direct shipments. Musa, Arnaout, and Jung (2010) and Ma et al. (2011) integrate direct shipments into cross-docking networks, using integer programming to balance transport costs and scheduling requirements. In the food sector, Mohammadi, Barzinpour, and Teimoury (2020) emphasize the value of hybrid models for managing shelf-life and optimizing the trade-offs between consolidation and speed.

However, within the domain of two-echelon vehicle routing problems (2E-VRPs), where flows typically pass from plants through distribution centers to customers, only a few studies allow flexible combinations with direct routes. The model of Zhou et al. (2024) is one of the few to incorporate direct deliveries alongside time windows and synchronization constraints. Still, broader research on hybrid distribution, especially in large-scale or cross-country networks, with real-world data, remains limited, presenting a significant opportunity for further exploration.

2.2 Key Performance Indicators in Hybrid Distribution

Evaluating hybrid distribution strategies requires clear performance metrics. Economic performance is commonly assessed through total costs, which include transportation, establishment, and handling costs (Fishani et al., 2022; Mohammadi, Barzinpour, and Teimoury, 2020). Environmental performance is increasingly measured through emissions or energy consumption, as supply chains aim for more sustainable operations (Fishani et al., 2022; Hosseini-Nasab et al., 2023). Key indicators found in hybrid distribution studies include:

- Transportation costs: includes truck mileage, load factors, and fuel use.
- Carbon emissions: derived from routing distance and vehicle type.
- Inventory and handling cost: influenced by the number and role of the DCs.
- Service level metrics: include delivery lead tmie and fulfillment rates.

Most existing work emphasizes cost minimization (Musa, Arnaout, and Jung, 2010; Ma et al., 2011), with limited integration of customer-centric or resilience indicators. Zhou et al. (2024) introduces time windows and direct delivery prioritization, demonstrating how hybrid approaches may enhance service while optimizing logistics KPIs. However, empirical assessments of trade-offs between cost, emissions, and service in hybrid models remain scarce.

2.3 Modeling Approaches for Hybrid Distribution Networks

Modeling hybrid distribution networks requires combining multiple logistics structures and decision layers. The Vehicle Routing Problem (VRP) forms the basis of most modeling approaches. Two-echelon VRPs (2E-VRPs) extend classical VRP formulations by introducing satellites or DCs between plants and customers (Sluijk et al., 2023; Sitek and Wikarek, 2015). In most 2E-VRPs, all customer demand is served through a two-stage distribution flow. However, hybrid models could aim to relax this by allowing direct routes as well.

Mixed-Integer Linear Programming (MILP) is widely applied for solving such models due to its flexibility in encoding routing, capacity, time, and synchronization constraints (Azizi and Hu, 2020; Zhou et al., 2024). MILP formulations have successfully incorporated cross-docking (Dondo, Méndez, and Cerdá, 2011), direct delivery logic (Mohammadi, Barzinpour, and Teimoury, 2020), and even adaptive routing in urban settings (Song, Gu, and Huang, 2017). Nevertheless, most models still assume deterministic conditions and are designed for city logistics contexts.

2.4 Research Gap and Contribution

Hybrid distribution strategies, such as integrating direct shipments into two-echelon networks, are increasingly recognized for their potential to reduce lead times and logistics costs. However, the environmental implications of these strategies remain underexplored, particularly the trade-offs they introduce between cost efficiency and carbon emissions. This gap is especially pronounced at a regional cross-border scale in FMCG networks, where shipment volumes are fragmented and transport emissions substantial.

Existing optimization models in the literature predominantly focus on minimizing cost objectives, without explicitly capturing the structural tension between economic and environmental goals. Moreover, few studies investigate how hybrid routing decisions behave under real-world constraints such as heterogeneous shipment eligibility, varying vehicle fill rates, and demand volatility in cross-border supply chains.

This study addresses that gap by developing a flow-based two-echelon vehicle routing problem (2E-VRP) model that incorporates both direct and indirect flows, calibrated on real operational data from a real-world European FMCG network. A key contribution lies in the integration of emissions as an explicit performance dimension, enabling a dual-perspective evaluation of routing decisions.

Through cost-emission sensitivity analyses, the study reveals a consistent and policy-relevant insight: while direct shipments often minimize logistics costs, it leads to lower vehicle utilization and thus higher emissions per unit. Conversely, emission-optimized solutions may favor consolidated two-echelon routes in some cases, even when direct shipping would be economically justified. This result underscores a structural misalignment between cost-optimal and sustainability-optimal routing behavior.

By quantifying this trade-off and identifying the conditions under which environmental priorities shift routing outcomes, this research advances theoretical understanding of hybrid network dynamics. It also informs practice by demonstrating that cost-driven hybrid distribution policies may require additional operational constraints or incentive structures to align with broader sustainability goals.

3 METHODOLOGY

This study develops a flow-based adaptation of the two-echelon vehicle routing problem (2E-VRP) to explicitly quantify trade-offs between logistics cost and transport-related emissions in hybrid distribution networks. The model captures two routing strategies, conventional two-echelon flows and direct shipments, and is formulated as a Mixed-Integer Linear Programming (MILP) problem. It is calibrated on real operational data from a European FMCG network, enabling practical insights into sustainability-aware supply chain design.

3.1 Flow-Based 2E-VRP Formulation

The initial model minimizes the total logistics cost of fulfilling customer demand under operational constraints. Two distribution modes are modeled: (1) two-echelon shipments from plant $\rightarrow$ DC $\rightarrow$ customer, and (2) direct shipment from plant $\rightarrow$ customer. Binary variables represent shipment activations on each leg, and continuous variables assign shipped volumes.

Unlike classical 2E-VRP models, vehicle tours and return routing are not explicitly modeled. This simplification reflects the operational setup in the case study, where third-party logistics providers (3PLs) handle transportation and availability is assumed sufficient. Transport costs and emissions are incurred per full truckload (FTL) movement, and vehicle capacity constraints ensure feasible allocation. All indices, sets, parameters and decision variables in the model are given in Table 1.

Table 1: Overview of model notation for indices, parameters, and variables

Indices and Sets	
$p \in \mathcal{P}$	Set of origin nodes (plants)
$d \in \mathcal{D}$	Set of satellite nodes (DC)
$c \in \mathcal{C}$	Set of destination nodes (customers)
$N = \mathcal{P} \cup \mathcal{D} \cup \mathcal{C}$	Set of all nodes in the network
$v \in \mathcal{V}$	Set of vehicles assigned to first echelon (plant origin)
$w \in \mathcal{W}$	Set of vehicles assigned to second echelon (DC origin)
Parameters	
Q	Vehicle capacity (in FP)
$dist_{ij}$	Travel distance between node i and node j, where $(i, j) \in N \times N$
c_{ij}^{T}	Fixed transportation cost from origin i to destination j (in €/FTL shipment)
c_i^{load}	Average loading cost at origin node $i \in \mathcal{P} \cup \mathcal{D}$ (in €/FP)
c_j^{unload}	Average unloading cost at destination node $j \in \mathcal{D} \cup \mathcal{C}$ (in €/FP)
Decision Variables	
Binary routing:	
$x_{pdv} \in \{0, 1\}$	1 if vehicle v transports SKUs from plant p to DC d (first echelon)
$y_{dcw} \in \{0, 1\}$	1 if vehicle w transports SKUs from DC d to customer c (second echelon)
$z_{pcv} \in \{0, 1\}$	1 if vehicle v transports SKUs directly from plant p to customer c (DPS)
Continuous flow:	
$q_{pcv} \geq 0$	FP volume transported directly from plant p to customer c using vehicle v (DPS)
$\theta_{pdv} \geq 0$	FP volume transported from plant p to DC d using vehicle v (first echelon)
$\theta_{dcw} \geq 0$	FP volume transported from DC d to customer c using vehicle w (second echelon)

The logistics cost objective $\bar{Z}_{\text{cost}}$ includes fixed transport rates and variable handling costs per unit of flow, in this case given by floor positions (FP), accounting for origin destination pairs and shipment size eligibility (see Equation 3.1).

$$\begin{aligned}
\min \bar{Z}_{\text{cost}} = & \sum_{p\in\mathcal{P}}\sum_{d\in\mathcal{D}}\sum_{v\in\mathcal{V}} \Big[\underbrace{\theta_{pdv} \cdot c_p^{\text{load}}}_{\substack{\text{loading cost at}\\ \text{plant}}} + x_{pdv} \cdot \underbrace{\left(\frac{\theta_{pdv}}{Q}\right) \cdot c_{pd}^{\text{T}}}_{\substack{\text{transport cost}\\ \text{plant}\to\text{DC}}} + \underbrace{\theta_{pdv} \cdot c_d^{\text{unload}}}_{\substack{\text{unloading cost at}\\ \text{DC}}} \Big] \\
& + \sum_{d\in\mathcal{D}}\sum_{c\in\mathcal{C}}\sum_{w\in\mathcal{W}} \Big[\underbrace{\theta_{dcw} \cdot c_d^{\text{load}}}_{\substack{\text{loading cost at}\\ \text{DC}}} + y_{dcw} \cdot \underbrace{\left(\frac{\theta_{dcw}}{Q}\right) \cdot c_{dc}^{\text{T}}}_{\substack{\text{transport cost}\\ \text{DC}\to\text{customer}}} \Big] \\
& + \sum_{p\in\mathcal{P}}\sum_{c\in\mathcal{C}}\sum_{v\in\mathcal{V}} \Big[\underbrace{q_{pcv} \cdot c_p^{\text{load}}}_{\substack{\text{loading cost at}\\ \text{plant}}} + z_{pcv} \cdot \underbrace{c_{pc}^{\text{T}}}_{\substack{\text{transport cost}\\ \text{plant}\to\text{customer}}} \Big]
\end{aligned} \tag{3.1}$$

The environmental objective $\bar{Z}_{\text{emissions}}$ captures transport-related emissions based on route distance, vehicle fill rates, and a fixed emissions factor per kilometer (see Equation 3.2).

$$\bar{Z}_{\text{emissions}} = \gamma \cdot \left(\sum_{p\in\mathcal{P}}\sum_{d\in\mathcal{D}}\sum_{v\in\mathcal{V}} \frac{\theta_{pdv}}{Q} \cdot x_{pdv} \cdot dist_{pd} + \sum_{d\in\mathcal{D}}\sum_{c\in\mathcal{C}}\sum_{w\in\mathcal{W}} \frac{\theta_{dcw}}{Q} \cdot y_{dcw} \cdot dist_{dc} + \sum_{p\in\mathcal{P}}\sum_{c\in\mathcal{C}}\sum_{v\in\mathcal{V}} z_{pcv} \cdot dist_{pc} \right) \tag{3.2}$$

Together, these formulations allow the model to reflect real-world routing constraints, cost structures, and carbon output, without requiring detailed vehicle routing or fleet scheduling assumptions. This level of abstraction is appropriate for strategic-level planning, where the focus lies on shipment allocation and network design rather than route sequencing.

