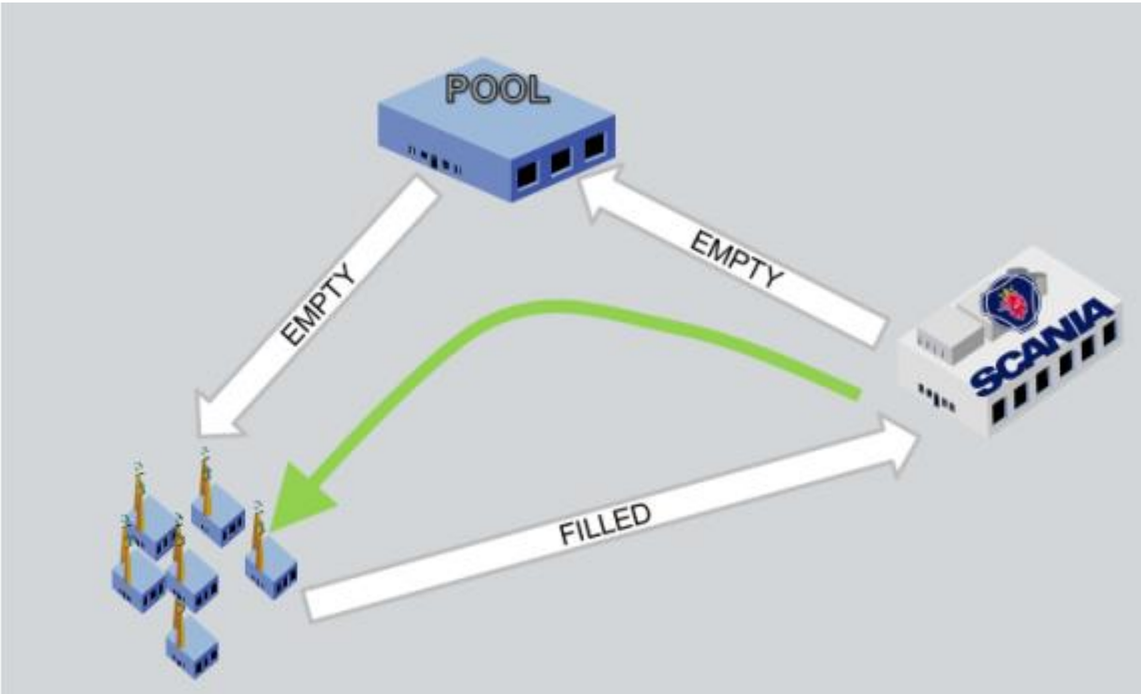

Allocation flow patterns and cost-aware routing in packaging networks

A data-driven decision support concept based on a case study at Scania

Master's thesis in Supply Chain Management

Adam Seraj
Viswanath Subramanian

DEPARTMENT OF TECHNOLOGY MANAGEMENT AND ECONOMICS
DIVISION OF SUPPLY AND OPERATION MANAGEMENT

CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026
www.chalmers.se

Allocation flow patterns and cost-aware routing in packaging networks

An empirical analysis of pool-to-pool movements and pool-to-supplier routing alternatives aimed at improving cost visibility and supporting cost-aware route selection in Scania's packaging network.

Adam Seraj

Viswanath Subramanian

Department of Technology Management and Economics
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026

Allocation flow patterns and cost-aware routing in packaging networks

Adam Seraj
Viswanath Subramanian

Department of Technology Management and Economics
Chalmers University of Technology
SE-412 96 Gothenburg Sweden
Telephone + 46 (0)31-772 1000

CONTENT PAGES

Contents

Abstract

Background – Scania a major manufacturer of heavy commercial vehicles, operates a European returnable packaging network composed of four interconnected pools that supply transport items to suppliers across Europe. Internal allocation orders between these pools generate substantial transport activity that the company's current planning department handles through manual decisions made under limited visibility of the underlying cost structure. In a network where pool-to-pool allocation movements account for approximately 11 per cent of the analysed transport loads, and where contracted prices vary substantially across destinations and shipment sizes, ensuring that direct and allocation-based routing decisions are made on a cost-aware basis is important for both operational efficiency and environmental performance.

Purpose – The thesis aims to develop a data-driven decision support concept that helps planners evaluate cost-aware transport routing options now of planning. The work has two interconnected aims. The first is to characterize the structural patterns of internal allocation flows between Scania's four packaging pools (Eschweiler, Angers, Järna and Oskarshamn) and to compare how pool-to-supplier transport costs differ across destinations and shipment sizes when historical and contracted pricing for standard, mega, and MSRT trailers are placed side by side are placed side by side. The second is to translate these insights into a prototype tool that demonstrates the analytical logic for a future production system to build on at the case company.

Method – The study combines quantitative analysis of historical transport data with engagement with the contracted rate sheet and qualitative input gathered through recurring discussions with transport planners at the case company. The empirical foundation consists of approximately 35,000 transport records covering March 2025 to March 2026 together with the company's rate sheet defining contracted prices across fourteen load meter brackets for more than 1,000 pool-to-destination combinations. Analysis was conducted at network, lane and shipment level. A spreadsheet-based prototype was developed iteratively to demonstrate how empirical findings can be operationalized at the level of the individual routing decision.

Results and Discussion – The analysis reveals that allocation flows between the four pools are highly asymmetric. A single corridor between Angers and Eschweiler carries close to half of all internal. Pool roles remain stable across every month of the studied period, indicating that the patterns are structural rather than transient and they always have differed demands in all the studied period of time frame. The comparison between historical and contracted pricing shows that historical weighted averages exceed rate sheet full truck prices on 83 per cent of comparable lanes, with average premiums ranging from 44 per cent at Järna to

135 per cent at Eschweiler. A separate allocation-order dataset reveals a boomerang pattern in which bilateral movements occur on the same corridor within the same week, indicating that approximately 14 per cent of the analysed allocation transport cost may be linked to potentially avoidable opposite-direction movements under a conservative minimum-overlap assumption. The decision support concept integrates these findings into a single planner facing tool that displays both pricing perspectives and ranks the several possible routing options for each supplier need and surfaces the information that is currently less visible at the moment of planning.

Recommendations for Scania – Based on the results of the analysis, Scania is recommended to develop the prototype into a production version built on a platform such as Power BI or internal Scania transport management tool (NILE) that connects directly to transport management and environmental reporting systems. A KPI capturing bilateral redundancy at the lane level should be introduced to make the boomerang pattern operationally visible. Cross-pool coordination should be embedded organizationally so that planners can see each other's activity rather than relying on individual initiative. Once route-level cost differences are made visible, Scania can use the results as input for future discussions with carriers and pool operators about pricing, responsibilities and incentive structures.

Keywords – Packaging logistics, allocation routing, decision support concept, returnable transport items, rate sheet pricing, multi-pool networks.

Acknowledgements

This thesis owes a great deal to several people whose contributions shaped what it became.

At Scania we want to thank our manager Felipe Garcia Tamosauskas, whose tips and feedback throughout the project helped shape the analytical direction of the work. Our mentors Marcus Söderberg and Kimia Mirzaei made the empirical part of the thesis possible by providing the data, the contractual material, and the operational knowledge of the packaging network that the analysis builds on. They also helped us understand how the patterns we saw in the data connected to the realities of daily planning. We are also grateful to the wider OA team at Scania whose discussions and willingness to share insight from their daily work gave the project a depth it could not have reached without them.

At Chalmers we want to thank our supervisor and examiner Violeta Roso who guided the academic side of the project. Her feedback at each stage helped us narrow a broad initial question into a focused contribution and her insistence on analytical clarity is reflected throughout the chapters.

Finally, we want to acknowledge each other. Writing a thesis together is a different kind of work from writing one alone. It requires that two people can disagree productively, can pick up where the other has left off and can keep moving when the analysis turns difficult. Through every phase of this project, we have stood by each other through the long evenings, the frustrating dead ends and the breakthroughs that made the work worth doing. The thesis that follows is genuinely the product of both of us.

1. Introduction

This chapter introduces the context, purpose of the study and describes the issues inherent in this project. At first the industrial background and the importance of efficient planning in complex logistics networks at companies are presented. This is followed by a discussion of the challenges that arise in connection with planning packaging flows and selecting cost-efficient routing alternatives across the pool network. The chapter continues by formulating the purpose of the study and the research questions that guide the analysis and results. To conclude the study's boundaries are defined and the overall structure of the report is presented to provide a clear overview of the work together with the project's boundaries. The chapter also introduces the decision support concept that emerges from the study to clarifies the role of the developed tool as a prototype that demonstrates how data-driven cost evaluation can support planners in their daily operational tasks.

1.1. Background

In modern industrial supply chains planning and coordinating logistical flows has become increasingly important. Manufacturing companies today operate in complex networks where materials, information and resources move continuously between production, suppliers and various logistics nodes. In such systems even relatively small planning decisions can affect transportation costs, resource utilization and delivery stability throughout the network. Research in supply chain management shows that transportation planning and network control are central factors in achieving both cost efficiency and operational stability in global supply chains (Rodrigue, 2020; Chopra & Meindl, 2019). At the same time demands for sustainability are increasing which means that companies increasingly need to understand how their logistical decisions affect emissions, energy use and resource consumption (McKinnon, Browne, Whiteing, & Piecyk, 2015).

A component that often plays a fundamental role in industrial supply chains is packaging. In many manufacturing systems, packaging is not a disposable product but a resource that circulates between different actors in the network. This is usually referred as returnable packaging systems. In such systems packaging is used several times and transported back to suppliers or storage points after being emptied in production. Studies in the field of packaging logistics show that a well-functioning packaging system can contribute to better transport utilization, lower material costs and reduced environmental impact by reusing packaging instead of replacing it with new ones (Hellström & Saghir, 2007; Twede & Clarke, 2004). At the same time shortcomings in planning and control of packaging flows can lead to the opposite effect where avoidable transport occurs and where a shortage of packaging will risk disrupting production.

In many industrial networks, the handling of returnable packaging is organized through central storage points or so called packaging pools. These packaging pools function as nodes where empty packaging is collected, stored and distributed to suppliers as needed. From a logistical perspective, such systems can be analysed as networks where flows occur between different nodes with varying demand and supply. Research shows that planning these flows often involves trade-offs between transportation costs, inventory levels and service levels (Simchi Levy, Kaminsky, & Simchi Levy, 2008). When demand and supply of packaging vary between different parts of the network, there is often a need for redistribution between storage points which in turn generates internal transport flows.

Scania is a global manufacturing company within heavy vehicles and industrial engines where suppliers in Europe deliver thousands of different components to the company's production units. In this system returnable packaging plays a central role because components are transported in specially adapted packaging that is owned and managed by Scania. This packaging always circulates between suppliers and production units which means that the availability of empty packaging must be planned and coordinated to ensure that suppliers can deliver components on time. The packaging flow can therefore be considered a parallel logistical system that enables the flow of materials to production.

Despite the strategic importance of packaging flows, planning is often based on established working methods and operational assessments rather than on systematic analysis of how different planning decisions affect transport performance and resource utilization. This means that internal redistributions of packaging between storage points can occur as a natural part of daily planning without there always being a clear picture of how these movements affect the overall efficiency of the network. Previous research has shown that data-driven analyses of logistics flows can help identify inefficient transport patterns and thereby create better decision-making information for planning and control (Tagaro et al., 2023; Domínguez Caamaño et al., 2017).

Against this background, there is a need to analyse how planning decisions within packaging flows impact transport costs, resource utilization and environmental impact in an industrial network. By studying how packaging moves between different nodes in the network and how planning decisions are made in practice, it may be feasible to identify improvement opportunities and develop more structured guidelines for planning such flows. This thesis takes such a data-driven approach by examining how internal allocation orders occur between four packaging pools in a real industrial network and by combining this with an analysis of the contracted transport pricing structure that governs deliveries from each pool to its surrounding supplier base. The combination of these two perspectives creates the foundation for a decision support concept that the study develops as one of its main outputs

1.2. Problem discussion

In Scania's packaging network empty packaging circulates continuously between supplier, pools and production sites around Europe. To ensure that suppliers can pack and deliver components to Scania production site, the right packaging type must be available in the right pool at the right time to be delivered to the supplier. When the supply in a pool is insufficient then there is a need for internal transfers between pools. These internal movements are created as allocation orders and planned in NILE by planners. Daily planning is largely based on manual decisions where the planner assesses availability as well as demand and upcoming delivery needs. In practice this means that internal transfers are often initiated when a specific need arises or when a surplus is identified in a pool. Since decisions are made continuously based on the current situation, internal transfers can occur as a natural part of the working method.

At the same time each allocation transport involves an extra movement in the network that affects both cost and transport work. At present, there is limited structured understanding of how these planning decisions affect the total transport flow and it is not permanently clear whether all internal movements are needed or whether some movements arise due to planning logic rather than actual needs. The cost dimension of these decisions is also more complex than it first appears. Transport pricing in such networks is rarely linear in shipment size, since contracted rates typically follow a bracket structure where the cost per load meter is high for small shipments and decreases until a full truck load top limit is reached. As a result, planners face decisions where the choice of source pool affects not only the allocation movement itself but also the downstream cost of delivering the packaging onward to the supplier. Without structured visibility into both the historical and the contracted side of these costs, it is difficult to evaluate whether a particular flow of routing is economical. Against this background there is a need to study in more detail how planning decisions within packaging flows affect transport work and resource utilization and how data-driven decision support could be designed to make the cost implications of allocation routing visible to the planner at the moment of the decision.

1.3. Purpose

The purpose of this study is to develop a data-driven decision support concept that helps planners evaluate cost-aware transport routing in a multi-pool packaging network. The study has three connected aims. The first is to understand the routing structure created by direct pool-to-supplier deliveries and allocation-based routes involving pool-to-pool movements. The second is to investigate how the cost and indicative CO_2 impact of these routing alternatives vary across destinations and shipment sizes when historical transport costs and contracted rate-card benchmarks are compared. The third is to translate this evaluation logic into a spreadsheet-based prototype that demonstrates how routing options can be evaluated and ranked at the moment of planning. The concept is presented as a prototype intended to serve as a basis for further development at the case company.

1.4. Research questions

The research questions are formulated to support the purpose of developing a decision support concept for cost-aware routing in Scania's packaging pool network. Since the project combines historical transport analysis, cost comparison and tool development, the questions are divided into three parts. The first question focuses on the structure of the routing alternatives, the second on how these alternatives can be evaluated in terms of cost and indicative CO_2 impact, and the third on how the evaluation logic can be translated into a planner-facing prototype. Before cost can be evaluated, it is necessary to understand which routing alternatives exist in the network and how they are shaped by pool-to-pool and pool-to-supplier transport relations. This motivates the first research question:

RQ1a. *What direct and allocation based routing alternatives exist between Scania's packaging pools and supplier postal-code zones?*

Once the routing alternatives have been identified, the next step is to evaluate how their cost and indicative environmental impact differ across destinations and shipment sizes. This requires comparison between historical transport cost, contracted rate-card benchmarks and estimated CO_2 values. This motivates the second research question:

RQ1b. *How do the estimated costs and indicative CO_2 impacts of these routing alternatives vary across destinations and shipment sizes when historical transport costs and contracted rate-card benchmarks are compared?*

To conclude the study aims not only to analyse the network but also to make the results usable for planners. The third question therefore concerns how the evaluation logic can be operationalised in a practical decision support tool:

RQ2. *How can the routing and cost evaluation logic be operationalised into a planner-facing decision support tool that ranks routing alternatives and improves cost visibility at the moment of planning?*

1.5. Scope and delimitations

The study covers the routing decisions involved in delivering packaging from the pool network to the suppliers it serves. The unit of analysis is the routing option rather than the individual allocation order. Direct deliveries from a single pool to a supplier are therefore treated alongside allocation routings where packaging is first moved between two pools before being dispatched onward. Both legs of these combined routings are within scope: the internal transfer between pools and the subsequent delivery from the dispatching pool to the supplier. Estimated CO_2 emissions are included as a parallel measure for each routing option. The study is limited to the four pools that form Scania's European packaging network: Eschweiler Pool in Germany, Angers Pool in France, Järna Pool and Oskarshamn Pool in Sweden. All transport relationships originating from these four pools are covered by the analysis. The empirical material refers to a 12 month period from March 2025 to March 2026.

The study draws on two parallel data sources. The first is a historical transport dataset that records actual shipments between pools and from pools to suppliers across the studied period, with associated load meters, costs and carrier information. The second is the contracted rate sheet provided by the company which specifies the official transport cost for each pool-to-supplier relation across a sequence of load meter sets. The combination of these two sources allows the study to examine both how transport costs have looked in practice and how they are structured contractually to compare the two perspectives where they differ.

The goal is to develop a decision support concept that helps planners evaluate the cost and CO_2 emissions of the routing options available for each supplier need. The ambition is instead to develop an analytical basis that can be used for future planning support, for example in the form of simpler guiding principle and limit values or practical decision rules that can be applied with limited resources. The study is data driven and aims to identify structural cost differences rather than to test a quantitative optimization model. The deliverable is a decision support concept implemented as a working prototype in spreadsheet form proposed to show the analytical logic to provide a basis for further development at the case company in a more scalable system environment.

Cost is the main starting point for the study while CO2 emission is treated as a secondary but still relevant viewpoint. However, the study does not aim to carry out a complete cost modelling of the entire network but rather to evaluate how transport planning decisions can be linked to optimize packaging flow. Furthermore, the study does not cover the entire packaging flow up to production. The internal management of production and effects in the production system are therefore not involved. The study also does not include real-time inventory calculations or forecast modelling and analysis of carrier agreements or detailed system implementation.

Some qualitative information from planners is also used to understand the data. This part is mainly used to understand the work process to provide context to observed deviations but does not constitute the main empirical basis of the study. There are some data related limitations, especially since complete follow-up of inventory balances is not available in the material used. This means that some underlying reasons for individual allocation decisions cannot be fully determined only via quantitative data.

1.6. Thesis structure

The remainder of this report is organized as follows. Chapter 2 establishes the frame of reference drawing on supply chain network theory, literature on returnable transport items, operational planning under uncertainty, organizational decision boundaries and transport cost structures. Chapter 3 describes the methodology including the research design for this project, the case study context, the evolution of study scope and the analytical approach. Chapter 4 presents the empirical findings which start with a description of the packaging network and data foundation and pool-level imbalances, pool-to-supplier cost structures and the comparison between historical and contracted pricing. Chapter 5 develops the analysis further by interpreting the patterns observed in Chapter 4 in light of the frame of reference and presenting the design of the decision support concept. Chapter 6 discusses the implications of the findings for planning practice and identifies the limits of the proposed concept. Chapter 7 concludes the study with a summary of contributions and recommendations for further development at the case company.

2. Frame of reference

The thoughts of this chapter are grounded in the practical problems identified in Chapter 1. Rather than adopting the conventional path of verifying established theories this chapter constructs its analytical framework based on existing supply chain literature, to examine the formation of allocation flows between packaging pools and the logic through which planning decisions shape these flows. All concepts introduced in each section of this chapter are developed to support subsequent analyses. First chapter 2.1 defines packaging pools as connected nodes within the supply chain network. Section 2.2 on internal redistribution explains inventory imbalance. Section 2.3 on planning logic explores operational transport decisions under uncertainty. Section 2.4 on boundaries of change fits in with theories of power and governance and then there is the point about transport cost structures covering how freight costs are built up empirically as well as organised contractually through rate sheet pricing enabling evaluation of the decision support concept developed during this study. The chapter concludes with a synthesis that integrates these concepts and demonstrates how they will be applied in the analysis to follow.

2.1. Supply chain networks and pool-based logistics

The packaging system studied in this project drives as a logistical network rather than a single transport process. So, to say various storage locations (pools) are connected to each other also to a wide range of suppliers afterwards the resulting flows form patterns that cannot be described by examining one transport at a time. This section therefore presents a network perspective on supply chains and combines it with literature on pool-based logistics for returnable transport packaging. The network perspective provides the structural language needed to describe pools, suppliers and transport links as connected elements of a wider system while the literature on pool-based logistics offers the operational concepts needed to understand how returnable packaging is managed across numerous locations. Together these two flows of literature form the basis for the analysis of allocation flows in the later chapters. The approach of seeing supply chains as networks has grown in support within the supply chain management literature. Carter, Rogers, and Choi (2015) assert that supply chain thinking should move beyond the notion of a linear sequence of buyer-supplier relationships and consider the supply chain as a complex network where the characteristics of the system are emergent from the relationship structure of its components. A similar viewpoint is adopted by Bellamy and Basole (2013) in their literature review of supply chain network research where they consider connection and node interdependence as core structural characteristics and key determinants of outcomes. This is particularly relevant for the scenario examined within this thesis, as the four packaging pools are not distinct entities but interdependent nodes with interconnected behaviours that contribute to empty packaging movement. The frame of the broader network can be narrowed down to that of pool-based logistics, the stream of literature which describes packaging reuse in closed loop systems where pooling points are established for storage, redirection and return (Glock, 2017; Hellström & Johansson, 2010; Kroon & Vrijens, 1995). This forms the basis for understanding pools as having structural, logistic value as something more than passive storage units, which is detail in subsequent subsections on node structure and the various redistributive activities.

2.1.1. The network perspective on supply chains

The network approach to supply chains has been substituted for the more traditional notion of the linear chain. Carter, Rogers, and Choi (2015), refer to this shift as taking the perspective of the supply chain as a network of firms interconnected by multiple flows often dependent on one another, instead of a sequence of dyadic buyer-supplier interactions. They describe the supply chain as bounded but internally complex and claim that the way the pieces are connected determines how the overall system responds. For example, an explanation for a cost, delay or imbalance at any point within a system might not be found in the actual shipment or transaction itself, but in the overall connectivity. A similar conclusion is made by Bellamy and Basole (2013), who present a systematic literature review of research applying the network approach to supply chains. They demonstrate that the network approach has been applied to investigate a number of issues, such as supplier relationships, information flow, and risk propagation, and also two related concepts that run through the different network research studies. These concepts are topology, the pattern of connections between nodes in a network, and interdependence between nodes whereby changes at one node in the network may ripple through to influence other nodes even if the changes at those individual nodes do not seem directly affected by what happens at the original node. These two concepts provide the terminology with which the packaging system will be described: the network of pools, the suppliers that provide their product and the transport relationships that link them together will be treated as nodes and connections of a network system whose overall behaviour, including the generation of allocation flows, will be determined by the pattern of the links.

2.1.2. Pools as buffering and redistribution nodes

In a logistical network not, all nodes have the same role. Some nodes function primarily as transit points where goods pass through with little change in form or storage time, while others actively hold buffer or redistribute resources to compensate for variation in supply and demand. Rodrigue (2020) describes this distinction in terms of the functional position of nodes within the wider transport system, where certain locations are designed to absorb sequential and spatial mismatches between where resources become available and where they are needed. Simchi-Levi, Kaminsky and Simchi-Levi (2008) similarly note that holding inventory at intermediate locations is one of the standard mechanisms for handling uncertainty and variability in supply networks, since stock at such locations can be used to smooth the gap between supply timing and demand timing across the system. The pools studied in this thesis correspond closely to such buffering and redistribution nodes. Each pool stores empty packaging that has been returned from production and forwards it to suppliers when needed, which means the pool absorbs short term mismatches between when packaging becomes available and when it is required at a supplier location. The literature on returnable transport items frames this function in similar terms. Glock (2017) describes pools as central locations in closed loop systems where reusable packaging is collected, held and dispatched. While Hellström and Johansson (2010) distinguish between different control strategies for such systems and show that the way a pool is operated has direct consequences for inventory levels, transport movements and operating costs. A central implication of this literature is that pools are not passive storage points but active management nodes whose configuration shapes the rest of the network. When a pool runs short of a particular packaging type, the response is not limited to local adjustment, since the shortage can be addressed by moving stock from another pool. This redistribution function is the conceptual basis for understanding allocation orders, which are the focus of the empirical analysis in the later chapters of this thesis.

2.1.3. Structural asymmetries and routing dependencies

The operation of a logistics network is not determined only by what each node does in isolation but also by the way in which each node is linked to all others. As Bellamy and Basole (2013) highlight, while two identical networks in terms of the number of nodes could achieve vastly different outcomes given differences in network topology since the connectivity structure determines which flow combinations are possible which require one or more intermediate hops, and which are not possible at all. This structural nature of networks has become central to research in supply chain network structure; writers have found it to be instrumental in processes such as risk propagation, variability in lead time and the responsiveness of the system to perturbations (Bellamy & Basole, 2013; Carter et al., 2015).

The nature of the network structure that is particularly interesting here is that which introduces structural asymmetries some nodes will have several direct links and others only one or a few limited links. The result of such structural features is that some flow paths will be channelled along others, and others will need to take a longer, indirect route to an intermediate point if the required direct link doesn't exist. The results of this are more than just increased distance: When a destination can only be reached via one intermediate location, that intermediate node becomes a structural bottleneck as decisions on the flow to the less connected node must pass through it. This creates a routing dependency and directly links the operational performance of what appear to be separate entities. These structural conditions have great importance to the following analysis since the four packaging pools are not linked symmetrically to one another, some pools have very robust connections while others are linked by a very restricted set of connections.