To ensure consistency with the per-unit cost formulation ($/Q$) in the two-echelon flows, the model assumes that first-echelon shuttles between plant and DC, as well as second-echelon deliveries from DC to customers, are always consolidated to full truckloads. In contrast, direct plant-to-customer shipments are modeled at full fixed cost per trip, since these flows are not assumed to be filled to FTL and this is highly relevant for the optimization. Furthermore, the model is subject to a broader set of constraints defined in the extended thesis report, including flow conservation, capacity restrictions, shipment activation logic, and customer demand fulfillment.

3.2 Sensitivity Analysis: Emissions-Oriented Routing Behavior

To systematically analyze the tension between cost efficiency and environmental sustainability, the model incorporates a multi-objective optimization approach. While minimizing logistics costs is a primary operational goal, reducing CO_2 emissions is increasingly prioritized in both corporate policy and regulation. The resulting trade-offs are not trivial: lower-cost routes may involve suboptimal vehicle utilization and higher emissions, whereas environmentally favorable strategies may increase operational costs.

To reflect this dual objective, the model evaluates two established methods that balance economic and environmental performance: a scalarization approach that combines both objectives using a weighted sum, and an ε-constraint method that minimizes emissions under budget constraints. These techniques enable systematic exploration of trade-offs and offer insight into how routing decisions shift under different sustainability priorities.

- **Scalarization:** A carbon price parameter λ (€/kg CO_2) is introduced to compute a weighted sum objective:

  $$\bar{Z}_{\text{total}} = \bar{Z}_{\text{cost}} + \lambda \cdot \bar{Z}_{\text{emissions}} \tag{3.3}$$

  By varying λ, the model explores Pareto-optimal trade-offs between cost and emissions. This approach reflects internalization of carbon pricing in logistics decision-making and aligns with recent studies on sustainability-aware optimization.
- **ε-constraint method:** The emissions objective is minimized directly as the sole objective, while logistics costs are constrained to remain within a specified upper bound. This approach reveals which routes and strategies are favored when emissions reduction is prioritized and cost becomes a constraint, without the influence of subjective weight settings.

Both methods are applied to multiple demand scenarios and policy configurations to assess routing behavior across a wide solution space. Together, they provide complementary perspectives on how different forms of environmental prioritization, implicit (via pricing) and explicit (via constraints), affect supply chain design and operational choices.

3.3 Case Application and Experimental Setup

The model is applied to a distribution network in an international European region, involving two production plants, one distribution center (DC), and a subset of high-volume destinations eligible for direct plant shipments. Historical shipment data and transport contracts inform the model parameters, including demand volumes, transport costs, and emissions per route.

Before introducing the environmental analyses described in subsection 3.2, the model was first solved under the cost-minimization framework (Equation 3.1) to understand baseline routing behavior.

3.3.1 Baseline Behavior: Break-Even Analysis

To explore when direct plant shipments become economically preferable, a break-even analysis is conducted. For each plant–destination pair, shipment volumes are gradually increased to determine the minimum threshold at which the model switches from two-echelon routing to direct shipment.

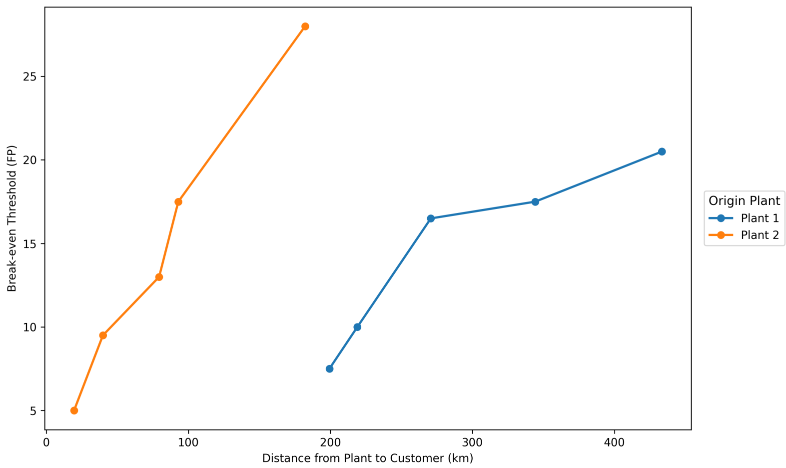

Figure 1: Break-even Threshold vs. Direct Distance to Customer

The resulting break-even thresholds as seen in Figure 1, reflect underlying cost trade-offs: longer transport distances require higher volumes to offset fixed full-truckload costs. The graph shows clear linear trends: customers located farther from a plant generally require higher shipment volumes to justify direct delivery. This indicates that uniform eligibility rules for direct shipments may be suboptimal. Instead, data-driven, route-specific thresholds could more effectively align operational decisions with cost-efficiency.

The trend for Plant 1, which serves customers at generally longer distances, appears less steep. This suggests that at higher distances, the marginal cost savings from direct shipments become more uniform, meaning that even modest increases in volume can justify switching to direct delivery. In contrast, for shorter distances (orange line), the threshold rises more steeply with distance, indicating stronger sensitivity to volume in determining cost-effectiveness.

These findings suggest that uniform eligibility rules for direct shipments may be suboptimal. Instead, route-specific, data-driven thresholds could better align operational decisions with cost-efficiency. Moreover, they provide context for the sensitivity analysis: although direct shipments may be cost-optimal at these break-even points, their typically lower vehicle utilization increases emissions per unit shipped, revealing a trade-off between economic and environmental objectives.

3.3.2 Sensitivity Analysis Design

To examine this trade-off in greater depth, experiments are conducted by combining four demand scenarios within internal policy configurations. Each experiment is solved using both sensitivity approaches described in subsection 3.2, enabling systematic exploration of how routing behavior shifts under varying levels of environmental prioritization. The resulting solutions form Pareto frontiers and trade-off curves that could identify thresholds where emissions objectives begin to override purely cost-driven decisions.

4 RESULTS

This section presents the results of two sensitivity analyses performed under the configuration in which no full-truckload (FTL) threshold is imposed. In this setup, the model has full flexibility to select shipment volumes that minimize cost and/or emissions, allowing a clearer view of how environmental objectives affect routing behavior when volume constraints are removed.

4.1 Scalarization: Cost-Emission Pareto Frontier

The scalarization approach reveals that increasing the carbon price parameter λ in its monetized form, leads to higher total logistics costs, as expected. However, routing structures remain largely unchanged, indicating that optimal flows are maintained unless environmental incentives reach a substantial threshold. Even at prices as high as €4/kg CO_2, changes in routing remain limited, suggesting a structural misalignment between cost-efficient and low-emission routes in the current network.

The scalarization analysis reveals that increasing the carbon price parameter λ (in €/kg CO_2) leads to a linear rise in total logistics costs, as expected. However, the routing structure remains largely unchanged across scenarios. As shown in Figure 2, emissions remain nearly constant despite rising costs, indicating that environmental pricing must reach very high levels before shifting route choices. Only in scenario S4 is a small emissions decrease observed at $\lambda = 0.13$, caused by a marginal shift toward a more emission-efficient route.

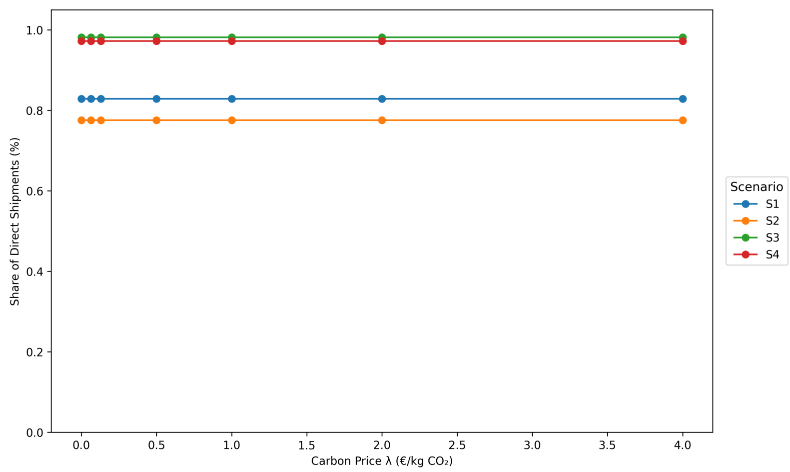

Figure 2: Pareto frontier of total logistics cost vs. CO_2 emissions across all scenarios

Total logistics cost is expressed as a percentage increase relative to the baseline ($\lambda = 0$), enabling a consistent cost comparison across scenarios. The results in Figure 2 highlight a structural misalignment between cost-optimal and emission-optimal routing decisions: pricing emissions alone does not meaningfully influence route selection under current network conditions.

4.2 ε-Constraint Method: Emissions vs. Share of Direct Shipments

The ε-constraint method reveals a different pattern. When total CO_2 emissions are minimized directly (subject to a maximum allowable cost increase, the ε-constraint), the model exhibits notable changes in routing behavior. As shown in Figure 3, the share of direct shipments declines across all scenarios as emissions are pushed downward. This is because direct flows tend to operate with lower vehicle fill rates, making them less emission-efficient compared to two-echelon flows consolidated through the DC.

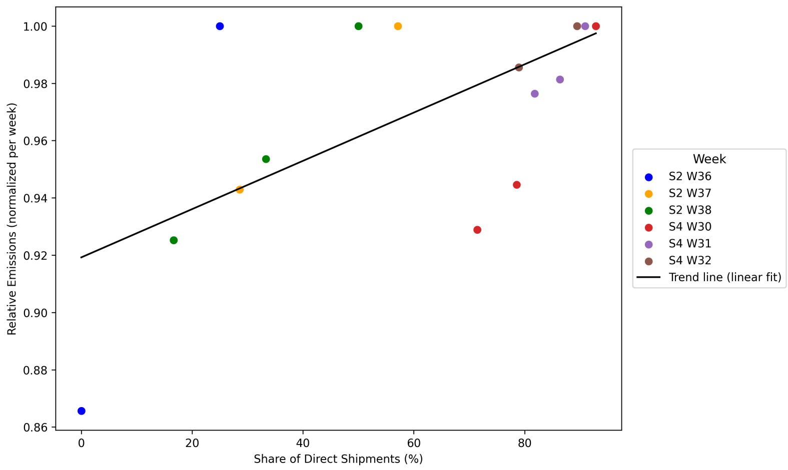

Figure 3: Relative CO_2 emissions vs. direct shipment share across the aggregated Pareto front

As cost constraints become stricter, the model increasingly shifts to centralized routing structures with higher utilization, thereby reducing emissions. The inverse relationship between emissions and direct shipment share seen in this analysis suggests that limiting emissions through direct constraints or utilization-based incentives is more effective than relying on carbon pricing alone.

5 DISCUSSION

5.1 Threshold Effects and Sustainability Misalignment

The results reveal a persistent trade-off between cost efficiency and environmental performance in hybrid distribution networks. Under the model configuration where no full truckload (FTL) threshold is imposed and shipment volumes are unconstrained, the model was free to choose the most efficient routing structure based on cost and/or emissions. Even in this flexible setting, expanding direct shipments does not align with sustainability goals when emissions are explicitly minimized.

The ε-constraint analysis shows that direct shipments are systematically deprioritized in favor of two-echelon flows when emissions are minimized. This shift arises because direct routes in cost-optimal solutions can suffer from low vehicle utilization, which increases emissions per shipped unit compared to consolidated deliveries via the DC. Notably, even without FTL constraints, the environmental cost of direct routes outweighs their logistics savings in many cases. This highlights a saturation point beyond which further direct plant shipping adoption may be environmentally counterproductive, particularly for small or fragmented flows.

These findings underscore the limitations of cost-driven routing in sustainability-oriented supply chains. They also demonstrate that two-echelon routing retains its environmental value, not only for low-volume or short-haul customers, but even for some high-volume routes where vehicle fill rates can be better optimized through central consolidation.

5.2 Strategic Implications for Sustainable Logistics

The scalarization results show that carbon pricing, in its current form, has limited influence on routing decisions. Even at elevated carbon prices (up to €4/kg CO_2), the model maintains cost-optimal routes with minimal emissions reduction, signaling a structural misalignment between economic and environmental objectives. This inertia reflects the weak sensitivity of current network configurations to environmental pricing alone.

To address this, more targeted environmental interventions are needed, such as route-level emission constraints, minimum VFR requirements, or differentiated service policies based on environmental performance. Adaptive eligibility rules that assess the real-time trade-off between cost and emissions could further improve alignment between direct shipment decisions and sustainability goals.

Overall, sustainable logistics design should move beyond static eligibility thresholds or pricing signals alone. Integrating emissions directly into the operational optimization logic, whether through hard constraints, dual objectives, or adaptive policies, offers a more effective strategy for balancing economic and environmental priorities in hybrid distribution networks.

5.3 Scientific Contribution

This research provides a detailed quantitative assessment of hybrid distribution networks that integrate direct plant shipments with two-echelon flows under flexible vehicle fill conditions. By systematically varying emission constraints and carbon prices in a fully unconstrained threshold configuration, the study offers novel insights into how routing structures respond to environmental incentives.

The dual application of scalarization and ε-constraint methods reveals that vehicle utilization, rather than eligibility thresholds, is the key driver of emissions performance. This challenges conventional assumptions that direct shipment strategies will always yield greener outcomes and highlights the nuanced trade-offs between cost, emissions, and service. The study contributes methodologically by applying multi-objective optimization and route-level emission accounting in a hybrid logistics setting, relevant for both academic and applied contexts.

6 CONCLUSION AND FURTHER RESEARCH RECOMMENDATIONS

This study has demonstrated that hybrid distribution systems can improve logistics performance and service levels, especially for high-volume, long-distance routes, while reducing DC-based inventory costs. However, these benefits are conditional: direct shipping does not guarantee emission savings unless vehicle fill rates remain high. Environmental performance is often superior in consolidated two-echelon flows, even in cost-neutral conditions.

From a strategic perspective, fixed thresholds for direct shipping eligibility appear suboptimal. A more dynamic, route-specific approach to different customers, accounting for volume, distance, and emission intensity, would better align operational efficiency with sustainability goals. Additionally, carbon pricing in its current form is insufficient to trigger meaningful routing shifts, indicating the need for more direct or constraint-based interventions.

Future research could extend this work by:

- Incorporating stochastic demand and delivery variability to test direct shipping robustness under uncertainty;
- Expanding the model to include return flows, intermodal options, or vehicle routing considerations;
- Applying the framework to different sectors or geographies to assess generalizability;
- Investigating behavioral or organizational barriers to implementing dynamic direct policies in practice.

Together, these directions would deepen the understanding of how to operationalize hybrid, sustainable logistics at scale.

REFERENCES

Azizi, Vahid and Guiping Hu (Aug. 2020). "Multi-product pickup and delivery supply chain design with location-routing and direct shipment". In: *International Journal of Production Economics* 226, p. 107648. ISSN: 0925-5273. DOI: 10.1016/J.IJPE.2020.107648.

Dondo, Rodolfo, Carlos A. Méndez, and Jaime Cerdá (Dec. 2011). "The multi-echelon vehicle routing problem with cross docking in supply chain management". In: *Computers & Chemical Engineering* 35 (12), pp. 3002–3024. ISSN: 0098-1354. DOI: 10.1016/J.COMPCHEMENG.2011.03.028.

Fishani, Behzad Aghaei et al. (May 2022). "Multi-objective location-allocation-routing problem of perishable multi-product supply chain with direct shipment and open routing possibilities under sustainability". In: *Concurrency and Computation: Practice and Experience* 34 (11). ISSN: 15320634. DOI: 10.1002/cpe.6860.

Hosseini-Nasab, Hasan et al. (Mar. 2023). "Transportation cost reduction using cross-docks linking". In: *Journal of Engineering Research (Kuwait)* 11 (1). Case study. ISSN: 23071885. DOI: 10.1016/j.jer.2023.100015.

Ma, Hong et al. (2011). "Crossdocking distribution networks with setup cost and time window constraint". In: *Omega* 39 (1), pp. 64–72. ISSN: 03050483. DOI: 10.1016/j.omega.2010.03.001.

Mohammadi, Zahra, Farnaz Barzinpour, and Ebrahim Teimoury (June 2020). "Designing sustainable supply chain network by considering direct and indirect shipment: Evidence from food industry". In: *Decision Science Letters* 9 (3), pp. 323–336. ISSN: 19295812. DOI: 10.5267/j.dsl.2020.5.003.

Musa, Rami, Jean Paul Arnaout, and Hosang Jung (Aug. 2010). "Ant colony optimization algorithm to solve for the transportation problem of cross-docking network". In: *Computers and Industrial Engineering* 59 (1), pp. 85–92. ISSN: 03608352. DOI: 10.1016/j.cie.2010.03.002.

Sitek, Paweł and Jarosław Wikarek (2015). "A hybrid approach to the optimization of multiechelon systems". In: *Mathematical Problems in Engineering* 2015. ISSN: 15635147. DOI: 10.1155/2015/925675.

Sluijk, Natasja et al. (Feb. 2023). *Two-echelon vehicle routing problems: A literature review*. DOI: 10.1016/j.ejor.2022.02.022.

Song, Liang, Hao Gu, and Hejiao Huang (May 2017). "A lower bound for the adaptive two-echelon capacitated vehicle routing problem". In: *Journal of Combinatorial Optimization* 33 (4), pp. 1145–1167. ISSN: 15732886. DOI: 10.1007/s10878-016-0028-6.

Zhou, Saiqi et al. (June 2024). "Two-echelon vehicle routing problem with direct deliveries and access time windows". In: *Expert Systems with Applications* 244. ISSN: 09574174. DOI: 10.1016/j.eswa.2023.121150.

B

Transport Data

Table B.1: Overview of Transportation Parameters in the Current State

Parameter	Value / Range	Explanation
Transport costs		
ISF cost: Plant 1 to DC 1		Fixed transport cost for the shuttle Plant 1 to DC 1
ISF cost: Plant 2 to DC 1		Fixed transport cost for the shuttle Plant 2 to DC 1
CF cost from DC 1 to customer		Range fixed CF rate to a customer-region based on historical data (see Table B.2)
CF cost from Plant 1 to customer		Range fixed CF rate to a customer-region based on historical data (see Table B.2)
CF cost from plant Plant 2 to customer		Range fixed CF rate to customer-region based on historical data (see Table B.2)
Handling and storage costs		
Loading cost (Inbound) per FP at plant		Standardized loading cost per Floor Position (FP) at the plant locations
Loading cost (Inbound) per FP at DC		Standardized loading cost per Floor Position (FP) at DC 1
Unloading cost (Outbound) per FP at DC		Standardized unloading cost per Floor Position (FP) at DC 1
Average storage cost per pallet at DC 1 (3PL)		Cost for holding 1 pallet (B1 or B2) at DC 1 owned by the 3PL
Average storage cost per pallet at plant (P&G owned)		Cost for holding pallets at the plant (P&G owned)
Other known parameters		
Vehicle capacity		Homogeneous vehicles on all routes (besides the different hauliers)
FTL threshold for DPS – customer		Below this, no DPS possible from this origin
FTL threshold from DC – Belgium customer		Below this, LTL is usually more cost-effective
FTL threshold from DC – Netherlands customer		Higher threshold than Belgium due to longer distance from DC 1

[1] Based on known average handling cost at Plant 2 DC for a B1 or B2 pallet, multiplied unit by 1.5 to reflect cost per FP
[2] Calculated as the average of handling one B2 pallet and two B1 pallets per FP.
[3] Calculated as the average of handling one B2 pallet and two B1 pallets per FP.