2.2. Pool imbalances and the role of internal redistribution

While the previous section established that pools function as buffering nodes within a connected network it did not explain why the demand placed on these pools varies across the network or why this variation gives rise to internal movements between them. Both questions are central to this project since allocation orders are by definition a response to a mismatch between where packaging is available and where it is needed in the network. The literature on supply chain coordination provides quite a few explanations for why such mismatches happen in multi-location systems also why redistribution between locations becomes essential in practice. First explanation concerns the geographical distribution of demand, in networks where suppliers are spread across regions the pull on each pool is formed by the volume and frequency of orders placed by suppliers in its catchment area and these patterns are infrequently symmetric across the network. Pålsson and Hellström (2016) describe packaging logistics as a system where service levels at one location depend on coordination across the wider network and they claim that a regional view of inventory balance is necessary because demand side variation explains directly into uneven pull-on different parts of the system. When some regions place consistently heavier demand on their local pool than others, the affected pool depletes faster than its peers, even when the total volume of packaging in the network would in principle be sufficient to cover all demand.

A second explanation concerns the asymmetric circulation of returnable transport items. Glock (2017) and Hellström and Johansson (2010) point out that closed loop systems for reusable packaging do not generally return to the same location they were shipped from. Empty packaging is collected at production sites and routed back to whichever pool is operationally convenient while suppliers draw packaging from the pool that serves their region. The result is a structural drift in stock between pools that is not caused by planning errors but by the way the system is built. Over time some pools accumulate more inflow than outflow, while others accumulate the opposite pattern, and the system depends on internal movements to bring the distribution back into balance.

A third explanation concerns variation over time. Demand for individual packaging types is rarely constant, and short-term changes in supplier needs can overtake the buffering capacity of a single pool. The supply chain coordination literature treats internal redistribution between locations as one of the standard responses to such variability, since moving stock between buffering nodes is faster and less disruptive than altering production or procurement (Simchi-Levi et al., 2008). Internal redistribution is therefore not a sign that the system is malfunctioning but a built-in mechanism for handling the kind of variation that any multi location buffering system will face. The relevant question for analysis is then not whether redistribution occurs but whether the volume and pattern of redistribution is proportionate to the underlying imbalance or whether some redistribution movements arise for reasons that lie outside the structural conditions described in this section. This distinction frames the empirical investigation of allocation orders conducted in the later chapters of this thesis.

2.3. Operational planning decisions under uncertainty

As claimed in the former sections, imbalancing is a natural phenomenon in pool networks, and internal redistribution is an intrinsic part of their functioning. However, they failed to explain at which level this redistribution actually takes place. In particular at the level of the concrete planning decisions of those who are managing the actual flow on a day-to-day basis. Every order for allocation in a packaging network is the product of a concrete decision, during which a planner decides whether a particular movement should be triggered, from which pool to source from and how to prioritize the move over other ongoing movements. Literature on operational planning offers a lens through which to explain the environment of these decisions and why they may not be in proportion to the actual imbalancing occurring in the network. A prominent feature of operational planning is that planning is conducted under incomplete information. Simon (1955) first coined the concept of bounded rationality to define situations in which the decision maker cannot observe all workable solutions and therefore has to rely on simplifying heuristics and judgment based on current information. This concept has subsequently been adapted by research in operations and supply chain management which reasoned that decisions made during planning for the complex flow of goods must often be heuristics rather than formal optimization due to pressures of time, partial visibility and a continuous stream of new information that needs to be processed quickly (Gigerenzer & Gaissmaier, 2011; van der Vaart & van Donk, 2008). When planning a reallocation decision in a packaging network, this suggests that the planner will not typically have complete information on all stock levels, upcoming returns and required supplier quantities across all pools and therefore will decide based on the available information visible in the planning system at that moment.

A second characteristic of operational planning that is highly relevant in this context is the distinction between reactive and proactive planning. While a reactive move is planned and executed when an imbalance is visible (i.e., When a specific pool is approaching a stock-out), an anticipatory move is made prior to a predicted issue (i.e., When a shortage is likely in the near future). For some time, the supply chain coordination literature has observed that planners are naturally biased towards reactive decision making because the cost associated with premature actions is more clearly visible than those of acting too late, and also because reactive decisions can more easily be justified at a later stage (van der Vaart & van Donk, 2008). As such, movements in the distribution system tend to occur when a deficit is about to become a problem rather than when a deviation from the expected distribution of stock is detected.

Thirdly, the role of experience and rules of thumb plays an important part in planning decisions made at an operational level. Where no formal planning rules are provided, or the rules provided are not sufficiently detailed, the planner relies on her/his own informal rules and reasoning. Based on previous experience, when a certain situation arises, the planner may perform the same actions as performed when facing an identical situation. Klein (2008) calls this the concept of recognition-primed decision making; experienced decision makers classify a situation as similar to a situation encountered earlier and react in the same way as they have been successful doing before. This provides an efficient way of planning but can result in a few implications for the present study. First, in such cases it is very difficult to reconstruct why a certain decision was made and thus, why the corresponding order was generated. Secondly, in the absence of explicit, formalized decision rules, two planners facing an identical situation will not necessarily arrive at the same decision. These three aspects, incomplete information, reactive decision making and rules of thumb, serve as the foundation for the subsequent analysis. They also motivate the type of analytical support that this study develops. When operational decisions are shaped by partial information, time pressure and informal rules, a structured decision support concept can play a useful role by making the cost consequences of alternative actions visible at the moment of planning. The literature on decision support systems describes such tools as analytical aids that complement rather than replace human judgement, and as instruments that help decision makers see the implications of choices that would otherwise rely on intuition or experience alone (Power, 2008; Arnott & Pervan, 2014). The decision support concept developed in this study is positioned within this tradition, as a prototype that demonstrates how routing options for allocation orders can be structured and ranked at the moment of decision.

2.4. Power, governance and decision boundaries in supply relationships

In supply chain networks involving multiple actors the distribution of power between buyers and suppliers is rarely equal. This asymmetry has important implications for how decisions are made, who is able to influence operational arrangements, and which changes can realistically be implemented. Several scholars argue that power relations in supply chains are not only a relational issue but a structural condition that shapes coordination patterns, adaptation processes and operational outcomes over time (Benton & Maloni, 2005; Nyaga et al., 2013). When one actor controls critical resources, information or access that other actors depend on the behaviour of the wider relationship is influenced in ways that may not be visible from an entirely logistical perspective. For this reason, power needs to be understood together with governance and decision authority. Governance mechanisms do not only coordinate transactions between actors they also reflect how control is distributed within the relationship and where the boundaries of action are located in practice. A theoretical understanding of buyer supplier power, contractual governance and operational decision boundaries is therefore necessary in order to explain why some forms of change remain possible within a given supply relationship, while others are constrained by the structure of control itself (Lacoste & Johnsen, 2015; Nyaga et al., 2013).

2.4.1. Power asymmetry in buyer & supplier relationships

Power asymmetry in buyer & supplier relationships arises when one party has a clearer advantage over the other and thus a greater opportunity to influence the terms, behaviours, and outcomes of the relationship. Such an imbalance is not an unusual allowance but rather a recurring feature in supply chains where the actors differ in size, expertise, degree of dependence, switching costs or control over important resources. Nyaga et al. show that relationships between buyers and suppliers often develop under precisely such asymmetric conditions and that this affects both cooperative behaviour and operational results in the relationship. The weaker party then becomes more vulnerable, not only to opportunism but also to having to adapt to requirements that it itself has not been involved in shaping (Nyaga et al., 2013).

In the industrial literature, power is described not only as formal control but also as the ability to get another actor to act in a certain way because this actor is dependent on the relationship. Benton and Maloni show in their study from the automotive industry that buyers in many cases have a structurally stronger position than their suppliers and that this imbalance affects how the relationship is experienced and how the supplier acts. The study also shows that different power bases produce different effects. When power is used in a way that is perceived as coercive or one-sided, the relationship risks being characterized by lower satisfaction and weaker long-term sustainability. This means that power imbalance is not only about who is in charge but also about how the relationship is shaped over time through the stronger party's way of exercising its influence (Benton & Maloni, 2005).

Lacoste and Johnsen provide a slightly different perspective by showing that power difference is not completely static even in relationships where the supplier initially has a weak position. Their study of packaging suppliers and strategic customers describes how low power on the part of the supplier can be met by the supplier gradually building up a more process-oriented offering and thereby creating increased mutual dependence. The interesting thing here is that power is not only understood as something that follows from volume or size, but also as something that can change when one party becomes more important to the other party's processes. At the same time, the starting point remains that the party that initially controls the central conditions of the relationship also has a clear advantage in how the interaction is organized and developed (Lacoste & Johnsen, 2015).

Lacoste and Johnsen show that a supplier with a weak starting position does not necessarily remain stuck in a low-power position. By developing more process-oriented and performance-related offerings, the supplier can in some cases create increased interdependence and thereby build up a form of countervailing power. This is illustrated in Figure 1.

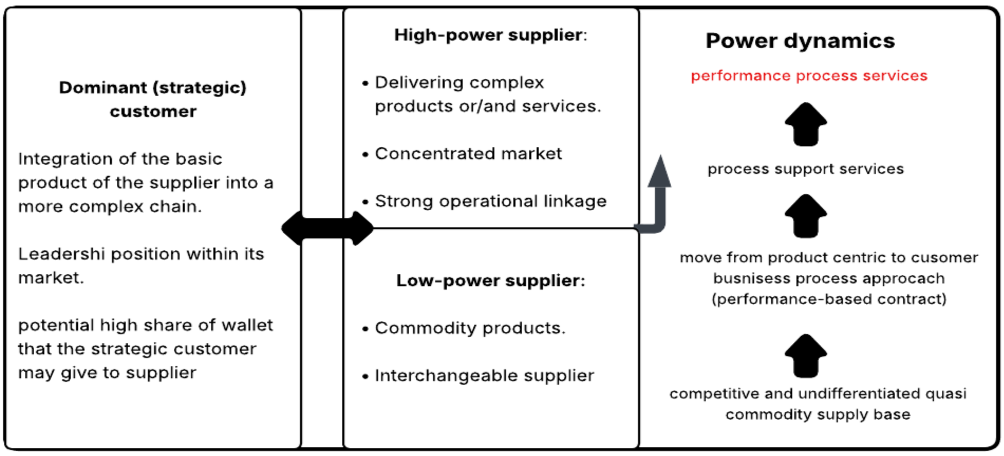

Figure 1. How a low power supplier may develop countervailing power through process support. Adapted from Lacoste and

The figure shows how power imbalances do not need to be understood as completely static. Instead, it illustrates how a supplier can try to move away from a pure commodity position by linking its offering more closely to the customer's processes and performance. In this way, the relationship becomes less one-sidedly driven by customer dominance and more characterized by mutual dependence.

Together, these studies point to the need to understand buyer-supplier relationships as relationships where influence is unevenly distributed from the outset. This inequality affects not only negotiation and cooperation, but also which actors can drive change and who are primarily expected to adapt. Power asymmetry therefore becomes an important theoretical starting point for understanding why certain operational patterns can persist even when they appear less rational from a purely logistical perspective.

2.4.2. Governance structures and contractual constraints

Governance in buyer-supplier relationships is about how the exchange is organized, what rules apply and how responsibilities are distributed between the parties. It therefore encompasses more than formal contracts in the narrow sense. Governance also concerns which processes govern cooperation, how follow-up takes place and which mechanisms are used to ensure that one party delivers according to the other party's requirements. In relationships where one party has a stronger position governance takes on a particularly important role since it not only coordinates flow but also clarifies who has the right to define conditions, standards and working methods in the relationship (Nyaga et al., 2013).

Nyaga et al. show that governance in asymmetric relationships can be understood through different power bases some of which are more direct and formal than others. Mediated power such as legal and coercive power is based on a party being able to refer to contracts, formal requirements or sanctions to influence the other party's behaviour. Non-mediated power such as expert and referent power is instead based more on competence, legitimacy and trust. The important point here is that governance is not only about whether a contract exists, but also about what type of control is actually used in the relationship. Formal governance can thus collaborate with more informal forms of influence, and the quality of the relationship affects which forms of governance are possible or effective in practice (Nyaga et al., 2013).

Lacoste and Johnsen contribute a more process-oriented perspective by showing that governance can also change when the relationship develops from a pure product exchange to a more integrated process approach. When the supplier becomes more involved in the customer’s business, for example through process support services or performance-related services, the relationship becomes less easy to control solely through standardized purchasing terms. This does not mean that the power imbalance disappears, but it does mean that the function of the contract changes. Instead of just regulating the delivery of a product, governance can start to be more about process responsibility, problem solving and joint performance. This also makes it clear that contractual constraints not only set limits on what can be done but also shape the type of improvements that can be actualized within the relationship at all (Lacoste & Johnsen, 2015).

Benton and Maloni correspondingly show that the buyer's influence in a supply chain is not only exercised through price or volume, but also through the structures that frame the relationship. Contractual and organizational arrangements can create a situation where the supplier is expected to adapt to given operational conditions, even when these conditions are not directly negotiated from time to time. Governance then becomes a way of stabilizing the relationship, but also a way of consolidating a certain power structure. This is particularly relevant in relationships where the supplier is dependent on the customer for continued business, since the formal framework can then have a strong impact even on everyday decisions and operational priorities (Benton & Maloni, 2005).

From a theoretical perspective this means that operational improvements in supply relationships cannot be understood as completely free technical or logistical choices. They are influenced by the governance structure on which the relationship is based, what obligations are already formalized, and how decision-making power is distributed between the parties. Contractual frameworks therefore function not only as background conditions, but as part of the structure that determines which changes are possible to drive and which in practice lie outside the operational scope of action.

2.4.3. operational decisions within limited decision authority

Operational decisions in buyer-supplier relationships are rarely made in a completely open decision-making room. Even when daily planning looks local and practical, it often takes place within a framework that has already been set by some other level in the relationship. This means that the actor working closest to the flow does not necessarily have a mandate to influence the more fundamental conditions behind the decision. Instead, much of the operational work is about adapting to given requirements, managing deviations and making everyday life work within already established limitations. Nyaga et al. show that behaviour in asymmetrical relationships is strongly shaped by how power is used, and that the weaker party often orients its actions according to what is accepted and possible rather than according to what would be optimal from a purely operational perspective (Nyaga et al., 2013).

This also means that many operational choices are not really full choices, but more limited action alternatives. For example, one party can influence the order, prioritization or implementation, but not necessarily the rules that determine why these particular options exist. Benton and Maloni show that power-driven relationships affect how suppliers act and how they perceive their own scope for action. When the purchasing party has a clear advantage, important frameworks tend to be already defined, which means that the other party works more reactively than formatively. This means that operational decisions often become short-term adaptive rather than strategically transformative, even when inefficiencies can actually be seen in the relationship (Benton & Maloni, 2005).

Lacoste and Johnsen also show that limited decision-making mandates do not always mean total passivity. A supplier with a weaker starting position can in some cases create greater scope for action by going deeper into the customer's processes and becoming more important to how the business works. When the relationship moves away from a simple product exchange and instead becomes more process-oriented, some operational decisions can gain greater weight and be linked more closely to joint performance. But even then, the scope for action changes gradually and not automatically. It is built through the relationship, not outside of it, and is still dependent on how much the stronger party allows responsibility and influence to be shifted in practice (Lacoste & Johnsen, 2015).

2.5. Transport cost drivers and performance measurement

In the transport literature cost and performance measurement are often treated as two parts of the same management issue. It is not enough to know what a transport costs, but measures are also required that make it possible to follow up on how the transport process actually works overtime. Domínguez Caamaño, García Arca, Fernández González and Prado (2017) describe KPIs as measures used to manage and control processes. At the same time, their literature analysis shows that transport-related indicators often revolve around lead time, delivery precision, transport cost and resource utilisation. They also show that such measures are not normally used in isolation, but as part of broader systems for performance measurement.

A similar approach is found in Gunasekaran, Patel and McGaughey (2004), who emphasize that performance measurement in supply chains needs to be linked to both planning, control and improvement work. Their framework is based on the fact that different measures are associated with different decision-making levels, strategic, tactical and operational, and that they must therefore be chosen based on the type of control they are to support. Against this background, cost drivers and performance measurement are treated here as closely linked concepts. The cost structure helps to explain what affects the financial outcomes of transport, while performance measures provide a framework for monitoring, comparing and interpreting these outcomes in practice.

2.5.1. Cost drivers in freight transport

Transport expenses are not described as a single homogeneous cost item but as a result of several different cost components that operate simultaneously. Izadi, Nabipour and Titidezh divide road transport costs into three broad categories, operational costs, value of time and external costs. Operational costs include both fixed and variable parts. Fixed costs do not change directly with how much the vehicle is used and include, for example, insurance, license fees, depreciation and various overhead costs. Variable costs, on the other hand, depend on actual use and include, among other things, fuel, tires, maintenance and repairs. In addition, they also highlight that time and reliability have their own economic value, since faster and more reliable transport in many cases motivates a higher willingness to pay on the part of the transport buyer. In this way, it becomes clear that the cost picture in freight transport is shaped both by physical resource consumption and by how the transport performance is experienced in terms of time and service quality (Izadi et al., 2020).

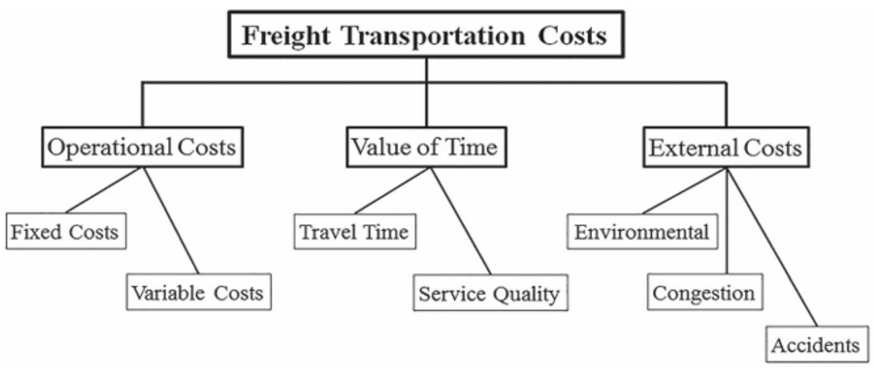

Figure 2. Overall model structure of road freight transportation costs. Reprinted from Izadi, A., Nabipour, M., & Titidezh, O. 2019.

As shown in Figure 2, freight transportation costs are not limited to direct operating expenses but also include time related factors and external costs. Conversely in operational road freight settings, the most immediate cost implications are often linked to the internal cost structure of transport execution, especially vehicle use, labour, fuel and capacity utilization. The figure therefore provides a broad theoretical basis, while the following discussion narrows the focus to those cost drivers that are most relevant for lane based and full truck load transport analysis.

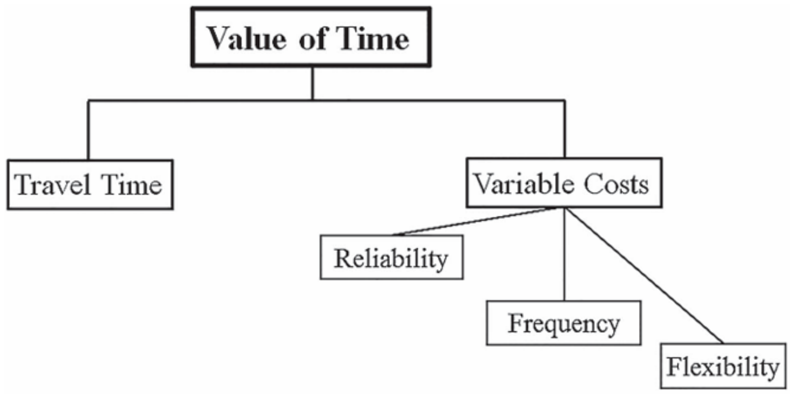

Figure 3. Model sub-structure of value of time in road freight transportation. Reprinted from Izadi, A., Nabipour, M., and Titidezh, O., 2020

Figure 3 clarifies that value of time in freight transport is not limited to travel time itself but also includes service-related dimensions such as reliability, frequency and flexibility. This is important because transport decisions are often shaped not only by direct operating cost, but also by how well a transport solution supports planning stability and service performance. The figure therefore helps explain why time related factors can influence transport choice even when they do not appear directly as a separate cost item in the shipment data.

This is important because it shows that cost drivers in road transport operate at different levels simultaneously. Some costs arise at the vehicle level, for example through usage time, mileage and wear and tear, while other costs are affected by how well the transport fits into a larger logistical context. For full truck load transports, this becomes quite clear. When a vehicle is in practice booked as a whole unit, it is not only the load itself that determines the cost but also how much of the vehicle's capacity is used and how efficiently it can be used on the current relationship. Against this background, weight and load meter become relevant cost proxies, as they capture two central dimensions of capacity utilization, namely how heavy the load is and how much physical space it takes up in the vehicle. They therefore provide a reasonable basis for analysing cost variation between shipments in an FTL network, even if they do not capture the entire cost picture on their own (Izadi et al., 2020).

At the same time, it is not enough to look only at the physical properties of the cargo. Zofío, Condeco Melhorado, Maroto Sánchez and Gutiérrez show that generalized transport costs are also affected by economic and infrastructural conditions that are linked to the transport relationship itself. In their model, the generalized transport cost is divided into economic costs and network-related variables, especially distance and time within the infrastructure. An important contribution from this literature is therefore that transport costs vary between different origin-destination relationships even when the cargo is similar, since different lanes have different geographical and infrastructural conditions. The cost is therefore shaped not only by how much goods are shipped, but also by which relationship the goods are shipped on and what temporal and spatial conditions apply there (Zofío et al., 2014).

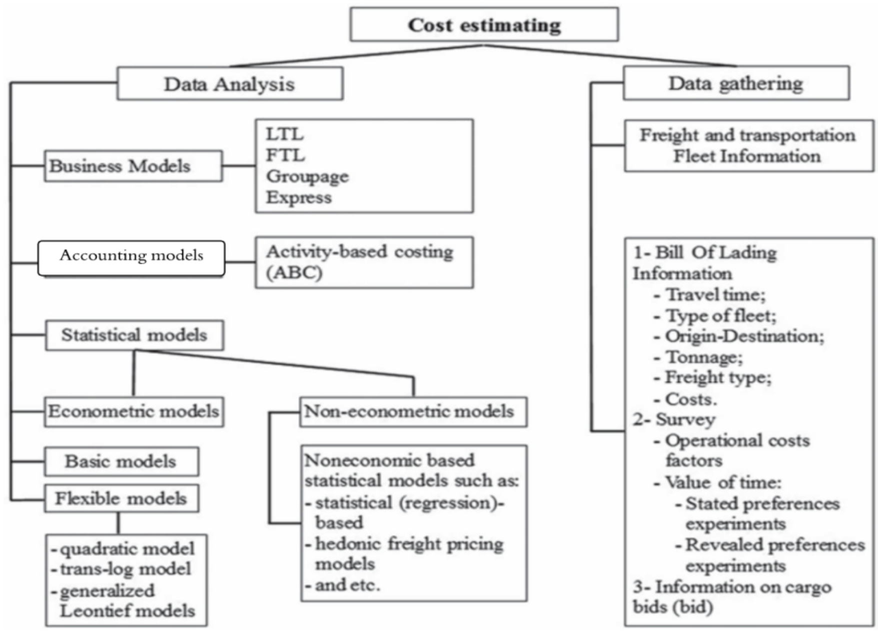

Figure 4. Overview of transport cost estimating approaches and main categories of input data. Reprinted from Izadi, A., Nabipour, M., and

Figure 4 shows that transport cost estimation is usually based on a combination of model structure and available transport data. What is important in this context is that several of the variables in the figure such as origin destination, travel time, tonnage and freight type, correspond closely to the kinds of factors that literature identifies as central cost drivers in road freight transport. The figure also makes clear that transport cost is not estimated from one isolated variable but from a set of operational and route related characteristics. This supports the use of lane level analysis together with shipment proxies such as weight and load meters when analysing cost patterns in a recurring transport network.

For the analysis of recurring internal transports, this becomes quite relevant. If the same type of packaging moves between different lanes, cost differences cannot be assumed to be due only to the weight or volume of the packaging. Part of the variation also lies in the structure of the routes, for example transport distance, travel time and how the infrastructure supports or limits the flow. Therefore, the lane level becomes an important analytical level when understanding transport costs from a network perspective. Weight and load meters can be used to represent how much capacity a particular shipment requires, but these measures need to be interpreted together with the relationship's own conditions to provide a more accurate picture of the cost drivers. It is precisely in the combination of capacity utilization and lane-specific conditions that the most relevant cost logic transports can be understood (Izadi et al., 2020; Zofío et al., 2014).