Table B.2: CF Transport Costs (€/FTL) from Each Origin to Destination Region. Green = known, Orange = estimated

Destination Region	Region Name	DC 1	Plant 1	Plant 2
	Drenthe			
	Flevoland			
	Friesland			
	Gelderland			
	Groningen			
	Limburg (NL)			
	Noord-Brabant			
	Noord-Holland			
	Overijssel			
	Utrecht			
	Zuid-Holland			
	Antwerpen			
	Brabant (BE)			
	Henegouwen			
	Limburg (BE)			
	Luik			
	Luxemburg (BE)			
	Namen			
	Oost-Vlaanderen			
	West-Vlaanderen			

Table B.3: Customer–Origin Distance Table (DC 1, Plant 1, Plant 2)

Customer (Region)	d_{ij} DC 1 (km)	d_{ij} Plant 1 (km)	d_{ij} Plant 2 (km)

C

Statistics Demand Data

Table C.1: Summary Statistics for DPS-eligible Customers (per Customer-Category Combination)

#	Customer*	Category	Mean (FP)	Std	CV
1	Customer 6	Category Y			
1	Customer 6	Category X			
2	Customer 7	Category Y			
2	Customer 7	Category X			
3	Customer 8	Category Y			
3	Customer 8	Category X			
4	Customer 9	Category Y			
4	Customer 9	Category X			
5	Customer 10	Category Y			
6	Customer 11	Category Y			
7	Customer 1	Category Y			
8	Customer 2	Category X			
9	Customer 3	Category Y			
9	Customer 3	Category X			
10	Customer 4	Category Y			
10	Customer 4	Category X			
11	Customer 12	Category Y			
11	Customer 12	Category X			
12	Customer 5	Category Y			
12	Customer 5	Category X			
13	Customer 13	Category Y			
14	Customer 14	Category Y			

Table C.2: Explanation of Key Demand Statistics

Metric	Description
Mean	Average weekly demand in Floor Positions (FP) for each customer-category-origin combination over the studied period.
Std	Standard Deviation: Measures the variability in weekly demand; higher values indicate greater fluctuations week to week.
CV	Coefficient of Variation: Calculated as $\text{CV} = \frac{\text{std}}{\text{mean}}$; it normalizes demand variability, where a lower CV indicates more stable demand, favorable for DPS planning.

Table C.3: FTL Consistency of DPS-eligible Customers (weeks ≥ 28 FP) - aggregated per Customer-Category

#	Customer	Category	FTL weeks	Share FTL weeks	DPS-eligible
1	Customer 6	Category Y			Yes
1	Customer 6	Category X			Yes
2	Customer 7	Category Y			Yes
2	Customer 7	Category X			Yes
3	Customer 8	Category Y			Yes
3	Customer 8	Category X			Yes
4	Customer 9	Category Y			Yes
4	Customer 9	Category X			Yes
5	Customer 10	Category Y			Yes
6	Customer 11	Category Y			Yes
7	Customer 1	Category Y			Yes
8	Customer 2	Category X			Yes
9	Customer 3	Category Y			Yes
9	Customer 3	Category X			Yes
10	Customer 4	Category Y			Yes
10	Customer 4	Category X			Yes
11	Customer 12	Category Y			Yes
11	Customer 12	Category X			Yes
12	Customer 5	Category Y			Yes
12	Customer 5	Category X			Yes
13	Customer 13	Category Y			Yes
14	Customer 14	Category Y			Yes

Table C.4: Average Autocorrelation ACF (lag 1–3) for Each DPS Candidate

Customer	Category	ACF
Customer 1	Cat X&Y	-0.191
Customer 1	Non-Cat X&Y	0.118
Customer 2	Cat X&Y	0.238
Customer 2	Non-Cat X&Y	0.067
Customer 3	Cat X&Y	-0.029
Customer 3	Non-Cat X&Y	0.247
Customer 4	Cat X&Y	0.017
Customer 4	Non-Cat X&Y	-0.133
Customer 5	Cat X&Y	-0.042
Customer 5	Non-Cat X&Y	0.056

Table C.5: ADF Stationarity Test Results for DPS Candidate Demand Series

Customer	Category	ADF	p-value	Stat.
Customer 1	Cat X&Y	-6.381	0.000	True
Customer 1	Other	-4.469	0.000	True
Customer 2	Cat X&Y	-0.576	0.876	False
Customer 2	Other	0.180	0.971	False
Customer 3	Cat X&Y	-5.740	0.000	True
Customer 3	Other	-2.307	0.170	False
Customer 4	Cat X&Y	-2.974	0.037	True
Customer 4	Other	-3.162	0.022	True
Customer 5	Cat X&Y	-4.899	0.000	True
Customer 5	Other	-5.032	0.000	True

D

Model Output for Data Validation

Table D.1: Detailed Cost Breakdown for Selected 2E Shipments in Validation Run

Plant	DC	Customer	Vol.	PD Dist.	DC Dist.	Transp. Cost	Load P. Cost	Load DC Cost	Unload DC Cost	Total Cost
Plant 1	DC 1	Customer 1								
Plant 1	DC 1	Customer 1								
Plant 1	DC 1	Customer 3								

Table D.2: Detailed Cost Breakdown for Selected DPS Shipments in Validation Run

Plant	DC	Customer	Vol.	PC Dist.	Transp. Dist.	Load P. Cost	Load DC Cost	Unload DC Cost	Total Cost
Plant 2	None	Customer 1							
Plant 2	None	Customer 3							