Based on this theory freight transport cost can therefore be understood as both resource-dependent and relationship-dependent. It is resource-dependent because vehicles, time and capacity are used and consumed in the transport itself. But it is also relationship-dependent, because each lane has its own cost structure that is shaped by geographical, temporal and infrastructural conditions. The theoretical conclusion is therefore that cost analysis in an FTL context should be based on the variables that best capture both of these levels. Weight and load meter are then particularly useful because they reflect the shipment's requirements for transport capacity, while lane analysis is needed to capture the structural differences that cause the cost to vary between different transport relationships, even when the load itself looks similar. However, the cost drivers discussed in this subsection describe the underlying determinants of freight cost rather than the pricing mechanisms through which these costs are actually translated into a price paid by the shipper. In practice, freight cost and freight price are related but not identical, since pricing is shaped by contractual structures, market conditions and the way capacity is allocated between shippers and carriers. For an analysis that aims to evaluate transport decisions from a planner's perspective, the pricing mechanism is often as important as the underlying cost driver, since it is the price that enters the decision at the moment of planning. The next subsection therefore introduces the pricing structure of freight transport and the role of rate sheet in establishing contracted prices for repeated transport relations.

2.5.2. Pricing structure and rate sheets in freight transport

Freight transport pricing is rarely a linear function of shipment size. In practice, the price paid by a shipper for moving a shipment varies non-linearly with how much capacity the shipment occupies, because the underlying transport operation involves both fixed components, such as driver time, vehicle deployment and lane setup, and variable components, such as fuel and capacity consumption. The result is a pricing structure where the cost per unit of capacity tends to be high for small shipments and decreases as the shipment grows, until a full truck load ceiling is reached and the price effectively becomes flat. This pattern is well documented in the literature on truckload and less-than-truckload transport, where it is generally understood that the unit economics of freight differ substantially between full truck and partial truck shipments (Caplice & Sheffi, 2003; Caplice, 2007).

The non-linear shape of the pricing curve has important consequences for how transport decisions can be evaluated. A linear approximation, in which the cost of a shipment is calculated as a single average rate per load meter multiplied by the shipment size, will systematically misrepresent the cost of small shipments and the cost of shipments close to a full truck. For very small shipments, the linear model understates the true cost because it ignores the higher unit prices that apply at low-capacity utilization. For shipments at or near full truck capacity, the linear model overstates the cost because it ignores the price ceiling that applies once a full truck has been booked. The practical implication is that an analytical model relying solely on average rates may give planners a misleading view of which routing options are economically attractive, particularly when the routing choice affects whether a given shipment is consolidated into a full truck or moved as a smaller load.

In repeated commercial transport, the pricing structure is typically formalized through a rate sheet, which is a contractual document that specifies the price the carrier will charge for transport between defined origin and destination points across a series of capacity brackets. Caplice (2007) describes such contracts as the dominant mechanism by which shippers and carriers structure their long-term transport relations, since they provide cost predictability for the shipper and demand stability for the carrier. A rate sheet is therefore not a marginal pricing instrument but the primary basis on which the shipper's transport budget is constructed, and it represents the cost that will actually be incurred when a shipment is dispatched under the agreement.

The distinction between historical and contracted pricing follows from this. Historical pricing describes what shipments have actually cost in the past, derived from observed transport data, and is therefore an empirical measure of average behaviour. Contracted pricing, expressed through the rate sheet, describes what shipments should cost under the agreement, and is therefore a normative measure of expected behaviour. The two perspectives are related but not identical, since historical pricing reflects the actual mix of shipment sizes, lanes and conditions that occurred during the observation period, while contracted pricing reflects the structural pricing schedule that governs each individual shipment. Differences between the two can arise when the actual mix of shipments is dominated by a particular bracket, when the data contain non-standard movements, or when the rate sheet has been renegotiated during the period in question. For the analysis carried out in this thesis, both perspectives are used in parallel, since they answer different questions: historical pricing characterizes what has happened, while contracted pricing characterizes what is expected to happen under the agreed terms.

2.5.3. KPI use in transport planning and control

KPIs in transport planning and control are not just about following up on results subsequently but rather about creating and examining a basis for ongoing decisions and prioritizing corrective actions. In a well-functioning control system, performance measurement should help the organization formulate goals, compare planned and actual outcomes besides identify when and where actions need to be taken. Gunasekaran, Patel, and McGaughey emphasize that performance measures affect strategic, tactical and operational planning at the same time, in addition they must therefore be designed to support control at multiple levels in the supply chain, not just local optimization in a single function (Gunasekaran et al., 2004). At the same time, Piotrowicz and Cuthbertson show that companies in practice often use KPI systems through frameworks such as the Balanced Scorecard, SCOR and some similar structural models precisely to connect observing with control and decision-making across multiple organizational levels (Piotrowicz & Cuthbertson, 2015).

Based on Gunasekaran et al., a well-functioning KPI system in transport planning can be described as a system where the measures are clearly linked to the organization's goals where they are distributed between strategic, tactical and operational levels where both financial and non-financial dimensions are taken into account. The authors emphasize that companies often use too many measures which makes management more diffuse rather than more precise. Instead, a smaller number of critical measures should be selected, i.e. those that truly capture the essence of the business's performance. They also believe that performance measures must be understandable for different decision-making levels and related to supply chain activities such as planning, sourcing, production and delivery. For transport planning, this becomes particularly relevant because decisions about for example lead time, delivery reliability, information flow and distribution cost cannot be understood in isolation but must be seen as parts of a larger management system. Such reasoning provides an important theoretical basis for later showing that planning that is mainly based on individual experience and situational assessment risks becoming difficult to standardize, compare and improve over time. More particularly, Gunasekaran et al. highlight that performance measures in the planning phase should support control of things like order lead time, total cycle time, information processing cost, customer inquiry time and forecast accuracy. In the delivery distribution phase measures such as delivery reliability, flexibility in the distribution system and total distribution cost become especially important. The key here is not only which measures are used but that they are linked to the right decision level. Strategic measures should support long-term priorities and competitiveness, tactical measures should help resource allocation and planning choices, while operational measures should make it possible to follow the day-to-day execution. This is relevant for transport planning because many decisions about internal transport in practice lie on the borderline between tactical and operational levels where planners need to both react quickly and at the same time take into account network effects, cost consequences and service impact (Gunasekaran et al., 2004).

Piotrowicz and Cuthbertson complement this more normative picture with empirical results on how companies actually work with performance measurement in supply chains. Their study shows that the most common approaches are the Balanced Scorecard, process-based measurement and SCOR while the most common ways of grouping measures are according to plan/source/make/deliver and according to strategic, tactical and operational levels. The study also shows that users value certain characteristics particularly highly in a KPI system: the measures should be linked to the fulfilment of customer needs, be easy to understand at different management levels, be clearly related to the company's strategy and preferably provide a forward-looking perspective. In addition, it is highlighted that common goals, visibility between the company's own performance and the performance of the supply chain, and trust between actors, are important for KPIs to actually have a governing function and not just become reporting measures. This is important for transport planning, as a KPI system that lacks a clear link to decisions, responsibilities and goals is easily reduced to ex post monitoring, rather than a support for planning and control. When it comes to which KPIs in practice appear to be most central, Piotrowicz and Cuthbertson show that economic measures dominate, especially cost-related and customer service-oriented indicators. Among the most highly valued quality measures are on-time delivery, customer satisfaction, order fill rate and product or service availability. Within the efficiency dimension, distribution costs, total costs, transport costs and loading capacity utilization rank highest while lead time, product lateness, stock-outs and forecast accuracy are prominent within responsiveness. This is particularly useful in this study, as it provides empirical support for cost as a legitimate and central key measure in transport planning, not only from an internal controlling perspective but also in the broader supply chain literature. At the same time the results show that although fill rate is a recognized measure in the supply chain context, it is primarily associated with service and availability dimensions, rather than it alone capturing the most important consequences of internal transportation decisions (Piotrowicz & Cuthbertson, 2015).

It is important not to interpret this dominance of economic measures as suggesting that environmental KPIs are unimportant. Piotrowicz and Cuthbertson show that environmental measures especially CO_2 emissions from transport processes are already part of companies' monitoring systems, although they are generally used less frequently plus are valued less than economic measures. The sources also point out that social and environmental measures tend to gain more weight when they are supported by legal requirements or clear external requirements. For this study, this makes CO_2 a reasonable KPI together with cost as both measures capture the consequences of allocation movements at the network level, albeit from two different logics: economic efficiency and environmental burden, respectively (Piotrowicz & Cuthbertson, 2015).

Adjacent to this background, it becomes clear why fill rate, despite its established position in the literature can be given a more limited role in the final delimitation of the study. In Piotrowicz and Cuthbertson's study, order fill rate appears as an important quality measure but not as the most central efficiency measure for transport management. In the current Scania case, this is even clearer, since complete history of stock levels and actual consumption at suppliers is not available in the available extract, while cost and CO_2 can be linked to BI data and transport statements at load level. This means that fill rate in this project has lower analytical precision as a main measure and can rather function as a complementary or supporting perspective than a primary optimization variable. Thus, the choice of cost and CO_2 as main KPIs becomes methodologically more robust, while fill rate can still be recognized as theoretically relevant in a broader supply chain context. The literature shows that KPI use in transport planning and control should be based on a limited but well-considered selection of measures that are strategically attached, operationally useful and understandable to those who actually plan the flows. Gunasekaran et al. provide a normative model for how such a system should be designed, while Piotrowicz and Cuthbertson show how companies in practice tend to prioritize cost, transport efficiency and customer service, and increasingly also emission measures. For this study, this means that cost and CO_2 appear as the most reasonable main measures for analysing and controlling internal allocation flows, while fill rate should be understood as a secondary measure whose importance is context-dependent and strongly dependent on data availability and the type of planning situation (Gunasekaran et al., 2004; Piotrowicz & Cuthbertson, 2015).

2.5.4. Relevance and limits of operational KPIs in allocation analysis

Operational KPIs have a clear function in transport planning because they make it possible to follow up on how well actual outcomes match planned goals and by this means create a basis for management and corrective action. In the classic performance measurement literature this logic is based on the fact that the output of processes can be measured, compared with defined standards and used to identify deviations between planned and actual outcomes. Gunasekaran et al. also emphasize that operational measures particularly require accurate data, since they are used to assess the results of decisions at a low organizational level and thus assume that it is possible to link observed outcomes to concrete actions in planning (Gunasekaran et al., 2004).

The literature shows that performance measures in supply chains have fundamental limitations that become particularly evident in complex networks. Shepherd and Gunter emphasize that in practice many supply chain procedures measure internal logistical performance rather than overall function of the entire supply chain and these risks losing the systems perspective and encouraging local optimization instead. They also show that the majority of measures used are quantitative which means that they often capture observable outcomes but not necessarily the relational, organizational, or situational factors that underlie a particular decision (Shepherd & Gunter, 2006). For allocation analysis, this is important because an internal transport between two pools is not only a transport outcome but also a result of a specific decision situation. For an operational KPI to function as a direct diagnostic tool, it is therefore not only necessary that the transport can be observed afterwards, but also that it is possible to reconstruct the decision basis that preceded the order. In practice this means that one needs to know the available stock levels at the time of the decision, the current need, possible alternative solutions, and the independence of action that the planner actually had. If this context is missing, the KPI becomes primarily a measure of consistency, not a sure measure of the appropriateness or necessity of the decision. This particular problem is central to this study. The available material lacks complete follow-up of inventory balances over time, and several underlying reasons for individual allocation decisions cannot be determined solely through the quantitative data.

This means that traditional KPI-based evaluation cannot be used to determine whether each individual allocation was right or wrong since the necessary decision environment is not fully observable afterwards. Conversely this does not mean that KPIs are irrelevant, but rather that their analytical role changes. In such a data situation, operational KPIs are better suited to describing consequences and recurring outcomes at a more aggregated level than to diagnosing each order. For example, cost and CO_2 can still be used to make visible which flows are resource-intensive or environmentally burdensome, but these measures cannot, on their own, explain why a certain allocation decision was made at a certain time. This limitation is structural, not just methodological, since it follows from the fact that performance measures require that lower-level data can be transformed into meaningful information in their decision-making context, something that Gunasekaran et al. explicitly highlight as crucial in planning and control systems (Gunasekaran et al., 2004).

2.6. Synthesis of the frame of reference

This framework is not intended as a collection of delimited literature reviews but as a set of concepts that together form the analytical lens used in the later chapters. To clarify how the different parts are related, they are tied together here around the practical problem that underlies the study, namely the question of why internal allocation orders arise between Scania's packaging pools and how these movements can be understood and influenced at an operational level. The starting point is the network perspective taken in section 2.1. by viewing Scania's packaging system as a network of interconnected nodes rather than a series of individual transports, it becomes possible to explain and understand why certain movements occur. Carter, Rogers and Choi (2015) and Bellamy and Basole (2013) contribute here with a structural vocabulary that makes it possible to talk about pools as nodes, about transports as links and about asymmetries in the structure of the network. The literature on returnable transport items complements this with a more specific understanding of pools as active nodes with a buffering and redistributing function (Glock, 2017; Hellström & Johansson, 2010; Kroon & Vrijens, 1995). This basic framework is necessary because the analysis cannot be interpreted as individual allocation orders without an understanding of the system in which they occur.

Section 2.2 builds on this foundation and explains why imbalances between pools a structural feature are rather than a sign of failure in the system. Geographically distributed demand, asymmetric circulation of returnable packaging and variation over time together create conditions where internal reallocation becomes necessary (Glock, 2017; Pålsson & Hellström, 2016; Simchi-Levi et al., 2008). This is an important point for this study because it means that the question is not whether allocations should occur or not, but which allocations reflect the structural conditions of the network, and which ones possibly arise for other reasons, and which are the best choice of pool. This distinction returns directly in the analysis where patterns and deviations in the data material are interpreted against this background.

Section 2.3 shifts the focus from the system level to the individual decisions that actually create the allocation orders. The concept of bounded rationality (Simon, 1955) together with literature on reactive and anticipatory decision making and on experience-based logic (Klein, 2008; van der Vaart & van Donk, 2008) provides a language for understanding how planners act when they have incomplete information, time pressure and a continuous stream of new needs to manage. It is at this level that the practical contribution of the study takes shape because the patterns identified in the data are interpreted by how planners actually work and not by an assumption of optimal action. It is also this level that the developed decision support addresses. The developed concept is positioned within the decision support literature, which describes such tools as analytical aids that complement rather than replace human judgement (Power, 2008; Arnott & Pervan, 2014).

Section 2.4 takes the perspective a step further and places operational decision making in a broader organizational context. Power asymmetries, governance structures and contractual constraints in the relationship between Scania and its transporters and suppliers largely determine the types of changes that can actually be implemented within the framework of this project (Benton & Maloni, 2005; Lacoste & Johnsen, 2015; Nyaga et al., 2013). This section has a somewhat different function than the others. It is used less to interpret individual data observations and more to explain why certain improvement proposals that emerged early in the project, such as changes in transport arrangements or renegotiation of freight contracts, later fell outside the final delimitation of the study. It thus constitutes a connection between the current focus of the study and the broader problem picture that was originally addressed and which is described in more detail in chapter 3.

Section 2.5 supplements the frame of reference with an operational dimension by explaining how transport costs are built up, how freight prices are structured, and how performance is measured in transport planning. The cost driver literature shows that the cost of a freight transport is determined by both the characteristics of the cargo and the structural conditions of the relationship (Izadi et al., 2020; Zofío et al., 2014). The pricing literature complements this by describing how these costs are translated into prices through non-linear pricing structures and contractual rate sheets, with the result that historical averages and contracted pricing can give substantively different views of the same lane (Caplice & Sheffi, 2003; Caplice, 2007). The KPI literature provides the methodological support for using cost and CO_2 as primary outcome measures while explaining why other measures, such as load factor, play a more limited role in this study (Gunasekaran et al., 2004; Piotrowicz & Cuthbertson, 2015; Shepherd & Günter, 2006). Together, these strands make it possible to translate observed patterns into measurable consequences and to evaluate routing options both empirically and contractually, which is a prerequisite for the decision support concept developed in the study.

Together the five areas form a coherent lens. The network perspective explains where and why movements occur, the imbalance section justifies why internal redistribution is necessary, the planning logic explains how individual decisions are shaped, the power and governance perspective frames what can actually be changed, the cost, pricing and KPI literature provides both the metrics and the pricing perspectives that make it possible to evaluate the results. The idea is that the reader should be able to enter chapters 5 and 6 with the concepts already in place, and that each observation in the analysis should be able to be linked back to at least one of the areas above. In this way, the frame of reference functions not only as a background but as an active tool for interpretation throughout the study.

3. Methodology

This chapter describes how the study was conducted afterward the methodological choices behind the analysis for this project. It explains the research design used how data was collected from both transport records and contracted rate sheets, how this data was prepared and processed. The chapter correspondingly explains the assumptions besides limitations of the project and how the study's reliability and validity were delivered. The goal is to provide a clear picture of how the work was performed so that the analysis results can be understood in their methodological context. The chapter describes how AI based tools were used as supporting aids through the project in line with the principle of methodological transparency.

3.1. Research design

In this study, the Scania packaging network serves as the case for analysis. The focus is on how internal transport decisions for empty packaging emerge between different pools and how these choices impact material flows in the network. A case study fits well for this investigation because it allows an in-depth look at organizational processes and operational decisions in a real industrial setting (Rashid et al., 2019).

This research is structured as a case study of an industrial logistics network. Case studies are useful when the goal is to examine complex events in their real-world context especially when the link between an event and its environment is not clearly defined. This method allows for a detailed search of a system to combines various data sources to deepen understanding of the processes operate in practice (Crowe et al., 2011). A case study design was selected over a survey, experiment or purely mathematical modelling approach because the purpose of the study was to understand and support a real operational planning problem in its organisational context. A survey would not have provided the transport-level cost and routing detail needed for the tool, and an experiment would not have reflected the actual planning environment. A formal optimisation model was also not selected as the main method because the available data did not include all operational constraints required for full optimisation, such as live inventory, carrier availability, pool capacity and supplier timing. The case study approach therefore allowed the project to combine real transport data, contracted pricing and planner input in a way that matched the practical problem.

The study is data driven. It relies on two complementary empirical sources: a historical transport dataset covering shipments between pools and from pools to suppliers across a twelve-month period and the contracted rate sheet which specifies official transport prices for each pool to supplier relation. By examining the actual orders together with the contractual pricing structure, the study makes it possible to characterize both how transport flows have looked in practice and how they are organized at the level of agreed conditions. This data-driven approach is the methodological core of the study, since logistics networks are often understood as complex systems where patterns emerge from the interaction of many interconnected flows and decisions (Bellamy & Basole, 2013). The research design can therefore be described as an industrial case study combined with a data-driven analysis that builds toward the design of a decision support concept for cost-aware routing. The combination of operational data, contractual pricing and understanding of planning practice allows for an examination of how internal transport decisions are made and what implications they have for routing choices in the network.

3.2. Case study context

Scania was chosen as a case because the company offers study objects in industrial logistics. The current case provides access to an operation where internal transport decisions are not peripheral but directly linked to how material supply works in practice. This makes the case particularly relevant for a study that wants to investigate how planning decisions are formed, how it is linked to the supply chain and how they are translated into operational flows and what consequences they have in the network. Another reason is Scania works well as a case study is that its operation has high practical relevance which offer good access to real data. The study uses data and work processes from a current industrial setting not hypothetical scenarios or simplified examples. This allows for analysis in a real planning context where decisions are made with actual limitations related to availability, time needs and transportation logistics. This circumstance is valuable for understanding the specific operation and for discussing wider topics about planning in logistics networks.

Scania was furthermore chosen for the reason that the company is in a situation where there is a concrete need to better understand internal transport and how these can be improved. The case study is therefore based on a real business issue rather than a purely academic problem formulation. This strengthens the relevance of the study, since the analysis is linked to an area where the results are expected to contribute to the continued development of planning principles and decision-making basis at the company. At the same time, the case provides a sufficient degree of complexity to be analytically interesting. It is not about an isolated transport issue but about an environment where several functions, actors and information flows affect how decisions are made. Scania is therefore a good fit as a case in a study that wants to examine a logistical problem in its organizational and operational context.

3.3. Evolution of the study scope

The original scope of the project was substantially broader than the final definition used in the study. In the first project definition the focus was on packaging order planning in a larger network perspective where both transports from pool to supplier and internal allocation orders between pools stayed to be analysed. The objective was to investigate how numerous planning decisions affected key figures such as cost, CO_2, fill rate and to eventually contribute with data for improved planning rules and future automated planning. The original focus was therefore based on a holistic perspective on packaging planning and was close to the first assignment description that was developed for the project.

During the first part of the project, work also began in line with this broader aim. In the beginning an initial review of available data, some data cleaning and simpler pivot analyses were carried out to create an overview of both normal and allocation flows, relationships and possible analysis dimensions. However, even at this early stage it enhanced clear that the broad scope presented a considerable methodological and practical challenge. The amount of data was large and the variation between flows was broad, and the number of possible analysis variables quickly became intense. At the same time, it turned out that the problem boundary in the original project definition was relatively uncertain. There stood some possible areas for improvement, but it was difficult to determine which parts were actually within the framework of what could be influenced in the study and which parts slightly concerned more comprehensive strategic or organizational decisions.

After a month when the work was in the middle of the analysis phase it became ever clearer that the project risked becoming too wide in relation to the time and knowledge that was available. This was first noticed by the students themselves in dialogue with the supervisor at Chalmers. In these discussions it surfaced that the initial focus probably needed to be delimited in order for the study to be carried out in a reasonable and methodologically sustainable way. The issue was then raised with Scania in a joint meeting where the continued focus was discussed based on the experiences and difficulties that had occurred in the first phase of the project. The final change in scope can thus be understood as a joint decision but was initiated in practice by the need for a clearer boundary being first identified in the work between the students and the supervisor.

There were several reasons why the original scope was no longer considered reasonable. A first reason was the large amount of data and the high degree of variation in the extensive material. In particular the analysis of the normal flow, i.e. transports from pool to supplier required significantly more data and a significantly more comprehensive analysis approach than had first been presumed. The number of relationships was large and the variation between supplier flows was high in addition the connection between transport outcomes and concrete planning decisions was difficult to trace in a clear way. This meant that the work risked being spread over too many levels of analysis at the same time without leading to a sufficiently sharp result. A second reason was that the possible solutions that began to be identified in the broader analysis in several cases lay outside the area that Scania wanted the project to affect. At an early stage for example thoughts were raised that cost reduction could be achieved by steering larger volumes differently between pools and suppliers or by simply changing how certain normal flows was planned. Other possible improvement ideas concerned changes in transport arrangements such as milk run principles to deliver to multiple suppliers at same truck or issues related to carrier agreements and contract structure. However, these types of procedures were not considered to be within the scope of Scania's objective for this study. Correspondingly it became clear that parts of the original work would have required knowledge of more advanced modelling and probably also a more developed mathematical optimization approach which was beyond the practical possibilities of the project.

The change in scope was therefore primarily a necessary reprioritization resulting from the project's limitations rather than a natural academic fine tuning of the research questions. It was not that the broader theme lacked relevance but that it was too relevant and extensive to be treated in appropriate depth within the framework of a thesis project. The broad problem, analysing both normal flow and allocation orders as part of a larger whole turned out to require more time, greater data access and a different degree of methodological specialization than was realistic in the project. The supervisor's assessment was also that a clearer delimitation would make the study more reasonable based on the timeframe.

It is also important to note that the shift in scope did not arise because allocation orders described for the largest share of costs in the network. On the contrary it appeared early on that the greater part of the total costs lay in the normal flow between pool and supplier. The final focus on transport orders should therefore not be understood as meaning that these transports were most important in a strict economic sense but rather as representing a more manageable and more concrete sub-problem within the larger whole. Within the planning group there was also a clear perception that there were preventable allocation orders and that precisely these internal movements could be an area where meaningful deviations could actually be identified. Allocation orders were also already included as part of the first assignment description which made it possible to narrow the project without completely abandoning the original problem statement.

When the study was delimited, it became possible to direct the analysis towards a clearer and more operational focus specifically internal movements between pools and how these could be understood. This change afforded better conditions for analysing deviations more concretely for following certain standard packaging through the network and for linking internal movements to observable planning decision. It also made it possible to combine the quantitative analysis of allocation orders with a more qualitative understanding of how planners' reason in practice, especially in situations where the outcome in the data seemed to deviate from what was anticipated. In this way a clearer analytical line was created between data, planning logic and improvement suggestions.

Several parts of the original project idea were thus removed and toned down in the final study. This applies in particular to things related to fill rate as a main indicator and logic for automation in addition a broader analysis approach for the entire normal flow. For example, fill rate was more relevant in relation to pool to supplier transports where LTL and other variations in fill rate were central, while allocation orders were largely carried out as FTL transports. The idea of developing more direct system logic or automation in NILE was also left out of the final study as the thesis was more focused on developing an analytical basis for future planning support than on building a complete technical solution. At the same time some elements from the wider stage survived and stayed important even after the redefinition. This particularly concerned the understanding of the planner's way of working and the overall logic behind order placement. The early phase also provided a necessary background to how allocation orders must be understood in relation to the normal flow especially when the analysis was carried out at product level. Some pool-to-supplier transports still needed to be considered in the continued study to create context. In this way the early work was not exhausted but functioned as an exploratory phase that helped to identify which sub-problem it was possible to study in more depth.