E
Computational Results

E.1. Results Cost and Environmental Impact per Scenario

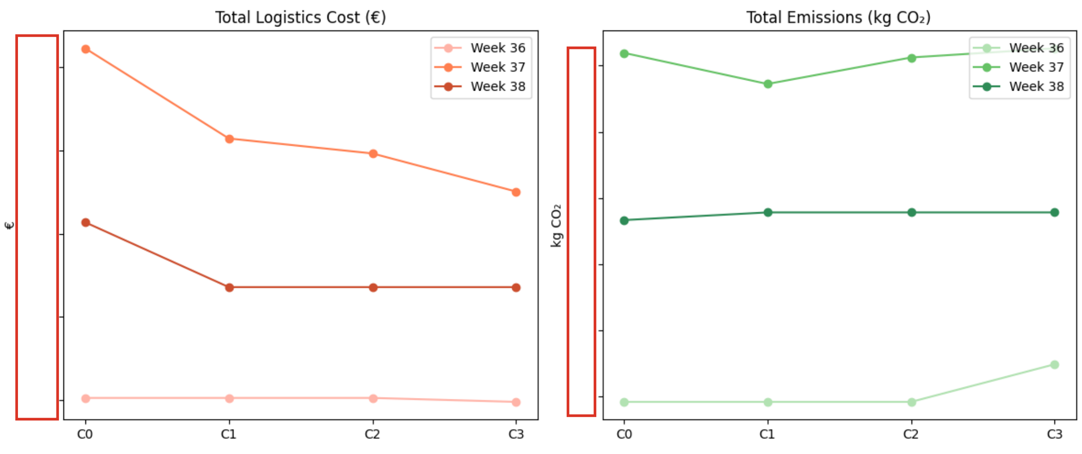

Figure E.1: Total Logistics Cost and CO_2 Emissions – S1 // C0-C3 // Week 36-38

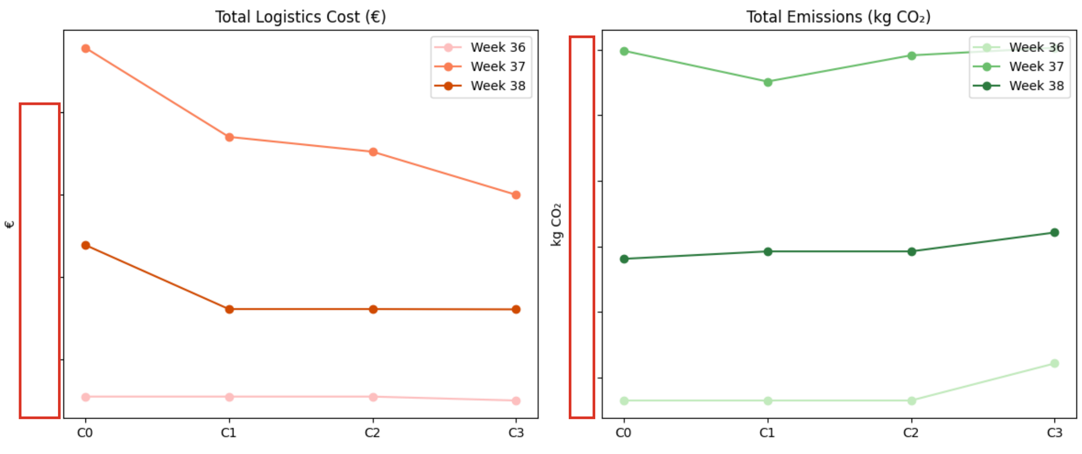

Figure E.2: Total Logistics Cost and CO_2 Emissions – S2 // C0-C3 // Week 36-38

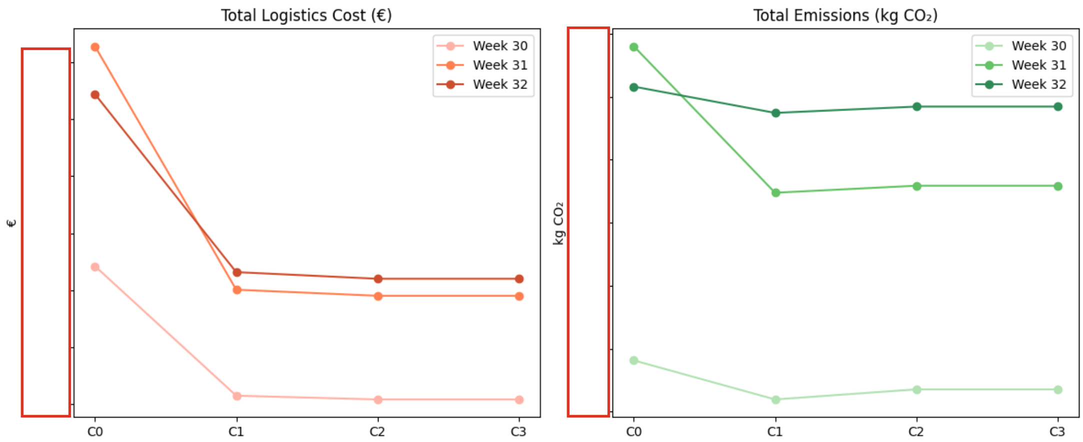

Figure E.3: Total Logistics Cost and CO_2 Emissions – S3 // C0-C3 // Week 30-32

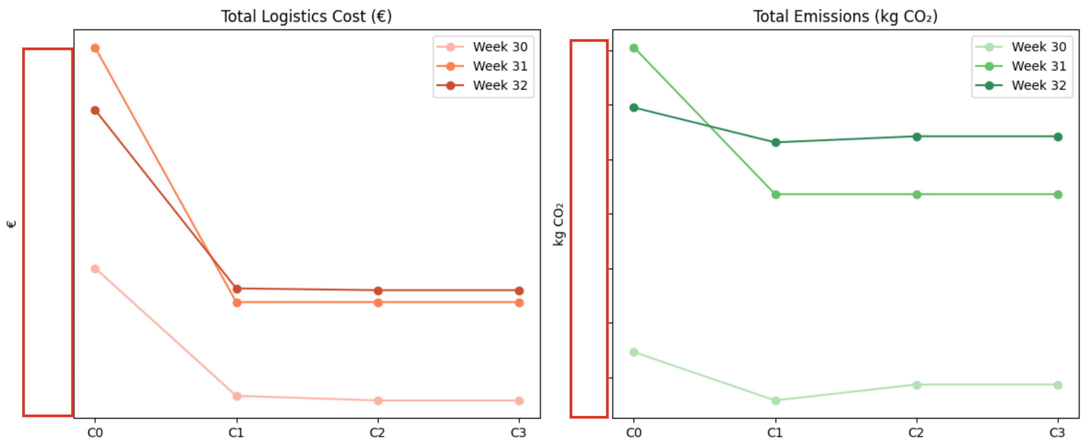

Figure E.4: Total Logistics Cost and CO_2 Emissions – S4 // C0-C3 // Week 30-32

E.2. Pairwise Comparison of Logistics Cost and Emissions Results

Table E.1: Pairwise Comparison of Configurations vs Baseline C0

		Logistics Cost		Emissions	
Scenario	**Comparison**	Mean % diff	Std % diff	Mean % diff	Std % diff
S1	C1 vs C0	-12.1%	10.5%	-1.2%	4.0%
	C2 vs C0	-13.0%	11.3%	0.4%	1.5%
	C3 vs C0	-16.3%	12.8%	7.4%	10.4%
S2	C1 vs C0	-11.2%	9.8%	-1.1%	3.7%
	C2 vs C0	-12.1%	10.5%	0.4%	1.4%
	C3 vs C0	-15.4%	12.4%	7.7%	7.5%
S3	C1 vs C0	-41.6%	4.5%	-10.9%	8.1%
	C2 vs C0	-42.9%	4.4%	-9.5%	8.3%
	C3 vs C0	-42.9%	4.4%	-9.5%	8.3%
S4	C1 vs C0	-42.7%	5.0%	-12.6%	7.7%
	C2 vs C0	-43.3%	4.8%	-11.0%	8.5%
	C3 vs C0	-43.3%	4.8%	-11.0%	8.5%
Total Avg	C1 vs C0	**-26.4%**	**7.5%**	**-7.0%**	**5.4%**
	C2 vs C0	**-27.3%**	**7.8%**	**-4.9%**	**4.9%**
	C3 vs C0	**-29.7%**	**8.1%**	**-1.1%**	**8.2%**

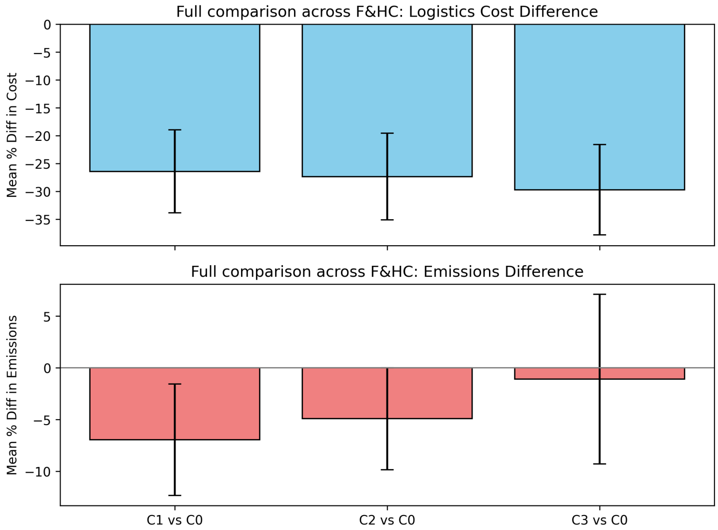

Figure E.5: Full Comparison Across Both Categories of Logistics Cost and Emissions Savings over the Configurations

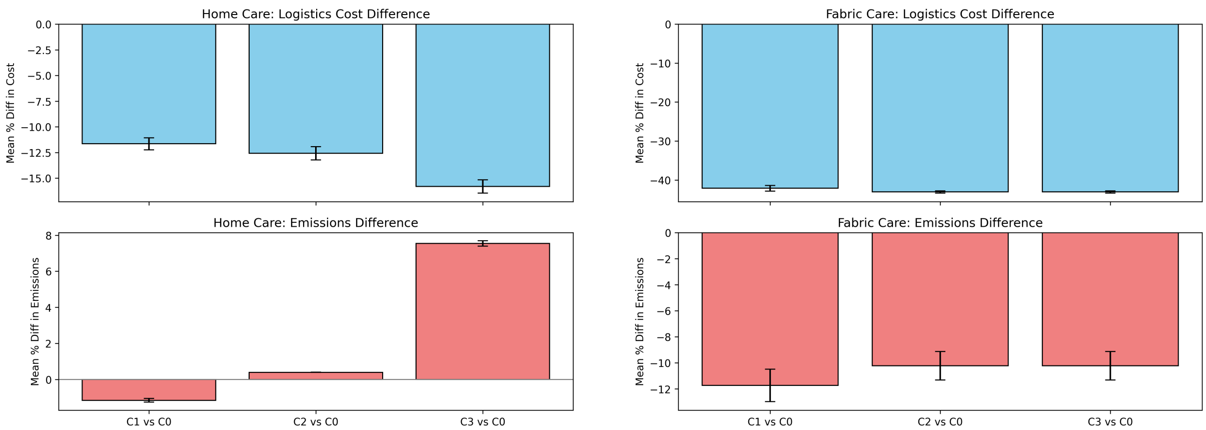

(a) Category X: Cost and Emissions Comparison

(b) Category Y: Cost and Emissions Comparison

Figure E.6: Pairwise Comparison of Logistics Cost and Emissions for Category X and Category Y

E.3. Results Service Level and Network Effects per Scenario

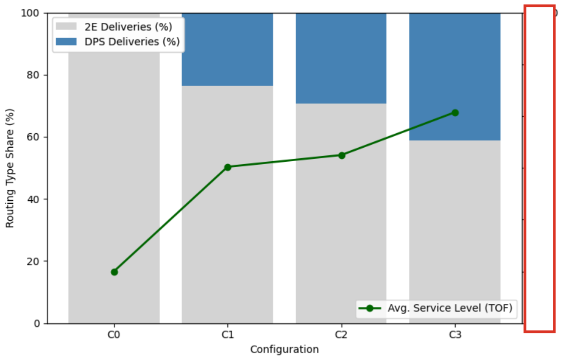

Figure E.7: Service Level and Delivery Type Split – S1 C0–C3, Average over Week 36–38

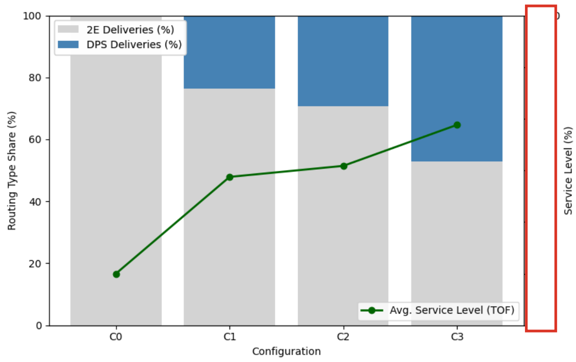

Figure E.8: Service Level and Delivery Type Split – S2 C0–C3, Average over Week 36–38

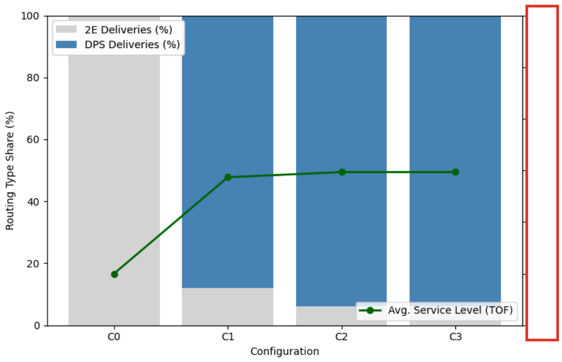

Figure E.9: Service Level and Delivery Type Split – S3 C0–C3, Average over Week 30–32

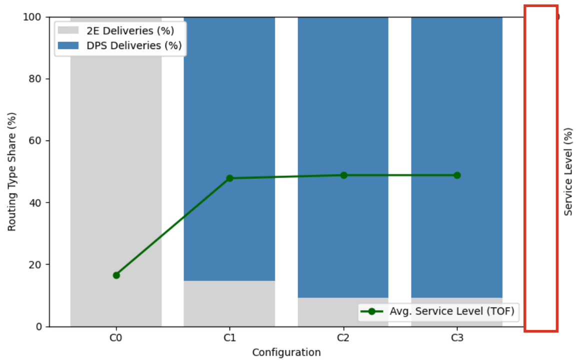

Figure E.10: Service Level and Delivery Type Split – S4 C0–C3, Average over Week 30–32