In methodological terms this development can therefore be described as a regular reformulation of the study's focus based on empirical and practical experiences from the first phase of the project. The original direction provided an important initial understanding of network then its complexity but also showed that not all parts could be considered with the same degree of depth within the framework of a degree project. The final delimitation towards optimizing transports orders therefore did not mean that the study changed subject, but a more precise object of analysis was chosen within the same overall problem area. This enabled a study with clearer boundaries, better analytical sharpness and a bigger chance of leading to concrete and useful conclusions.

A further methodological development took place later in the project, when the contracted rate sheet became available from Scania. The addition of this data source did not represent a change in scope but a refinement of the analytical foundation since it made it possible to evaluate transport costs from two complementary perspectives, the historical and the contractual to use both perspectives in the design of the decision support concept developed in the tool.

3.4. Data collection

The empirical material in the study consists of four data sources. The first is the historical transport dataset, exported from the case company's transport management system. This dataset contains load level records of shipments between pools and from pools to suppliers across the period from March 2025 to March 2026, with associated variables such as origin and destination, carrier, trailer type, weight, load meters, cost in euros and date of transport. The raw transport dataset contains 35,375 shipment records of which 34,999 contain valid non-zero transport cost values after data cleaning; rows with zero or missing cost were retained for context but excluded from cost-rate calculations where relevant.

The second source is the contracted rate sheet provided by Scania, delivered in spreadsheet form. The rate sheet specifies the official transport price for each pool to supplier relation across fourteen load meter brackets, ranging from 0.5 load meters up to a full truck load capacity of 13.6 load meters. The rate sheet covers a total of approximately 1,090 pools to destination combinations, distributed unevenly across the four pools, with Eschweiler, Järna and Oskarshamn covering broad European destination ranges and Angers covering a smaller, predominantly western European destination set.

The third source consists of distance and emission-estimation inputs used to calculate indicative CO_2 values in the prototype. The CO_2 values are therefore used as indicative comparisons between routing alternatives rather than as certified emission-reporting figures.

The fourth source is qualitative information gathered through recurring meetings and conversations with planners and supervisors at the company. These meetings were not formal interviews, but operational discussions held throughout the project, in which planners explained the working logic behind routing decisions, clarified the meaning of specific data fields and helped validate the interpretation of cost and route information. The qualitative input was not the main empirical basis of the study but provided the operational context necessary to read the quantitative material correctly.

3.5. Data preparation and processing

Before the analysis could be conducted, the raw data sources were prepared and processed in several steps. The first step was a manual review of the data structure, including the meaning of each variable the units used and the relationship between the different fields in each dataset. This step also identified records that were incomplete, duplicated or otherwise unsuitable for analysis which were flagged in a dedicated column called Valid for Analysis. Records that were flagged as invalid were retained in the dataset for traceability but excluded from the analytical operations.

The second step was the standardisation of identifiers across data sources. Pool names were harmonised to a consistent set of four labels, Eschweiler, Angers, Järna and Oskarshamn, since the raw data contained variations such as company names and historical naming conventions. Postal codes were standardised into a single format consisting of a country prefix followed by a two-digit area code, which made it possible to link transport records to the corresponding rate sheet entries.

The third step was the creation of derived variables required for the analysis. A LaneKey variable was constructed by concatenating the sender and receiver postal codes which made it possible to group transports by pool to destination relation. A pool-to-pool flag was created to identify shipments where both sender and receiver were among the four pools, separating internal allocation movements from pool to supplier deliveries. A WeekNumber variable was added to support analysis of temporal patterns and a cost per load meter variable was calculated by dividing the recorded cost by the recorded load meters for each shipment. For shipments where the load meter value was missing or zero the cost per load meter variable was left undefined to prevent calculation errors.

The fourth step was the integration of the rate sheet into the analytical workflow. Each pool to destination combination in the rate sheet was indexed by a unique key which made it possible to retrieve the contracted price for any combination of pool destination and shipment size through a structured lookup. The bracket structure was implemented in such a way that any shipment size between two defined brackets was matched to the higher of the two, reflecting how contracted pricing applies in practice where partial brackets are not separately priced. The result of the preparation process was an analytical dataset where each transport record could be cross referenced with both its observed historical cost and the corresponding contracted rate sheet price, which formed the basis for the empirical analysis presented in the following chapter.

3.6. Analytical approach

The analytical approach combines four complementary perspectives applied to the prepared dataset. The first perspective is a network level analysis in which the four pools treated as nodes and the transport relations between them as connections. At this level, combined measures are calculated for each pool including the total number of loads sent and received, the total volume in load meters, the net flow position, the share of activity associated with each pool. This perspective supports the analysis of pool level imbalances and the characterisation of each pool's functional role in the network.

The second perspective is a lane level analysis in which each pool to pool or pool to supplier relation is treated as a unit of analysis. At the lane level the dataset is grouped by sender and receiver, aggregate measures are computed for each lane including the number of loads, total load meters, total cost, average cost per load and average cost per load meter. The lane level perspective supports the analysis of mutual flow, directional cost asymmetries and the comparative economics of alternative routing options.

The third perspective is a comparative cost analysis, in which the historical weighted average price per load meter is compared with the corresponding contracted rate sheet price for each relevant pool to destination combination. The comparison is performed at the level of individual brackets since the rate sheet price varies non-linearly with shipment size. This perspective supports the identification of structural differences between observed pricing and contracted pricing, and it provides the analytical basis for the dual cost methodology implemented in the decision support concept.

The fourth perspective is a deviation and pattern analysis, recurring movements and uncommon observations are identified through systematic filtering and grouping of the dataset. Examples include lanes with disproportionate volumes, relations where the mutual balance is extreme in one direction, periods with sudden changes in activity in addition pool to supplier combinations where the historical and contracted pricing diverge substantially. This perspective links the descriptive and comparative analyses to the underlying planning logic discussed in Chapter 2 to provides input for the design of the decision support concept.

The four perspectives were not treated as separate analyses but as interconnected views of the same dataset. Findings from one perspective were used to inform the others for example when a network level observation about asymmetric pool participation directed the lane level analysis toward specific relations of interest. The combined analytical structure was implemented through pivot tables, conditional formulas, lookup functions and aggregation logic in spreadsheet form, which made it possible to test specific hypotheses iteratively as new patterns emerged in the material. This analytical approach was chosen because the study required both descriptive understanding and practical route evaluation. Network level and lane level analysis were needed to understand the transport structure while comparative cost analysis was needed to connect historical spending with contracted price logic. The approach was chosen instead of a full mathematical optimisation model because the objective was not to calculate a globally optimal network plan but to develop a decision support logic that planners could understand and use as a basis for route selection.

3.7. Scenario and rule evaluation

The decision support concept developed in this study is implemented as a working prototype in spreadsheet form. The choice of spreadsheet as the development environment was deliberate. It allowed the design logic to be tested and refined rapidly against the actual data, it produced an artefact that could be inspected and modified by the case company without requiring specialised software and it kept the implementation close to the analytical workflow described in the previous section. The prototype is not intended as a production ready system but as a demonstrator of the underlying logic that the case company can develop further in a more scalable system environment.

The structure of the prototype combines three functional layers. The input layer accepts a destination postal-code zone and a shipment size in load metres from the planner. The calculation layer evaluates eleven predefined routing alternatives, consisting of four direct pool-to-supplier routes and seven allocation-based routes where stock is first moved between two pools before being dispatched to the supplier. For each available alternative, the prototype estimates total route cost using historical weighted average €/LM values for the relevant allocation and delivery legs, supplemented by rate-card-based estimates where historical observations are unavailable. The output layer presents the routing options ranked by estimated total cost and shows indicative CO_2 values together with rate-card audit information for direct-route benchmarks and trailer-type comparison. The prototype does not calculate a complete official contracted rate-card cost for every two-leg allocation route, which is why the rate-card information is used primarily for benchmarking and auditability.

The development of the prototype followed an iterative process in which each functional component was first drafted then experimented against expected outcomes for a set of representative input combinations and finally refined based on the observed behaviour. The bracket lookup logic in particular required careful assessment, since the rate sheet defines prices at discrete load meter levels rather than as a continuous function and the prototype needs to apply the correct bracket for any input size. The testing process included edge cases at the boundaries between brackets, at the full truck load ceiling and for pool to destination combinations not covered by the rate sheet, in order to ensure that the logic responded consistently across the expected range of planning situations.

The conceptual evaluation of the prototype consisted of two complementary checks. The first was an internal consistency check in which outputs of the prototype were compared with manual calculations performed independently on the same input combinations to ensure that the prototype produced the expected results. The second was a practical evaluation with planners who reviewed the structure of the tool, the interpretation of the outputs and the suitability of the ranking logic. The feedback from these reviews was used to refine the display logic and the choice of which variables to present to the user. The prototype as it now stands is therefore the result of an iterative process where data, contractual structure, calculation logic and user feedback have been combined to produce a functional concept that can serve as a basis for further development at the case company.

3.8. Use of AI tools in the research process

In the interest of methodological transparency this section describes the use of artificial intelligence tools during the research and writing process. Large language models, accessed through standard chat-based interfaces were used as supporting aids at several stages of the project. The use of these tools was supplementary in character, and all analytical decisions, data interpretations and final content choices remain the responsibility of the authors. AI tools were used in three main ways.

First, they assisted in the development of the spreadsheet-based prototype, particularly in drafting and reviewing the formulas required for the lookup logic, the rate sheet bracket selection and the cost calculation layer. Each formula generated with AI assistance was tested manually against expected outputs and validated against the underlying data before being adopted in the working tool which means that no formula was incorporated into the prototype without prior human verification.

Second, AI tools were used to support the identification of relevant academic literature. When concepts emerged in the analysis that required theoretical grounding, AI based search and summarization tools surfaced potentially relevant references. All identified sources were subsequently retrieved in their original form, read by the authors and independently checked against the citations used in the report. No claim or reference in the thesis is based on AI generated summaries alone and every cited source has been verified against its original publication.

Third AI tools were used to support the structuring and drafting of the report. This included reviewing the organization of chapters, suggesting alternative phrasings for difficult passages, identifying inconsistencies between sections and providing feedback on clarity. Drafts produced with AI support were subsequently revised by the authors who retained full editorial control over the final wording, the argumentative structure and the conclusions drawn from the empirical material. The disclosure of AI tool use follows from the general principle that academic transparency requires not only an honest account of the data and methods used in a study but also an honest account of the tools used to support the work itself. The authors take full responsibility for the content presented in this thesis.

3.9. Validity and reliability

The quality of the study has been improved by the analysis being based on multiple data sources and by the same selection criteria being used regularly throughout the work. The empirical material has been drawn from the transport management system at the case company and complemented by the contracted rate sheet provided by Scania. The combination of historical transport records and contracted pricing made it possible to cross reference the same set of transport relations from two independent perspectives, which is a form of triangulation that strengthens the credibility of the cost related findings. To ensure that the data was reasonable specific order lines have been checked manually and a clear data cleaning has been carried out before the analysis began. Only the four current pools the selected time period have been included which has provided a stable basis for comparisons and pattern identification.

The validity has also been supported by the order type being clearly classified in the data. It has thus been possible to delimit allocation orders from normal flow in a reliable way. At the same time the quantitative data has not been interpreted in isolation. In order to understand what different movements actually mean in practice the analysis has been continuously discussed with people at Scania especially in weekly meetings with mentors and at recurring reconciliations with supervisory functions. The combination of order data, planner responses and supplementary files made it feasible to understand the results in a more operational context.

Reliability has been managed by conducting the analysis systematically and by checking the results through repeated filtering and comparisons in the data material. The analysis steps have also been documented in files and working papers that made the process observable and possible to follow afterwards. Conversely there are certain limitations that affect the reliability of the study. Complete follow-up of stock balances has not been available and in some cases clear information about the reason behind an order is missing. There are also records where certain fields are unclear or where order numbers are missing. This means that not all individual decisions can be fully explained using the data alone. The results should therefore primarily be understood as a credible analysis of patterns within the studied case rather than as a complete explanation of each individual transport.

3.10. Methodological limitations

As in most case studies some methodological limitations impact the interpretation of the results. The study is based on a single industrial case and closely aligned with the prevailing planning environment in Scania's packaging network. This means that the results provide an in-depth understanding of this particular context but cannot be easily transferred to other companies, networks, or types of logistics systems. The ambition has therefore not been to attain statistical generalization but to construct a well-founded analysis of a concrete planning problem in its actual context. Another limitation lies in the sequential delimitation. The analysis covers a twelve-month period from March 2025 to March 2026 which provides a stable annual cycle of activity but does not capture multi-year trends or longer-term changes in planning logic. This provides a clear and manageable data frame but at the same time means that seasonal variations or longer development and longer-term changes in planning logic are not fully captured. The patterns identified in the study should therefore be understood as representative of the period examined rather than as a complete expression of all deviations over time.

The study is also limited by its focus on standard packaging and on the packaging numbers that emerged most clearly in the material. This delimitation has been necessary to enable a deeper analysis but means that other parts of the packaging flow have been left out like special packaging, disposable packaging and locally used packaging types can therefore affect the overall picture in practice without being visible in the final analysis. Correspondingly the study has primarily focused on allocation orders, although certain pool to supplier flows have been used as support for the interpretation. The broader transport structure is therefore only partially addressed.

Another limitation is that the study is mainly based on historical business data. Such data is valuable for identifying patterns but does not always provide full comprehension into why an individual decision was made at a certain time. In several cases direct information about the reason for the order, the planner's intention or the complete stock status at the time of the decision is missing. This means that the analysis in some parts must be based on explanation of observed outcomes rather than a complete reconstruction of the decision process. It should also be emphasized that the study does not develop or test an appropriate optimization model. The recommendations that can be formulated based on the analysis should therefore be seen primarily as an analytical basis for further work not as fully validated solutions. This is a conscious methodological limitation, but it also means that the study stops at pointing out possible directions for improvement rather than proving the effect of specific measures.

4 Empirical findings

This chapter presents the empirical findings of the study. It begins by describing the setting in which the routing decisions take place, including the structure of Scania's packaging network, the role of pool-to-pool movements and the data on which the route evaluation logic is built. This foundation is necessary because the route-ranking logic cannot be interpreted without an understanding of the four pools, their connections, the planning environment in which routing choices are made, and the main data sources used in the analysis: historical transport records, contracted rate sheet data and distance-based CO_2 estimation inputs. The chapter then moves from network context to empirical results, organised around pool-to-pool movements, pool roles, pool-to-supplier cost structures, the comparison between historical and contracted pricing and the information gaps that motivate the developed decision support prototype. The intention is to keep description and analysis in the same chapter so that each empirical observation can be read directly against the operational context in which it arises.

4.1. The packaging network and data foundation

Scania's packaging system is structured as a regional European network with four main pools that function as warehouses and supply points for empty packaging. These pools are geographically located to support suppliers in different parts of Europe. The four pools are referred to in this thesis as Eschweiler Pool in Germany, Angers Pool in France, Järna Pool and Oskarshamn Pool in Sweden, based on their geographical location rather than the operating company at each site. Basically, the pools have almost the same function. They store, handle and forward empty packaging to suppliers. At the same time, there are some structural differences between them. The pool in Eschweiler is the largest and has the capacity to handle significantly higher volumes than the others. The pool in Angers is the smallest and has a more limited capacity. Järna and Oskershamn are approximately equal in size and in practice function as regional supply points for Swedish supplier flows.

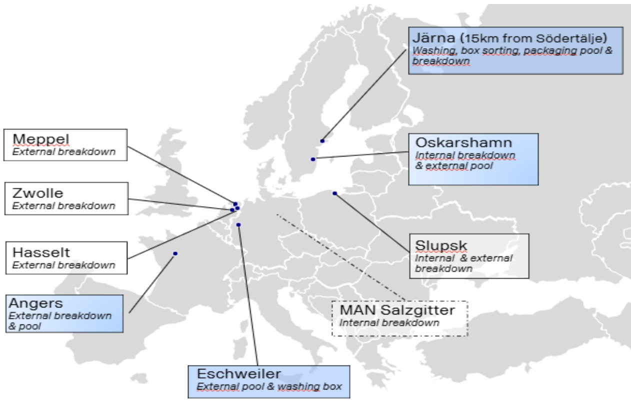

Figure 5 – Pools and breakdown map

The network is not completely symmetrically interconnected. Packaging can be moved between most pools, but Angers is only directly connected to Eschweiler and not to the Swedish pools. This creates a structural limitation in how flows can be balanced between France and Sweden. In practice, this means that any imbalances in Angers often need to be corrected via Eschweiler. Angers functions more as a structural source pool with greater outflow than inflow, which means that internal movements from Angers to Eschweiler occur relatively frequently. Eschweiler can in turn function as a redistribution point further in the network. Transports between the pools take place exclusively by truck, but of different sizes depending on the region, and in most cases, these are planned as full trucks FTL. Transporters differ between regions, which means that the network is also operationally divided into geographical zone.

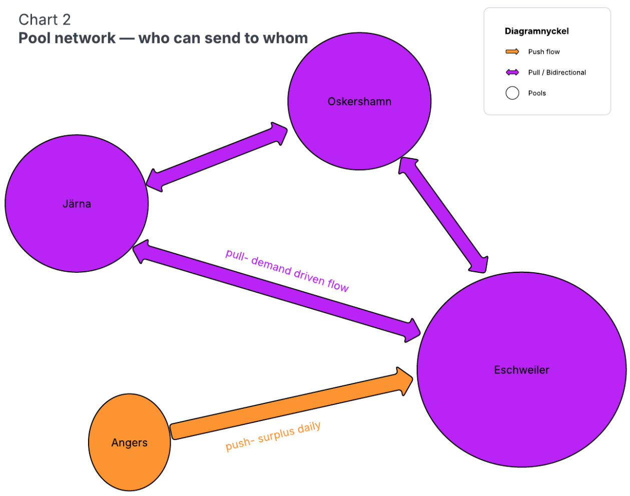

Figure 6 – Pool network

The pools are primarily intended to be regional suppliers. Suppliers in each geographical area should normally be served from a nearby pool. At the same time, deviations from this logic arise depending on the availability of packaging. This study focuses on transport routing and cost evaluation rather than on detailed inventory optimisation per packaging type. The system is mainly controlled by manual decisions in the planning tool operated by Scania. There are no strictly automated rules that force the choice of a particular pool. Instead, the planners make assessments based on the supplier's location and current stock balances. Alongside this background the network can be understood as a semi-centralised pool network where routing choices are shaped by regional supplier demand, packaging availability, pool capacity and the cost of both direct and allocation-based transport alternatives. Pool-to-pool flows are therefore not an exception but an inherent part of the structure.

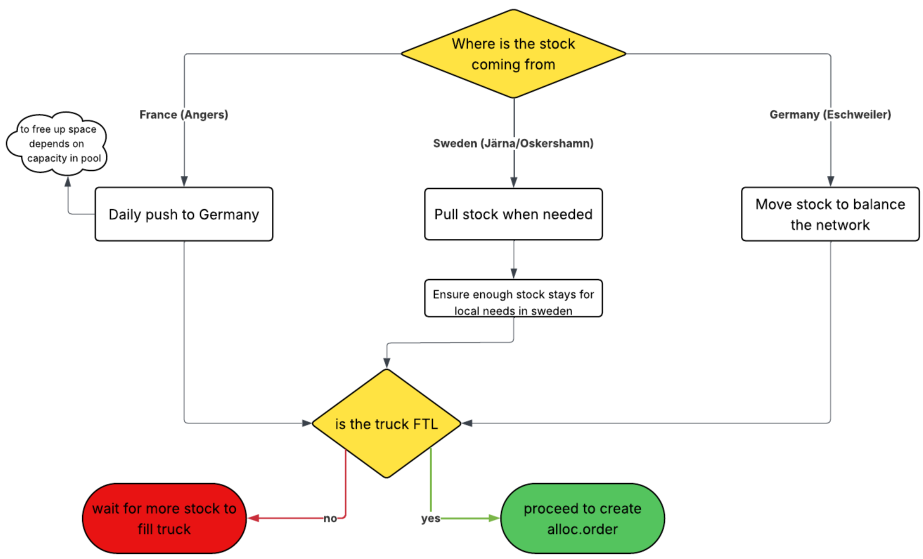

Figure 7 – Main allocation process from trigger to closure

The pools are organized to be able to store and distribute empty packaging to suppliers in different parts of Europe. Each pool holds a range of standard packaging that is normally available in all pools, but the quantity per packaging type varies. In practice, it is therefore important to keep track of the balance per packaging number in each pool, since planning is often about exactly which types are available, and not just the total volume. An allocation order can be described in this context as an internal order in NILE where empty packaging is moved between two pools. The purpose of such an order is to change the stock position in the network, so that a pool can cover an upcoming need, or so that balances become more balanced between the pools. Allocation thus differs from supplier deliveries in that the recipient is another pool and not an external party. Allocations are primarily triggered when there is a risk of shortage of a certain packaging type in a pool. This can occur when supplier needs are registered and planning shows that the pool that would normally serve the area does not have sufficient balance. It can also be triggered by a more general imbalance, where one pool has a surplus of a packaging type while another pool repeatedly ends up in deficit. In some cases, the need may arise at short notice, for example if demand changes rapidly or if planning needs to ensure delivery capacity close to time. Based on the current working method, allocations appear to be largely reactive, as they are often initiated when a concrete need is visible in the order flow and a pool lacks a balance. At the same time, there is a planning horizon where planners can act in advance, if it is seen early on that a certain packaging type will become critical in a pool. In that situation, allocation becomes more preventive but still linked to expected shortages rather than to a fixed rebalancing plan. The decision-making process is mainly manual. Suppliers enter their needs in NILE, but decisions about which pool packaging to take from, and whether an internal transfer should be created, are made by Scania's planners. NILE is then used as a planning tool and source of information for order status and availability, but the prioritization and choice of action itself is largely based on the planners' judgment and experience.

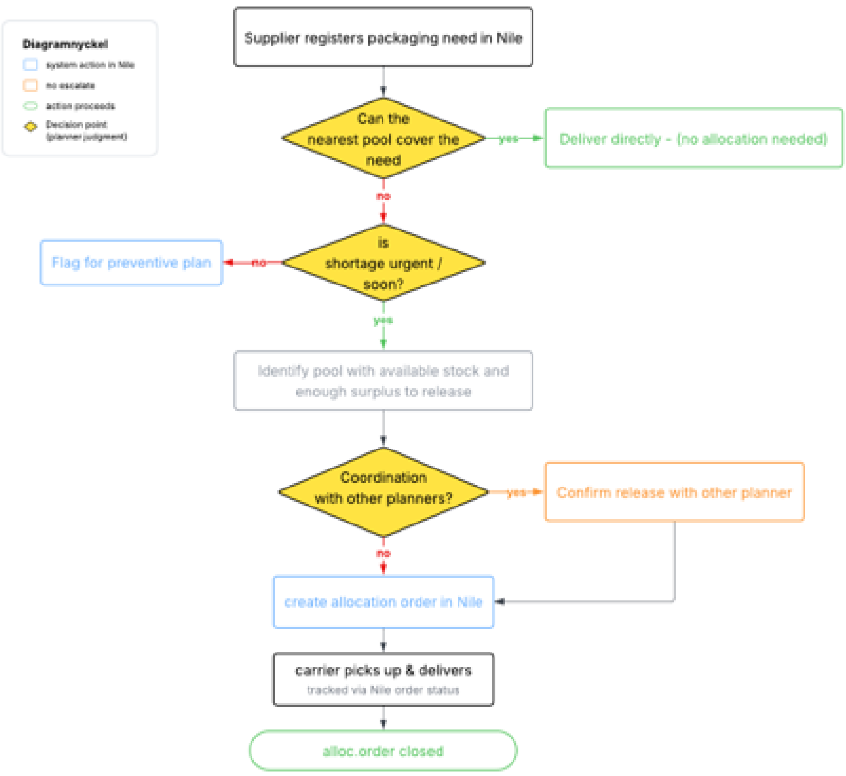

Figure 8 – Planner decision tree

An allocation order starts when the planners see a risk of shortage of a certain packaging type in a pool, or when an incoming need cannot be covered from the pool that is normally used. The planner checks the situation in NILE and decides whether an internal move is needed, then the sending and receiving pools are selected and the order is created in the system with packaging number, quantity and planned pickup date, in some cases several lines are linked to the same load number. The carrier carries out pickup and delivery, and the order is followed up in NILE through order status and planned and actual pickup date, until it can be closed when the receiving pool has confirmed receipt and the stock balance has been updated.