E.4. Pairwise Comparison of Service Level Results

Table E.2: Pairwise Comparison of Configurations vs Baseline C0 for Service Level

Scenario	**Comparison**	**Service Level** Mean % diff	Std % diff
S1	C1 - C0	1.0%	0.9%
	C2 - C0	1.2%	1.0%
	C3 - C0	1.5%	0.6%
S2	C1 - C0	1.0%	0.9%
	C2 - C0	1.1%	0.9%
	C3 - C0	1.5%	0.6%
S3	C1 - C0	1.0%	0.0%
	C2 - C0	1.0%	0.0%
	C3 - C0	1.0%	0.0%
S4	C1 - C0	1.0%	0.0%
	C2 - C0	1.0%	0.0%
	C3 - C0	1.0%	0.0%
Total Avg	C1 - C0	**1.0%**	**0.6%**
	C2 - C0	**1.1%**	**0.6%**
	C3 - C0	**1.2%**	**0.5%**

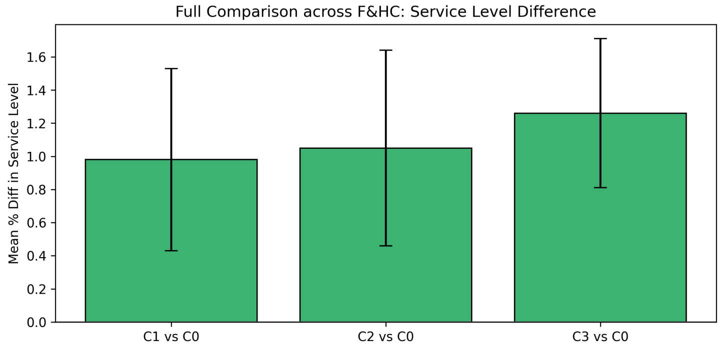

Figure E.11: Full Comparison Across Both Categories of Service Level Change over the Configurations

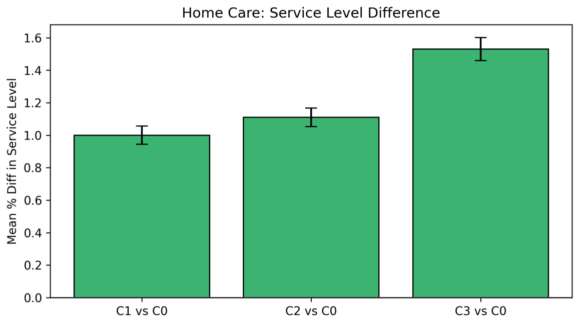

(a) Category X: Service Level Difference

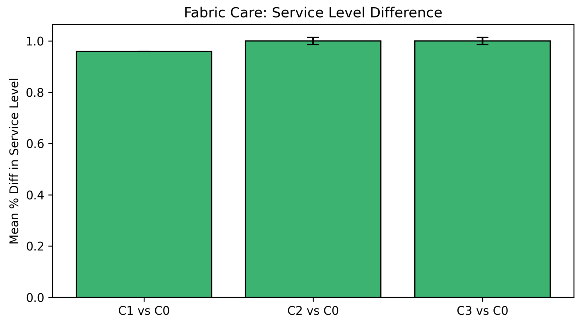

(b) Category Y: Service Level Difference

Figure E.12: Pairwise Comparison of Service Level Differences for Category X and Category Y

E.5. Detailed Results Demand Scenario 1

Table E.3: Demand – Category X Scenario S1 – Top 5 Customers (Weeks 36–38)

Plant	Customer	Week 36 (FP)	Week 37 (FP)	Week 38 (FP)
Plant 1	Customer 1			
Plant 1	Customer 2			
Plant 1	Customer 3			
Plant 1	Customer 4			
Plant 1	Customer 5			
	Total			

Table E.4: Results – Category X Scenario S1 (Production only) for Weeks 36–38

Configuration	Metric	Week 36	Week 37	Week 38
C0	Total logistic costs			
	Total distance			
	Total emissions			
	2E/DPS shipment-ratio			
	Total service level			
C1	Total logistic costs			
	Total distance			
	Total emissions			
	2E/DPS shipment-ratio			
	Total service level			
C2*(DPS Threshold =)*	Total logistic costs			
	Total distance			
	Total emissions			
	2E/DPS shipment-ratio			
	Total service level			
C3*(DPS Threshold = 0)*	Total logistic costs			
	Total distance			
	Total emissions			
	2E/DPS shipment-ratio			
	Total service level			

Table E.5: Detailed Cost Breakdown Example: S1 // C1 // Week 37

Plant	**DC**	**Cust.**	**Vol.**	**PC** Dist.	**PD** Dist.	**DC** Dist.	**Transp.** Cost	**Load P.** Cost	**Load DC** Cost	**Unload DC** Cost	**Total** Cost
Plant 1	None	Customer 3									
Plant 1	None	Customer 5									
Plant 1	DC 1	Customer 1									
Plant 1	DC 1	Customer 2									
Plant 1	DC 1	Customer 3									
Plant 1	DC 1	Customer 4									
Plant 1	DC 1	Customer 5									

E.6. Detailed Results Demand Scenario 2

Table E.6: Demand – Category X Scenario S2 – Top 5 Customers (Weeks 36–38)

Plant	Customer	Week 36 (FP)	Week 37 (FP)	Week 38 (FP)
Plant 1	Customer 1			
Plant 1	Customer 2			
Plant 1	Customer 3			
Plant 1	Customer 4			
Plant 1	Customer 5			
	Total			

Table E.7: Results – Category X Scenario S2 (Production + Storage) for Weeks 36–38

Configuration	Metric	Week 36	Week 37	Week 38
C0	Total logistic costs			
	Total distance			
	Total emissions			
	2E/DPS shipment-ratio			
	Total service level			
C1	Total logistic costs			
	Total distance			
	Total emissions			
	2E/DPS shipment-ratio			
	Total service level			
C2*(DPS Threshold =)*	Total logistic costs			
	Total distance			
	Total emissions			
	2E/DPS shipment-ratio			
	Total service level			
C3*(DPS Threshold = 0)*	Total logistic costs			
	Total distance			
	Total emissions			
	2E/DPS shipment-ratio			
	Total service level			

E.7. Detailed Results Demand Scenario 3

Table E.8: Demand – Category Y Scenario S3 – Top 5 Customers (Weeks 30–32)

Plant	Customer	Week 30 (FP)	Week 31 (FP)	Week 32 (FP)
Plant 2	Customer 1			
Plant 2	Customer 2			
Plant 2	Customer 3			
Plant 2	Customer 4			
Plant 2	Customer 5			
	Total			

Table E.9: Results – Category Y Scenario S3 (Production only) for Weeks 30-32

Configuration	Metric	Week 30	Week 31	Week 32
C0	Total logistic costs			
	Total distance			
	Total emissions			
	2E/DPS shipment-ratio			
	Total service level			
C1	Total logistic costs			
	Total distance			
	Total emissions			
	2E/DPS shipment-ratio			
	Total service level			
C2*(DPS Threshold =)*	Total logistic costs			
	Total distance			
	Total emissions			
	2E/DPS shipment-ratio			
	Total service level			
C3*(DPS Threshold = 0)*	Total logistic costs			
	Total distance			
	Total emissions			
	2E/DPS shipment-ratio			
	Total service level			

Table E.10: Detailed Cost Breakdown Example: S3 // C1 // Week 31

Plant	**DC**	**Cust.**	**Vol.**	**PC** Dist.	**PD** Dist.	**DC** Dist.	**Transp.** Cost	**Load P.** Cost	**Load DC** Cost	**Unload DC** Cost	**Total** Cost
Plant 2	None	Customer 1									
Plant 2	None	Customer 1									
Plant 2	None	Customer 3									
Plant 2	None	Customer 3									
Plant 2	None	Customer 3									
Plant 2	None	Customer 3									
Plant 2	None	Customer 3									
Plant 2	None	Customer 3									
Plant 2	None	Customer 3									
Plant 2	None	Customer 3									
Plant 2	None	Customer 4									
Plant 2	None	Customer 4									
Plant 2	None	Customer 4									
Plant 2	None	Customer 4									
Plant 2	None	Customer 5									
Plant 2	None	Customer 5									
Plant 2	None	Customer 5									
Plant 2	None	Customer 5									
Plant 2	DC 1	Customer 1									
Plant 2	DC 1	Customer 2									

E.8. Detailed Results Demand Scenario 4

Table E.11: Demand – Category Y Scenario S4 – Top 5 Customers (Weeks 30–32)

Plant	Customer	Week 30 (FP)	Week 31 (FP)	Week 32 (FP)
Plant 2	Customer 1			
Plant 2	Customer 2			
Plant 2	Customer 3			
Plant 2	Customer 4			
Plant 2	Customer 5			
	Total			

Table E.12: Results – Category Y Scenario S4 (Production + Storage) for Weeks 30-32

Configuration	Metric	Week 30	Week 31	Week 32
C0	Total logistic costs			
	Total distance			
	Total emissions			
	2E/DPS shipment-ratio			
	Total service level			
C1	Total logistic costs			
	Total distance			
	Total emissions			
	2E/DPS shipment-ratio			
	Total service level			
C2(*DPS Threshold =*)	Total logistic costs			
	Total distance			
	Total emissions			
	2E/DPS shipment-ratio			
	Total service level			
C3(*DPS Threshold = 0)*	Total logistic costs			
	Total distance			
	Total emissions			
	2E/DPS shipment-ratio			
	Total service level			

E.9. Numerical Results ε-constraint Analysis

Table E.13: Results of the ε-constraint Analysis for S2 and S4, Minimizing Emissions under Varying Cost Caps

Scenario	ε (€)	Cost (€)	Emissions (kg CO_2)	2E / DPS Count
S2 (Week 36)	Initial Cost Optimization Solution			
S2 (Week 37)	Initial Cost Optimization Solution			
S2 (Week 38)	Initial Cost Optimization Solution			
S4 (Week 30)	Initial Cost Optimization Solution			
S4 (Week 31)	Initial Cost Optimization Solution			
S4 (Week 32)	Initial Cost Optimization Solution			