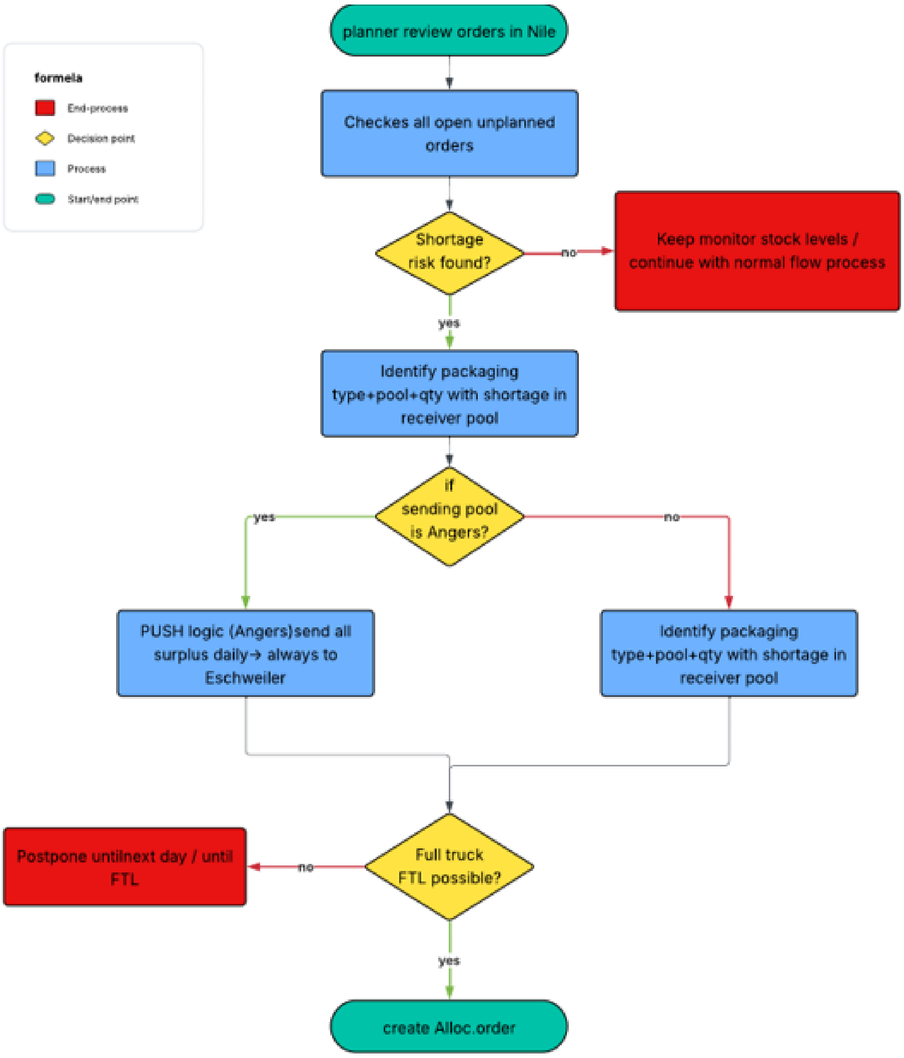

Figure 9 – Main allocation process from trigger to

In practical terms, the planning of allocation orders is performed on a continuous basis rather than according to a fixed weekly or daily planning horizon. When creating allocation orders, planners review all open unplanned orders in NILE, regardless of date or delivery-based date, and assess what needs to be repositioned between pools. For allocations from the French pool in Angers, internal transfers are created on a daily basis, and surplus volumes are systematically pushed towards the German pool in Eschweiler. These movements are always planned as full truck loads. If a full truck cannot be created, the allocation is postponed until sufficient volume is available the following day. For the Swedish pools, the logic is more demand driven. Packaging is pulled based on normal flow requirements and upcoming supplier needs.

When deciding what can be moved, planners evaluate how much must remain in the sending pool in order to secure its expected demand. In some cases, coordination with other planners is required to confirm whether certain packaging numbers can be released. In situations where surplus is identified in the French pool, all excess must be transferred to avoid overstock. Additionally, planners may receive input from colleagues regarding packaging types that require repositioning even if no immediate open order exists in the receiving pool. When it comes to thresholds and rules, according to the information available, there do not appear to be clearly defined automatic limit values in the system, such as the minimum quantity per allocation order or fixed levels for when rebalancing must occur. However, the outcome is affected by practical conditions. Pool-to-pool transports are most often planned as full trucks, which means that in many cases you need to collect volume or combine needs to create feasible flows, especially when several packaging types are to be moved simultaneously. Allocation possibilities are also affected by structural differences between the pools. The largest pool often acts as a natural balancing point, both due to higher capacity and because it is connected to more flows. This means that some internal movements in practice become more predetermined, since not all paths exist between all pools. In the network there is also a difference in how moves are initiated depending on the role of the pools in practice. From Angers there is largely a push towards Eschweiler, i.e. packaging is moved from there to the German pool. From the other pools the logic is more pull, where packaging is pulled in when needed to meet their normal flows and planned deliveries. This affects which relationships become more common in the internal allocation orders and how the decisions are justified.

4.2. Historical pool-to-pool movements as routing context

This section provides a collective overview of the historical pool-to-pool movements observed in the transport data. The purpose is to establish the overall scale and role of pool-to-pool movements before later sections examine how these movements affect routing alternatives and total route cost. All figures presented here are derived from the historical transport dataset described in section 4.1, filtered to movements where both sender and receiver are among the four pools and where the record is valid for analysis.

Across the twelve-month observation period from March 2025 to March 2026, the dataset contains 3,033 pools to pool transport loads. Together these loads carry a total of approximately 41,400 load meters, which corresponds to roughly 3,040 full truck equivalents at a 13.6 LM trailer capacity, meaning that almost every allocation is planned and executed at or very close to full truck capacity. This confirms the operational pattern described in section 4.1, where pool to pool movements is systematically set up as full truck loads and postponed when sufficient volume cannot be assembled. Pool-to-pool movement is therefore a substantial recurring component of the routing system and must be considered when evaluating route alternatives from the pool network to suppliers. The orders are not distributed evenly across the period. Monthly volumes range from around 130 loads in July 2025 to approximately 320 loads in April 2025, with most other months falling between 180 and 270 loads. The clearest pattern is a marked summer dip in July which is consistent with the typical European reduction in industrial activity during the summer period. Activity recovers in August, peaks again briefly in September and October, and stabilises around 200 loads per month from November onward through the end of the period. The general pattern is similar across the four pools but most pronounced for Angers, where the steady push toward Eschweiler creates a near constant baseline of activity that contributes between 100 and 150 loads to almost every month of the dataset.

Pool participation in the allocation activity is markedly asymmetric. As sender, Angers Pool alone accounts for around 48 per cent of the 3,033 loads, followed by Eschweiler Pool with roughly 20 per cent, Järna Pool with 17 per cent and Oskarshamn Pool with 16 per cent. As recipient, the distribution is more concentrated. Eschweiler Pool receives close to 66 per cent of all pool-to-pool loads, while Järna Pool receives 19 per cent, Oskarshamn Pool 12 per cent and Angers Pool only 3 per cent. This confirms the role of Eschweiler Pool as the central redistribution hub of the network and reflects the structural condition described in section 4.1, where Angers Pool is connected mainly to Eschweiler Pool and where the German pool serves as the natural balancing point for the wider network.

The dominant pool to pool relation is Angers to Eschweiler, with 1,442 loads, representing close to half of all observed allocations. The second most frequent relation is Eschweiler to Järna at 399 loads, followed by Järna to Eschweiler at 303 loads and Oskarshamn to Eschweiler at 265 loads. Between Järna and Oskarshamn, the data show 185 and 165 loads in each direction respectively, while Eschweiler to Oskarshamn contributes another 184 loads. The remaining three relations, all directed toward Angers from the other three pools, together account for fewer than 100 loads. This distribution shows that allocation activity is structured around a small set of recurrent corridors, with the Angers to Eschweiler axis being the single most important by a wide margin and with Eschweiler functioning as the connecting node for most of the other major relations.

The trailer mixes across all pool-to-pool loads are dominated by Mega trailers, which account for around 61 per cent of the loads. Standard trailers represent about 27 per cent of the loads and MSRT trailers about 12 per cent. The use of each trailer type is concentrated on specific relations, which is examined in more detail in section 4.3. At the aggregate level, the strong share of Mega trailers reflects the operational logic of treating pool to pool transport as a high-volume backbone where the goal is to consolidate volume into full trucks of the largest available capacity.

At an aggregate level, the pool-to-pool movements represent a recurring and operationally significant part of the transport activity observed during the studied period. Although detailed cost values are omitted for confidentiality reasons, the data show that these movements create a material cost base and therefore need to be considered when evaluating routing alternatives and planning-support requirements. With substantial variation across relations and trailer types. Indicative CO_2 values associated with these movements are estimated from distance and shipment size and are treated as a complementary comparison metric alongside cost. These aggregate figures provide the historical cost and volume context for interpreting the pool-to-supplier routing structure and the later comparison between historical and contracted pricing.

4.3. Pool-to-pool flow structure

While the previous section examined the aggregate scale and the share of activity associated with each pool, it did not look at how the flows are organised between pairs of pools. This section therefore examines the route characteristics of the ten pool-to-pool relations from four complementary angles: bilateral balance, unit cost, trailer type and average load size. Taken together these dimensions show that the flows are not a homogeneous set of similar movements but a structured collection of corridors with sharply different characteristics.

The most pronounced bilateral imbalance in the network occurs on the Angers to Eschweiler axis. Over the twelve-month period, 1,442 loads moved from Angers Pool toward Eschweiler Pool against only 20 loads in the opposite direction, giving a load ratio of approximately 72 to

1 and a net flow of around 18,900 load meters toward Eschweiler. No other relation in the network approaches this level of asymmetry. The result is that the Angers axis functions in practice as a one-way redistribution channel rather than a bidirectional exchange, and it accounts on its own for the largest single share of allocation activity in the system.

The three corridors that connect Eschweiler Pool to the Swedish pools and the Swedish pools to each other show much more moderate imbalances. On the Eschweiler to Järna relation, 399 loads moved south and 303 loads moved north, a ratio of around 1.32 leaving a net of approximately 1,400 load meters toward Järna. The Eschweiler to Oskarshamn relation shows the opposite direction at 184 against 265 loads, a ratio of 0.69 with a net of around 1,000 load meters toward Eschweiler. The Järna and Oskarshamn corridor is the most balanced of the three, with 185 and 165 loads in each direction respectively and a small net of around 250 load meters toward Järna. None of these three corridors approaches the asymmetry of the Angers axis, and together they describe a German and Swedish balancing pattern where movements in both directions are of broadly similar magnitude.

A different pattern is visible on the flows directed toward Angers Pool. Each of the other three pools sends some packaging to Angers across the period, with 41 loads from Oskarshamn, 29 loads from Järna and 20 loads from Eschweiler. In none of these three relations does Angers send anything in return. Combined with the dominant flow from Angers to Eschweiler, this creates an unusual structural pattern in which Angers receives small volumes from the rest of the network while pushing a much larger and uniform outflow toward Eschweiler. The pool therefore functions as a recipient of small ad hoc balancing movements and as a source for one large structural channel, but it does not function as a bidirectional redistribution node.

The unit cost of transport varies substantially between the pool-to-pool relations. This means that allocation legs cannot be treated as interchangeable when routing alternatives are evaluated. Some internal corridors show comparatively low historical cost per load metre, while others carry substantially higher unit rates. A particularly important finding is that costs are not always symmetric in both directions on the same corridor. Movements from Eschweiler toward the Swedish pools are consistently more expensive per load metre than the corresponding movements in the opposite direction. Since average shipment sizes are broadly comparable in both directions, this asymmetry is better explained by the underlying lane rates and carrier pricing conditions than by load utilisation alone. The Angers–Eschweiler axis also illustrates that distance is not the only driver of unit cost, since a long and frequently used corridor can still appear comparatively cost-efficient when the transport setup and lane rate are favourable. For the decision-support tool, this reinforces the need to evaluate each routing alternative using lane-specific cost information rather than assuming that the shortest or most direct route is always the cheapest option.

The trailer types used on each relation show a strong pattern of operational specialisation. The Angers to Eschweiler axis are operated almost exclusively with Mega trailers, with around 99 per cent of the 1,442 loads using Mega capacity. The two northbound relations into Eschweiler from Järna and Oskarshamn, in contrast, are operated almost entirely with Standard trailers, also at around 99 per cent. The two corridors connecting the Swedish pools to each other are operated almost exclusively with MSRT trailers in both directions. The Eschweiler to Järna lane is the only major relation with a mixed trailer composition, using around 62 per cent Mega and 38 per cent Standard, while the Eschweiler to Oskarshamn lane is more uniform at 94 per cent Mega. The small flows directed toward Angers from any pool are operated entirely with Standard trailers. This distribution shows that each lane has a dedicated transport setup, with the trailer choice tied to the underlying transport contract rather than selected flexibly per load.

An operational implication is that switching trailer type at the load level is not generally available as an adjustment mechanism in the current network.

Average load size varies in line with the trailer mix. For most relations, the average load uses between 13.1 and 13.5 load meters, which is consistent with full Mega and Standard trailer capacity at the 13.6 LM benchmark used in the rate sheet. The two MSRT lanes between Järna and Oskarshamn are the clear exception. The average load uses around 14.9 load meters in the Järna to Oskarshamn direction and around 18.2 load meters in the Oskarshamn to Järna direction, with the latter substantially above the standard benchmark. This reflects the higher load meter capacity of the MSRT trailer compared to a standard Mega configuration and shows that the Swedish internal lanes carry more volume per movement than any other relation in the network, which partly explains why these lanes also show the lowest costs per load meter.

The structural observations together suggest that the ten pool to pool relations can be grouped into three operationally distinct categories. The first is the Angers push axis, characterised by extreme directional asymmetry, the largest single volume in the network, low per load meter cost and exclusive use of Mega trailers. The second consists of the Eschweiler and Swedish corridors, comprising the relations between Eschweiler Pool and Järna Pool and between Eschweiler Pool and Oskarshamn Pool. These corridors are characterised by moderate flow volumes, near balanced bilateral activity and a pronounced directional cost asymmetry in which the southbound direction is substantially more expensive than the northbound. The third category comprises the Swedish connector lanes between Järna Pool and Oskarshamn Pool, characterized by smaller flow volumes among the major relations, exclusive use of MSRT trailers the highest load meters per movement and the lowest per load meter costs in the dataset. The remaining one-way flows toward Angers Pool sit outside these three categories and represent occasional balancing movements rather than recurring route categories. The pool level imbalances that emerge from these flow categories are examined in the next section.

4.4. Pool roles and routing context

While the previous section examined the structure of flows between pairs of pools, this section turns to the question of imbalance at the level of individual pools. For each of the four pools, the focus is on whether the pool mainly acts as a sender, receiver or balanced node, and how this role affects the routing alternatives available in the network. The aim is to identify how the operational role of each pool influences the route choices that the decision support tool later evaluates.

Across the twelve-month period, the net flow position of each pool differs sharply. Angers Pool is by far the largest net sender, with 1,442 loads sent and only 90 received, leaving a positive net of around 1,350 loads or approximately 18,000 load meters. Eschweiler Pool is the corresponding net receiver, with 2,010 loads received against 603 sent, a negative net of around 1,400 loads or approximately 18,600 load meters. The two Swedish pools sit much closer to balance. Oskarshamn Pool is a moderate net sender with 471 loads sent and 369 received, leaving a positive net of approximately 1,800 load meters. Järna Pool sits slightly on the receiving side, with 517 loads sent and 564 received and a small negative net of approximately 1,300 load meters. The four pools therefore occupy three distinct positions in the network: a strong net sender, a strong net receiver and two pools that sit close to balance with small net flows in opposite directions.

The imbalance profile of Angers Pool is the most pronounced in the network. All 1,442 loads sent from Angers during the period went to a single destination, Eschweiler Pool, while the smaller number of loads received in Angers came from all three other pools in fairly even shares. This means that Angers operates as a one-way source pool rather than as a redistribution node. It does not balance its outflow against inflow from the same partner, nor does it spread its outflow across multiple destinations. The imbalance at Angers is therefore not the result of a complex pattern with several partners but of a single recurring redistribution movement toward Eschweiler.

Eschweiler Pool shows the opposite pattern but with greater complexity. It is the largest net receiver in the network, and its inflow is heavily concentrated. Around 72 per cent of Eschweiler's incoming loads come from Angers, with the remainder split fairly evenly between Järna at 15 per cent and Oskarshamn at 13 per cent. On the outbound side, Eschweiler distributes its smaller outflow across all three other pools, with 66 per cent going to Järna, 31 per cent to Oskarshamn and 3 per cent to Angers. This combination of concentrated inflow and distributed outflow gives Eschweiler the characteristics of a redistribution hub, where packaging arriving mainly from one source is forwarded onward to other destinations within the network. The substantial net imbalance of around 18,600 load meters in favour of Eschweiler indicates that the pool absorbs considerably more than it forwards to the other pools, reflecting its dual role as both a balancing point in the internal network and as a supply base for European supplier deliveries that originate from this pool.

The two Swedish pools occupy intermediate positions, but they are not identical. Oskarshamn Pool is a consistent moderate net sender across all of its connections. It sends more to Eschweiler than it receives, more to Järna than it receives, and small flows to Angers without any return, producing a positive net of around 1,800 load meters over the period. Järna Pool sits slightly on the receiving side. Its outflows are similar in volume to its inflows, but the small net of around 1,300 load meters falls in the inbound direction. Järna also has the most concentrated inflow pattern in the network, with around 71 per cent of its incoming loads coming from Eschweiler and the remainder from Oskarshamn. Taken together, Järna and Oskarshamn balance each other to some degree, with the moderate surplus at Oskarshamn flowing partly toward Järna and partly toward Eschweiler, while Järna's small deficit is filled mainly through inflows from the same two pools.

The imbalances observed at the pool level are not concentrated in particular months but persist throughout the studied period. Angers Pool shows a positive net flow in every month from March 2025 to March 2026, ranging from around 700 load meters in July 2025, the summer dip month identified in section 4.2, to over 2,000 load meters in September 2025. Eschweiler Pool shows a negative net flow of corresponding magnitude in every single month of the period, ranging between around 1,000 and 2,300 load meters. The two Swedish pools fluctuate more from month to month. Järna varies between small positive and small negative net flows without a clear directional trend, while Oskarshamn shows a small but recurring positive net flow that becomes more pronounced from late summer onward. This stability indicates that the imbalances of Angers and Eschweiler are structural rather than transient, while the smaller imbalances of Järna and Oskarshamn fluctuate around their long-term positions without changing direction.

The way each pool's imbalance is distributed across its connections also differs. For Angers, the entire imbalance is concentrated on a single relation, since the outflow goes exclusively toward Eschweiler. The same applies in mirror image to Eschweiler, where the Angers axis alone produces more net inflow than the pool's total net imbalance, meaning that the other corridors connecting Eschweiler to the Swedish pools partially offset rather than reinforce the dominant pattern. For Oskarshamn, the imbalance is spread more evenly across the network, with positive net flows toward all three partner pools in similar proportions. For Järna, the small net inflow comes mainly from Eschweiler, while the corridor toward Oskarshamn is close to balanced. These results suggest that the routing context is strongly shaped by the Angers–Eschweiler axis, while the Swedish pools contribute a smaller and more distributed layer of secondary pool-to-pool movements.

Taken together, the pool level evidence supports a clear functional categorization of the network. Angers Pool operates as a structural source, sending packaging in one consistent direction and receiving very little in return. Eschweiler Pool operates as the central hub, absorbing the largest single inflow and partially redistributing it to the other pools while retaining a substantial net surplus that reflects its position as a supply base for European supplier deliveries. Oskarshamn Pool operates as a secondary source with a smaller and more spread-out outflow pattern. Järna Pool operates as the most balanced node in the network, with broadly similar flows in and out and a small net inflow from Eschweiler. These observations are consistent with the conceptual framing introduced in section 2.2 where imbalances in pool networks were treated as structural features arising from geographical demand asymmetry and the asymmetric circulation of returnable transport items rather than from planning failures. The next section moves from pool-to-pool movements to pool-to-supplier routing where the wider cost structure of the network becomes visible and where the contracted rate sheet complements the historical cost perspective.

4.5. Pool-to-supplier routing and cost structure

This section moves the analysis outward to the flows that the pools direct toward suppliers in the wider European supply base. The purpose is twofold. First, pool-to-supplier deliveries represent the final delivery leg that the routing prototype evaluates for each supplier destination. Understanding the pool-to-supplier cost structure is therefore necessary for evaluating both direct routes and allocation-based routes. Second, the cost structure of pool-to-supplier deliveries determines the total cost of each routing option, since allocation-based routes combine an internal pool-to-pool leg with a final delivery leg to the supplier. This section presents the volume, geographic coverage, cost and trailer characteristics of pool to supplier flows across the four pools.

Across the twelve-month period, the four pools dispatched 21,042 loads to suppliers, carrying a total of approximately 221,000 loads. Pool to supplier activity is therefore around seven times larger than the internal allocation activity in load count, around five times larger in load meters and around four and a half times larger in total cost. The average pool to supplier load uses around 10.5 load meters, which is lower than the average pool to pool load of 13.7 load meters reported in section 4.2. This reflects the fact that deliveries to individual suppliers are often partial truck loads adjusted to the specific volume that the supplier needs, while pool to pool movements are systematically consolidated into full trucks before being dispatched.

The four pools differ in how much they contribute to the pool to supplier activity. Eschweiler Pool is by far the largest contributor, dispatching 9,541 loads, which represents around 45 per cent of all pool to supplier deliveries in the period. Järna Pool and Oskarshamn Pool dispatch broadly similar volumes, with 4,948 and 5,221 loads respectively, together accounting for around half of the total. Angers Pool is the smallest contributor by a wide margin, dispatching only 1,332 loads, which corresponds to around 6 per cent of the activity. This distribution mirrors the imbalance role each pool plays in the internal network. The pool that receives the largest net inflow internally, Eschweiler, also dispatches the largest volume to suppliers, while the pool that sends the largest internal net outflow, Angers, dispatches the smallest external volume. Pool-to-pool movements and pool-to-supplier deliveries are therefore closely linked and should be evaluated as complete routing alternatives rather than as separate decisions.

The four pools also differ markedly in the geographic breadth of their supplier base. Eschweiler Pool serves 272 unique postal destinations across 25 European countries, the broadest reach in the network. Järna Pool and Oskarshamn Pool serve smaller but still substantial supplier bases, with 209 and 165 unique postal destinations and around 20 countries each. Angers Pool is the most constrained, with only 61 unique postal destinations distributed across 5 countries. This narrower coverage reflects the regional positioning of Angers as a pool serving western European suppliers rather than acting as a continental hub. The asymmetry in geographic coverage has direct consequences for the routing decisions examined later in the study, since the availability and cost of direct routes differ across pools and supplier destinations.

The supplier base of each pool reflects a clear geographic specialisation. Eschweiler Pool delivers primarily to Germany, which accounts for around 43 per cent of its loads, followed by the Netherlands, Italy and the Czech Republic at between 7 and 14 per cent each. Angers Pool delivers primarily to France at around 57 per cent and Spain at around 31 per cent, with smaller volumes to the Netherlands and Portugal. Järna Pool delivers around half of its loads to Swedish suppliers and the remainder mainly to Poland, Estonia, the Czech Republic and Germany. Oskarshamn Pool follows a similar pattern, with around 46 per cent of loads to Swedish suppliers and a substantially larger share to Poland at around 29 per cent. These specialisations are consistent with the role of each pool as a regional supply point, and they explain the structural distinction between the broad European reach of the German pool and the more focused regional reach of the Swedish and French pools.

Cost characteristics differ noticeably between the pools as well. The historical cost per load metre is not uniform across the network, since each pool serves a different mix of destinations, distances, trailer types and carrier agreements. Pools with broader and more geographically dispersed destination coverage tend to face more heterogeneous transport conditions, while pools serving more concentrated regional flows can often operate under more stable cost conditions. Shipment size also matters, since smaller or less consolidated deliveries generally create higher unit costs than shipments closer to full-truck utilisation. This means that differences in unit cost between pools should not be interpreted simply as differences in pool efficiency. Instead, they reflect the operational role of each pool in the wider routing network. For the decision-support tool, this reinforces the need to evaluate costs at the level of the specific routing alternative, destination and shipment size, rather than relying on average pool-level cost assumptions.

The trailer mix on pool to supplier loads differs sharply from the internal pool to pool flows analysed in the previous sections. Standard trailers account for around 74 per cent of all pool to supplier loads, MSRT trailers for around 22 per cent and Mega trailers for only around 4 per cent. This is essentially the inverse of the internal flows, where Mega trailers dominate. The pattern is consistent with the smaller average load size on supplier deliveries, since Standard trailers are better suited to partial truck loads and to dispersed delivery patterns where multiple suppliers may need to be reached. The very low share of Mega trailers indicates that consolidated full truck deliveries to single suppliers are uncommon in the network and that the bulk of supplier service is provided through smaller, more flexible transport arrangements.