E.10. Data Demand-Based Break-Even Analysis

Table E.14: Overview of Distances and Break-even DPS Thresholds by Customer and Origin

Cust	Origin	Plant–Cust (km)	Plant–DC (km)	DC–Cust (km)	Threshold τ (FP)
Customer 1	Plant 2				
Customer 1	Plant 1				
Customer 5	Plant 2				
Customer 5	Plant 1				
Customer 2	Plant 2				
Customer 2	Plant 1				
Customer 3	Plant 2				
Customer 3	Plant 1				
Customer 4	Plant 2				
Customer 4	Plant 1				

F

Python Model

The mathematical flow-based 2E-VRP model described in chapter 5 was implemented in Python using Visual Studio Code. This appendix presents the main code blocks that form the core of the implementation, highlighting the most relevant components used in model formulation and optimization. Non-essential sections (e.g. scenario testing, extended analyses, and visualization scripts) are excluded from this appendix to maintain clarity and focus on core model components.

```
"""
VS Code Editor
Thesis Tessa van der Hulst
"""
import pandas as pd
import numpy as np
import matplotlib.pyplot as plt
import seaborn as sns
from IPython.display import display, Markdown
import os
os.environ['XPRESS'] = r"c:\Users\vanderhulst.tp\xpauth.xpr"
import xpress as xp
from itertools import combinations
from xpress import *
```

Data Preprocessing

```
# Load and preprocess raw demand data
df_raw = pd.read_csv("demand_full_2.csv", skiprows=1)
df_raw.columns = df_raw.columns.str.strip()

# Reshape from wide to long format
df_long = df_raw.melt(
    id_vars=['Customer Ship To Name', 'Category', 'Origin Plant', 'First Storage Location'],
    var_name='Week',
    value_name='Demand_FP'
).rename(columns={
    'Customer Ship To Name': 'CustomerID',
    'Origin Plant': 'Origin',
    'First Storage Location': 'StorageLoc'
})

# Clean and convert columns
df_long[['CustomerID', 'Category', 'Origin', 'StorageLoc']] = df_long[['CustomerID', 'Category',
↪ 'Origin', 'StorageLoc']].ffill()
df_long['Week'] = df_long['Week'].astype(int)
df_long['Demand_FP'] = pd.to_numeric(df_long['Demand_FP'], errors='coerce')
df_long.dropna(subset=['Demand_FP'], inplace=True)

```

```python
# Flag FTL-eligible shipments
df_long['FTL_eligible'] = df_long['Demand_FP'] >=
# Select data for specific weeks (e.g., week 37)
selected_weeks = [37]
df_weeks = df_long[df_long['Week'].isin(selected_weeks)].copy()
# Aggregate demand per week per origin-customer combination
df_grouped = (
    df_weeks
    .groupby(['CustomerID', 'Category', 'Origin', 'StorageLoc', 'Week'])['Demand_FP']
    .sum()
    .reset_index()
)
# Total demand across selected weeks
df_final = (
    df_grouped
    .groupby(['CustomerID', 'Category', 'Origin', 'StorageLoc'])['Demand_FP']
    .sum()
    .reset_index()
    .sort_values(by=['CustomerID', 'Category', 'Origin', 'StorageLoc'])
)
# Average demand per week across ALL weeks
df_avg_allweeks = (
    df_long
    .groupby(['CustomerID', 'Category', 'Origin', 'StorageLoc'])['Demand_FP']
    .mean()
    .reset_index()
    .rename(columns={'Demand_FP': 'Avg_Demand_FP_per_week_allweeks'})
)
# Average demand per week across SELECTED weeks
df_avg = (
    df_grouped
    .groupby(['CustomerID', 'Category', 'Origin', 'StorageLoc'])['Demand_FP']
    .mean()
    .reset_index()
    .rename(columns={'Demand_FP': 'Avg_Demand_FP_per_week'})
)
# Load cost data and create cost parameter dictionaries
df_costs = pd.read_csv("csv_fixedtransportcosts.csv")
c_pd_ISF = {(r['From'], r['To']): r['Cost'] for _, r in df_costs[df_costs['Type'] ==
↪   'plant_to_dc'].iterrows()}
c_dc_CF = {(r['From'], r['To']): r['Cost'] for _, r in df_costs[df_costs['Type'] ==
↪   'dc_to_customer'].iterrows()}
c_pc_CF = {(r['From'], r['To']): r['Cost'] for _, r in df_costs[df_costs['Type'] ==
↪   'plant_to_customer'].iterrows()}
# Aggregate total demand per customer and storage location, excluding DC 1
q_pc_c2c4 = (
    df_final[df_final['StorageLoc'] != 'DC 1']
    .groupby(['CustomerID', 'StorageLoc'])['Demand_FP']
    .sum()
    .reset_index()
)
```

Model Formulation: Sets, Parameters, Distance Matrix, Decision Variables

```python
# --- Demand dictionary: {(plant, customer): demand} ---
q_pc_small = {
    (row['StorageLoc'], row['CustomerID']): row['Demand_FP']
    for _, row in q_pc_c2c4.iterrows()
}

# --- Define sets ---
C = sorted(set(c for (_, c) in q_pc_small))         # Customers
P = sorted(set(p for (p, _) in q_pc_small))[1:]      # Plants (e.g., [1:] for Plant 1)
D = sorted(set(d for (d, _) in c_dc_CF))             # Distribution Centers

```

```
# --- Filter demand to selected plants only ---
q_pc_small = {
    (p, c): q for (p, c), q in q_pc_small.items()
    if p in P
}

# --- Vehicle sets ---
V = list(range(50))        # Plant-based vehicles
W = list(range(50, 100))   # DC-based vehicles

# --- Filter transport cost dictionaries based on selected customers ---
c_pc_CF_small = {(p, c): cost for (p, c), cost in c_pc_CF.items() if c in C}
c_dc_CF_small = {(d, c): cost for (d, c), cost in c_dc_CF.items() if c in C}
c_pd_ISF_small = c_pd_ISF.copy()  # Use full plant→DC cost dict

# --- Create optimization model ---
prob = xp.problem()

# --- Define global parameters ---
Q =                             # Vehicle capacity in FP
c_load_plant = {p:  for p in P}
c_load_dc = {d:  for d in D}
c_unload = {d:  for d in D}
lambda_emission = 0.065        # €/kg CO2
gamma = 0.90                   # kg CO2 per km

# --- Load and clean distance matrix ---
distance_matrix_all = pd.read_csv("full_distance_matrix.csv", index_col=0)
distance_matrix_all.index = distance_matrix_all.index.astype(str).str.strip()
distance_matrix_all.columns = distance_matrix_all.columns.astype(str).str.strip()
distance_matrix_all.rename(
    index=lambda x: x.replace("Plant", "").replace("DC", ""),
    columns=lambda x: x.replace("Plant", "").replace("DC", ""),
    inplace=True
)

# --- Build full distance lookup ---
distance_lookup = {
    (i, j): distance_matrix_all.loc[i, j]
    for i in distance_matrix_all.index
    for j in distance_matrix_all.columns
    if not pd.isna(distance_matrix_all.loc[i, j])
}

# --- Distance dictionaries per leg ---
d_pc = {(p, c): distance_lookup[(p, c)] for (p, c) in c_pc_CF_small if (p, c) in
↪  distance_lookup}
d_pd = {(p, d): distance_lookup[(p, d)] for (p, d) in c_pd_ISF_small if (p, d) in
↪  distance_lookup}
d_dc = {(d, c): distance_lookup[(d, c)] for (d, c) in c_dc_CF_small if (d, c) in
↪  distance_lookup}

# --- Binary routing variables ---
# x_{pdv}: Plant → DC (first leg of 2E)
x = {(p, d, v): xp.var(name=f"x_{p}_{d}_v{v}", vartype=xp.binary)
     for (p, d) in c_pd_ISF_small for v in V}
# z_{pcv}: Direct Plant → Customer (DPS)
z = {(p, c, v): xp.var(name=f"z_{p}_{c}_v{v}", vartype=xp.binary)
     for (p, c) in c_pc_CF_small for v in V}
# y_{dcw}: DC → Customer (second leg of 2E)
y = {(d, c, w): xp.var(name=f"y_{d}_{c}_w{w}", vartype=xp.binary)
     for (d, c) in c_dc_CF_small for w in W}

# --- Continuous flow variables ---
# q_{pcv}: DPS flow in FP
```

```
q = {(p, c, v): xp.var(name=f"q_{p}_{c}_v{v}", lb=0)
      for (p, c) in c_pc_CF_small for v in V}
# _{pdv}: First leg of 2E (Plant → DC)
theta_pd = {(p, d, v): xp.var(name=f"theta_pd_{p}_{d}_v{v}", lb=0)
             for (p, d) in c_pd_ISF_small for v in V}
# _{dcw}: Second leg of 2E (DC → Customer)
theta_dc = {(d, c, w): xp.var(name=f"theta_dc_{d}_{c}_w{w}", lb=0)
             for (d, c) in c_dc_CF_small for w in W}
```

Model Formulation: Objective Function, Constraints

```
# Add variables to the problem
prob.addVariable(
    list(x.values()) +
    list(y.values()) +
    list(z.values()) +
    list(q.values()) +
    list(theta_pd.values()) +
    list(theta_dc.values())
)

# OBJECTIVE FUNCTION: LOGISTICS COST
objective_terms = []
# Two-echelon: plant → DC
for (p, d, v), flow in theta_pd.items():
    cost_transport = c_pd_ISF_small.get((p, d), 0)
    load_p = c_load_plant.get(p, 0)
    unload_d = c_unload.get(d, 0)

    objective_terms.append(flow * load_p)
    objective_terms.append((flow / Q) * cost_transport * x[(p, d, v)])
    objective_terms.append(flow * unload_d)
# Two-echelon: DC → customer
for (d, c, w), flow in theta_dc.items():
    cost_transport = c_dc_CF_small.get((d, c), 0)
    load_d = c_load_dc.get(d, 0)

    objective_terms.append(flow * load_d)
    objective_terms.append((flow / Q) * cost_transport * y[(d, c, w)])
# DPS: direct plant → customer
for (p, c, v), flow in q.items():
    cost_transport = c_pc_CF_small.get((p, c), 0)
    load_p = c_load_plant.get(p, 0)

    objective_terms.append(flow * load_p)
    objective_terms.append(z[(p, c, v)] * cost_transport)

# Set the objective
prob.setObjective(xp.Sum(objective_terms), sense=xp.minimize)

# OBJECTIVE FUNCTION: EMISSIONS
objective_emissions = []
for (p, d, v), flow in theta_pd.items():
    if (p, d) in d_pd:
        dist = d_pd[(p, d)]
        objective_emissions.append((flow / Q) * dist * gamma * x[(p, d, v)])
for (d, c, w), flow in theta_dc.items():
    if (d, c) in d_dc:
        dist = d_dc[(d, c)]
        objective_emissions.append((flow / Q) * dist * gamma * y[(d, c, w)])
for (p, c, v) in z:
    if (p, c) in d_pc:
        dist = d_pc[(p, c)]
        objective_emissions.append(gamma * dist * z[(p, c, v)])