Taken together the pool to supplier analysis shows that the four pools serve operationally distinct supplier bases that vary in geographic breadth, volume and cost. Eschweiler Pool functions as a continental hub with the broadest coverage, the highest volume and the highest unit costs. Angers Pool functions as a constrained regional pool serving a limited western European base. Järna Pool and Oskarshamn Pool function as Swedish regional pools with extended reach into Eastern Europe with Oskarshamn showing the lowest unit costs in the network. The routing and cost structure established here forms the empirical base for the next section, where historical pool-to-supplier costs are compared with contracted rate sheet pricing. This comparison is central to the prototype because it allows routing alternatives to be evaluated against both observed historical spending and contracted cost benchmarks.

4.6. Pricing structure: historical versus contracted rates

This section turns to the second pricing perspective introduced in section 4.1 the contracted rate sheet and compares it systematically with the historical pricing. The purpose of the comparison is not to evaluate which perspective is more correct, since the two measure different things, but to make visible how they relate to each other across the network and to identify where they diverge. The comparison is also important for the study because the decision support prototype combines historical weighted average costs with contracted rate sheet benchmarks, and planners need to understand the role of each perspective when interpreting the output.

The contracted rate sheet from the case company defines the official transport price for each pool to supplier combination across fourteen load meter brackets, ranging from 0.5 load meters up to a full truck load of 13.6 load meters. The complete rate sheet contains 1,088 pools to destination entries, distributed unevenly across the four pools. Eschweiler Pool covers 303 destinations in the rate sheet, while Järna Pool and Oskarshamn Pool each cover 362 destinations. Angers Pool covers only 61 destinations. This distribution mirrors the geographic coverage of each pool's historical activity described in section 4.5, with the smaller and more regional Angers Pool having the narrowest contractual footprint. For each pool to destination combination, the rate sheet therefore defines fourteen prices that together describe how the cost of transport on that lane changes with shipment size.

A second feature of the rate sheet is that the price effectively plateaus at the full truck load level. Across the rate sheet entries, the price at the 12-load meter bracket is typically equal to or very close to the price at the 13.6 load meter bracket, meaning that the contracted price stops increasing once a full truck has been booked. As an example, on the Eschweiler Pool to DE-58 lane, the contracted price is the same at the 12 LM and 13.6 LM brackets, meaning that adding additional volume within this range does not increase the contracted cost. This full truck top limit has practical implications for decision making since the marginal cost of loading additional packaging into a planned truck is effectively zero up to the truck's capacity limit. Planners who can consolidate volume into a single truck therefore obtain a more favourable per unit cost than planners who dispatch the same total volume across multiple partial loads.

When the rate sheet is matched against the historical activity recorded in the transport extract, 656 of the 1,088 rate sheet entries also appear in the historical data as pool to supplier lanes that have actually been used during the studied period. The remaining 432 rate sheet entries correspond to lanes that are contractually available but have not generated any recorded transport activity in the twelve-month period, which suggests that the rate sheet covers a wider potential supplier base than is currently active. In addition, 54 historical pools to destination combinations appear in the transport data but not in the rate sheet, which corresponds to deliveries that have taken place outside the standard contracted lanes. These cases are less common but indicate that not all transport activity flows through pricing structures defined in the rate sheet.

Across the 656 lanes that appear in both data sources, a systematic pattern emerges when the historical weighted average price per load meter is compared with the rate sheet price per load meter at the full truck bracket. The historical weighted average exceeds the rate sheet price at full truck in 83 per cent of the comparable lanes, with an average difference of approximately 78 euros per load meter, corresponding to about 87 per cent above the full truck reference. The pattern is not uniform across the four pools. For Eschweiler Pool, the historical weighted average exceeds the full truck rate in 93 per cent of cases, with an average premium of around 135 per cent. For Angers Pool, the premium reaches 97 per cent of cases and approximately 121 per cent on average. For Järna Pool, the historical exceeds the full truck rate in 78 per cent of cases with an average premium of around 44 per cent, while Oskarshamn Pool shows the same direction in 70 per cent of cases with an average premium of around 53 per cent.

The systematic gap between the historical weighted average and the rate sheet price at full truck does not indicate that the rate sheet is incorrect or that the historical data is unrepresentative. The two measures answer different questions. The historical weighted average reflects the actual blend of shipment sizes that has occurred on each lane, including the smaller loads that carry a higher per unit premium according to the bracket structure. The rate sheet price at the full truck bracket reflects only the cost that would apply if every load on the lane were a full truck. The systematic difference therefore reflects the fact that pool to supplier shipments are not generally full trucks. Section 4.5 showed that the average pool to supplier load is around 10.5 load meters, below the 13.6 load meter benchmark, with substantial variation across pools. Pools that ship a higher share of smaller loads, in particular Eschweiler Pool with an average load of around 9.8 load meters and a broad European destination mix, show the largest gaps. The implication is that the historical weighted average cannot be used as a substitute for the rate sheet price at a specific shipment size, since the historical metric blends across all sizes while the rate sheet distinguishes between them.

Taken together, the comparison shows that the two pricing perspectives provide complementary rather than competing information. The historical perspective describes what transport has cost in practice on each lane, including the effect of the actual mix of shipment sizes. The contracted perspective describes what transport should cost for any given shipment size under the agreed terms. For a planner who needs to evaluate a specific routing option for a specific shipment size, the rate sheet provides a contracted benchmark that accounts for the relevant load-metre bracket. The historical perspective adds context by showing how the lane has performed across the wider mix of shipments. The decision support concept developed in the study is built on this complementarity. It displays historical cost references and contracted rate sheet benchmarks in parallel, allowing the planner to compare observed lane costs with the contracted cost level for the selected shipment size. The design and structure of the decision support concept are presented in the next section.

4.7. Decision chain misalignments and the decision support concept

A planner who is about to dispatch packaging from one pool to a supplier rarely has all the information needed to evaluate the choice in front of them. The empirical findings documented earlier in the chapter make this concrete. Pool-to-pool movements are asymmetric, pool-to-supplier costs vary across destinations, historical weighted average costs can differ substantially from contracted rate sheet benchmarks, and the rate sheet itself prices each shipment differently depending on its load-metre bracket. None of these conditions are visible in the daily planning environment described in section 4.1. What emerges is a structural mismatch between the information that planners actually use and the information they would need in order to make cost aware decisions in real time.

This mismatch can be understood as a decision chain that runs from supplier need to executed transport through several steps. At each step, information is added or transformed, but it does not consistently return to the planner in a form that supports route-level decision-making. Three specific misalignments emerge from the empirical findings presented throughout this chapter: limited route-level cost visibility, limited access to contracted rate sheet benchmarks in the planning context, and the separation between direct delivery and allocation-based routing decisions. These misalignments explain why planners may lack a complete view of the cost consequences of a routing choice at the moment the decision is made. Each misalignment is described in turn below before the structure of the decision support concept is presented.

The first misalignment relates to the visibility of cost at the moment of the allocation decision. Unit costs vary by a factor of five between pool-to-pool corridors, as documented in section 4.3, and directional cost asymmetries can reach close to 90 per cent on the corridors connecting Eschweiler to the Swedish pools. To this, section 4.6 added that the cost of a specific shipment depends on its size, since the rate sheet prices small shipments at multiples of the per unit cost of a full truck. A planner who chooses a source pool or routing alternative based mainly on availability and operational experience has limited visibility of whether the chosen route is more expensive than another feasible alternative, since route-level cost comparison is not presented directly in the current planning environment.

A second misalignment involves the absence of contracted pricing in the planning context. Where any cost information is available to planners, it tends to be based on historical averages or aggregate budget figures rather than on the contracted rate sheet. Section 4.6 demonstrated that these two perspectives can differ substantially for individual lanes, and that the difference is systematic rather than random. A planner relying only on historical averages will overestimate the cost of full truck shipments and underestimate the cost of small ones, with no way to distinguish between the two cases. The contracted rate sheet resolves this ambiguity but is not currently integrated into the operational decision environment.

The third misalignment is structural rather than informational. Direct delivery decisions and allocation-based routing decisions are often considered separately in the planning workflow, even though they can represent alternative ways of fulfilling the same supplier need. Choosing a routing alternative determines whether the supplier need is fulfilled through a direct pool-to-supplier delivery or through a combined route that includes both a pool-to-pool leg and a final delivery leg. The total cost of meeting a supplier need therefore depends on both legs of the routing, but the current planning logic does not require these legs to be evaluated together. As a result, a route that appears operationally suitable from one perspective may still be more expensive than another available routing alternative when the full route cost is evaluated.

To respond to these three misalignments, the study develops a decision support concept that brings the analytical material presented throughout this chapter into a single planner facing artefact. The concept is implemented as a working prototype in spreadsheet form, as described in section 3.7, and is built around the routing decision that planners face at the moment a supplier need has been registered. Given an input consisting of a destination postal-code zone and a shipment size in load metres, the prototype evaluates eleven predefined routing alternatives included in the tool. Four of these are direct deliveries from each of the four pools to the supplier. The remaining seven are allocation routings, in which packaging is first moved from one pool to another and dispatched onward to the supplier from there. For each routing alternative the prototype estimates total route cost using historical weighted average EUR/LM values for the relevant transport legs, supplemented by rate sheet-based estimates where historical observations are unavailable, and ranks the available alternatives from cheapest to most expensive.

The design of the prototype reflects each of the three misalignments directly. Cost is made visible in the decision interface itself, which removes the first misalignment by presenting the cost of each routing option at the moment the planner is choosing between them. Historical cost references and contracted rate sheet benchmarks are displayed in parallel where applicable, which responds to the second misalignment by giving planners visibility of both observed spending and contracted price levels. The full routing option, including both the internal allocation leg and the supplier delivery leg, is presented as one combined evaluation, which addresses the third misalignment by linking the two decisions that are currently treated separately. Additional outputs include an indicative CO_2 estimate calculated from distance, shipment size and a fixed emission factor, together with route descriptions and rate sheet audit information that improve transparency and traceability.

It is important to be clear about what the prototype is and what it is not. It runs on static historical data and contractual pricing rather than on live system inputs. It does not connect to inventory or capacity information at the pool level, nor to the operational order management system that planners use today. Its purpose is to demonstrate the analytical logic that a future system would need to embed, and to provide the case company with a concrete reference for how such a system could be structured. The expectation is that the case company will use the prototype as a basis for the development of a more scalable system that can be integrated with planning tools and operational data sources, retain the dual pricing logic and the routing comparison structure, and extend the concept with features that lie outside the scope of a spreadsheet implementation, such as real time inventory checks, automated alerts and integration with the rate sheet maintenance workflow. In this sense, the prototype carries the empirical and analytical contribution of the study into a single usable form while leaving the operational implementation as the next step for the organization.

Taken together, the chapter has moved from the structure of the pool network and its historical transport movements to pool-to-supplier routing costs, the comparison between historical and contracted pricing, and finally the decision support prototype that brings these elements together. The empirical observations presented across the chapter are not isolated findings but mutually reinforcing elements of the same underlying analysis. The asymmetries in pool-to-pool movements, the different roles of the four pools, the variation in cost per load metre across routing alternatives, and the gap between historical and contracted pricing all point toward the same conclusion: planners need route-level cost visibility when choosing between direct and allocation-based routing alternatives.

5 Analysis

Chapter 5 interprets these findings from the perspective of routing decision support. The analysis is organised around three layers: the routing logic of the pool network, the cost logic of route selection, and the planning logic behind routing choices. These layers are then carried together in the established excel spreadsheet prototype which estimates and ranks direct and allocation-based routing alternatives for a selected destination and shipment size.

5.1. The routing logic of the pool network

Two findings from Chapter 4 are especially important for the routing prototype. First, pool-to-pool movements are not random background activity but recurring transport legs that can form part of a complete route to a supplier. Second, the four pools do not play identical roles in the network. Some pools mainly act as source pools, some as receiving or redistribution points, and some as regional delivery bases. The analytical question is therefore not only why these movements occur, but how they should be represented when planners compare routing alternatives. The pool network does not allow every theoretical combination of origin, intermediate pool and supplier destination. In the prototype this is treated by assessing a predefined set of operationally relevant routing alternatives rather than by solving an unrestricted network optimisation problem. The tool includes four direct pool-to-supplier routes and seven allocation-based routes, reflecting the pool-to-pool legs for which operationally relevant data and routing logic exist. This means that the topology of the network becomes a practical constraint in the decision support tool: the planner compares the routing alternatives that are available in the model, not every possible route that could theoretically be imagined.

The structure of this routing is not to be regarded as a simple planner mistake. Returnable packaging networks are driven by the differences in availability of packaging, by the emergence of supplier demand and by the existence of contracted transport connections between pools. The prototype is not designed to optimize the packaging inventory across the whole network. Instead, it makes visible the routing consequences of these structural conditions by showing how different direct and allocation-based alternatives affect cost for a given destination and shipment size.

Eschweiler's role in the network is still important, but in the context of the prototype it is mainly relevant because it appears as a direct delivery pool and as an intermediate pool in several allocation-based routes. This double duty affects the route evaluation. Sometimes a supplier can be served directly from Eschweiler, sometimes a route will first move packaging from another pool to Eschweiler before delivering it on to the supplier. The tool thus evaluates the total route cost without considering the pool-to-pool leg and the supplier delivery leg as separate planning decisions.

The Swedish pool connections have a different routing logic. Geographically, Järna and Oskarshamn are closer to each other than to the continental pools, and the historical allocation rates between the two are lower than several of the longer pool-to-pool alternatives. This implies that a route based on allocation is not necessarily inefficient. Sometimes even a route that has a pool-to-pool leg can be competitive if the allocation leg is cheap and the final delivery leg from the receiving pool is good.

The empirical stability found in the historical data supports the use of historical transport rates as planning references, but it should not be overemphasized as a forecasting model. The prototype uses historically weighted average costs to model observed transport behavior and contracted rate sheet data to provide benchmark prices. It does not forecast demand for the future, inventory availability, or supplier order timing. The routing logic is thus to be interpreted as decision support structure based on available historical and contractual information.

The implication for the study is that the discussion of improvement should not be about simply cutting allocation flows. Sometimes allocation routes are operationally necessary and sometimes they are cost competitive. The opportunity for improvement lies instead in making the full route cost visible before a routing decision is made. A planner should be able to compare a direct pool to supplier route with allocation-based alternatives and see what the allocation leg, the delivery leg, the shipment size, and the contracted price benchmark are due to the total cost. This brings us to the cost logic of route choice.

5.2. The cost logic of route selection

Transport cost in this transports network is not a single number. It depends on the route the shipment size, the historical cost level of each transport leg and the contracted rate sheet bracket that applies to the shipment. This is why the prototype does not evaluate routing alternatives only by geographical proximity or by whether a route is direct or allocation based. Instead, each alternative is evaluated as a cost combination of the relevant allocation leg, delivery leg and load-metre input.

The non-linear shape of the rate sheet reflects a basic feature of road freight economics. A truckload transport has substantial fixed components: the driver's time, the lane setup, the deployment of a vehicle from one location to another. These costs apply whether the truck is fully loaded or only partially loaded. As the load grows, the same fixed costs are spread over more cargo and the cost per unit of cargo decreases. Once the truck is full, no additional volume can be added without booking a second truck, and the price effectively plateaus. Caplice and Sheffi (2003) describe this pattern as the standard form of truckload pricing, and the bracket structure used in the Scania rate sheet is one of the common ways of operationalising it (Caplice, 2007). The pattern is not a peculiarity of this particular carrier or this particular contract. It is how freight pricing works.

Once this is recognized the difference between historical weighted average costs and contracted rate sheet costs becomes easier to clarify. The historical weighted average per load meter is the actual cost of a lane across the actual shipments that have run on it. The rate sheet on the other hand gives a contracted benchmark for a certain pool destination and load-meter bracket. So, these two measures are asking different questions. The historical cost is useful for route ranking and for understanding observed spending. The rate sheet is useful for benchmarking contracts.

The variation across pools should therefore be interpreted as a consequence of route coverage, shipment-size mix and supplier geography rather than as a simple ranking of pool efficiency. A pool that serves a broad and dispersed supplier base will naturally include more partial loads and more heterogeneous lane costs. A pool with a more concentrated supplier base may show a lower weighted average cost per load metre because a larger share of its loads is closer to full-truck capacity. For the prototype, this means that route cost must be evaluated for the specific destination and shipment size rather than inferred from the general role of the pool. This has a direct consequence for how the prototype should be interpreted. The route ranking is primarily based on estimated route cost derived from historical weighted average €/LM values for the relevant transport legs, supplemented by rate sheet-based estimates where historical observations are unavailable. The Rate Sheet Audit sheet then provides a contracted benchmark for direct pool-to-supplier routes and trailer-type comparisons. The two perspectives should therefore be read together: the historical cost estimate supports route comparison, while the rate sheet supports auditability and contracted-cost visibility.

A second feature of the cost structure deserves separate attention. The Allocation Matrix shows that the historical €/LM rate can differ substantially by direction on the same pool-to-pool corridor. For example, the Eschweiler-to-Swedish-pool directions are considerably more expensive per load metre than the corresponding Swedish pool to Eschweiler directions in the prototype's historical allocation rates. This asymmetry is not caused by distance, since the physical corridor is the same in both directions. It reflects directional pricing and freight-market conditions. For the planner, the practical implication is that the choice of routing alternative determines not only which distance is travelled, but also which directional rate is applied.

From a performance measurement standpoint, the analysis exposes a gap that connects to the broader literature on supply chain KPIs. Gunasekaran, Patel and McGaughey (2004) argue that operational KPIs are useful only when they are visible at the level where decisions are made. In the present case, transport cost and environmental indicators may be available in reporting contexts, but they are not presented as a side-by-side comparison of routing alternatives at the moment of planning. The decision support prototype addresses this gap by repositioning route-level cost information, contracted rate sheet benchmarks and indicative CO_2 estimates closer to the routing decision itself.

The cost logic of the network is therefore more complex than a flat €/LM comparison would suggest. It includes load-metre brackets, directional allocation rates, differences between observed historical spending and contracted price benchmarks, and variation in the availability of historical data across lanes. These conditions explain why a planner-facing tool is needed: the cost of a routing alternative cannot be inferred reliably from distance, source pool or allocation status alone. The next analytical layer therefore considers the planning logic behind routing choices.

5.3. The decision logic behind routing choices

The planning situation can be interpreted through the lens of bounded rationality. Planners make routing-related decisions under time pressure, with incomplete cost visibility and without a single interface that compares direct and allocation-based alternatives side by side. In such a setting, it is unrealistic to expect every alternative route, cost bracket and allocation leg to be evaluated manually for each shipment. The absence of integrated route-level cost information therefore encourages planners to rely on operational experience and feasible routing choices rather than systematic cost comparison. This is not a failure of planner competence; it is a rational response to the information conditions under which the decisions are made.

Gigerenzer and Gaissmaier (2011) show that operational decision makers often develop simplified rules that work well in their specific context. In a packaging pool environment, such a rule might be to choose a feasible pool with available stock or to prefer a familiar route that has worked before. This type of heuristic can be operationally effective, but it does not reveal whether another direct or allocation-based route would have produced a lower total cost for the same supplier need. The role of the prototype is therefore not to replace planner judgement, but to add cost and benchmark information to the set of cues available when the routing decision is made.

A second decision mechanism is recognition-based judgement. As described by Klein (2008), experienced operators often recognise situations as similar to previous cases and apply responses that have worked before. In Scania's context, this may help planners act quickly when similar destinations, pools or supplier needs recur. However, recognition-based judgement is limited when the underlying cost structure changes with shipment size, LM bracket, directional allocation rate or contracted price benchmark. The prototype addresses this limitation by translating these cost variables into a ranked route comparison.

The same logic applies to routing decisions. If the planning environment primarily stages operational possibility while cost outcomes are only visible later in reporting the safer decision is often to choose a route that solves the immediate need. This can be reasonable from a service perspective, but it does not guarantee that the chosen route is cost efficient. The decision support prototype therefore focuses on making the cost consequences of feasible routing alternatives visible before the transport is finished.

The planning logic therefore supports the central design choice of the prototype. If planners are expected to make cost-aware routing decisions, route-level cost information must be visible at the point where those decisions are made. Asking planners to choose the lowest-cost route without presenting comparable route costs, contracted benchmarks and shipment-size effects means asking them to optimise on information that is not available in the decision interface.

This is the gap that the decision support literature addresses. Power (2008) describes decision support systems as analytical tools that present relevant information at the moment of decision and complement human judgement rather than replace it. Arnott and Pervan (2014) argue that effective decision support must fit the cognitive and organisational situation of its users. The prototype developed in this study follows this logic. It does not automate the routing decision, and it does not replace planner judgement. Instead, it changes what is visible when the planner evaluates a route: estimated route cost, direct and allocation-based alternatives, contracted rate sheet benchmarks where applicable, and indicative CO_2 impact. Whether this changes planning behaviour in practice is outside the scope of the thesis, but the analysis suggests that cost-aware decisions require cost-aware information conditions.

5.4. Integrating the three logics through the decision support concept

The established prototype is to understand as the point where the routing, cost and planning analyses become a single working object. The routing analysis defines which direct and allocation-based alternatives should be evaluated. The cost analysis defines how historical €/LM value, allocation rates and contracted rate sheet benchmarks should be clarified. The planning analysis explains why this information needs to be presented in a planner-facing interface rather than only in retrospective reports. The prototype brings these elements together in one route-ranking tool.

The routing layer of the prototype makes the network structure operational. For each selected supplier postal code zone and shipment size the tool evaluates eleven predefined routing alternatives: four direct pool-to-supplier routes and seven allocation-based routes that combine a pool-to-pool movement with a final supplier delivery. The planner can therefore see which alternatives are available which require an intermediate pool and how the total cost changes when the route includes both an allocation leg and a delivery leg. The network structure becomes an input to the decision rather than a hidden constraint.

The cost layer of the prototype translates the pricing logic into a practical route comparison. Estimated route cost is calculated from historical weighted average €/LM values for the relevant allocation and delivery legs, with rate sheet-based estimates used where historical observations are unavailable. In parallel, the Rate sheet Audit sheet provides official contracted benchmarks for direct pool-to-supplier routes and trailer type comparisons. This distinction is important: the prototype does not calculate a complete official contracted rate sheet cost for every two-leg allocation route. Instead, it combines historical route-cost estimation with contracted-cost transparency where the rate sheet supports it.

The planning layer explains why this information must be presented. The prototype positions route descriptions estimated route costs, indicative CO_2 values and rate sheet audit information in the same workbook, with the main Lookup Tool ranking available alternatives from cheapest to most expensive. Cost is therefore not only analysed after transport execution but also translated into a route comparison view that can support the planning decision before execution.

A concrete illustration from the latest prototype is the lookup for postal-code zone DE-74 with a shipment size of 13 LM. In this case, the tool ranks Eschweiler Pool → Supplier as the lowest-cost available route, with an estimated historical route cost of €1,125.49. The same lookup shows alternative direct routes from Oskarshamn and Järna, as well as allocation-based alternatives such as Järna → Oskarshamn → Supplier and Angers → Eschweiler → Supplier. The Rate Sheet Audit sheet shows that the official contracted direct-route benchmark for Eschweiler Pool to DE-74 at the applied 13.6 LM bracket is €904.04, while the corresponding direct-route benchmarks are €1,431.09 from Oskarshamn and €1,741.25 from Järna. This illustrates the central function of the prototype: it ranks feasible routing alternatives by estimated cost while also making contracted direct-route benchmarks visible for comparison.

The prototype also includes an indicative CO_2 estimate alongside cost. This estimate is calculated from estimated road distance, shipment size in load metres and a fixed emission factor. It is therefore intended as a directional comparison between routing alternatives, not as certified emission reporting. The tool does not automatically optimise a combined cost-and-emissions objective. Instead, it surfaces the approximate CO_2 implication next to the estimated route cost, allowing planners to see whether the cheapest route also appears favourable from an environmental perspective or whether there may be a trade-off.

The prototype is implemented as a spreadsheet-based decision support tool rather than as a production system. This is a deliberate methodological choice. The aim is to demonstrate how historical transport costs, allocation rates, contracted rate sheet benchmarks and indicative CO_2 estimates can be combined into a working planner-facing route-ranking interface. The prototype runs on static data, evaluates predefined routing alternatives and does not connect to live inventory, order or capacity systems. Its contribution is therefore not system integration, but the operationalisation of the route evaluation logic in a form that planners and stakeholders can inspect, test and further develop.

The integration of routing, cost and planning logic in the prototype completes the substantive analysis of the study. The analysis no longer points only to historical transport behaviour but to a practical way of supporting future routing decisions. What remains is to state in compact form. What this analysis shows in relation to the research questions.

5.5. Answering the research questions

RQ1a asks what direct and allocation-based routing alternatives exist between Scania's packaging pools and supplier postal-code zones, and how these alternatives are shaped by the historical transport structure. The analysis shows that the routing option, rather than the allocation order alone, should be treated as the unit of evaluation. A route may be direct from one pool to a supplier, or it may combine a pool-to-pool allocation leg with a final supplier delivery leg. The empirical findings also show that the four pools do not play identical roles in the network and that some pool-to-pool relations are more structurally important than others. These findings justify why the prototype evaluates predefined routing alternatives rather than treating all theoretical routes as equally relevant.

RQ1b asks how the estimated costs and indicative CO_2 impacts of the routing alternatives vary across destinations and shipment sizes when historical transport costs and contracted rate-card benchmarks are compared. The analysis shows that historical weighted average €/LM values and contracted rate-card costs answer different questions. Historical costs describe observed spending on a lane or transport leg, while the rate card provides contracted benchmarks for specific pool-to-destination relations and load-metre brackets. In the prototype, this distinction is operationalised by using historical €/LM values to estimate route costs for the relevant legs, while the Rate Card Audit sheet provides official contracted benchmarks for direct routes and trailer-type comparisons. The indicative CO_2 values add a secondary comparison dimension but are not treated as certified emission reporting.

RQ2 asks how the routing and cost evaluation logic can be operationalised into a planner facing decision support tool. The answer is the spreadsheet-based Pool Routing Optimizer prototype. Given a destination postal-code zone and shipment size in load metres, the tool evaluates eleven predefined routing alternatives: four direct pool-to-supplier routes and seven allocation-based routes. For each available alternative, the prototype estimates total route cost from the relevant allocation and delivery rates and ranks the alternatives from cheapest to most expensive. It also presents indicative CO_2 estimates and rate-card audit information for contracted direct-route benchmarks. The prototype is not a mathematical optimisation model, and it does not connect to live inventory or order systems. Its contribution is to operationalise the routing evaluation logic in a planner-facing interface that improves cost visibility at the moment of planning.

These answers complete the analytical argument of the study. RQ1a explains how the routing alternatives are structured in the pool network. RQ1b explains how these alternatives can be evaluated using historical cost, contracted rate-card benchmarks and indicative CO_2 information. RQ2 shows how this evaluation logic can be translated into a working decision support prototype. Together, the three questions connect the empirical analysis to the practical tool developed in the thesis.

6 Discussion

This chapter discusses what those findings mean for planning practice for the case company and for the wider supply chain literature. The discussion focuses on the value of cost-aware routing decision support, the practical conditions required for implementation, the contribution of the study, the way the project scope evolved, and the limitations that define how the results should be interpreted. The discussion is structured through the frame of reference developed in Chapter 2. The network and pool logistics literature is used to interpret why routing alternatives are structurally constrained by the design of the pool network. The literature on operational planning and bounded rationality explains why route-level information must be visible at the moment of decision. The transport cost, rate-card and KPI literature explains why historical costs, contracted benchmarks and indicative CO_2 values need to be shown as separate but related perspectives. Finally, the governance literature frames which recommendations can be acted on operationally and which require broader organisational or contractual decisions.

6.1. Implications for cost-aware routing decision support

The main implication of the analysis is not that allocation flows should be reduced as a category. This interpretation is consistent with the literature on returnable transport item systems, where internal redistribution is treated as a normal response to spatial imbalance and demand variation rather than as an inefficiency in itself (Glock, 2017; Hellström & Johansson, 2010; Pålsson & Hellström, 2016). Allocation-based routes can therefore be operationally necessary and may in some cases be cost competitive. The central issue is that planners need visibility of the full routing alternative before choosing how a supplier need should be served. A direct pool-to-supplier route and an allocation-based route are not separate planning questions from a cost perspective; they are alternative ways of fulfilling the same need. The developed prototype addresses this by comparing complete routing alternatives and ranking them by estimated route cost for a selected destination and shipment size, which is aligned with the decision-support view that analytical tools should complement rather than replace human judgement (Power, 2008; Arnott & Pervan, 2014).

The historical analysis still matters but it should be taken as diagnostic context rather than as the main function of the tool. Circular or bidirectional movements indicate that limited cross-pool visibility can create avoidable transport work, but individual cases may still be justified by packaging availability, supplier timing, service requirements or packaging-type compatibility. For that reason, the thesis should not claim that every observed circular movement is automatically avoidable. The safer conclusion is that such observations strengthen the need for better route-level visibility and provide a possible direction for future development.

The strongest practical justification for the prototype follows from the planning literature: when planners work under incomplete information and time pressure, the relevant decision information must be available at the moment of choice rather than only in retrospective reports (Simon, 1955; Gigerenzer & Gaissmaier, 2011; Klein, 2008). The current version does not include live inventory, pending order visibility or alerts for circular flows. Its function is narrower and more concrete: it allows a planner to enter a destination postal-code zone and shipment size, evaluate predefined direct and allocation-based alternatives, and compare their estimated cost, contracted benchmarks where applicable and indicative CO_2 impact. Future versions could extend this logic by adding alerts for recent or pending opposite-direction movements, but that is not part of the current prototype.

The discussion also raises a broader measurement issue. Performance-measurement literature emphasises that KPIs are useful only when they are linked to the decision level, they are intended to support (Gunasekaran et al., 2004; Piotrowicz & Cuthbertson, 2015). If transport cost is mainly reviewed after execution and at an aggregated level, planners receive limited feedback on the cost consequences of individual routing choices. A useful performance measure should therefore connect more directly to the route decision itself. Instead of only tracking total transport spend, the company could monitor route-level cost deviations, differences between historical cost and contracted benchmarks, and repeated use of high-cost alternatives where lower-cost routes appear available. Such measures would align better with the purpose of the prototype than a KPI focused only on allocation volume.

In summary, the main discussion point is that cost-aware routing cannot depend only on planner experience, geographical intuition or whether stock is available in a given pool. The cost of a routing choice depends on the combination of allocation leg, final delivery leg, shipment size, historical cost level and contracted benchmark. The prototype makes these elements visible in one planner-facing view. This is the central practical implication of the study.

6.2. Implications for planning practice at the case company

The work produces an Excel-based prototype that demonstrates the analytical logic for cost-aware routing decision support. This follows the decision-support literature, which describes such tools as analytical aids that should present relevant information at the moment of decision while still leaving judgement with the user (Power, 2008; Arnott & Pervan, 2014). The prototype is not intended to remain as a stand-alone spreadsheet solution in the long term. Its practical value lies in showing what information planners need, how route alternatives can be compared, and how historical cost and contracted benchmarks can be brought closer to the planning decision.

The prototype was built to be inspected and tested, not to operate as a scalable production system. A production version would need to connect to live order data, updated rate sheet information, historical shipment records and relevant feasibility checks such as inventory and capacity. The important design principle is that the analytical logic should remain clear: the system should show the planner which routing alternatives are available, how the estimated cost is calculated, and how the result compares with contracted benchmarks where applicable. The spreadsheet prototype can therefore serve as a functional specification for a future implementation.

Information that does not reach the planner at the right moment has limited operational value. Three categories of information should be positioned closer to the routing decision. The first is historical transport cost, which provides the weighted average €/LM references used to estimate route cost. This reflects the transport-cost literature, where cost is understood as both resource-dependent and lane-dependent (Izadi et al., 2020; Zofío et al., 2014). The second is contracted rate-card information, which provides direct-route benchmarks and supports auditability. This is important because freight pricing is often structured through non-linear contractual rate cards rather than simple linear cost functions (Caplice & Sheffi, 2003; Caplice, 2007). The third is indicative CO_2 information, which in the prototype is calculated from estimated distance, shipment size and a fixed emission factor. These inputs already exist in different forms, but the prototype shows how they can be combined into one route-comparison view.

A practical implementation should also define ownership of route level performance follow-up. The tool creates value only if its output is connected to planning routines, procurement discussions and operational review. For example, repeated differences between historical cost and contracted benchmarks should be reviewed by someone with visibility across planning, transport procurement and pool operations. Similarly recurring use of high-cost routes should trigger discussion about whether the choice was operationally necessary or whether a lower-cost alternative could have been used. This type of follow-up requires ownership beyond a single pool or individual planner.

Not all implications of the analysis lie within the company's full control. The literature on buyer supplier relationships emphasizes that pool operators function as partners with their own commercial interests and their own contractual logic (Benton & Maloni, 2005; Nyaga et al., 2013). Lacoste and Johnsen (2015) describe how operational improvements in such relationships depend on negotiation and on agreement about which side benefits from a given change. A reduction in bilateral churning between Eschweiler and Järna, to take the most pronounced case in the data, would benefit the case company directly through lower transport cost. The pool operators on the receiving and dispatching side may have a different perspective on the same change depending on how their compensation is structured. The implication is not that improvements are blocked. It is that improvements need to be negotiated with the operators rather than imposed unilaterally. The decision support concept makes the cost effect of routing choices visible. Acting on that visibility is a separate organizational task and one that engages the company's commercial relationships with the pool operators. The prototype on its own changes no transport decision. Its value is that it provides a concrete analytical specification for how cost-aware routing decisions could be supported.

6.3. Contributions to the supply chain literature

The study does not claim to introduce a new theory of returnable packaging networks. Its contribution is more practical and methodological. Existing literature has already established the importance of returnable transport items, pool-based logistics and supply chain network structure (Kroon & Vrijens, 1995; Bellamy & Basole, 2013; Carter et al., 2015; Glock, 2017). What this study adds is a worked industrial example of how these themes can be combined with transport pricing, historical cost visibility and operational decision support in a working prototype for route evaluation. The contribution therefore lies in demonstrating how cost-aware routing logic can be operationalised in an industrial planning context.

The case contributes to the literature on returnable transport items by connecting network structure with planner facing decision support. The closed loop nature of returnable packaging creates planning challenges because packaging availability, supplier demand and transport connections do not always align, which is consistent with previous research on reusable packaging systems and pool management (Hellström & Johansson, 2010; Pålsson & Hellström, 2016; Glock, 2017). This study adds an applied example of how those conditions can be translated into route evaluation logic. Instead of only describing pool movements, the thesis shows how historical transport data and contracted pricing information can be transformed into a tool that supports routing decisions.

The study also contributes to decision-support research in supply chain networks. Decision-support literature emphasises that such tools should improve the information conditions under which decisions are made, rather than replace human judgement with full automation (Power, 2008; Arnott & Pervan, 2014). Many network models also simplify routing cost into distance- or volume-based approximations. The case shows why this can be insufficient in practice. In Scania's setting, the relevant decision is not simply which pool is geographically closest, but which complete routing alternative gives the most favourable cost outcome for a specific destination and shipment size. The prototype demonstrates how this type of decision can be supported without developing a full mathematical optimisation model.

The cost dimension is an important part of this contribution. Freight transport cost is shaped by both shipment characteristics and lane-specific conditions (Izadi et al., 2020; Zofío et al., 2014), while freight pricing is often governed by contractual rate-card structures rather than by a simple linear relationship between distance and volume (Caplice & Sheffi, 2003; Caplice, 2007). The prototype shows that route evaluation in a packaging pool network benefits from distinguishing between historical weighted average costs and contracted rate-card benchmarks. Historical €/LM values help estimate and compare route alternatives, while the rate card supports contracted-cost visibility and auditability. The study therefore contributes a practical example of how multiple pricing perspectives can be shown together without forcing them into a single blended measure. This is especially relevant in operational settings where planners need fast, transparent comparisons rather than a complex optimisation model that is difficult to interpret.

6.4. Reflecting on the project's evolution from broad to specific

Where the project ends are not where it began. The original brief from the company case was wider than what this thesis has examined. It covered the full sequence of suppliers' needs to delivered packaging, with interest in transport contracts, pool operator agreements, packaging type rationalization and overall cost reduction across the European network. The work that follows from that brief looks different from the work presented here.

The narrowing happened in two main steps. The first step recognised that several of the original questions were governance questions rather than analytical ones. Contract renegotiation, pool operator compensation and network redesign depend on commercial relationships, organisational mandates and governance structures that lie partly outside the planners' operational decision space (Benton & Maloni, 2005; Lacoste & Johnsen, 2015; Nyaga et al., 2013). An analytical thesis could inform such questions, but it could not resolve them. The second step recognised that operational routing decisions offered the most productive site for a practical contribution. This was where historical transport data, rate sheet information and planner needs could be combined into a decision support prototype.

Three areas were identified early in the project but fell outside the final scope. Transport contract renegotiation was one. The rate sheet analysis can inform procurement discussions, but renegotiation itself belongs to the procurement and commercial domain. A second area was pool operator compensation. The analysis can highlight where routing choices create cost consequences, but incentive design is an organisational and contractual question. A third area was packaging type rationalisation. The final prototype evaluates routing alternatives by destination and shipment size; it does not optimise the packaging portfolio itself. The narrowing was therefore necessary. By focusing on route evaluation and planner-facing decision support, the project produced a concrete tool rather than a broad but less actionable diagnosis.

What was left outside the scope still matters for future work. Contract management, operator incentives, live inventory visibility and packaging portfolio decisions all influence the practical value of routing optimisation. However, these questions require organisational decisions beyond the analytical scope of this thesis. The thesis should therefore be understood as one building block: it provides the route evaluation logic and a working prototype that can support further development within the case company.

6.5. Limitations of the study and its concept

The empirical base of the study is a single industrial case at one company. As case study literature notes, such designs are well suited for analysing complex operational problems in context, but they do not support statistical generalisation in the same way as broader quantitative studies (Crowe et al., 2011; Rashid et al., 2019). The results are shaped by Scania's specific pool structure, supplier geography, transport contracts, data availability and planning environment. The qualitative input from planners was gathered through recurring operational discussions rather than structured interviews, which means it provided important context but was not analysed as a separate interview dataset. In addition, the transport data and packaging-level information are not always available at the same level of granularity, which limits how precisely individual packaging behaviours can be traced through the full network.

Methodologically, the comparison between historical and contracted pricing rests on two assumptions. First, the historical cost values in the transport data are treated as valid representations of observed transport spend. Second, the rate sheet is treated as the relevant contracted benchmark for the lanes and shipment-size brackets included in the prototype. These assumptions are reasonable for the purpose of a decision support prototype, but they cannot fully replace invoice-level audit. Circular-flow and bidirectional-movement findings should also be interpreted as diagnostic opportunity indicators rather than confirmed avoidable savings, since operational constraints such as packaging compatibility, timing, stock availability and service requirements may justify individual movements. This is consistent with KPI literature, which warns that performance measures can identify recurring outcomes but cannot fully explain operational decisions unless the decision context is observable (Gunasekaran et al., 2004; Shepherd & Günter, 2006).

The decision support tool is a prototype rather than a production system. It uses static historical data and a static rate sheet snapshot rather than live updates. It evaluates one shipment at a time and does not include live inventory, pool capacity, supplier order timing or carrier availability. It also does not calculate a complete official contracted rate sheet cost for every two-leg allocation route. The rate sheet audit mainly supports direct-route benchmarking and trailer type comparison. A planner using the prototype must therefore verify operational feasibility before acting on a recommendation. These limitations are addressable in a production version, but they define the current prototype.

Another limitation concerns the future configuration of the pool network. The study is based on the current location and role of the four pools. If pool locations, transport contracts, supplier geography or planning rules change, the distance references, routing alternatives and cost relationships in the prototype will need to be updated. The findings should therefore be interpreted as valid for the current network structure and data foundation.

None of these limitations invalidates the findings. They define what the study can reasonably claim. Within these boundaries, the prototype demonstrates how historical cost data, allocation rates, contracted rate sheet benchmarks and indicative CO_2 estimates can be combined into a planner-facing route comparison. The case company receives a concrete artefact that can be tested and developed further, while the thesis contributes an applied example of cost-aware routing decision support in a packaging pool network. The conclusion now states these contributions formally and outlines possible next steps.

7 Conclusion

A thesis ends where its argument ends. The argument was closed in chapter 6. What remains is to state the contribution in a form that stands on its own; to set out the concrete path forward for the case company and to identify the questions that the work leaves open for future research.

7.1. The study's contribution

The study contributes a worked example of how the cost structure of a multi pool packaging network can be made visible at the level of the individual routing decision. The contribution rests on two pieces of work that the literature has rarely combined. The first is a documented empirical case at sufficient granularity to expose how network topology meets cost structure inside the cognition of a planner. The second is a prototype that operationalizes that intersection in a tool planners can actually use.

In substantive terms the work delivers four pieces of analytical material to the literature and to the case company. The flow patterns between four packaging pools were mapped across a full year of operations together with their structural asymmetries and their persistence over time. The contracted rate sheet was placed beside the historical pricing on the same lanes to expose where the two diverge and why. A third piece is the tracing of the decision chain down to the specific points where information about cost or routing falls out before reaching the planner. The decision support concept then brings these three threads together in a single planner facing tool that can be developed further at the case company.

What the study has shown can be stated plainly enough. Allocation routing in a multi pool returnable packaging network is not a problem of optimization but a problem of visibility. Once the cost structure of the alternatives is made visible at the moment of the decision the planner has what they need to act on it. Visibility is what the study has built. Acting on it is what the case company can now begin to do.

7.2. Implementation path forward for the case company

The prototype as it stands today is deliberately built in Excel. The choice made sense for proof-of-concept work where the priority was to demonstrate the analytical logic in a form that could be reviewed and refined quickly. For actual operational use the tool needs to leave this environment. The principal reason is that the prototype today runs on static data extracts that must be refreshed manually each time the underlying transport records or the rate sheet change. In a planning environment where allocation decisions are made daily, a tool that relies on yesterday's data carries the risk of giving advice based on conditions that are no longer held.

A platform such as Power BI fits the next development step for several reasons. The platform connects directly to the source systems that produce the data the prototype currently reads in static form. Once the connections have been built the figures presented to the planner refresh on the schedule of the underlying data rather than on the schedule of a manual export. The platform also handles the visualization layer without requiring custom software development, which lowers both the cost and the time needed to reach a usable version. A further advantage is that the analytical logic developed in the spreadsheet prototype can be transferred to the new environment in steps rather than rebuilt from scratch. The bracket lookup, the dual cost comparison, the routing enumeration: these are logic structures that translate naturally into a query and dashboard environment.

The first development step is to move the data layer from static exports to live connections. Three connections are required for the tool to function as a live decision aid. The transport extract should flow from the company's transport management system on a daily schedule so that historical figures stay current as new shipments accumulate. The rate sheet should connect to wherever procurement maintains it so that contractual changes propagate to the planning interface without manual intervention. The CO_2 data should connect to the environmental reporting system so that the environmental measure shown beside each routing option reflects the same values that appear in the company's emission reports.

Once the data layer is in place, the next focus is on functional extensions that the prototype cannot support in spreadsheet form. The most important of these is a forward-looking view that lets the planner see pending and recent loads on the corridor under consideration. This addresses the boomerang risk identified earlier by surfacing the information that planners in different pools currently cannot see about each other's activity. A second extension is a batch view that lets the planner evaluate several shipments at once rather than one at a time so that consolidation opportunities can be identified before separate allocations are dispatched. A third extension is the integration of inventory and capacity information at the pool level so that routing recommendations are filtered into options that are actually feasible given current stock conditions at the proposed source pool.

Beyond the functional extensions, a KPI layer would let the platform monitor the network rather than only supporting individual decisions. The most useful KPI to add is bilateral redundancy at the lane level which captures the boomerang pattern documented in section 6.1 in a form that can be tracked and acted on. A second useful KPI is the difference between the routing option a planner selected and the lowest cost option available at the time of the decision which gives a measure of cost efficiency at the network level. A third is the corresponding environmental measure which tracks how often planners choose options with higher CO_2 when lower CO_2 alternatives existed. Together these KPIs would create the feedback loop that the current measurement system does not provide, and that section 5.2 identified as missing from the planning environment.

None of these technical developments work without the organizational integration that has to accompany them. The tool needs to live in the planners' regular workflow rather than as a parallel system that competes for their attention. This means embedding the dashboard in the same screen environment where allocation orders are managed today. The tool also needs to be available to all planners across all four pools simultaneously so that cross corridor visibility becomes structural rather than dependent on individual initiative. A short training and onboarding period would help planners understand what the bracket structure means in practice and how to read the dual cost display when the two measures point in different directions.

The work outlined here is sequenced deliberately. The data layer is the foundation because nothing else functions without current data. The functional extensions come second because they multiply the value of the dashboard once the data is alive. The KPI layer comes third because it only becomes useful when it can run on live decisions made through the tool. Organizational integration runs in parallel with all of this because technology only delivers value when it sits where the decisions happen. A reasonable timeline for moving from the current prototype to a working production version is between six and twelve months depending on the rate of internal alignment and the availability of development resources at the company.

7.3. Directions for future research

Several research directions follow naturally from the work presented here. They divide it into questions that ask whether the patterns observed at Scania apply elsewhere and questions that ask how the methodology developed here translates to other industrial settings. The boomerang pattern documented in chapter 6 deserves comparative study. The data examined here came from a single company in a single European context which leaves open whether bilateral churning is a property of multi pool returnable packaging networks in general or a feature of the specific operational setup at Scania. A study that sampled allocation data from several different companies operating similar networks would be able to test this question. The same methodology used here could be applied to identify whether bilateral redundancy emerges at comparable rates in other settings or whether it depends on factors such as the structure of pool ownership, the coordination between planning teams or the configuration of supplier bases.

Beyond the question of generalizability, the dual cost methodology developed in chapter 4 offers a second avenue for further research. The decision support concept treats historical pricing and contracted pricing as complementary rather than competing measures of cost. Whether this approach generalizes to other freight contexts in which nonlinear pricing structures exist is an open question. Industries that come immediately to mind are automotive component logistics, fast moving consumer goods distribution and post pandemic ocean freight in which contracted and spot pricing diverge widely. The methodological question is whether dual pricing displays change decision behavior in those settings or whether decision makers default to one measure regardless of what is shown.

A third direction is the empirical study of how decision support tools change planner behavior in practice. The work here demonstrates that the cost structure of routing decisions can be made visible now of choice. The next question is whether making it visible shifts decisions toward lower cost routings. A longitudinal study at the case company that compared decision patterns before and after the tool is deployed would provide direct evidence on this. The cognitive science literature on naturalistic decision making suggests that habits formed under information scarce conditions can persist even when information becomes available which means the question is empirical rather than something that can be assumed.

A further direction concerns the network itself rather than the routing decisions made within it. The study took the four-pool network as a given and analyzed the flows that emerge from it. A natural extension is to invert the analytical question. Given the cost structure documented here the analytical question becomes whether a different network design would produce more favorable outcomes. The question links operational research on facility location with the empirical patterns observed in the present case and would benefit from explicitly considering the rate sheet structure that this thesis has shown matters for routing economics. The company itself may eventually face a network design decision when its supplier base or production volumes shift, and the analytical material developed here would form a useful starting point for that work.

A thesis is one move in a longer conversation. The contribution of this work is to show that allocation routing in multipool packaging networks is a richer analytical territory than it appears at first sight. The cost structure is nonlinear. The planning decisions are bound by information conditions. The patterns are stable but not optimal. The questions that remain about this territory are worth pursuing. The hope is that this thesis makes them slightly easier to ask.

REFERENCES

Arnott, D., & Pervan, G. (2014). A critical analysis of decision support systems research revisited: The rise of design science. Journal of Information Technology, 29(4), 269–293.

Bellamy, M. A., & Basole, R. C. (2013). Network analysis of supply chain systems: A systematic review and future research. Systems Engineering, 16(2), 235–249.

Benton, W. C., & Maloni, M. (2005). The influence of power-driven buyer/seller relationships on supply chain satisfaction. Journal of Operations Management, 23(1), 1–22.

Caplice, C. (2007). Electronic markets for truckload transportation. Production and Operations Management, 16(4), 423–436.

Caplice, C., & Sheffi, Y. (2003). Optimization-based procurement for transportation services. Journal of Business Logistics, 24(2), 109–128.

Carter, C. R., Rogers, D. S., & Choi, T. Y. (2015). Toward the theory of the supply chain. Journal of Supply Chain Management, 51(2), 89–97.

Chopra, S., & Meindl, P. (2019). Supply chain management strategy planning and operation 7th ed. Pearson.

Crowe, S., Cresswell, K., Robertson, A., Huby, G., Avery, A., Sheikh, A. 2011. The case study approach. BMC Medical Research Methodology, 11, 100.

Domínguez Caamaño, P., García Arca, J., Fernández González, A. J., & Prado Prado, J. C. (2017). Transport KPIs for supply chain improvement a literature analysis. In M. Amorim et al. Engineering Systems and Networks pp. 251 258. Springer.

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. Annual Review of Psychology, 62, 451–482.

Glock, C. H. (2017). Decision support models for managing returnable transport items in supply chains: A systematic literature review. International Journal of Production Economics, 183, 561–569.

Gunasekaran, A., Patel, C., & McGaughey, R. E. (2004). A framework for supply chain performance measurement. International Journal of Production Economics, 87(3), 333–347.

Hellström, D., & Johansson, O. (2010). The impact of control strategies on the management of returnable transport items. Transportation Research Part E: Logistics and Transportation Review, 46(6), 1128–1139.

Hellström, D., & Saghir, M. (2007). Packaging and logistics interactions in retail supply chains. Packaging Technology and Science, 20(3), 197–216.