```

```
Z_emissions = xp.Sum(objective_emissions)

# CONSTRAINTS
# This turns on and off the constraints
config = {
    "demand_constraints_per_plant": True,
    "flow_activ_vehiclecap_constraints": True,
    "ftl_enforcement": True,
    "dc_flow_balance": True,
    "single_task_per_vehicle": True,
    "limit_dps_to_one_customer": True
}

# Demand Satisfaction
if config["demand_constraints_per_plant"]:
    for (p, c), demand in q_pc_small.items():
        direct = xp.Sum(q[(p, c, v)] for v in V if (p, c, v) in q)
        indirect = xp.Sum(theta_dc[(d, c, w)] for d in D for w in W if (d, c, w) in theta_dc)
        prob.addConstraint(direct + indirect == demand)

# Flow Activation & Vehicle Capacity
if config["flow_activ_vehiclecap_constraints"]:
    for (p, c, v) in q:
        prob.addConstraint(q[(p, c, v)] <= Q * z[(p, c, v)])
    for (p, d, v) in theta_pd:
        prob.addConstraint(theta_pd[(p, d, v)] <= Q * x[(p, d, v)])
    for (d, c, w) in theta_dc:
        prob.addConstraint(theta_dc[(d, c, w)] <= Q * y[(d, c, w)])

# DPS Threshold Enforcement
DPS_THRESHOLD =
if config["ftl_enforcement"]:
    for (p, c, v) in z:
        demand = q_pc_small.get((p, c), 0)
        if demand > DPS_THRESHOLD:
            prob.addConstraint(q[(p, c, v)] <= Q * z[(p, c, v)])
            prob.addConstraint(q[(p, c, v)] >= DPS_THRESHOLD * z[(p, c, v)])
        else:
            prob.addConstraint(z[(p, c, v)] == 0)

# DC Flow Balance
if config["dc_flow_balance"]:
    for d in D:
        inflow = xp.Sum(theta_pd[(p, d, v)] for p in P for v in V if (p, d, v) in theta_pd)
        outflow = xp.Sum(theta_dc[(d, c, w)] for c in C for w in W if (d, c, w) in theta_dc)
        prob.addConstraint(inflow == outflow)

# Vehicle Task Exclusivity
if config["single_task_per_vehicle"]:
    for v in V:
        dps = xp.Sum(z[(p, c, v)] for p in P for c in C if (p, c, v) in z)
        twoe = xp.Sum(x[(p, d, v)] for p in P for d in D if (p, d, v) in x)
        prob.addConstraint(dps + twoe <= 1)
    for w in W:
        prob.addConstraint(xp.Sum(y[(d, c, w)] for d in D for c in C if (d, c, w) in y) <= 1)

# DPS: One Customer per Vehicle
if config["limit_dps_to_one_customer"]:
    for v in V:
        prob.addConstraint(xp.Sum(z[(p, c, v)] for p in P for c in C if (p, c, v) in z) <= 1)
```

Problem Solver & Output

```
# Z_total Objective: include emissions cost
emission_cost_expr = e * Z_emissions
```

```
total_cost_with_emissions = xp.Sum(objective_terms) + emission_cost_expr
prob.setObjective(total_cost_with_emissions, sense=xp.minimize)

# Solve
status = prob.solve()

# --- Main results ---
print(f"\n[Fixed Carbon Price]")
print(f"Total Cost (logistics): €{prob.getSolution(xp.Sum(objective_terms)):.2f}")
print(f"Total Emissions: {prob.getSolution(Z_emissions):.2f} kg CO ")
print(f"Emissions Cost: €{prob.getSolution(emission_cost_expr):.2f}")
print(f"Total Cost incl. Emissions: €{prob.getObjVal():.2f}")

# Summary trackers
dps_total_cost = 0
e2_total_cost = 0
dps_emissions_total = 0
e2_emissions_total = 0
dps_emissions_arc = 0
e2_emissions_arc = 0
total_distance = 0
dps_count = 0
e2_count = 0

logistics_cost = float(prob.getSolution(xp.Sum(objective_terms)))
emissions = float(prob.getSolution(Z_emissions))

print(f"\n Logistics-Only Cost: €{logistics_cost:,.2f}")
print(f" Emissions: {emissions:,.2f} kg CO ")

# DPS
print("\nDirect Plant → Customer (DPS) Routes:")
for (p, c, v) in q:
    flow = float(prob.getSolution(q[(p, c, v)]))
    if flow > 1e-3:
        dps_count += 1
        arc_used = float(prob.getSolution(z[(p, c, v)]))
        dist = d_pc.get((p, c), 0)
        cost = arc_used * c_pc_CF_small.get((p, c), 0) + flow * c_load_plant.get(p, 0)
        emissions_flow = gamma * dist * (flow / Q)
        emissions_arc = gamma * dist

        dps_total_cost += cost
        dps_emissions_total += emissions_flow
        dps_emissions_arc += emissions_arc
        total_distance += dist

        print(f"  - Plant {p} → Customer {c} | Vehicle {v} | Flow: {flow:.1f} | "
              f"Cost: €{cost:.2f} | Emissions: {emissions_flow:.1f} kg (flow-based), "
              f"{emissions_arc:.1f} kg (arc-based) | Distance: {dist:.1f} km")

# 2E-routes
print("\nTwo-Echelon Routes (Plant → DC → Customer):")
for (p, d, v) in theta_pd:
    flow_pd = float(prob.getSolution(theta_pd[(p, d, v)]))
    if flow_pd > 1e-3:
        for (d2, c, w) in theta_dc:
            if d2 == d:
                flow_dc = float(prob.getSolution(theta_dc[(d, c, w)]))
                if flow_dc > 1e-3:
                    e2_count += 1
                    cost = (
                        flow_pd * (c_load_plant.get(p, 0) + c_unload.get(d, 0)) +
                        (flow_pd / Q) * c_pd_ISF_small.get((p, d), 0) +
                        flow_dc * c_load_dc.get(d, 0) +
```

```
                        (flow_dc / Q) * c_dc_CF_small.get((d, c), 0)
                    )
                    dist_pd = d_pd.get((p, d), 0)
                    dist_dc = d_dc.get((d, c), 0)
                    segment_distance = dist_pd + dist_dc
                    emissions_flow = gamma * segment_distance * (flow_dc / Q)
                    emissions_arc = gamma * segment_distance

                    e2_total_cost += cost
                    e2_emissions_total += emissions_flow
                    e2_emissions_arc += emissions_arc
                    total_distance += segment_distance

                    print(f"  - Plant {p} → DC {d} → Customer {c} | Vehicles {v}/{w} | "
                          f"Flow: {flow_dc:.1f} | Cost: €{cost:.2f} | Emissions: {emissions_flow:.1f} kg (flow), "
                          f"{emissions_arc:.1f} kg (arc) | Distance: {segment_distance:.1f} km")

print("\nSummary:")
print(f"   DPS Total Cost: €{dps_total_cost:,.2f}")
print(f"   2E Total Cost: €{e2_total_cost:,.2f}")
print(f"   DPS Emissions: {dps_emissions_total:.1f} kg (flow), {dps_emissions_arc:.1f} kg (arc)")
print(f"   2E Emissions: {e2_emissions_total:.1f} kg (flow), {e2_emissions_arc:.1f} kg (arc)")
print(f"   Total Emissions (hybrid): {e2_emissions_total + dps_emissions_arc:.1f} kg")
print(f"   Total Emissions (arc-based): {e2_emissions_arc + dps_emissions_arc:.1f} kg")
print(f"   Total Distance Traveled: {total_distance:.1f} km")
print(f"   # DPS Shipments: {dps_count}")
print(f"   # 2E Shipments: {e2_count}")

# SERVICE LEVEL PERFORMANCE
# --- Estimated Service Level Calculation ---

# Historical service level per origin
SL_Plant 2 =
SL_Plant 1 =
SL_DC 1 =

# Initialize shipped volumes per source
Vol_Plant 2 = 0
Vol_Plant 1 = 0
Vol_DC 1 = 0

# DPS volumes (from plants)
for (p, c, v), var in q.items():
    flow = prob.getSolution(var)
    if flow > 1e-3:
        if p == "Plant 2":
            Vol_Plant 2 += flow
        elif p == "Plant 1":
            Vol_Plant 1 += flow
# Two-echelon volumes (from DC 1)
for (d, c, w), var in theta_dc.items():
    flow = prob.getSolution(var)
    if flow > 1e-3 and d == "DC 1":
        Vol_DC 1 += flow
# Total shipped volume
total_shipped = Vol_Plant 2 + Vol_Plant 1 + Vol_DC 1

# Weighted service level calculation
if total_shipped > 0:
    SL_estimated = (
        SL_Plant 2 * Vol_Plant 2 +
        SL_Plant 1 * Vol_Plant 1 +
        SL_DC 1 * Vol_DC 1
```

```
    ) / total_shipped
3
    print(f"\nEstimated Service Level (weighted by origin): {SL_estimated:.4f}")
    print(f"    Volume from Plant 2:   {Vol_Plant 2:.1f}")
    print(f"    Volume from Plant 1: {Vol_Plant 1:.1f}")
    print(f"    Volume from DC 1 DC: {Vol_DC 1:.1f}")
else:
    print("\nNo shipments executed - cannot estimate service level.")
```