Izadi, A., Nabipour, M., & Titidezh, O. (2020). Cost models and cost factors of road freight transportation: A literature review and model structure. Fuzzy Information and Engineering, 11, 257–278.

Klein, G. (2008). Naturalistic decision making. Human Factors, 50(3), 456–460.

Kroon, L., & Vrijens, G. (1995). Returnable containers: An example of reverse logistics. International Journal of Physical Distribution & Logistics Management, 25(2), 56–68.

Lacoste, S., & Johnsen, R. E. (2015). Supplier–customer relationships: A case study of power dynamics. Journal of Purchasing and Supply Management, 21(4), 229–240.

McKinnon, A., Browne, M., Whiteing, A., & Piecyk, M. (2015). Green logistics improving the environmental sustainability of logistics 3rd ed. Kogan Page.

Nyaga, G. N., Lynch, D. F., Marshall, D., & Ambrose, E. (2013). Power asymmetry, adaptation and collaboration in dyadic relationships involving a powerful partner. Journal of Supply Chain Management, 49(3), 42–65.

Piotrowicz, W., & Cuthbertson, R. (2015). Performance measurement and metrics in supply chains: An exploratory study. International Journal of Productivity and Performance Management, 64(8), 1068–1091.

Power, D. J. (2008). Decision support systems: A historical overview. In F. Burstein & C. W. Holsapple (Eds.), Handbook on decision support systems 1 (pp. 121–140). Springer.

Pålsson, H., & Hellström, D. (2016). Packaging logistics in supply chain practice: Current state, trade-offs and improvement potential. International Journal of Logistics: Research and Applications, 19(5), 351–368.

Rashid, Y., Rashid, A., Warraich, M. A., Sabir, S. S., Waseem, A. 2019. Case study method. A step by step guide for business researchers. International Journal of Qualitative Methods, 18, 1–13.

Rodrigue, J. P. (2020). The geography of transport systems 5th ed. Routledge.

Shepherd, C., & Günter, H. (2006). Measuring supply chain performance: Current research and future directions. International Journal of Productivity and Performance Management, 55(3/4), 242–258.

Simchi Levy, D., Kaminsky, P., & Simchi Levy, E. (2008). Designing and managing the supply chain 3rd ed. McGraw Hill.

Simon, H. A. (1955). A behavioral model of rational choice. The Quarterly Journal of Economics, 69(1), 99–118.

Tagaro, J. C., Valda, D. T. S., Villa, S. E., & Yasuda, M. D. (2023). Logistics optimization a literature review of techniques for streamlining land transportation in supply chain operations.s

Van der Vaart, T., & van Donk, D. P. (2008). A critical review of survey-based research in supply chain integration. International Journal of Production Economics, 111(1), 42–55

Zofío, J. L., Condeço-Melhorado, A. M., Maroto, A., & Gutierrez, J. (2014). Generalized transport costs and index numbers: A geographical analysis of economic and infrastructure fundamentals. Transportation Research Part A: Policy and Practice, 67, 141–157.

APPENDICES

Appendix A — List of abbreviations

Abbreviation	Full term
AI	Artificial Intelligence
CO_2	Carbon dioxide
EP	Empty Packaging
FTL	Full Truck Load
KPI	Key Performance Indicator
LM	Loading Meter
LTL	Less Than Truckload
MSRT	Mega Swedish Road Train
NILE	Scania's internal planning system for packaging orders
RC	Rate sheets
RTI	Returnable Transport Items
SCM	Supply Chain Management

Appendix B — Tool structure

Tool prototype section

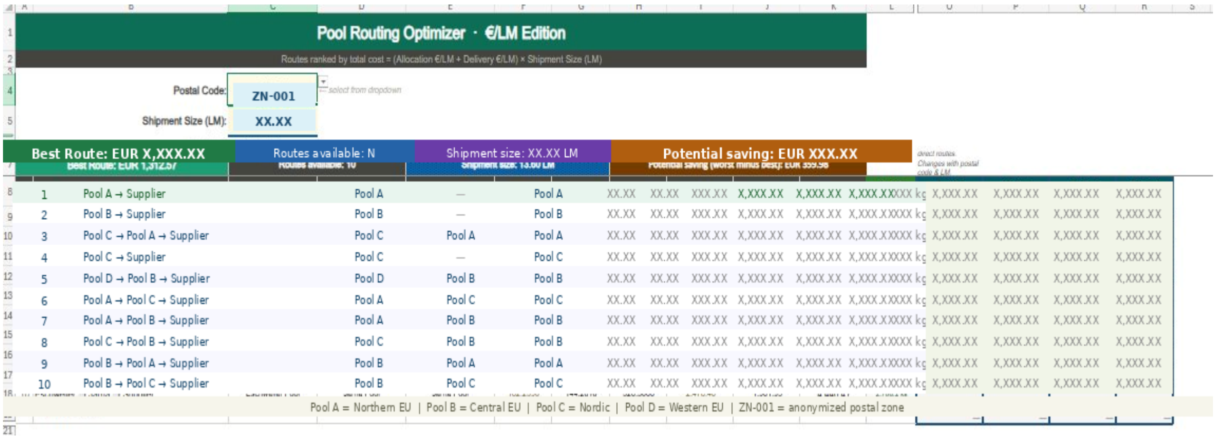

Tool architecture

Rate sheet integration

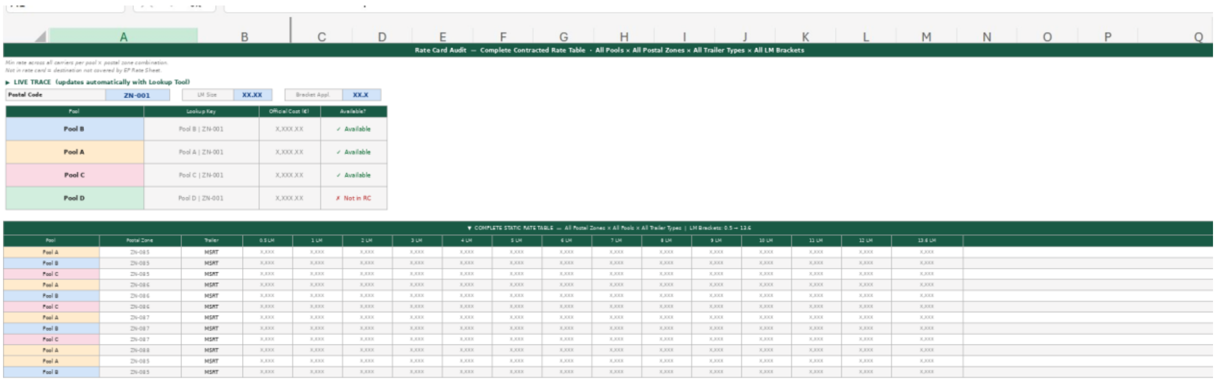

Rate Card Audit — Complete Contracted Rate Table · All Pools × All Postal Zones × All Trailer Types × All LM Brackets

Min rate across all carriers per pool × postal zone combination.
Not in rate card = destination not covered by EP Rate Sheet.

▶ LIVE TRACE (updates automatically with Lookup Tool)

Postal Code	ZN-001	LM Size	XX.XX	Bracket Appl.	XX.X

Pool	Lookup Key	Official Cost (€)	Available?
Pool B	Pool B | ZN-001	X,XXX.XX	✓ Available
Pool A	Pool A | ZN-001	X,XXX.XX	✓ Available
Pool C	Pool C | ZN-001	X,XXX.XX	✓ Available
Pool D	Pool D | ZN-001	X,XXX.XX	✗ Not in RC

▼ COMPLETE STATIC RATE TABLE — All Postal Zones × All Pools × All Trailer Types | LM Brackets: 0.5 → 13.6

Pool	Postal Zone	Trailer	0.5 LM	1 LM	2 LM	3 LM	4 LM	5 LM	6 LM	7 LM	8 LM	9 LM	10 LM	11 LM	12 LM	13.6 LM
Pool A	ZN-085	MSRT	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX
Pool B	ZN-085	MSRT	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX
Pool C	ZN-085	MSRT	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX
Pool A	ZN-086	MSRT	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX
Pool B	ZN-086	MSRT	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX
Pool C	ZN-086	MSRT	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX
Pool A	ZN-087	MSRT	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX
Pool B	ZN-087	MSRT	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX
Pool C	ZN-087	MSRT	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX
Pool A	ZN-088	MSRT	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX
Pool A	ZN-085	MSRT	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX
Pool B	ZN-085	MSRT	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX	X,XXX

Historical rate data

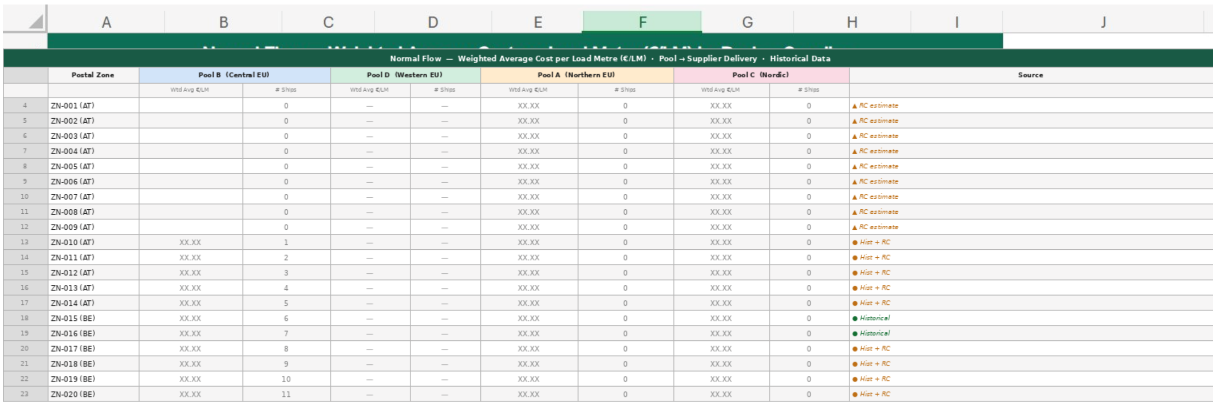

Normal Flow — Weighted Average Cost per Load Metre (€/LM) · Pool → Supplier Delivery · Historical Data

	Postal Zone	Pool B (Central EU)		Pool D (Western EU)		Pool A (Northern EU)		Pool C (Nordic)		Source
		Wtd Avg €/LM	# Ships	Wtd Avg €/LM	# Ships	Wtd Avg €/LM	# Ships	Wtd Avg €/LM	# Ships	
4	ZN-001 (AT)		0	—	—	XX.XX	0	XX.XX	0	▲ RC estimate
5	ZN-002 (AT)		0	—	—	XX.XX	0	XX.XX	0	▲ RC estimate
6	ZN-003 (AT)		0	—	—	XX.XX	0	XX.XX	0	▲ RC estimate
7	ZN-004 (AT)		0	—	—	XX.XX	0	XX.XX	0	▲ RC estimate
8	ZN-005 (AT)		0	—	—	XX.XX	0	XX.XX	0	▲ RC estimate
9	ZN-006 (AT)		0	—	—	XX.XX	0	XX.XX	0	▲ RC estimate
10	ZN-007 (AT)		0	—	—	XX.XX	0	XX.XX	0	▲ RC estimate
11	ZN-008 (AT)		0	—	—	XX.XX	0	XX.XX	0	▲ RC estimate
12	ZN-009 (AT)		0	—	—	XX.XX	0	XX.XX	0	▲ RC estimate
13	ZN-010 (AT)	XX.XX	1	—	—	XX.XX	0	XX.XX	0	● Hist + RC
14	ZN-011 (AT)	XX.XX	2	—	—	XX.XX	0	XX.XX	0	● Hist + RC
15	ZN-012 (AT)	XX.XX	3	—	—	XX.XX	0	XX.XX	0	● Hist + RC
16	ZN-013 (AT)	XX.XX	4	—	—	XX.XX	0	XX.XX	0	● Hist + RC
17	ZN-014 (AT)	XX.XX	5	—	—	XX.XX	0	XX.XX	0	● Hist + RC
18	ZN-015 (BE)	XX.XX	6	—	—	XX.XX	0	XX.XX	0	● Historical
19	ZN-016 (BE)	XX.XX	7	—	—	XX.XX	0	XX.XX	0	● Historical
20	ZN-017 (BE)	XX.XX	8	—	—	XX.XX	0	XX.XX	0	● Hist + RC
21	ZN-018 (BE)	XX.XX	9	—	—	XX.XX	0	XX.XX	0	● Hist + RC
22	ZN-019 (BE)	XX.XX	10	—	—	XX.XX	0	XX.XX	0	● Hist + RC
23	ZN-020 (BE)	XX.XX	11	—	—	XX.XX	0	XX.XX	0	● Hist + RC

Pool-to-Pool Rates

Allocation Matrix — Weighted Average €/LM to Move Stock Between Pools (FROM row → TO column)

FROM ↓ · TO →	Pool B (Central EU)		Pool D (Western EU)		Pool A (Northern EU)		Pool C (Nordic)	
	Wtd Avg €/LM	# Ships	Wtd Avg €/LM	# Ships	Wtd Avg €/LM	# Ships	Wtd Avg €/LM	# Ships
Pool B	— (same pool)		N/A		XX.XX	NNN	XX.XX	NNN
Pool D	XX.XX	NNN	— (same pool)		N/A		N/A	
Pool A	XX.XX	NNN	N/A		— (same pool)		XX.XX	NNN
Pool C	XX.XX	NNN	N/A		XX.XX	NNN	— (same pool)	

Rates = SUM(Historical Cost) ÷ SUM(Load Metres) across all pool-to-pool shipments. Mar 2025–Mar 2026. N/A = no historical allocation data for that corridor.

CO₂ Methodology

Distance Reference (km) — Pool Routing Optimizer

Distance Reference | Estimated Road Distance (km) | Method: Haversine × 1.30 | Source: Pool OPS + ISO centroids

Distance Reference | Road Distance (km) from each Pool to each Postal Code | Method: Haversine × 1.30 | Source: Pool OPS + ISO centroids

	Postal Code	Pool B (km)	Pool D (km)	Pool A (km)	Pool C (km)
2	ZN-001 (AT)	XXX	XXX	X,XXX	X,XXX
3	ZN-002 (AT)	XXX	XXX	X,XXX	X,XXX
4	ZN-003 (AT)	XXX	XXX	X,XXX	X,XXX
5	ZN-004 (AT)	XXX	XXX	X,XXX	X,XXX
6	ZN-005 (AT)	XXX	XXX	X,XXX	X,XXX
7	ZN-006 (AT)	XXX	XXX	X,XXX	X,XXX
8	ZN-007 (AT)	XXX	XXX	X,XXX	X,XXX
9	ZN-008 (AT)	XXX	XXX	X,XXX	X,XXX
10	ZN-009 (AT)	XXX	XXX	X,XXX	X,XXX
11	ZN-010 (AT)	XXX	XXX	X,XXX	X,XXX
12	ZN-011 (AT)	XXX	XXX	X,XXX	X,XXX
13	ZN-012 (AT)	XXX	XXX	X,XXX	X,XXX
14	ZN-013 (AT)	XXX	XXX	X,XXX	X,XXX
15	ZN-014 (AT)	XXX	XXX	X,XXX	X,XXX
16	ZN-015 (BE)	XXX	XXX	X,XXX	X,XXX
17	ZN-016 (BE)	XXX	XXX	X,XXX	X,XXX
18	ZN-017 (BE)	XXX	XXX	X,XXX	X,XXX
19	ZN-018 (BE)	XXX	XXX	X,XXX	X,XXX
20	ZN-019 (BE)	XXX	XXX	X,XXX	X,XXX
21	ZN-020 (BE)	XXX	XXX	X,XXX	X,XXX
22	ZN-021 (BE)	XXX	XXX	X,XXX	X,XXX

Pool-to-Pool Distances (km)

From Pool	To Pool	Distance (km)
Pool B	Pool A	XXX km
Pool B	Pool C	X,XXX km
Pool D	Pool B	XXX km
Pool A	Pool B	X,XXX km
Pool A	Pool C	XXX km
Pool C	Pool B	X,XXX km
Pool C	Pool A	XXX km

Emission Factor

0.067 kg CO_2 per LM per km (GLEC Framework v3 / EN 16258)

Basis: Euro VI truck, well-to-wheel, 70% load factor

Method: Haversine straight-line × 1.30 road correction, country centroid

Note: Indicative estimates only — not for precise CO_2 reporting

Data Foundation

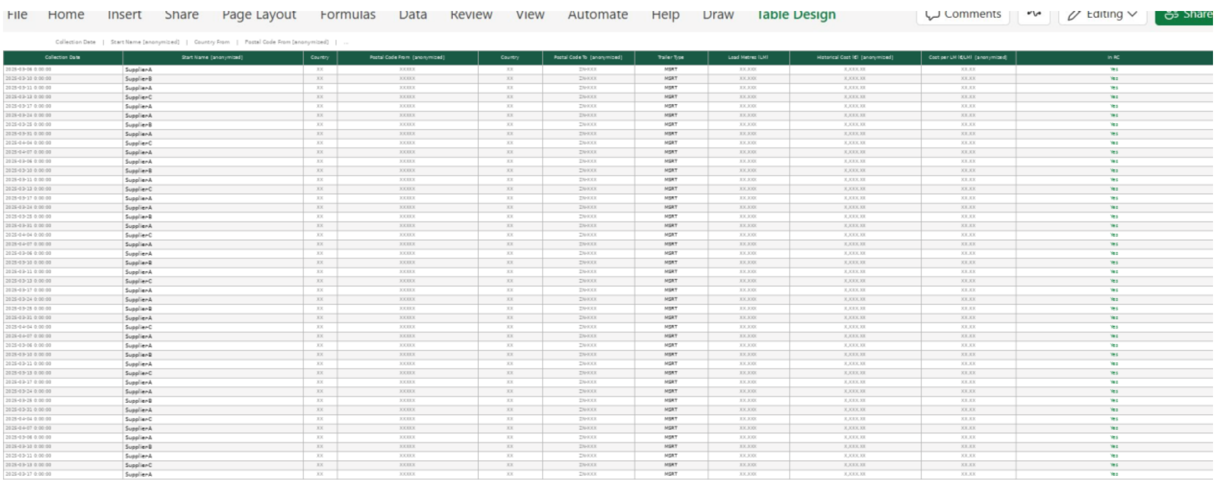

Guide Sheet

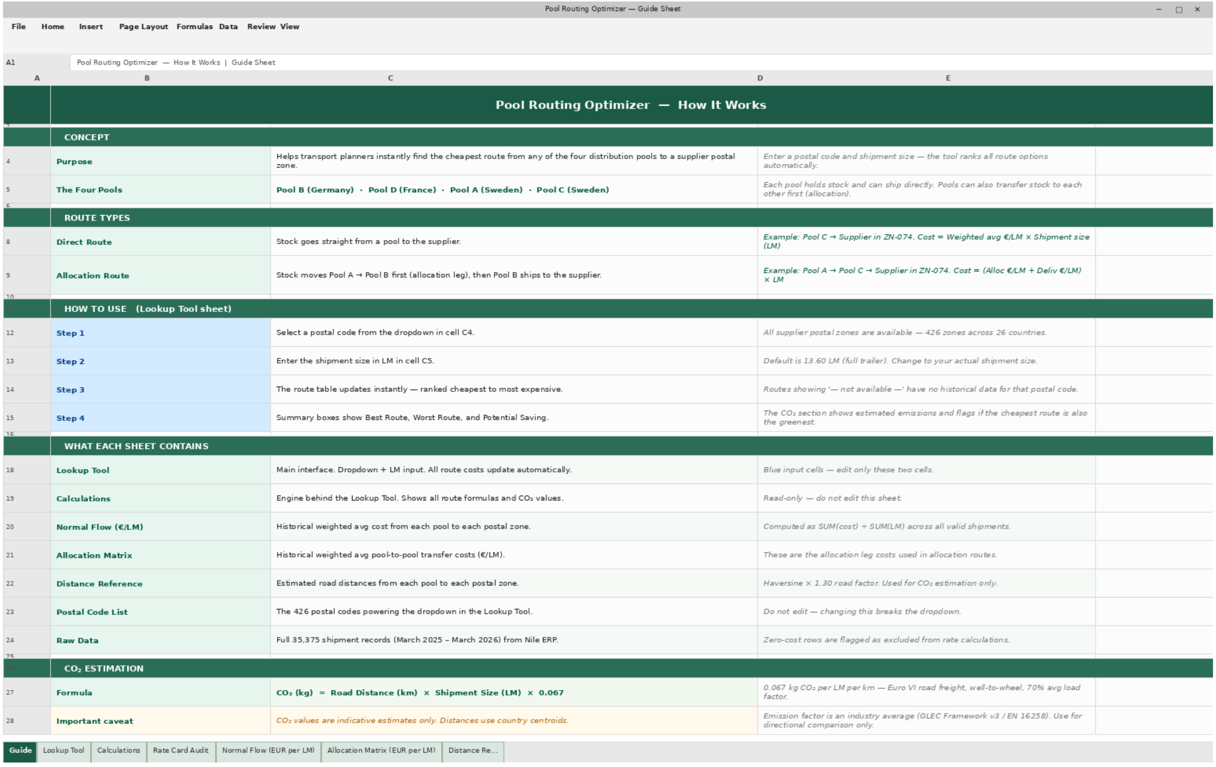

Pool Routing Optimizer — Guide Sheet

File Home Insert Page Layout Formulas Data Review View

A1 | Pool Routing Optimizer — How It Works | Guide Sheet

	A	B	C	D
1	**Pool Routing Optimizer — How It Works**			
3	**CONCEPT**			
4	**Purpose**	Helps transport planners instantly find the cheapest route from any of the four distribution pools to a supplier postal zone.	*Enter a postal code and shipment size — the tool ranks all route options automatically.*	
5	**The Four Pools**	**Pool B (Germany) · Pool D (France) · Pool A (Sweden) · Pool C (Sweden)**	*Each pool holds stock and can ship directly. Pools can also transfer stock to each other first (allocation).*	
7	**ROUTE TYPES**			
8	**Direct Route**	Stock goes straight from a pool to the supplier.	*Example: Pool C → Supplier in ZN-074. Cost = Weighted avg €/LM × Shipment size (LM)*	
9	**Allocation Route**	Stock moves Pool A → Pool B first (allocation leg), then Pool B ships to the supplier.	*Example: Pool A → Pool C → Supplier in ZN-074. Cost = (Alloc €/LM + Deliv €/LM) × LM*	
11	**HOW TO USE (Lookup Tool sheet)**			
12	**Step 1**	Select a postal code from the dropdown in cell C4.	*All supplier postal zones are available — 426 zones across 26 countries.*	
13	**Step 2**	Enter the shipment size in LM in cell C5.	*Default is 13.60 LM (full trailer). Change to your actual shipment size.*	
14	**Step 3**	The route table updates instantly — ranked cheapest to most expensive.	*Routes showing '— not available —' have no historical data for that postal code.*	
15	**Step 4**	Summary boxes show Best Route, Worst Route, and Potential Saving.	*The CO_2 section shows estimated emissions and flags if the cheapest route is also the greenest.*	
17	**WHAT EACH SHEET CONTAINS**			
18	**Lookup Tool**	Main interface. Dropdown + LM input. All route costs update automatically.	*Blue input cells — edit only these two cells.*	
19	**Calculations**	Engine behind the Lookup Tool. Shows all route formulas and CO_2 values.	*Read-only — do not edit this sheet.*	
20	**Normal Flow (€/LM)**	Historical weighted avg cost from each pool to each postal zone.	*Computed as SUM(cost) ÷ SUM(LM) across all valid shipments.*	
21	**Allocation Matrix**	Historical weighted avg pool-to-pool transfer costs (€/LM).	*These are the allocation leg costs used in allocation routes.*	
22	**Distance Reference**	Estimated road distances from each pool to each postal zone.	*Haversine × 1.30 road factor. Used for CO_2 estimation only.*	
23	**Postal Code List**	The 426 postal codes powering the dropdown in the Lookup Tool.	*Do not edit — changing this breaks the dropdown.*	
24	**Raw Data**	Full 35,375 shipment records (March 2025 – March 2026) from Nile ERP.	*Zero-cost rows are flagged as excluded from rate calculations.*	
26	**CO_2 ESTIMATION**			
27	**Formula**	**CO_2 (kg) = Road Distance (km) × Shipment Size (LM) × 0.067**	*0.067 kg CO_2 per LM per km — Euro VI road freight, well-to-wheel, 70% avg load factor.*	
28	**Important caveat**	*CO_2 values are indicative estimates only. Distances use country centroids.*	*Emission factor is an industry average (GLEC Framework v3 / EN 16258). Use for directional comparison only.*	

Guide | Lookup Tool | Calculations | Rate Card Audit | Normal Flow (EUR per LM) | Allocation Matrix (EUR per LM) | Distance Re...

www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY